Cisco Secure Firewall Management Center Administration Guide, 7.3

First Published: 2022-11-29

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

CONTENTS

PART I

Getting Started

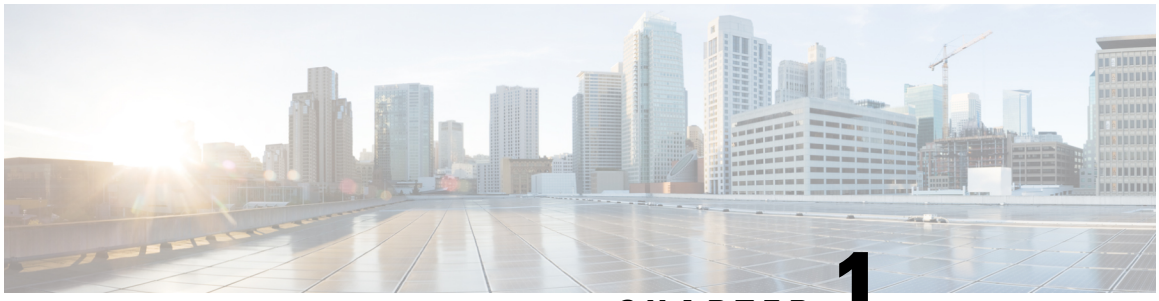

CHAPTER 1

Management Center Overview

This guide applies to an *on-premises* Secure Firewall Management Center, either as your primary manager or as an analytics-only manager. When using the Cisco Security Cloud Control (Security Cloud Control) cloud-delivered management center as your primary manager, you can use an on-prem management center for analytics. Do not use this guide for Security Cloud Control management; see Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Cisco Security Cloud Control.

The Secure 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. You should use the management center if you want a multi-device manager, and you require all features on the threat defense. The 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 management center for analytics only, then the on-prem management center does not support policy configuration or upgrading. Some chapters and procedures in this guide related might not apply to devices whose primary manager is Security Cloud Control.

For the management center used as the primary manager: The management center is not compatible with other managers because the management center owns the threat defense configuration, and you are not allowed to configure the threat defense directly, bypassing the management center.

Quick Start: Basic Setup

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

Installing and Performing Initial Setup on Physical Appliances

Procedure

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

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

    Cisco Secure Firewall Management Center Getting Started Guides
- **Threat Defense managed devices**
  - Cisco Firepower 1010 Getting Started Guide
  - Cisco Firepower 1100 Getting Started Guide
  - 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

Deploying Virtual Appliances

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

Procedure

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

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

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

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

Logging In for the First Time

Before logging in to a new management center for the first time, prepare the appliance as described in Installing and Performing Initial Setup on Physical Appliances, on page 2 or Deploying Virtual Appliances, on page 2.

The first time that you log in to a new management center (or a management center newly restored to factory defaults), use the **admin** account for either the CLI or the web interface and follow the instructions in the *Cisco Secure Firewall Management Center Getting Started Guide* for your management center model. When you complete the initial configuration process, the following aspects of your system will be configured:

- The passwords for the two **admin** accounts (one for web interface access and the other for CLI access) will be set to the same value, complying with strong password requirements as described in Guidelines and Limitations for User Accounts for Management Center, on page 124. The system synchronizes the passwords for the two **admin** accounts only during the initial configuration process. If you change the password for either **admin** account thereafter, they will no longer be the same and the strong password requirement can be removed from the web interface **admin** account. (See Add or Edit an Internal User, on page 125.)
- The following network settings the management center uses for network communication through its management interface (eth0) will be set to default values or values you supply:
  - Fully qualified domain name (`<hostname>.<domain>`)
  - Boot protocol for IPv4 configuration (DHCP or Static/Manual)
  - IPv4 address
  - Network mask
  - Gateway

- DNS Servers
- NTP Servers

Values for these settings can be viewed and changed through the management center web interface; see Modify Management Center Management Interfaces, on page 81 and Time Synchronization, on page 104 for more information.

- As part of the initial configuration, the system schedules weekly GeoDB updates. We recommend you review this task and make changes if necessary, as described in Schedule GeoDB Updates, on page 228.
- As part of the initial configuration, the system schedules weekly downloads. We recommend you review this task and make changes if necessary, as described in Automating Software Downloads, on page 497.

Important This task only downloads the updates. It is your responsibility to install any updates this task downloads.

- As part of the initial configuration, the system schedules weekly configuration-only management center backups (locally stored). We recommend you review this task and make changes if necessary, as described in Schedule Management Center Backups, on page 489.
- As part of the initial configuration, the system downloads and installs the latest VDB. To keep the system up to date, we recommend you schedule recurring updates as described in Vulnerability Database Update Automation, on page 499.
- As part of the initial configuration, the system schedules daily intrusion rule updates. We recommend you review this task and make changes if necessary, as described in Schedule Intrusion Rule Updates, on page 231.

On completion of 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 Specifying Your Home Page, on page 205.)

When you complete the initial configuration, begin controlling and analyzing traffic by configuring the basic policies as described in Setting Up Basic Policies and Configurations, on page 4.

Setting Up Basic Policies and Configurations

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

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

Before you begin

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

Procedure

Step 1 Set a time zone for this account as described in Setting Your Default Time Zone, on page 210.

Step 2 If needed, add licenses as described in Licenses, on page 253.

Step 3 Add managed devices to your deployment as described in *Add a Device to the 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 threat defense devices.
- *Interface Overview* in the Cisco Secure Firewall Management Center Device Configuration Guide, to configure interfaces on the 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.

- In most cases, Cisco suggests setting the **Balanced Security and Connectivity** intrusion policy as your default action. For more information, see *Access Control Policy Default Action* 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, see About Connection Logging, on page 713.

Step 6 Apply the system-provided default health policy as described in Apply a Health Policy, on page 374.

Step 7 Customize a few of your system configuration settings:

- If you want to allow inbound connections for a service (for example, SNMP or the syslog), modify the ports in the access list as described in Configure an Access List, on page 47.
- Understand and consider editing your database event limits as described in Configuring Database Event Limits, on page 59.
- If you want to change the display language, edit the language setting as described in Set the Language for the Web Interface, on page 73.
- If your organization restricts network access using a proxy server, edit your proxy settings as described in Modify Management Center Management Interfaces, on page 81.

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 traffic on your network. In most cases, Cisco suggests restricting discovery to the addresses in RFC 1918.

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 management center, see Add or Edit an Internal User, on page 125.
- If you want to use LDAP or RADIUS external authentication to allow access to the management center, see Configure External Authentication for the Management Center, on page 127.

Step 10 Deploy configuration changes; see the Cisco Secure Firewall Management Center Device Configuration Guide.

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

Although the management center can manage devices running previous versions (as specified in the compatibility matrix available at Cisco Secure Firewall Threat Defense Compatibility Guide), this guide only includes features supported on the *latest* version of device software.

For features that are only supported on old device versions, refer to the guide that matches your version.

Threat Defense Devices

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

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

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

Compatibility

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

Features

These tables list some commonly used features.

Appliance and System Management Features

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

If you want to...	Configure...	As described in...
Manage user accounts for logging in to your Secure Firewall devices	Device authentication	Users, on page 117 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	About Health Monitoring, on page 359
Back up data on your appliance	Backup and restore	Backup/Restore, on page 453
Upgrade to a new version	System updates	Cisco Secure Firewall Threat Defense Upgrade Guide for Management Center Cisco Secure Firewall Threat Defense Release Notes
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 223
Apply licenses in order to take advantage of license-controlled functionality	Smart licensing	About Licenses, on page 253
Ensure continuity of appliance operations	Managed device high availability and/or management center high availability	*About Secure Firewall Threat Defense "High Availability chapter"* in the Cisco Secure Firewall Management Center Device Configuration Guide About Management Center High Availability, on page 301

If you want to...	Configure...	As described in...
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
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 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 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	Introduction to Multitenancy Using Domains, on page 213
View and manage appliance configuration using a REST API client	REST API and REST API Explorer	REST API Preferences, on page 90 *Secure Firewall Mangement Center REST API Quick Start Guide*
Troubleshoot issues	N/A	Troubleshooting, on page 425

Features for Detecting, Preventing, and Processing Potential Threats

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

If you want to...	Configure...	As described in...
Inspect, log, and take action on network traffic	Access control policy, the parent of several other policies	*Introduction to Access Control* in the Cisco Secure Firewall Management Center Device Configuration Guide

If you want to...	Configure...	As described in...
Block or monitor connections to or from IP addresses, URLs, and/or domain names	Security Intelligence within your access control policy	*About Security Intelligence* in the Cisco Secure Firewall Management Center Device Configuration Guide
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

If you want to...	Configure...	As described in...
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
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 617
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 425

Integration with External Tools

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

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 1003 *Firepower System Remediation API Guide*
Stream event data from a management center to a custom-developed client application	eStreamer integration	eStreamer Server Streaming, on page 637 *Secure Firewall Mangement Center Event Streamer Integration Guide*
Query database tables on a management center using a third-party client	External database access	External Database Access, on page 62 *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 using external event data storage tools and other data resources	Integration with external event analysis tools	Event Analysis Using External Tools, on page 617

If you want to...	Configure...	As described in...
Troubleshoot issues	N/A	Troubleshooting, on page 425

Search the Management Center

You can use the global search feature to quickly locate and navigate to elements of your Secure Firewall Management Center configuration.

Note This feature is supported in Light and Dusk themes only. To change the theme, see Change the Web Interface Appearance, on page 205.

You can search the management center configuration for the following entities:

- Names of web interface pages in top-level menus. (See Search for Web Interface Menu Options, on page 14.)
- For certain policy types:
  - Policy names
  - Policy descriptions
  - Rule names
  - Rule comments

  (See Search for Policies, on page 14.)
- For certain object types:
  - Object names
  - Object descriptions
  - Configured values

  (See Search for Objects, on page 16 .)
- How To walkthroughs.

  The search returns a list of walkthroughs that contain the search term, with links to each. (See Search for How To Walkthroughs, on page 20.)

Keep the following in mind when using global search:

- When you open the global search tool, the most recent ten searches appear in a history list below the search text box. You can select an item from this list to re-execute a search.
- When you type a search expression, the interface replaces the search history with search results that update as you type your search; you do not need to press Enter to execute the search.
- You can navigate the history list or the search results using the mouse or the keyboard arrow keys and the Enter key. Pressing the Enter key selects the currently highlighted item in the search results. In the

case of results for web interface pages, this causes the management center interface to display the highlighted page. For objects and policies, this displays details about the found entity.

- Search is not case-sensitive.
- You can use the following 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; that is, global search does not return policies or objects that reference objects where a search term is found. However, you can determine which policies or objects reference many found objects by viewing the **Usages** tab for the found object 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 management center. If global search fails to return something you are expecting to find, try refining your search, try using the search or filter tool that appears at the top of many GUI pages, or try some of the configuration-specific search features 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

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 see 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.

In the example screenshot below, the deployment consists of three domains at three levels: Global, Domain1, and SubDomainA. The user, whose current domain is Domain1, has entered a search for the string “example” in both ancestor and child domains.

Figure 1: Example of Global Search in a Multidomain Environment

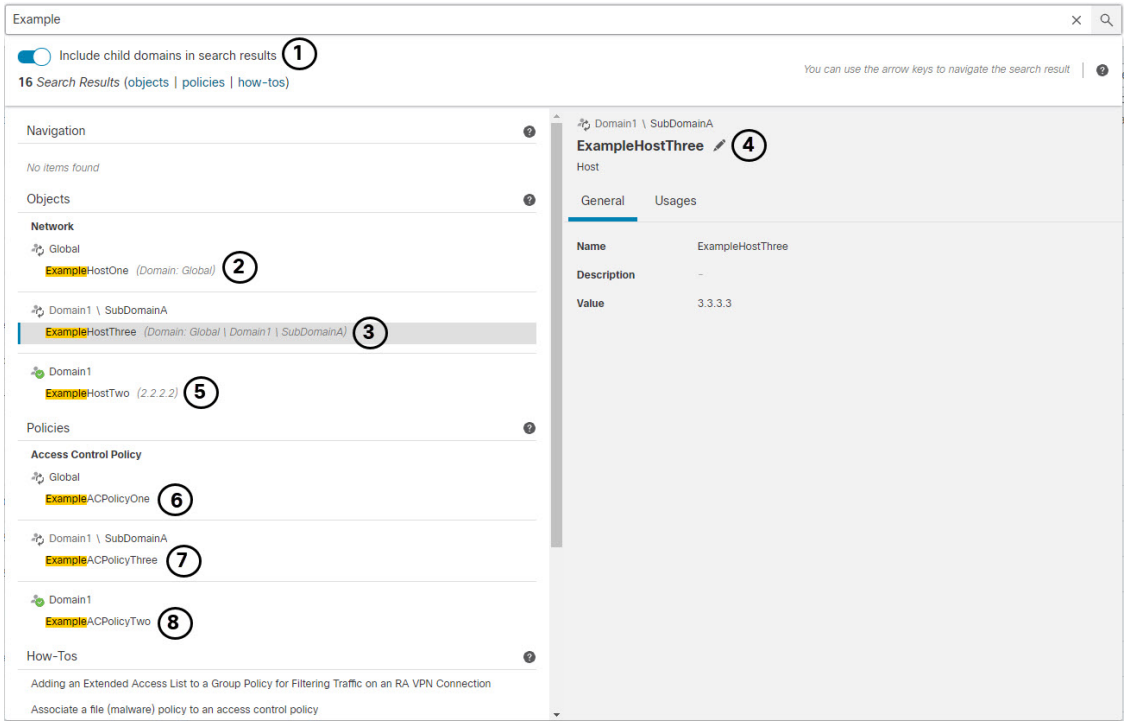

1	The user has chosen to search child domains (SubDomainA) as well as the current domain (Domain1) and its ancestor (Global).	2	A matching network object ExampleHostOne defined in the parent domain Global is displayed with the domain name, and the **External Domain** ( ) icon indicating the user must 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

You can search to find locations of pages in the top-level menus of the web interface. For example, to view or configure Quality of Service settings, search for `QoS`.

Before you begin

This feature is not available in the Classic theme. To change the theme, see Change the Web Interface Appearance, on page 205.

Procedure

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

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

Step 2 Enter one or more letters of the name of the menu option you seek. Search results appear below 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 go to a page listed under **Navigation**, click the menu path in the search results list.

Search for Policies

The following table indicates which policy types you can search for by name:

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

In Scope	Out of Scope
	DNS Policy Malware & 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 polices whose names match the search term, as well as access control policies using rules whose name or comments match the search term. If you see an access control policy in the search result list whose name does not match the search, the match was made on the name or comments for a rule configured within the policy.

Important Global search returns the top results for your search expression determined by its relevance to the most commonly used configuration entities in the 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, see Search the Management Center.

Before you begin

This feature is not available in the Classic theme. To change the theme, see Change the Web Interface Appearance, on page 205.

Procedure

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

- In the menu bar at the top of the 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 appear below 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 see 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 do the following:

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, and in the details pane click**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.

Search for Objects

The following 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 • Secure Client File • Custom Attribute

Global search returns objects whose names or description match the search term, as well as objects with configured values that match the search term. If you see an object in the search result list whose name does not match the search, the match was made on the description or a configured value within the object.

Important Global search returns the top results for your search expression determined by its relevance to the most commonly used configuration entities in the 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, see Search the Management Center.

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

- IPv4 and IPv6 address information, including the following 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, `192.164.1.1-192.168.1.5` 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 `192.168.1.0/24`, `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, see Change the Web Interface Appearance, on page 205.

Procedure

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

- In the menu bar at the top of the 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 appear below the text box and update as you type; you do not 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 see objects in those descendant domains.

Step 4 Search results appear divided by category. In a multidomain deployment, within the **Objects** category the search results are grouped by the domains within which found objects are defined. Under the **Objects** category you can do the following:

To:	Do this:
View search results for a single object type.	Click on 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 polices 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

You can search for How To walkthroughs that address tasks of interest. For example, to find walkthroughs that describe device set up procedures, you can search for the term "device."

Procedure

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

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

Step 2 Enter a search term associated with a task for which you would like to see a walkthrough. Search results appear below 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, see Online Help, How To, and Documentation, on page 23.

Switching Domains on the Secure Firewall Management Center

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

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

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

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

When you switch to a domain, the system displays:

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

Procedure

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

The Context Menu

Certain pages in the 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.

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 620.

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.

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.

Sharing Data with Cisco

You can opt to share data with Cisco using the following features:

- Cisco Success Network

  See Configure Management Center to Share Usage Metrics and Statistics with Cisco , on page 44

- Web analytics

  See Web Analytics, on page 112

Online Help, How To, and Documentation

You can reach the online help from the web interface:

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

How To is a widget that provides walkthroughs to navigate through tasks on the management center. The walkthroughs guide you to perform the steps required to achieve a task by taking you through each step, one after the other irrespective of the various UI screens that you may have to navigate, 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** > **How-Tos**.

Note The walkthroughs are generally available for all UI pages, and are not user role sensitive. However, depending on the privileges of the user, some of the menu items will not appear on the management center interface. Thereby, the walkthroughs will not execute on such pages.

The following walkthroughs are available on management center:

For a list of feature walkthroughs supported in the management center, see Feature Walkthroughs Supported in Secure Firewall Management Center.

You can find additional documentation using the documentation roadmap:

Navigating the Cisco Secure Firewall Threat Defense Documentation.

User Guides on Cisco.com

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

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

Secure Firewall Management Center

- Secure Firewall Management Center hardware appliances:

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

- Secure Firewall Management Center Virtual appliances:

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

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

- Secure Firewall Threat Defense software:

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

- Secure Firewall Threat Defense Virtual:

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

- Firepower 1000 series:

  https://www.cisco.com/c/en/us/support/security/firepower-1000-series/tsd-products-support-series-home.html

- Firepower 2100 series:

  https://www.cisco.com/c/en/us/support/security/firepower-2100-series/tsd-products-support-series-home.html

- Secure Firewall 3100:

  https://www.cisco.com/c/en/us/support/security/secure-firewall-3100-series/series.html

- Firepower 4100 series:

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

- Firepower 9300:

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

- ISA 3000:

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

License Statements in the Documentation

The License statement at the beginning of a section indicates which Classic or Smart license you must assign to a managed device 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 Defense license.

For more information about licenses, see About Licenses, on page 253.

Related Topics

About Licenses, on page 253

Supported Devices Statements in the Documentation

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

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

Access Statements in the Documentation

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

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

IP Address Conventions

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

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

Additional Resources

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

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

Note Some of the videos, technical notes, and reference material in the Firewalls Community points to older versions of the management center. Your version of the 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.

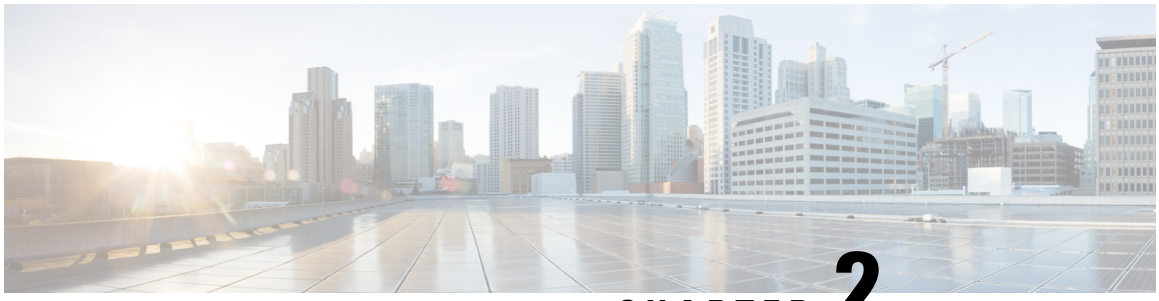

CHAPTER 2

Logging into the Management Center

The following topics describe how to log into the system:

User Accounts

You must provide a username and password to obtain local access to the web interface or CLI on 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 management center, all CLI users can use the `expert` command. The threat defense and 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 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 management center CLI provides a single **admin** user who has access to all commands. The features 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.

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

Caution All 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.

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

Secure Firewall Management Centers

Secure Firewall Management Centers support the following user account types:

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

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 the following user account types:

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

The 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 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 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 management center documentation..

Related Topics

Add or Edit an Internal User, on page 125

Web Interface Considerations

- If your organization uses Common Access Cards (CACs) for authentication, external users authenticated with LDAP can use CAC credentials to obtain access to the web interface of an appliance.
- The menus and menu options listed at the top of the default home page are based on the privileges for your user account. However, the links on the default home page include options that span the range of user account privileges. If you click a link that requires different privileges from those granted to your account, the system displays a warning message and logs the activity.
- Some processes that take a significant amount of time may cause your web browser to display a message that a script has become unresponsive. If this occurs, make sure you allow the script to continue until it finishes.

Related Topics

Specifying Your Home Page, on page 205

Session Timeout

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

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

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

System > **Configuration** > **Shell Timeout**

Related Topics

Configure Session Timeouts, on page 102
Configure SAML Single Sign-On, on page 144

Logging Into the Secure Firewall Management Center Web Interface

Note This task applies to internal users and external users authenticated by LDAP or RADIUS servers. For SSO login, see Logging Into the Management Center Web Interface Using SSO, on page 32.

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 management centers share the same IP address:

- Each management center can support only one login session at a time.
- To access different 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 125.

Procedure

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

Related Topics

Session Timeout, on page 30

Logging Into the Management Center Web Interface Using SSO

The 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 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 log in as a different user.

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

- Each management center can support only one login session at a time.
- To access different 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 management center for SSO access. See Configure SAML Single Sign-On, on page 144.
- If you do not have access to the web interface, contact your system administrator to configure your account at the SSO IdP.

Procedure

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

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

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

Step 3 The system responds in one of two ways:

- If you are already logged into the SSO federation, the management center default home page appears.

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

Related Topics

Session Timeout, on page 30
Configure SAML Single Sign-On, on page 144

Logging Into the Secure Firewall Management Center with CAC Credentials

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

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

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

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

Before you begin

- If you do not have access to the web interface, contact your system administrator to modify your account privileges, or log in as a user with Administrator access and modify the privileges for the account.
- Create user accounts as described in the Add or Edit an Internal User, on page 125.
- Configure CAC authentication and authorization as described in Configure Common Access Card Authentication with LDAP, on page 143.

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 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**.

Related Topics

Configure Common Access Card Authentication with LDAP, on page 143
Session Timeout, on page 30
SSO Guidelines for the Management Center, on page 145

Logging Into the Management Center Command Line Interface

The **admin** CLI user and certain custom external users can log into the 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 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 Logging In for the First Time, on page 3.

Procedure

Step 1 Use the **admin** user name and password to connect to the 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.

View Your Last Login

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.

Logging Out of the Management Center Web Interface

When you are no longer actively using the 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. Logging out ends your web session and ensures that no one can use the interface with your credentials.

Note If you are logging out of an SSO session at the management center, when you log out the system redirects your browser to the SSO IdP for your organization. To ensure management center security and prevent others from accessing the 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 management center, the system redirects you to the SSO IdP for your organization. Log out at the IdP to ensure management center security.

Related Topics

Session Timeout, on page 30

History for Logging into the Management Center

Feature	Minimum Management Center	Minimum 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 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: management center
Automatic CLI access for the management center	6.5	Any	When you use SSH to log into the 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 management center. As a consequence of deprecating this option, the virtual management center no longer displays the **System** > **Configuration** > **Console Configuration** page, which still appears on physical 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 management center	6.3	Any	New/Modified screens: New check box available to administrators in management center web interface: **Enable CLI Access** on the **System** > (⚙) > **Configuration** > **Console Configuration** page. • Checked: Logging into the management center using SSH accesses the CLI. • Unchecked: Logging into 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: management center

PART II

System Settings

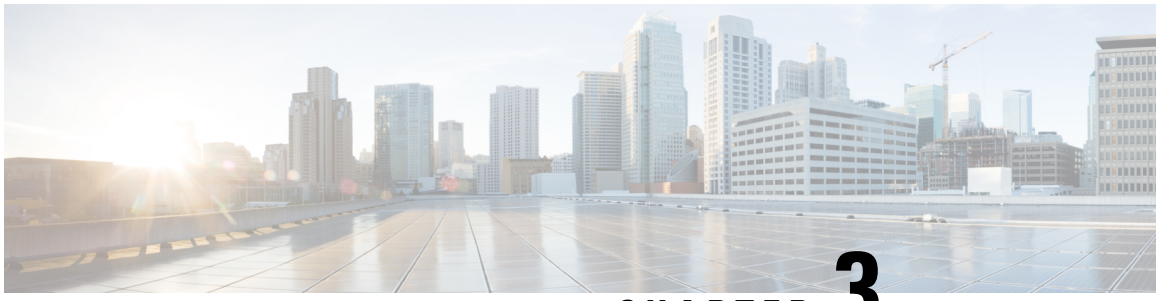

CHAPTER 3

System Configuration

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

Integrate Management Center with the Cisco Security Cloud

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

Use your Cisco Security Cloud Control account to authorize and register the management center with Cisco Security Cloud. This integration brings your firewall deployment onboard to the Cisco cloud tenancy, providing capabilities such as:

- Establish a consistent policy across management centers.
- Implement Zero-Touch Provisioning of the threat defense devices.
- Send events to the cloud and use Cisco Security Cloud services to enrich your threat hunts and investigations.
- Get a centralized view of inventory across management centers.

For more information about onboarding a management center to Security Cloud Control, refer to Onboard an On-Prem Management Center.

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

Enable SecureX Integration

Integrate the management center with SecureX to onboard both the management center and its managed devices to a Security Cloud Control tenant. When the management center is onboarded to Security Cloud Control, you can view its managed devices, view managed network objects, and cross-launch to the management center UI to manage associated devices and objects.

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 management center with Cisco Security Cloud. If you do not already have a Security Cloud Control tenant, request for a tenant or create one 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 Firewall in Security Cloud Control and Cisco XDR Tenant Accounts.
- Your management center must be between version 7.0.2 and 7.0.x, or version 7.2 and later to perform this task.

Procedure

Step 1 In the management center, choose **Integration** > **SecureX**.

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

Note

- The regional cloud you choose here is also used for the Cisco Success Network and Cisco Support Diagnostics capabilities. This setting also governs the cloud region for the Secure Network Analytics cloud using Security Analytics and Logging (SaaS).
- If you have already registered the management center with Smart License, the region selected by default will correspond to your Smart Licensing region. In this case, you don't have to change the region.

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

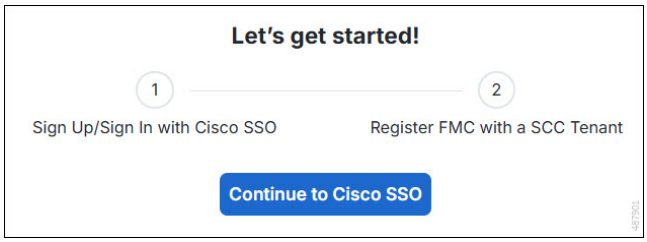

Step 5 Log in to your Security Cloud Control account.

Figure 3: Security Cloud Control Sign On

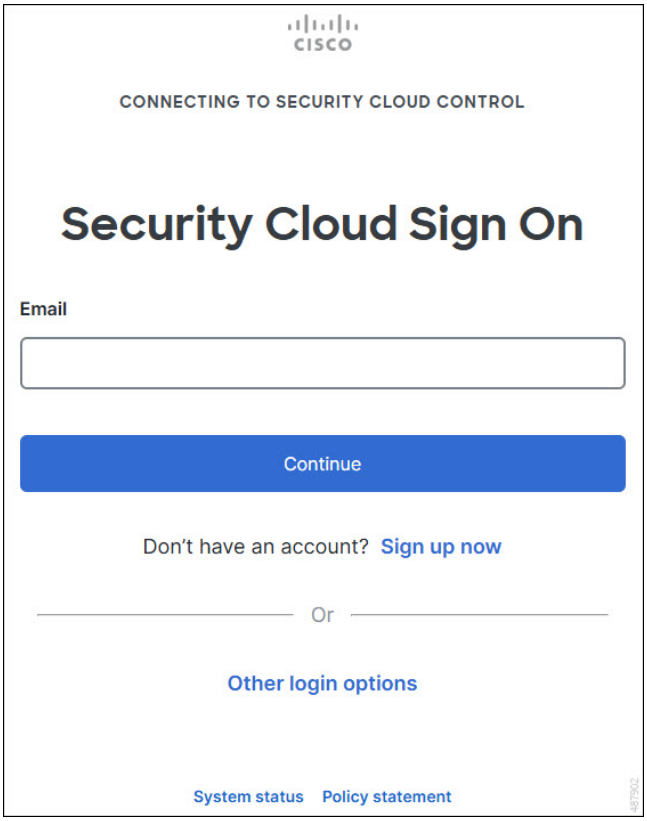

If you do not have a Security Cloud Sign On account to log in to Security Cloud Control and you want to create one, click **Sign up now** in the **Security Cloud Sign On** page. See Create a New Cisco Security Cloud Sign On Account.

Step 6 Choose a Security Cloud Control tenant that you want to use for this integration. The 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

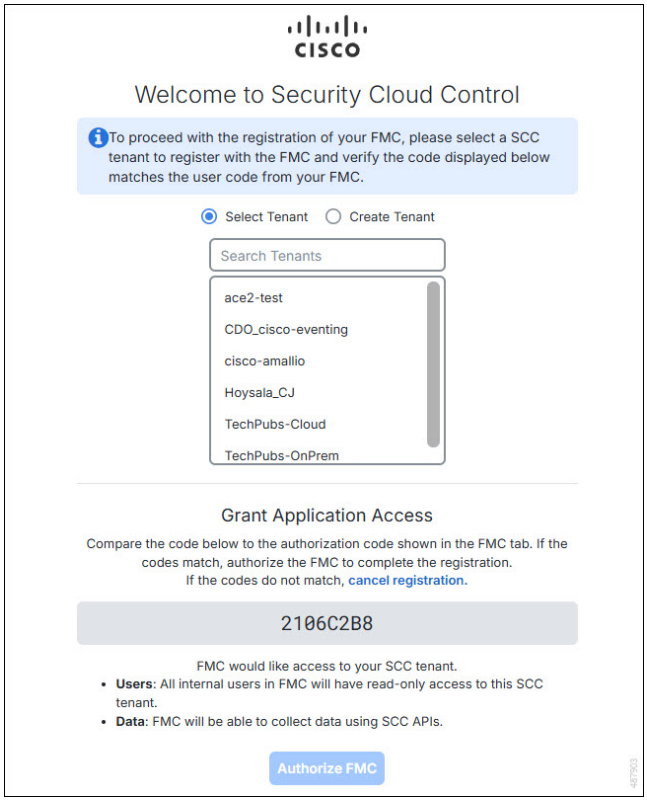

If you do not already have a Security Cloud Control tenant or if you want to use a new tenant for this integration, create a new tenant. See Request a Security Cloud Control Tenant for more information.

Step 7 Verify that the code displayed in the Security Cloud Control login page matches the code provided by the management center.

Figure 5: Verification Code in Management Center

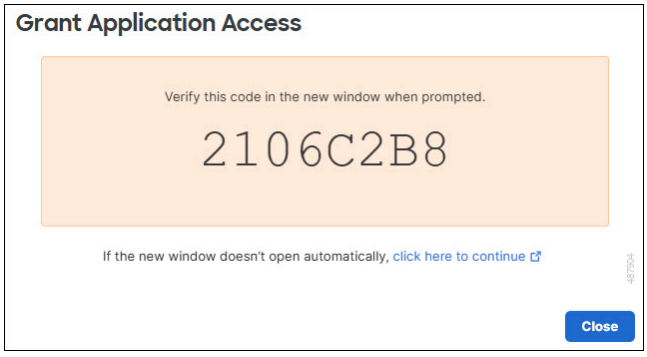

Step 8 Click **Authorize FMC**.

Step 9 In the management center, configure the following:

- **Event Configuration** : Enable this setting to allow threat defense devices to send events directly to the cloud. The event types configured on this page can be used for multiple integrations, where applicable, and enabled. For more information, see Enable Sending Events to the Cisco Security Cloud, on page 618.
- **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 Management Center to Share Usage Metrics and Statistics with Cisco , on page 44 and Configure Management Center to Share Device Health Data with Cisco, on page 45.

Step 10 Click **Save**.

Configure Management Center to Share Usage Metrics and Statistics with Cisco

Cisco Success Network is a cloud service that enables the management center to establish a secure connection to Cisco cloud and stream usage information and statistics. Streaming this telemetry provides a mechanism to select data of interest from the threat defense device and send it in a structured format to remote management stations for the following reasons:

- 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, see Cisco Success Network Telemetry Data Collected from Cisco Secure Firewall Management Center Devices.

The 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 management center and the 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 management center from the Cisco cloud.

You can enable Cisco Success Network when you register the management center with the Smart Software Manager.

Note

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

Before you begin

Enable SecureX integration or register your management center with the Smart License to perform this task.

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**.

Configure Management Center to Share Device Health Data with Cisco

Cisco Support Diagnostics is a user-enabled cloud-based TAC support service. When enabled, the management center and the managed devices establish a secure connection with the Cisco cloud to stream system health-related information.

Cisco Support Diagnostics provides an enhanced user experience during troubleshooting by allowing Cisco TAC to securely collect essential data from your device during the resolution of a TAC case. Moreover, Cisco periodically collects health data, and processes this data using an automated problem-detection system to notify you of issues if any. While data collection service during the resolution of a TAC case is available for all users with support contracts, the notification service is available only to users with specific service contracts.

Cisco Support Diagnostics allows both threat defense devices and the management center to establish and maintain secure connections with the Cisco cloud. The management center sends the collected data to the regional cloud selected on the **SecureX Integration** page.

You can turn off this connection at any time by disabling both Cisco Success Network and Cisco Support Diagnostics, which disconnect these features from the Cisco cloud.

Administrators can view a sample data set collected from the management center by following the steps in Producing Troubleshooting Files for Specific System Functions.

Before you begin

Enable SecureX integration or register your management center with the Smart License to perform this task.

Procedure

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

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

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

Step 3 Click **Save**.

Requirements and Prerequisites for the System Configuration

Model Support

Management Center

Supported Domains

Global

User Roles

Admin

Manage the Secure Firewall Management Center System Configuration

The system configuration identifies basic settings for the management center.

Procedure

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

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

Access List

You can limit access to the management center by IP address and port. By default, the following ports are enabled for any IP address:

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

You can also add access to poll for SNMP information over port 161. Because SNMP is disabled by default, you must first enable SNMP before you can add SNMP access rules. For more information, see Configure SNMP Polling, on page 100.

⚠

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

Configure an Access List

This access list does not control external database access. See Enabling External Access to the Database, on page 63.

Caution If you delete access for the IP address that you are currently using to connect to the management center, and there is no entry for "`IP=any port=443`", you will lose access when you save.

Before you begin

By default, the access list includes rules for HTTPS and SSH. To add SNMP rules to the access list, you must first enable SNMP. For more information, see Configure SNMP Polling, on page 100.

Procedure

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

Step 2 (Optional) Click **SNMP** to configure SNMP if you want to add SNMP rules to the access list. By default, SNMP is disabled; see Configure SNMP Polling, on page 100.

Step 3 Click **Access List**.

Step 4 To add access for one or more IP addresses, click **Add Rules**.

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

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

Step 7 Click **Add**.

Step 8 Click **Save**.

Related Topics

IP Address Conventions, on page 25

Access Control Preferences

Configure access control preferences on **System** (⚙) > **Configuration** > **Access Control Preferences**.

Requiring Comments on Rule Changes

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

Audit Log

The management center records user activity in read-only audit logs. You can review audit log data in several ways:

- Use the web interface: Audit and Syslog, on page 405.

  Audit logs are presented in a standard event view where you can view, sort, and filter audit log messages based on any item in the audit view. You can easily delete and report on audit information and you can view detailed reports of the changes that users make.

- Stream audit log messages to the syslog: Stream Audit Logs to Syslog, on page 48.
- Stream audit log messages to an HTTP server: Stream Audit Logs to an HTTP Server, on page 50.

Streaming audit log data to an external server allows you to conserve space on the 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 ; see Audit Log Certificate, on page 51.

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

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 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.

Before you begin

Make sure the management center can communicate with the syslog server. When you save your configuration, the system uses ICMP/ARP and TCP SYN packets to verify that the syslog server is reachable. Then, the system by default uses port 514/UDP to stream audit logs. If you secure the channel (optional, see Audit Log Certificate, on page 51), you must manually configure port 1470 for TCP.

Procedure

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

Step 2 Click **Audit Log**.

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

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

Option	Description
Host	The IP address or the fully qualified name of the syslog server to which you will send audit logs. 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 555. For example, choose AUDIT.
Severity	The severity of the message. Choose a severity described in Syslog Severity Levels, on page 556.
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 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**.

Stream Audit Logs to an HTTP Server

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

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

Where the local date, time, and originating hostname precede the bracketed optional tag, and the sending appliance 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; see Audit Log Certificate, on page 51.

Procedure

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

Step 2 Click **Audit Log**.

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

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

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

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

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

Step 6 Click **Save**.

Audit Log Certificate

You can use Transport Layer Security (TLS) certificates to secure communications between the management center and a trusted audit log server.

Client Certificates (Required)

Generate a certificate signing request (CSR), submit it to a Certificate Authority (CA) for signing, then import the signed certificate onto the management center. Use the local system configuration: Obtain a Signed Audit Log Client Certificate for the Management Center, on page 52 and Import an Audit Log Client Certificate into the Management Center, on page 53.

Server Certificates (Optional)

For additional security, we recommend you require mutual authentication between the management center and the audit log server. To accomplish this, load one or more certificate revocation lists (CRLs). You cannot stream audit logs to servers with revoked certificates listed in those CRLs.

Secure Firewall supports CRLs encoded in Distinguished Encoding Rules (DER) format. Note that these are the same CRLs that the system uses to validate HTTPS client certificates for the management center web interface.

Use the local system configuration: Require Valid Audit Log Server Certificates, on page 53.

Securely Stream Audit Logs

If you stream the audit log to a trusted HTTP server or syslog server, you can use Transport Layer Security (TLS) certificates to secure the channel between the management center and the server. You must generate a unique client certificate for each appliance you want to audit.

Before you begin

See ramifications of requiring client and server certificates at Audit Log Certificate, on page 51.

Procedure

Step 1 Obtain and install a signed client certificate on the management center:

a) Obtain a Signed Audit Log Client Certificate for the Management Center, on page 52:

Generate a Certificate Signing Request (CSR) from the 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 management center and the audit log server, the client certificate must be signed by the same CA that signed the server certificate to be used for the connection.

b) After you receive the signed certificate from the certificate authority, import it into the management center. See Import an Audit Log Client Certificate into the Management Center, on page 53.

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

See Require Valid Audit Log Server Certificates, on page 53.

Step 3 Configure audit log streaming if you have not yet done so.

See Stream Audit Logs to Syslog, on page 48 or Stream Audit Logs to an HTTP Server, on page 50.

Obtain a Signed Audit Log Client Certificate for the Management Center

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

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

Before you begin

Keep the following 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.

Procedure

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

Step 2 Click **Audit Log Certificate**.

Step 3 Click **Generate New CSR**.

Step 4 Enter a country code in the **Country Name (two-letter code)** field.

Step 5 Enter a state or province postal abbreviation in the **State or Province** field.

Step 6 Enter a **Locality or City**.

Step 7 Enter an **Organization** name.

Step 8 Enter an **Organizational Unit (Department)** name.

Step 9 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

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

Step 10 Click **Generate**.

Step 11 Open a new blank file with a text editor.

Step 12 Copy the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines, and paste it into a blank text file.

Step 13 Save the file as *clientname*.csr, where *clientname* is the name of the appliance where you plan to use the certificate.

Step 14 Click **Close**.

What to do next

- Submit the certificate signing request to the certificate authority that you selected using the guidelines in the "Before You Begin" section of this procedure.
- When you receive the signed certificate, import it to the appliance; see Import an Audit Log Client Certificate into the Management Center, on page 53.

Import an Audit Log Client Certificate into the Management Center

In the management center high availability setup, you *must* use the active peer.

Before you begin

- Obtain a Signed Audit Log Client Certificate for the Management Center, on page 52.
- Make sure you are importing the signed certificate for the correct management center.
- If the signing authority that generated the certificate requires you to trust an intermediate CA, be prepared to provide the necessary certificate chain (or certificate path). The CA that signed the client certificate must be the same CA that signed any intermediate certificates in the certificate chain.

Procedure

Step 1 On the 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**.

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 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. See Audit Log Certificate, on page 51.
- Obtain and import the client certificate following the steps in Securely Stream Audit Logs, on page 51 and the topics referenced in that procedure.

Procedure

Step 1 On the 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 (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, the following 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, you have to add the issuing CA before adding the signing CA from the subordinate CA (Root CA), and so on.

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

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 (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, the following 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 the following 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.

Enabling fetching of the CRL creates a scheduled task for the client to regularly update (download) the CRL or CRLs. The CRL(s) are used for server certificate verification, where, the verification fails if there is a CRL from the CA specifying that the server certificate being verified has been revoked by the CA.

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

Repeat to add up to 25 CRLs.

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

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

Step 8 Click **Save**.

What to do next

(Optional) Set the frequency of CRL updates. See Configuring Certificate Revocation List Downloads, on page 491.

View the Audit Log Client Certificate on the Management Center

You can view the audit log client certificate only for the appliance that you are logged in to. In management center high availability pairs, you can view the certificate only on the active peer.

Procedure

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

Step 2 Click **Audit Log Certificate**.

Change Reconciliation

To monitor the changes that users make and ensure that they follow your organization's preferred standard, you can configure the system to send, via email, a detailed report of changes made over the past 24 hours. Whenever a user saves changes to the system configuration, a snapshot is taken of the changes. The change reconciliation report combines information from these snapshots to present a clear summary of recent system changes.

The following sample graphic displays a User section of an example change reconciliation report and lists both the previous value for each configuration and the value after changes. When users make multiple changes to the same configuration, the report lists summaries of each distinct change in chronological order, beginning with the most recent.

You can view changes made during the previous 24 hours.

Configuring Change Reconciliation

Before you begin

- Configure an email server to receive emailed reports of changes made to the system over a 24-hour period; see Configuring a Mail Relay Host and Notification Address, on page 62 for more information.

Procedure

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

Step 2 Click **Change Reconciliation**.

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

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

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

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

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

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

Step 8 Click **Save**.

Related Topics

Using the Audit Log to Examine Changes, on page 410

Change Reconciliation Options

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

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

Note The change reconciliation report does not include changes to threat defense interfaces and routing settings.

DNS Cache

You can configure the system to resolve IP addresses automatically on the event view pages. You can also configure basic properties for DNS caching performed by the appliance. Configuring DNS caching allows you to identify IP addresses you previously resolved without performing additional lookups. This can reduce the amount of traffic on your network and speed the display of event pages when IP address resolution is enabled.

Configuring DNS Cache Properties

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

Procedure

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

Step 2 Choose **DNS Cache**.

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

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

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

The default setting is 300 minutes (five hours).

Step 5 Click **Save**.

Related Topics

Configuring Event View Settings, on page 206

Dashboard

Dashboards provide you with at-a-glance views of current system status through the use of widgets: small, self-contained components that provide insight into different aspects of the system. The system is delivered with several predefined dashboard widgets.

You can configure the management center so that Custom Analysis widgets are enabled on the dashboard.

Related Topics

About Dashboards, on page 339

Enabling Custom Analysis Widgets for Dashboards

Use Custom Analysis dashboard widgets to create a visual representation of events based on a flexible, user-configurable query.

Procedure

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

Step 2 Click **Dashboard**.

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

Step 4 Click **Save**.

Related Topics

About Dashboards, on page 339

Database

To manage disk space, the 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 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; see Data Purge and Storage, on page 515.

Configuring Database Event Limits

Before you begin

- If you want to receive email notifications when events are pruned from the management center's database, you must configure an email server; see Configuring a Mail Relay Host and Notification Address, on page 62.

Procedure

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

Step 2 Choose **Database**.

Step 3 For each of the databases, enter the number of records you want to store.

For information on how many records each database can maintain, see Database Event Limits, on page 59.

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

Step 5 Click **Save**.

Database Event Limits

The following table lists the minimum and maximum number of records for each event type that you can store per management center.

Table 1: Database Event Limits

Event Type	Upper Limit	Lower Limit
Intrusion events	10 million (management center Virtual) 30 million (management center 1000, management center 1600,) 60 million (management center 2500, management center 2600, , FMCv 300) 300 million (management center 4500, management center 4600)	10,000
Discovery events	10 million (management center Virtual) 20 million (management center 2500, management center 2600, management center 4500, management center 4600, FMCv 300)	Zero (disables storage)

Event Type	Upper Limit	Lower Limit
Connection events Security-related connection events	50 million (management center Virtual) 100 million (management center 1000, management center 1600,) 300 million (management center 2500, management center 2600, , FMCv 300) 1 billion (management center 4500, management center 4600) Limit is shared between connection events and security-related connection 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-related connection, intrusion, file, and malware events are not stored on the management center. **Caution** Setting **Maximum Connection Events** to zero immediately purges existing connection events other than security-related connection events. See below for the effect of this setting on Maximum Flow Rate. These settings do not affect connection summaries.
Connection summaries (aggregated connection events)	50 million (management center Virtual) 100 million (management center 1000, management center 1600,) 300 million (management center 2500, management center 2600, , FMCv 300) 1 billion (management center 4500, management center 4600)	Zero (disables storage)
Correlation events and compliance allow list events	1 million (management center Virtual) 2 million (management center 2500, management center 2600, management center 4500, management center 4600, FMCv 300)	One
Malware events	10 million (management center Virtual, management center 1600,) 20 million (management center 2500, management center 2600, , management center 4500, management center 4600, FMCv 300)	10,000
File events	10 million (management center Virtual, management center 1600,) 20 million (management center 2500, management center 2600, , management center 4500, management center 4600, FMCv 300)	Zero (disables storage)
Health events	1 million	Zero (disables storage)
Audit records	100,000	One

Event Type	Upper Limit	Lower Limit
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
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 management center hardware model is specified in the **Platform Specifications** section of the 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-related connection events, intrusion, file, and malware events are not counted toward the maximum flow rate for your 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 Notification

Configure a mail host if you plan to:

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

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

Configuring a Mail Relay Host and Notification Address

Procedure

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

Step 2 Click **Email Notification**.

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

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

Typical ports include:

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

Step 5 Choose an **Encryption Method**:

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

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

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

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

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

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

Step 9 Click **Save**.

External Database Access

You can configure the management center to allow read-only access to its database by a third-party client. This allows you to query the database using SQL using any of the following:

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

Use the management center's system configuration to enable database access and create an access list that allows selected hosts to query the database. Note that this access list does not also control appliance access.

You can also download a package that contains the following:

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

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

Enabling External Access to the Database

Procedure

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

Step 2 Click **External Database Access**.

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

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

Note
In 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**.

Related Topics

IP Address Conventions, on page 25

HTTPS Certificates

Secure Sockets Layer (SSL)/TLS certificates enable management centers to establish an encrypted channel between the system and a web browser. A default certificate is included with all firewall devices, but it is not generated by a certificate authority (CA) trusted by any globally known CA. For this reason, consider replacing it with a custom certificate signed by a globally known or internally trusted CA.

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

Note HTTPS certificates are not supported on the management center REST API.

Default HTTPS Server Certificates

If you use the default server certificate provided with an appliance, do not configure the system to require a valid HTTPS client certificate for web interface access because the default server certificate is not signed by the CA that signs your client certificate.

The 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**.

Note that some Secure Firewall software upgrades can automatically renew the certificate. For more information, see the appropriate version of the Cisco Secure FirewallRelease Notes.

On the management center, you can renew the default certificate on the **System** (⚙) > **Configuration** > **HTTPS Certificate** page.

Custom HTTPS Server Certificates

You can use the management center web interface to generate a server certificate request based on your system information and the identification information you supply. You can use that request to sign a certificate if you have an internal certificate authority (CA) installed that is trusted by your browser. You can also send the resulting request to a certificate authority to request a server certificate. After you have a signed certificate from a certificate authority (CA), you can import it.

HTTPS Server Certificate Requirements

When you use HTTPS certificates to secure the connection between your web browser and the Secure Firewall appliance web interface, you must use certificates that comply with the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile (RFC 5280). When you import a server certificate to the appliance, the system rejects the certificate if it does not comply with version 3 (X.509 v3) of that standard.

Before importing an HTTPS server certificate, be certain it includes the following fields:

Certificate Field	Description
Version	Version of the encoded certificate. Use version `3`. See RFC 5280, section 4.1.2.1.
Serial number	A positive integer assigned to the certificate by the issuing CA. Issuer and serial number together uniquely identify the certificate. See RFC 5280, section 4.1.2.2.
Signature	Identifier for the algorithm used by the CA to sign the certificate. Must match the signatureAlgorithm field. See RFC 5280, section 4.1.2.3.
Issuer	Identifies the entity that signed and issued the certificate. See RFC 5280, section 4.1.2.4.
Validity	Interval during which the CA warrants that it will maintain information about the status of the certificate. See RFC 5280, section 4.1.2.5.
Subject	Identifies the entity associated with the public key stored in the subject public key field; must be an X.500 distinguished name (DN). See RFC 5280, section 4.1.2.6.
Subject Alternative Name	Domain names and IP addresses secured by the certificate. Subject Alternative Name is defined in section RFC 5280, section 4.2.1.6. We recommend you use 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. See RFC 5280, section 4.1.2.7.
Authority Key Identifier	Provides a means of identifying the public key corresponding to the private key used to sign a certificate. See RFC 5280, section 4.2.1.1.
Subject Key Identifier	Provides a means of identifying certificates that contain a particular public key. See RFC 5280, section 4.2.1.2.
Key Usage	Defines the purpose of the key contained in the certificates. See RFC 5280, section 4.2.1.3.
Basic Constraints	Identifies whether the certificate Subject is a CA, and the maximum depth of validation certification paths that include this certificate. See RFC 5280, section 4.2.1.9. For server certificates used in Secure Firewall appliances, use `critical CA:FALSE`.

Certificate Field	Description
Extended Key Usage extension	Indicates one or more purposes for which the certified public key may be used, in addition to or in place of the basic purposes indicated in the Key Usage extension. See RFC 5280, section 4.2.1.12. Be certain you import certificates that can be used as server certificates.
signatureAlgorithm	Identifier for the algorithm the CA used to sign the certificate. Must match the Signature field. See RFC 5280, section 4.1.1.2.
signatureValue	Digital signature. See RFC 5280, section 4.1.1.3.

HTTPS Client Certificates

You can restrict access to the system web server using client browser certificate checking. When you enable user certificates, the web server checks that a user's browser client has a valid user certificate selected. That user certificate must be generated by the same trusted certificate authority that is used for the server certificate. The browser cannot load the web interface under any of the following circumstances:

- The user selects a certificate in the browser that is not valid.
- The user selects a certificate in the browser that is not generated by the certificate authority that signed the server certificate.
- The user selects a certificate in the browser that is not generated by a certificate authority in the certificate chain on the device.

To verify client browser certificates, configure the system to use the online certificate status protocol (OCSP) or load one or more certificate revocation lists (CRLs). Using the OCSP, when the web server receives a connection request it communicates with the certificate authority to confirm the client certificate's validity before establishing the connection. If you configure the server to load one or more CRLs, the web server compares the client certificate against those listed in the CRLs. If a user selects a certificate that is listed in a CRL as a revoked certificate, the browser cannot load the web interface.

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

Viewing the Current HTTPS Server Certificate

Procedure

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

Step 2 Click **HTTPS Certificate**.

Generating an HTTPS Server Certificate Signing Request

If you install a certificate that is not signed by a globally known or internally trusted CA, the user's browser displays a security warning when they try to connect to the web interface.

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

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

Procedure

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

Step 2 Click **HTTPS Certificate**.

Step 3 Click **Generate New CSR**.

The following figure shows an example.

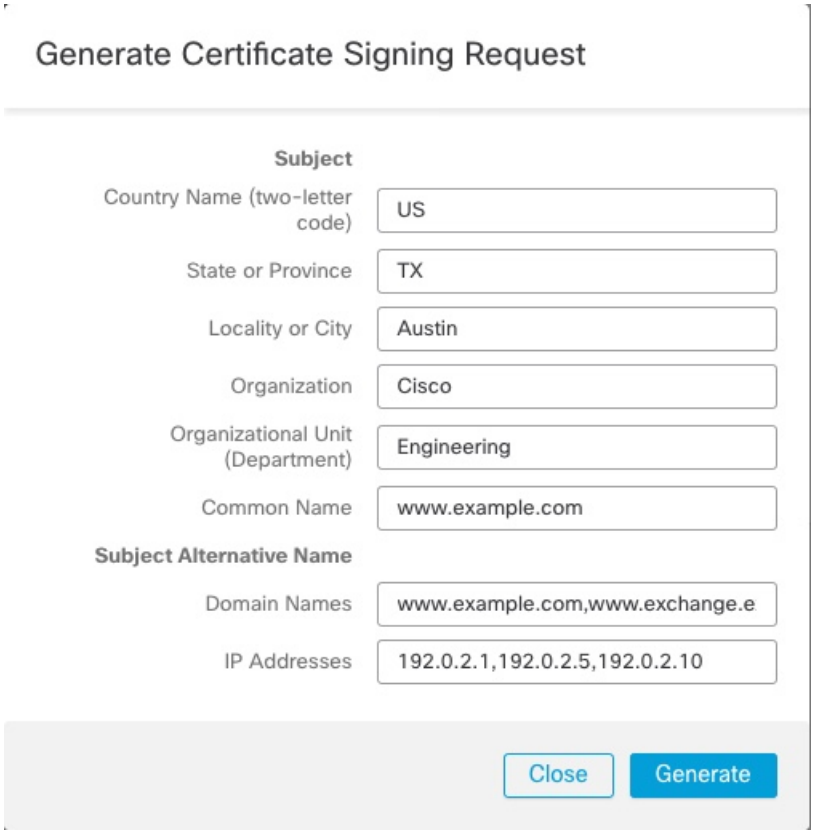

Step 4 Enter a country code in the **Country Name (two-letter code)** field.

Step 5 Enter a state or province postal abbreviation in the **State or Province** field.

Step 6 Enter a **Locality or City**.

Step 7 Enter an **Organization** name.

Step 8 Enter an **Organizational Unit (Department)** name.

Step 9 Enter the fully qualified domain name of the server for which you want to request a certificate in the **Common Name** field.

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

Step 10 To request a certificate that secures multiple domain names or IP addresses, enter the following 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 11 Click **Generate**.

Step 12 Open a text editor.

Step 13 Copy the entire block of text in the certificate request, including the `BEGIN CERTIFICATE REQUEST` and `END CERTIFICATE REQUEST` lines, and paste it into a blank text file.

Step 14 Save the file as *`servername`*`.csr`, where *servername* is the name of the server where you plan to use the certificate.

Step 15 Click **Close**.

What to do next

- Submit the certificate request to the certificate authority.
- When you receive the signed certificate, import it to the management center; see Importing HTTPS Server Certificates, on page 68.

Importing HTTPS Server Certificates

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

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

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

Caution The management center supports 4096-bit HTTPS certificates. If the certificate used by the management center was generated using a public server key larger than 4096 bits, you will not be able to log in to the Secure Firewall Management Center web interface. For more information about updating HTTPS Certificates to Version 6.0.0, see "Update Management Center HTTPS Certificates to Version 6.0" in *Firepower System Release Notes, Version 6.0*. If you generate or import an HTTPS Certificate and cannot log in to the management center web interface, contact Support.

Before you begin

- Generate a certificate signing request; see Generating an HTTPS Server Certificate Signing Request, on page 67.
- Upload the CSR file to the certificate authority where you want to request a certificate or use the CSR to create a self-signed certificate.
- Confirm that the certificate meets the requirements described in HTTPS Server Certificate Requirements, on page 64.

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 Generating an HTTPS Server Certificate Signing Request, on page 67), the system already has the private key and you need not enter one here.
- If you generated the Certificate Signing Request using some other means, you must supply the private key here. Open the private key file and copy the entire block of text, include the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Paste this text into the **Private Key** field.

Step 6 Open any required intermediate certificates, copy the entire block of text for each, and paste it into the **Certificate Chain** field. If you received a root certificate, paste it here. If you received an intermediate certificate, paste it below the root certificate. In both cases, copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines.

Step 7 Click **Save**.

Requiring Valid HTTPS Client Certificates

Use this procedure to require users connecting to the 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, to ensure that the list of revoked certificates stays current, you can create a scheduled task to update the CRLs. The system displays the most recent refresh of the CRLs.

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

Before you begin

- Import a server certificate signed by the same certificate authority that signed the client certificate to be used for the connection; see Importing HTTPS Server Certificates, on page 68.
- Import the server certificate chain if needed; see Importing HTTPS Server Certificates, on page 68.

Procedure

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

Step 2 Click **HTTPS Certificate**.

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

Step 4 You have three options:

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

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

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

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

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

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

Step 8 Click **Save**.

Related Topics

Configuring Certificate Revocation List Downloads, on page 491

Renewing the Default HTTPS Server Certificate

You can only view server certificates for the appliance you are logged in to.

Procedure

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

Step 2 Click **HTTPS Certificate**.

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 below 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**.

What to do next

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

Information

The **System** > **Configuration** page of the web interface includes the information listed in the table below. Unless otherwise noted, all fields are read-only.

Note See also the **Help** > **About** page, which includes similar but slightly different information.

Field	Description
Name	A descriptive name you assign to the 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 management center with Cisco XDR. If you change the name, all registered devices are marked out of date and deployment is required to push the new name to the devices.
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

Configure various intrusion policy preferences to monitor and track changes to the critical policies in your deployment.

Set Intrusion Policy Preferences

Configure the intrusion policy preferences.

Procedure

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

Step 2 Click **Intrusion Policy Preferences**.

Step 3 You have the following 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 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 enabled, the system retains the rule overrides in the new replacement rules that are added as part of the LSP update. On the management center menu bar, click **Notifications** > **Tasks** to view the notifications. This option is enabled by default.

Language

You can use the Language page to specify a different language for the web interface.

Set the Language for the Web Interface

The language you specify here is used for the web interface for every user. You can choose from:

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

Procedure

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

Step 2 Click **Language**.

Step 3 Choose the language you want to use.

Step 4 Click **Save**.

Login Banner

You can use the Login Banner page to specify session, login, or custom message banners for a security appliance or shared policy.

You can use ASCII characters and carriage returns to create a custom login banner. The system does not preserve tab spacing. If your login banner is too large or causes errors, Telnet or SSH sessions can fail when the system attempts to display the banner.

Customize the Login Banner

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**.

Management Interfaces

After setup, you can change the management network settings, including adding more management interfaces, hostname, search domains, DNS servers, and HTTP proxy on the management center.

About Management Center Management Interfaces

By default, the management center manages all devices on a single management interface. You can also perform initial setup on the management interface and log into the management center on this interface as an administrator. The management interface is also used to communicate with the Smart Licensing server, to download updates, and to perform other management functions.

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

About Device Management

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

By using the 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 management center

Note If you have a Security Cloud Control-managed device and are using the on-prem management center for analytics only, then the on-prem 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 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 management center to manage nearly every aspect of a device's behavior.

Note Although the 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 threat defense software are not available to these previous-release devices. Some management center features may be available for earlier versions.

The Management Connection

After you configure the device with the management center information and after you add the device to the management center, either the device or the management center can establish the management connection. Depending on initial setup:

- Either the device or the management center can initiate.
- Only the device can initiate.
- Only the management center can initiate.

Initiation always originates with eth0 on the management center or with the lowest-numbered management interface on the device. Additional management interfaces are tried if the connection is not established. Multiple management interfaces on the management center let you connect to discrete networks or to segregate management and event traffic. However, the initiator does not choose the best interface based on the routing table.

Make sure the management connection is stable, without excessive packet loss, with at least 5 Mbps throughput. By default, the management connection uses TCP port 8305 (this port is configurable). If you place another threat defense between devices and the management center, to prevent potential management disruption, be sure to 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 itself 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 Interfaces on the Management Center

The management center uses the eth0 interface for initial setup, HTTP access for administrators, management of devices, as well as other management functions such as licensing and updates.

You can also configure additional management interfaces. When the management center manages large numbers of devices on different networks, adding more management interfaces can improve throughput and performance. You can also use these interfaces for all other management functions. You might want to use each management interface for particular functions; for example, you might want to use one interface for HTTP administrator access and another for device management.

For device management, the management interface carries two separate traffic channels: the *management traffic channel* carries all internal traffic (such as inter-device traffic specific to managing the device), and the *event traffic channel* carries all event traffic (such as web events). You can optionally configure a separate event-only interface on the management center to handle event traffic; you can configure only one event interface. You must also always have a management interface for the management traffic channel. Event traffic can use a large amount of bandwidth, so separating event traffic from management traffic can improve the performance of the management center. For example, you can assign a 10 GigabitEthernet interface to be the event interface, if available, while using 1 GigabitEthernet interfaces for management. You might want to configure an event-only interface on a completely secure, private network while using the regular management interface on a network that includes Internet access, for example. Though you may use both management and event interfaces on the same network, we recommend that placing each interface on a separate network to avoid potential routing problems, including routing problems from other devices to the management center. Managed devices will send management traffic to the management center's management interface and event traffic to the management center's event-only interface. If the managed device cannot reach the event-only interface, then it will fall back to sending events to the management interface. However, the management connections cannot be made through the event-only interface.

Management connection initiation from the management center is always attempted first from eth0 and 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 47). Conversely, you cannot restrict an interface to *only* HTTP access; management interfaces always support device management (management traffic, event traffic, or both).

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

Management Interface Support Per Management Center Model

See the hardware installation guide for your model for the management interface locations.

See the following table for supported management interfaces on each management center model.

Table 2: Management Interface Support on the Management Center

Model	Management Interfaces
MC1000	eth0 (Default) eth1
MC2500, MC4500	eth0 (Default) eth1 eth2 eth3
MC1600, MC2600, MC4600	eth0 (Default) eth1 eth2 eth3 CIMC (Supported for Lights-Out Management only.)
Management Center Virtual	eth0 (Default)

Network Routes on Management Center Management Interfaces

Management interfaces (including event-only interfaces) support only static routes to reach remote networks. When you set up your management center, the setup process creates a default route to the gateway IP address that you specify. You cannot delete this route; you can only modify the gateway address.

You can configure multiple management interfaces on some platforms. The default route does not include an egress interface, so the interface chosen depends on the gateway address you specify, and which interface's network the gateway belongs to. In the case of multiple interfaces on the default network, the device uses the lower-numbered interface as the egress interface.

At least one static route is recommended per management interface to access remote networks. We recommend placing each interface on a separate network to avoid potential routing problems, including routing problems from other devices to the management center.

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.

NAT Environments

Network address translation (NAT) is a method of transmitting and receiving network traffic through a router that involves reassigning the source or destination IP address. The most common use for NAT is to allow private networks to communicate with the internet. Static NAT performs a 1:1 translation, which does not pose a problem for management center communication with devices, but port address translation (PAT) is more common. PAT lets you use a single public IP address and unique ports to access the public network; these ports are dynamically assigned as needed, so you cannot initiate a connection to a device behind a PAT router.

Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the management center specifies the device IP address when you add a device, and the device specifies the management center IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The management center and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

For example, you add a device to the management center, and you do not know the device IP address (for example, the device is behind a PAT router), so you specify only the NAT ID and the registration key on the management center; leave the IP address blank. On the device, you specify the management center IP address, the same NAT ID, and the same registration key. The device registers to the management center's IP address. At this point, the management center uses the NAT ID instead of IP address to authenticate the device.

Although the use of a NAT ID is most common for NAT environments, you might choose to use the NAT ID to simplify adding many devices to the management center. On the management center, specify a unique NAT ID for each device you want to add while leaving the IP address blank, and then on each device, specify both the management center IP address and the NAT ID. Note: The NAT ID must be unique per device.

The following example shows three devices behind a PAT IP address. In this case, specify a unique NAT ID per device on both the management center and the devices, and specify the management center IP address on the devices.

Figure 6: NAT ID for Managed Devices Behind PAT

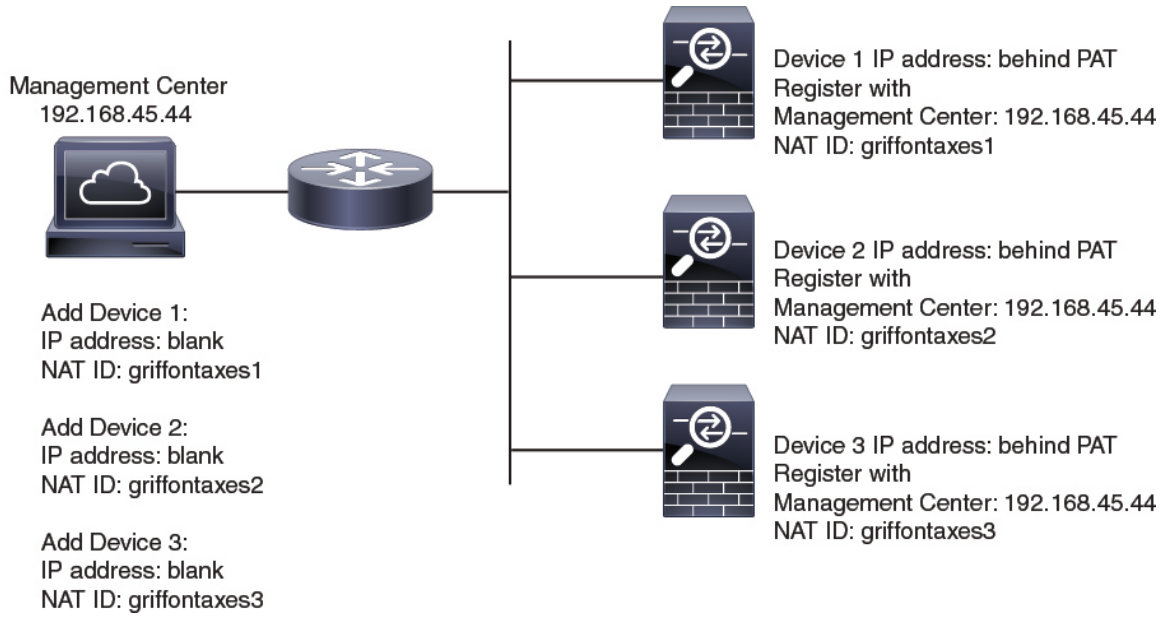

The following example shows the management center behind a PAT IP address. In this case, specify a unique NAT ID per device on both the management center and the devices, and specify the device IP addresses on the management center.

Figure 7: NAT ID for Management Center Behind PAT

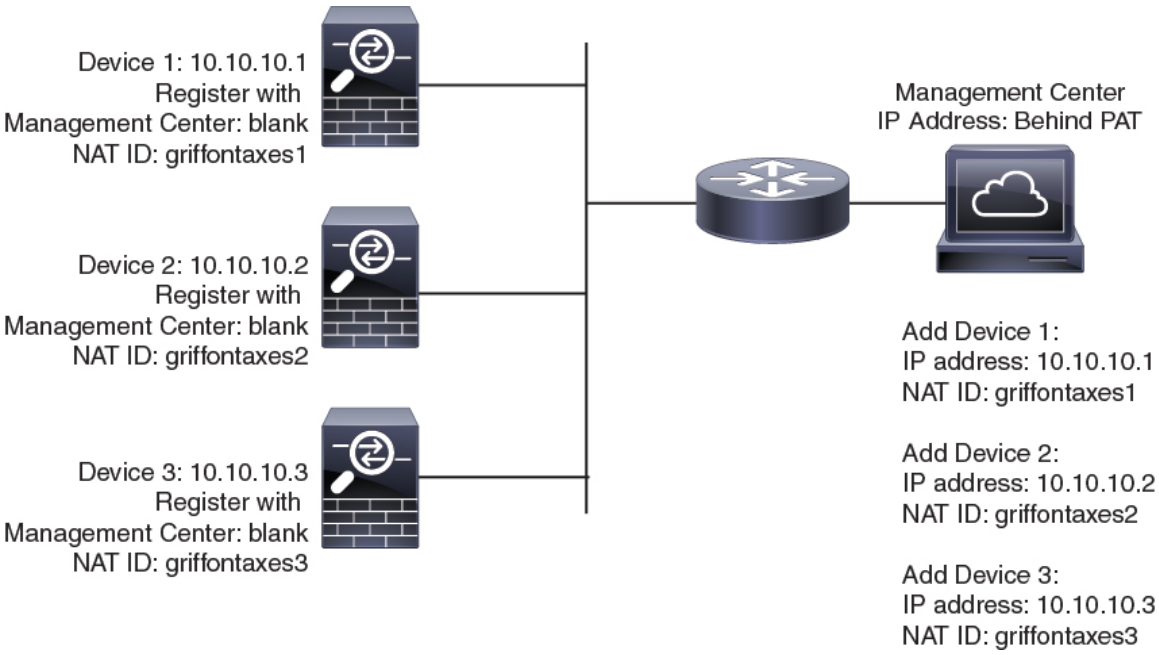

Management and Event Traffic Channel Examples

Note If you use a data interface for management on a threat defense, you cannot use separate management and event interfaces for that device.

The following example shows the management center and managed devices using only the default management interfaces.

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

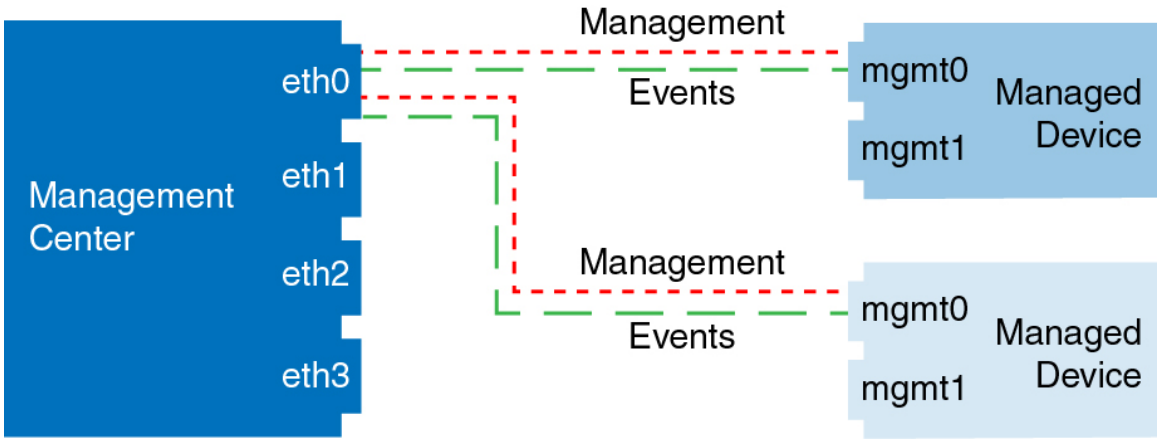

The following example shows the management center using separate management interfaces for devices; and each managed device using 1 management interface.

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

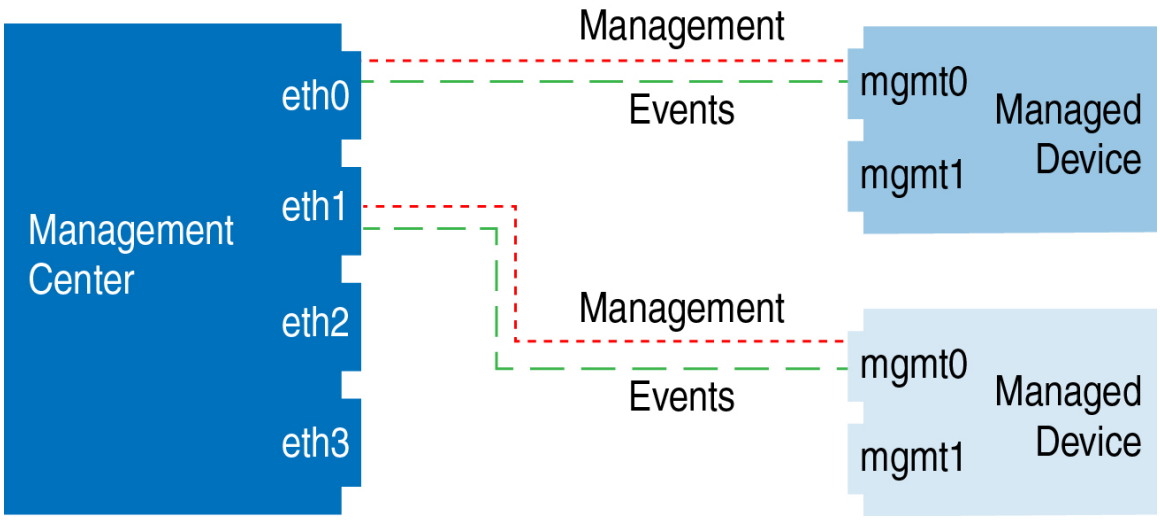

The following example shows the management center and managed devices using a separate event interface.

Figure 10: Separate Event Interface on the Secure Firewall Management Center and Managed Devices

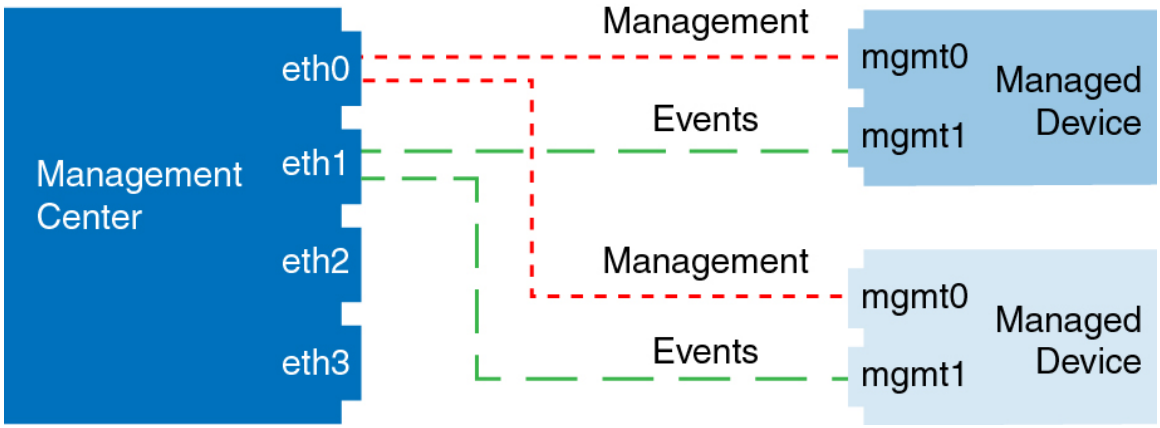

The following example shows a mix of multiple management interfaces and a separate event interface on the management center and a mix of managed devices using a separate event interface, or using a single management interface.

Figure 11: Mixed Management and Event Interface Usage

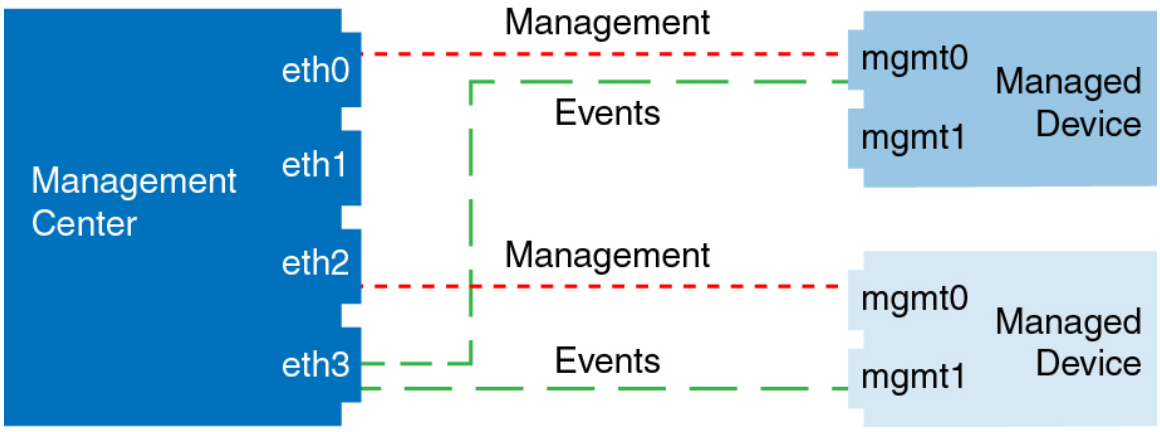

Modify Management Center Management Interfaces

Modify the management interface settings on the 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 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 management center IP address, then see *Edit the management center IP Address or Hostname on the Device* in the Cisco Secure Firewall Management Center Device Configuration Guide. If you change the 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 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 management center and you specified the NAT ID only. Even in other cases, we recommend keeping the management center IP address or hostname up to date for extra network resiliency.

In a high availability configuration, when you modify the management IP address of a registered device from the device CLI or from the management center, the secondary management center does not reflect the changes even after an high availability synchronization. To ensure that the secondary management center is also updated, switch roles between the two management centers, making the secondary management center as the active unit. Modify the management IP address of the registered device on the Device Management page of the now active management center.

If you modify the management IP address of one peer management center in a high availability configuration, the remote peer does not reflect the changes even after an high availability synchronization. To ensure that the remote peer management center is also updated, you must log in to the remote peer management center, navigate to **Integration** > **Other Integrations** > **High Availability** > **Peer Manager**, and then manually update the IP address of its peer manager. For more detailed instructions, see Change the IP Address of the Management Center in a High Availability Pair, on page 319.

Before you begin

- For information about how device management works, see *About Device Management Interfaces* in the Cisco Secure Firewall Management Center Device Configuration Guide.
- If you use a proxy:
  - Proxies that use NT LAN Manager (NTLM) authentication are not supported.
  - If you use or will use Smart Licensing, the proxy FQDN cannot have more than 64 characters.

Procedure

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

Step 2 In the **Interfaces** area, click **Edit** next to the interface that you want to configure.

All available interfaces are listed in this section. You cannot add more interfaces.

You can configure the following options on each management interface:

- **Enabled**—Enable the management interface. Do **not** disable the default eth0 management interface. Some processes require the eth0 interface.
- **Channels**—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 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 management center network configuration gets out of sync. To recover from a DHCP address change, connect to the management center (using the hostname or the new IP address) and navigate to **System** (⚙) > **Configuration** > **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, see Network Routes on Management Center Management Interfaces, on page 77.

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 the following settings for a static route:

- **Destination**—Set the destination address of the network to which you want to create a route.
- **Netmask** or **Prefix Length**—Set the netmask (IPv4) or prefix length (IPv6) for the network.
- **Interface**—Set the egress management interface.
- **Gateway**—Set the gateway IP address.

Step 4 In the **Shared Settings** area, set network parameters shared by all interfaces.

Note
If you selected **DHCP** for the eth0 interface, you cannot manually specify some shared settings derived from the DHCP server.

You can configure the following shared settings:

- **Hostname**—Set the 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 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 management center, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system**. The domains are used only on the management interface, or for commands that go through the management interface.
- **Primary DNS Server**, **Secondary DNS Server**, **Tertiary DNS Server**—Set the DNS servers to be used in order of preference.

- **Remote Management Port**—Set the remote management port for communication with managed devices. The 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 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 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.

See proxy requirements in the prerequisites to this topic.

a) Check the **Enabled** check box.
b) In the **HTTP Proxy** field, enter the IP address or fully-qualified domain name of your proxy server.

   See requirements in the prerequisites to this topic.

c) In the **Port** field, enter a port number.
d) Supply authentication credentials by choosing **Use Proxy Authentication**, and then provide a **User Name** and **Password**.

Step 7 Click **Save**.

Step 8 If you change the management center IP address, then see *Edit the management center IP Address or Hostname on the Device* in the Cisco Secure Firewall Management Center Device Configuration Guide.

If you change the 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 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 management center and you specified the NAT ID only. Even in other cases, we recommend keeping the management center IP address or hostname up to date for extra network resiliency.

Change Both Management Center and Threat Defense IP Addresses

You might want to change both management center and 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 ( ).

Figure 12: Disable Management

You are prompted to proceed with disabling management; click **Yes**.

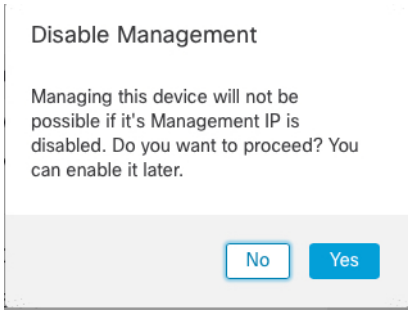

Step 2 Change the device IP address in the management center to the new device IP address.

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 **Remote Host Address** IP address and optional **Secondary Address** (when using a redundant data interface) or hostname by clicking **Edit** ( ).

Figure 13: Edit Management Address

b) In the **Management** dialog box, modify the name or IP address in the **Remote Host Address** field and the optional **Secondary Address** field, and click **Save**.

Figure 14: Management IP Address

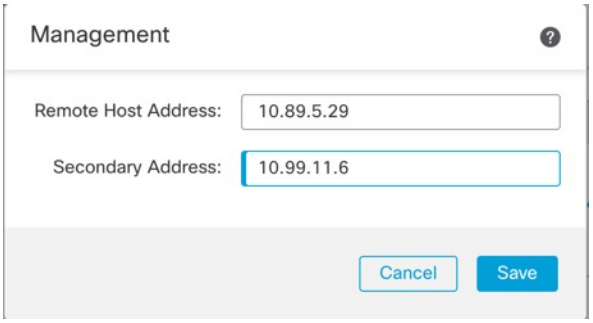

Step 3 Change the management center IP address.

Caution
Be careful when making changes to the management center interface to which you are connected; if you cannot re-connect because of a configuration error, you need to access the 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 threat defense CLI, view the 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 management center IP address or hostname.

configure manager edit *identifier* {**hostname** {*ip_address* | *hostname*} | **displayname** *display_name*}

If the 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 Reenable management by clicking the slider so it is enabled ( ).

For a high-availability pair or cluster, perform these steps on all units.

Figure 15: Enable Management Connection

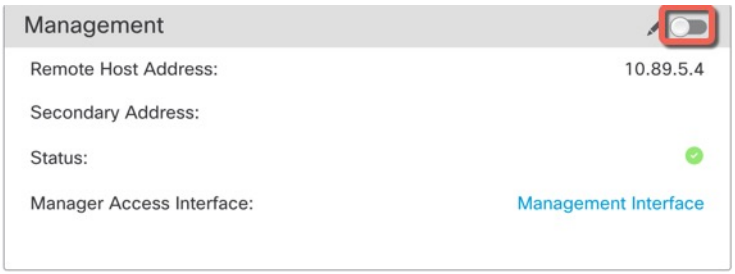

Step 7 (If using a data interface for manager access) Refresh the data interface settings in the management center.

For a high-availability pair, perform this step on both units.

a) Choose **Devices** > **Device Management** > **Device** > **Management** > **Manager Access - Configuration Details**, and click **Refresh**.
b) Choose **Devices** > **Device Management** > **Interfaces**, 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 management center, check the management connection status on the **Devices** > **Device Management** > **Device** > **Management** > **Manager Access - Configuration Details** > **Connection Status** page.

At the threat defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 16: Connection Status

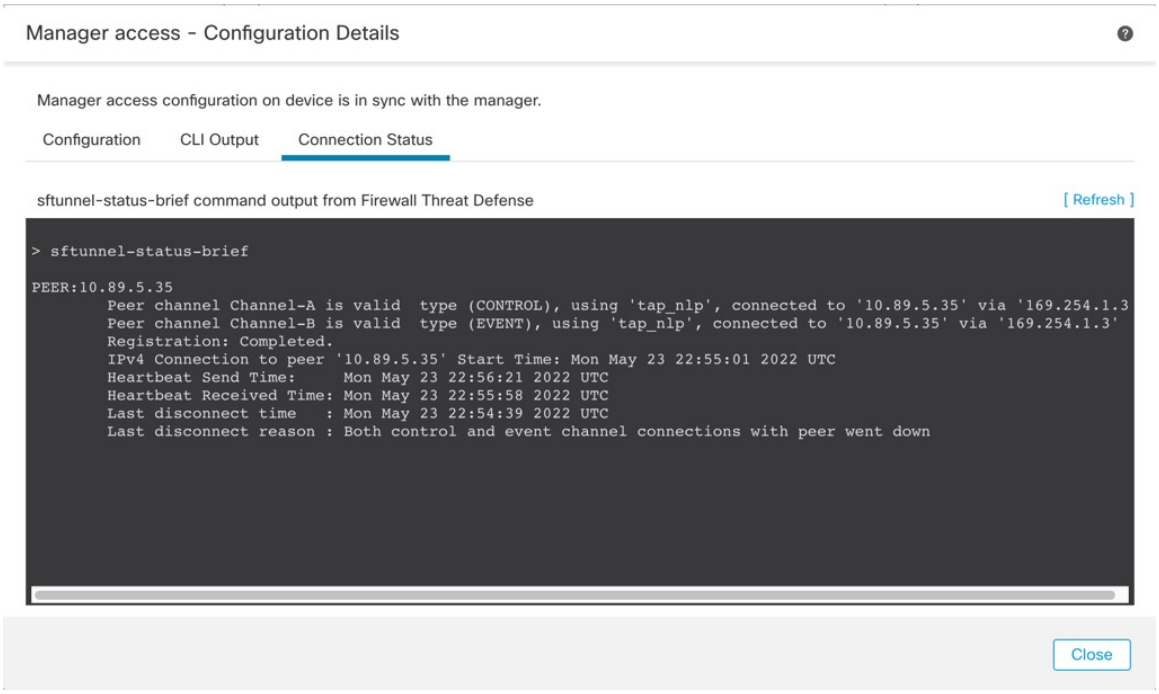

Step 9 (For a high-availability management center pair) Repeat configuration changes on the secondary management center.

a) Change the secondary 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 management center.
f) Reenable the management connection.

Network Analysis Policy Preferences

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.

Process

Use the web interface to control the shut down and restart of processes on the management center. You can:

- 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.

Shut Down or Restart the Management Center

Procedure

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

Step 2 Choose **Process**.

Step 3 Do one of the following:

Shut down	Click **Run Command** next to **Shutdown Management Center**.
Reboot	Click **Run Command** next to **Reboot Management Center**. **Note** Rebooting logs you out, and the system runs a database check that can take up to an hour to complete.
Restart the console	Click **Run Command** next to **Restart Management Center Console**. **Note** Restarting may cause deleted hosts to reappear in the network map.

REST API Preferences

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, see the Secure Firewall Management Center REST API Quick Start Guide.

Note HTTPS certificates are not supported on the management center REST API.

By default, the management center allows requests from applications using the REST API. You can configure the management center to block this access.

Enabling REST API Access

Note In deployments using the management center high availability, this feature is available only in the active management center.

Procedure

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

Step 2 Click **REST API Preferences**.

Step 3 To enable or disable REST API access to the 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
```

Remote Console Access Management

You can use a Linux system console for remote access on supported systems via either the VGA port (which is the default) or the serial port on the physical appliance. Use the Console Configuration page to choose the option most suitable to the physical layout of your organization's Secure Firewall deployment.

On supported physical-hardware-based systems, you can use Lights-Out Management (LOM) on a Serial Over LAN (SOL) connection to remotely monitor or manage the system without logging into the management interface of the system. You can perform limited tasks, such as viewing the chassis serial number or monitoring such conditions as fan speed and temperature, using a command line interface on an out-of-band management connection. The cable connection to support LOM varies by management center model:

- For management center models MC1600, MC2600, and MC4600, use a connection with the CIMC port to support LOM. See the *Cisco Firepower Management Center 1600, 2600, and 4600 Getting Started Guide* for more information.
- For all other management center hardware models, use a connection with the default (eth0) management port to support LOM. See the *Navigating the Cisco Secure Firewall Threat Defense Documentation Guide* for your hardware model.

You must enable LOM for both the system and the user you want to manage the system. After you enable the system and the user, you use a third-party Intelligent Platform Management Interface (IPMI) utility to access and manage your system.

Configuring Remote Console Settings on the System

You must be an Admin user to perform this procedure.

Before you begin

- Disable Spanning Tree Protocol (STP) on any third-party switching equipment connected to the device's management interface.
- If you plan to enable Lights-Out Management see the *Getting Started Guide* for your appliance for information about installing and using an Intelligent Platform Management Interface (IPMI) utility.

Procedure

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

Step 2 Click **Console Configuration**.

Step 3 Choose a remote console access option:

- Choose **VGA** to use the appliance's VGA port.
- Choose **Physical Serial Port** to use the appliance's serial port.
- Choose **Lights-Out Management** to use an SOL connection on the management center. (This may use the default management port or the CIMC port depending on your management center model. See the *Getting Started Guide* for your model for more information.)

Step 4 To configure LOM via SOL:

- Choose the address **Configuration** for the system (**DHCP** or **Manual**).
- If you chose manual configuration, enter the necessary IPv4 settings:
  - Enter the **IP Address** to be used for LOM.

    Note
    The LOM IP address must be different from and in the same subnet as the management center management interface IP address.

  - Enter the **Netmask** for the system.

- Enter the **Default Gateway** for the system.

Step 5 Click **Save**.

Step 6 The system displays the following warning: "You will have to reboot your system for these changes to take effect." Click **OK** to reboot now or **Cancel** to reboot later.

What to do next

- If you configured serial access, be sure the rear-panel serial port is connected to a local computer, terminal server, or other device that can support remote serial access over ethernet as described in the *Getting Started Guide* for your management center model.
- If you configured Lights-Out Management, enable a Lights-Out Management user; see Lights-Out Management User Access Configuration, on page 92.

Lights-Out Management User Access Configuration

You must explicitly grant Lights-Out Management permissions to users who use the feature. LOM users also have the following restrictions:

- You must assign the Administrator role to the user.
- The username may have up to 16 alphanumeric characters. Hyphens and longer 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 120. 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 management centers can have up to 13 LOM users.

Note that if you deactivate and then reactivate a user with LOM while that user is logged in, that 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 management centers. You must create different admin users with LOM enabled on the active 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 management center, the changes get synced to both the active and standby management centers and the active management center CIMC. The new password is not synced with the standby management center for CIMC login. To ensure that the standby management center is also updated, reset the CIMC login password for the local user on the standby management center.

Enabling Lights-Out Management User Access

You must be an Admin user to perform this procedure.

Use this task to grant LOM access to an existing user. To grant LOM access to a new user, see Add or Edit an Internal User, on page 125.

Procedure

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

Step 2 To grant LOM user access to an existing user, click **Edit** (✎) next to a user name in the list.

Step 3 Under **User Configuration**, enable the Administrator role.

Step 4 Check the **Allow Lights-Out Management Access** check box.

Step 5 Click **Save**.

Serial Over LAN Connection Configuration

You use a third-party IPMI utility on your computer to create a Serial Over LAN connection to the appliance. If your computer uses a Linux-like or Mac environment, use IPMItool; for Windows environments, you can use IPMIutil or IPMItool, depending on your Windows version.

Note Cisco recommends using IPMItool version 1.8.12 or greater.

Linux

IPMItool is standard with many distributions and is ready to use.

Mac

You must install IPMItool on a Mac. First, confirm that your Mac has Apple's XCode Developer tools installed, making sure that the optional components for command line development are installed (UNIX Development and System Tools in newer versions, or Command Line Support in older versions). Then you can install macports and the IPMItool. Use your favorite search engine for more information or try these sites:

```
https://developer.apple.com/technologies/tools/
http://www.macports.org/
http://github.com/ipmitool/ipmitool/
```

Windows

For Windows Versions 10 and greater with Windows Subsystem for Linux (WSL) enabled, as well as some older versions of Windows Server, you can use IPMItool. Otherwise, you must compile IPMIutil on your Windows system; you can use IPMIutil itself to compile. Use your favorite search engine for more information or try this site:

```
http://ipmiutil.sourceforge.net/man.html#ipmiutil
```

Understanding IPMI Utility Commands

Commands used for IPMI utilities are composed of segments as in the following example for IPMItool on Mac:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

where:

- `ipmitool` invokes the utility.
- `-I lanplus` specifies to use an encrypted IPMI v2.0 RMCP+ LAN Interface for the session.
- `-H IP_address` indicates the IP address you have configured for Lights-Out Management on the appliance you want to access.
- `-U user_name` is the name of an authorized remote session user.
- *`command`* is the name of the command you want to use.

Note Cisco recommends using IPMItool version 1.8.12 or greater.

The same command for IPMIutil on Windows looks like this:

```
ipmiutil command -V 4 -J 3 -N IP_address -Uuser_name
```

This command connects you to the command line on the appliance where you can log in as if you are physically present near the appliance. You may be prompted to enter a password.

Configuring Serial Over LAN with IPMItool

You must be an Admin user with LOM access to perform this procedure.

Procedure

Using IPMItool, enter the following command, and a password if prompted:

```
ipmitool -I lanplus -H IP_address -U user_name sol activate
```

Configuring Serial Over LAN with IPMIutil

You must be an Admin user with LOM access to perform this procedure.

Procedure

Using IPMIutil, enter the following command, and a password if prompted:

```
ipmiutil -J 3 -N IP_address -U username sol -a
```

Lights-Out Management Overview

Lights-Out Management (LOM) provides the ability to perform a limited set of actions over an SOL connection on the default (`eth0`) management interface without the need to log into the system. You use the command to create a SOL connection followed by one of the LOM commands. After the command is completed, the connection ends.

Caution In rare cases, if your computer is on a different subnet than the system's management interface and the system is configured for DHCP, attempting to access LOM features can fail. If this occurs, you can either disable and then re-enable LOM on the system, or use a computer on the same subnet as the system to ping its management interface. You should then be able to use LOM.

Caution Cisco is aware of a vulnerability inherent in the Intelligent Platform Management Interface (IPMI) standard (CVE-2013-4786). Enabling Lights-Out Management (LOM) on a system exposes this vulnerability. To mitigate this vulnerability, deploy your systems on a secure management network accessible only to trusted users and use a complex, non-dictionary-based password of the maximum supported length for your system and change it every three months. To prevent exposure to this vulnerability, do not enable LOM.

If all attempts to access your system have failed, you can use LOM to restart your system remotely. Note that if a system is restarted while the SOL connection is active, the LOM session may disconnect or time out.

Caution Do **not** restart your system unless it does not respond to any other attempts to restart. Remotely restarting does not gracefully reboot the system and you may lose data.

Table 3: Lights-Out Management Commands

IPMItool	IPMIutil	Description
(not applicable)	`-V 4`	Enables admin privileges for the IPMI session
`-I lanplus`	`-J 3`	Enables encryption for the IPMI session
`-H hostname/IP address`	`-N nodename/IP address`	Indicates the LOM IP address or hostname for the management center
`-U`	`-U`	Indicates the username of an authorized LOM acc
`sol activate`	`sol -a`	Starts the SOL session
`sol deactivate`	`sol -d`	Ends the SOL session
`chassis power cycle`	`power -c`	Restarts the appliance

IPMItool	IPMIutil	Description
`chassis power on`	`power -u`	Powers up the appliance
`chassis power off`	`power -d`	Powers down the appliance
`sdr`	`sensor`	Displays appliance information, such as fan speeds an temperatures

For example, to display a list of appliance information, the IPMItool command is:

```
ipmitool -I lanplus -H IP_address -U user_name sdr
```

Note Cisco recommends using IPMItool version 1.8.12 or greater.

The same command with the IPMIutil utility is:

```
ipmiutil sensor -V 4 -J 3 -N IP_address -U user_name
```

Configuring Lights-Out Management with IPMItool

You must be an Admin user with LOM access to perform this procedure.

Procedure

Enter the following command for IPMItool and a password if prompted:

```
ipmitool -I lanplus -H IP_address -U user_name command
```

Configuring Lights-Out Management with IPMIutil

You must be an Admin user with LOM access to perform this procedure.

Procedure

Enter the following command for IPMIutil and a password if prompted:

```
ipmiutil -J 3 -N IP_address -U username command
```

Remote Storage Device

You can store management center backups and reports locally, or mount one of the following systems:

- Network File System (NFS)
- Server Message Block (SMB)/Common Internet File System (CIFS)
- Secure Shell Filesystem (SSHFS)

You cannot send backups to one remote system and reports to another, but you can choose to send either to a remote system and store the other on the 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. See Database, on page 58.

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**.

Configure NFS for Remote Storage

The management center can store backups and reports on an NFS mount.

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**.

You can specify the version number of the NFS storage using the following format:

vers=*version*

For example, to select NFSv4, enter:

vers=4.0

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**.

Troubleshooting

When there is a random latency in the NFS connection with the firewall device, perform the following activities, and then contact Cisco TAC for troubleshooting:

- 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, see Troubleshoot Firepower File Generation Procedures.
- Collect the incoming and exiting traffic PCAP records. For information on the procedure, see Packet Capture Overview, on page 441.
- 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, see show snort counters.

Configure SMB for Remote Storage

The management center can store backups and reports on an SMB mount.

Before you begin

Ensure that your external remote storage system is functional and accessible from your 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 management center has ownership of and read/change access to the share location.

- To ensure security, you should install SMB 2.0 or greater.

Procedure

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

Step 2 Click **Remote Storage Device**.

Step 3 Choose **SMB** from the **Storage Type** drop-down list.

Step 4 Add the connection information:

- Enter the IPv4 address or hostname of the storage system in the **Host** field.
- Enter the share of your storage area in the **Share** field.
- Optionally, enter the domain name for the remote storage system in the **Domain** field.
- Enter the user name for the storage system in the **Username** field and the password for that user in the **Password** field.

Step 5 Optionally, check the **Use Advanced Options** check box and enter any required mount options in **Command Line Options**.

You can specify the version number of the SMB storage using the following 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

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**.

Configure SSH for Remote Storage

The management center can store backups and reports on an SSHFS mount.

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**.

SNMP

You can enable Simple Network Management Protocol (SNMP) polling. This feature supports use of versions 1, 2, and 3 of the SNMP protocol. This feature allows access to the standard management information base (MIB), which includes system details such as contact, administrative, location, service information, IP addressing and routing information, and transmission protocol usage statistics.

Note When selecting SNMP versions for the SNMP protocol, note that SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

Enabling SNMP polling does not cause the system to send SNMP traps; it only makes the information in the MIBs available for polling by your network management system.

Configure SNMP Polling

Before you begin

Add SNMP access for each computer you plan to use to poll the system. See Configure an Access List, on page 47.

Note The SNMP MIB contains information that could be used to attack your deployment. We recommend that you restrict your access list for SNMP access to the specific hosts that will be used to poll for the MIB. We also recommend you use SNMPv3 and use strong passwords for network management access.

Procedure

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

Step 2 Click **SNMP**.

Step 3 From the **SNMP Version** drop-down list, choose the SNMP version you want to use:

- **Version 1** or **Version 2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

  Note
  Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- **Version 3**: Click **Add User** to display the user definition page. SNMPv3 only supports read-only users and encryption with AES128.

Step 4 Enter a **Username**.

Step 5 Choose the protocol you want to use for authentication from the **Authentication Protocol** drop-down list.

Step 6 Enter the password required for authentication with the SNMP server in the **Authentication Password** field.

Step 7 Re-enter the authentication password in the **Verify Password** field.

Step 8 Choose the privacy protocol you want to use from the **Privacy Protocol** list, or choose **None** to not use a privacy protocol.

Step 9 Enter the SNMP privacy key required by the SNMP server in the **Privacy Password** field.

Step 10 Re-enter the privacy password in the **Verify Password** field.

Step 11 Click **Add**.

Step 12 Click **Save**.

Session Timeout

Unattended login sessions may be security risks. You can configure the amount of idle time before a user's login session times out due to inactivity.

Note that you can exempt specific web interface users from timeout, for scenarios where you plan to passively, securely monitor the system for long periods of time. Users with the Administrator role, whose complete access to menu options poses an extra risk if compromised, cannot be made exempt from session timeouts.

Configure Session Timeouts

Procedure

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

Step 2 Click **CLI Timeout**.

Step 3 Configure session timeouts:

- Web interface (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, see Add or Edit an Internal User, on page 125.

- CLI: Configure the **CLI Timeout (Minutes)** field. The default value is `0`; the maximum value is `1440` (24 hours).

Step 4 Click **Save**.

Time

Time settings are displayed on most pages in local time using the time zone you set on the Time Zone page in User Preferences (the default is America/New York), but are stored on the appliance using UTC time.

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

Procedure

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

Step 2 Click **Time**.

The current time is displayed using the time zone specified for your account in User Preferences.

If your appliance uses an NTP server: For information about the table entries, see NTP Server Status, on page 103.

NTP Server Status

If you are synchronizing time from an NTP server, you can view connection status on the **Time** page (choose **System > Configuration**).

Table 4: NTP Status

Column	Description
NTP Server	The IP address or name of the configured NTP server.
Status	The status of the NTP server time synchronization: • **Being Used** indicates that the appliance is synchronized with the NTP server. • **Available** indicates that the NTP server is available for use, but time is not yet synchronized. • **Not Available** indicates that the NTP server is in your configuration, but the NTP daemon is unable to use it. • **Pending** indicates that the NTP server is new or the NTP daemon was recently restarted. Over time, its value should change to **Being Used**, **Available**, or **Not Available**. • **Unknown** indicates that the status of the NTP server is unknown.
Authentication	The authentication status for communication between the 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

Synchronizing the system time on your Secure Firewall Management Center (management center) and its managed devices is essential to successful operation of your system. We recommend that you specify NTP servers during management center initial configuration, but you can use the information in this section to establish or change time sychronization settings after initial configuration is complete.

Use a Network Time Protocol (NTP) server to synchronize system time on the management center and all devices. The management center supports secure communications with NTP servers using MD5, SHA-1, or AES-128 CMAC symmetric key authentication; for system security, we recommend using this feature.

The management center can also be configured to connect solely with authenticated NTP servers; using this option improves security in a mixed-authentication environment, or when you migrate your system to different NTP servers. It is redundant to use this setting in an environment where all reachable NTP servers are authenticated.

Note If you specified an NTP server for the 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 management center and managed devices.

To synchronize time on management center and managed devices, see:

- Recommended: Synchronize Time on the Management Center with an NTP Server, on page 104

  This topic provides instructions for configuring your management center to synchronize with an NTP server or servers and includes links to instructions on configuring managed devices to synchronize with the same NTP server or servers.

- Otherwise: Synchronize Time Without Access to a Network NTP Server, on page 106

  This topic provides instructions for setting the time on your management center, configuring your management center to serve as an NTP server, and links to instructions on configuring managed devices to synchronize with the management center NTP server.

Synchronize Time on the Management Center with an NTP Server

Time synchronization among all of the components of your system is critically important.

The best way to ensure proper time synchronization between management center and all managed devices is to use an NTP server on your network.

The management center supports NTPv4.

You must have Admin or Network Admin privileges to do this procedure.

Before you begin

Note the following:

- If your management center and managed devices cannot access a network NTP server, do not use this procedure. Instead, see Synchronize Time Without Access to a Network NTP Server, on page 106.
- 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, see the documentation for your hardware model.

Caution If the 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.

Procedure

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

Step 2 Click **Time Synchronization**.

Step 3 If **Serve Time via NTP** is **Enabled**, choose **Disabled** to disable the management center as an NTP server.

Step 4 For the **Set My Clock** option, choose **Via NTP**.

Step 5 Click **Add**.

Step 6 In the **Add NTP Server** dialog box, enter the host name or IPv4 or IPv6 address of an NTP server.

Step 7 (Optional) To secure communication between your management center and the NTP server:

a) Select **MD5**, **SHA-1** or **AES-128 CMAC** from the **Key Type** drop-down list.

b) Enter an the corresponding MD5, SHA-1, or AES-128 CMAC **Key Number** and **Key Value** from the specified NTP server.

Step 8 Click **Add**.

Step 9 When only two NTP servers are configured, the offset difference between them becomes high. This results in the management center using the Local Time. Hence, we recommend that you configure atleast three NTP servers.

To add more NTP servers, repeat Steps 5 through 8.

Step 10 (Optional) To force the management center to use only an NTP server that successfully authenticates, check the **Use the authenticated NTP server only** check box.

Step 11 Click **Save**.

What to do next

Set managed devices to synchronize with the same NTP server or servers:

- Configure device platform settings: *Configure NTP Time Synchronization for Threat Defense* in the Cisco Secure Firewall Management Center Device Configuration Guide.

  Note that even if you force the 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

If your devices cannot directly reach the network NTP server, or your organization does not have a network NTP server, a physical-hardware management center can serve as an NTP server.

Important

- Do not use this procedure unless you have no other NTP server. Instead, use the procedure in Synchronize Time on the Management Center with an NTP Server, on page 104.
- Do not use a virtual management center as an NTP server.

To change the time manually **after** configuring the management center as an NTP server, you must disable the NTP option, change the time manually, and then re-enable the NTP option.

Procedure

Step 1 Manually set the system time on the 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 two 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 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 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 instructions:

For threat defense devices, see *Configure NTP Time Synchronization for Threat Defense* in the Cisco Secure Firewall Management Center Device Configuration Guide.

About Changing Time Synchronization Settings

- Your management center and its managed devices are heavily dependent on accurate time. The system clock is a system facility that maintains the time of the system. The system clock is set to Universal Coordinated Time (UTC), which is the primary time standard by which the world regulates clocks and time.

  DO NOT ATTEMPT TO CHANGE THE SYSTEM TIME. Changing the system time zone from UTC is NOT supported, and doing so will require you to reimage the device to recover from an unsupported state.

- If you configure the 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 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 management center as an NTP server, you must disable the NTP option, change the time manually, and then re-enable the NTP option.

UCAPL/CC Compliance

Your organization might be required to use only equipment and software complying with security standards established by the U.S. Department of Defense and global certification organizations. For more information about this setting, see Security Certifications Compliance Modes, on page 329.

User Configuration

Global User Configuration settings affect all users on the 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 applies to CLI access as well; the system maintains separate password lists for each form of access. Setting the limit to zero (the default) places no restrictions on password reuse. See Set Password Reuse Limit, on page 108.

- **Track Successful Logins**: The number of days that the system tracks successful logins to the 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. See Track Successful Logins, on page 109.
- **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 login attempts while the temporary lockout is in force:
  - 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 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. 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), the number of concurrent sessions is unlimited.

If you change the concurrent session limit to a value more restrictive, the system will not close any currently open sessions; it will, however, prevent new sessions beyond the number specified from being opened.

Set Password Reuse Limit

If you enable the **Password Reuse Limit**, the system keeps encrypted *password histories* for management center users. Users cannot reuse passwords in their histories. You can specify the number of stored passwords for each user, per access method (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.

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**.

Track Successful Logins

Use this procedure to enable tracking successful logins for each user for a specified number of days. When this tracking is enabled, the system displays the successful login count when users log into the web interface or the CLI.

Note If you lower the number of days, the system deletes records of older logins. If you then increase the limit, the system does not restore the count from those days. In that case, the reported number of successful logins may be temporarily lower than the actual number.

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**.

Enabling Temporary Lockouts

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.

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**.

Set Maximum Number of Concurrent Sessions

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.

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 change the concurrent session limit to a value more restrictive, the system will not close any currently open sessions; it will, however, prevent new sessions beyond the number specified from being opened.

Step 4 Click **Save**.

VMware Tools

VMware Tools is a suite of performance-enhancing utilities intended for virtual machines. These utilities allow you to make full use of the convenient features of VMware products. Secure Firewall virtual appliances running on VMware support the following plugins:

- guestInfo
- powerOps
- timeSync
- vmbackup

You can also enable VMware Tools on all supported versions of ESXi. For information on the full functionality of VMware Tools, see the VMware website (http://www.vmware.com/).

Enabling 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**.

Vulnerability Mapping

The system automatically maps vulnerabilities to a host IP address for any application protocol traffic received or sent from that address, when the server has an application ID in the discovery event database and the packet header for the traffic includes a vendor and version.

For any servers which do not include vendor or version information in their packets, you can configure whether the system associates vulnerabilities with server traffic for these vendor and versionless servers.

For example, a host serves SMTP traffic that does not have a vendor or version in the header. If you enable the SMTP server on the Vulnerability Mapping page of a system configuration, then save that configuration to the management center managing the device that detects the traffic, all vulnerabilities associated with SMTP servers are added to the host profile for the host.

Although detectors collect server information and add it to host profiles, the application protocol detectors will not be used for vulnerability mapping, because you cannot specify a vendor or version for a custom application protocol detector and cannot select the server for vulnerability mapping.

Mapping Vulnerabilities for Servers

This procedure requires any Smart License or the Protection classic license.

Procedure

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

Step 2 Choose **Vulnerability Mapping**.

Step 3 You have the following choices:

- To prevent vulnerabilities for a server from being mapped to hosts that receive application protocol traffic without vendor or version information, clear the check box for that server.
- To cause vulnerabilities for a server to be mapped to hosts that receive application protocol traffic without vendor or version information, check the check box for that server.

Tip
You can check or clear all check boxes at once using the check box next to **Enabled**.

Step 4 Click **Save**.

Web Analytics

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 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 the following procedure.

Procedure

Step 1 Choose **System > Configuration**.

Step 2 Click **Web Analytics**.

Step 3 Make your choice and click **Save**.

What to do next

(Optional) Determine whether to share data via the Configure Cisco Success Network Enrollment.

History for System Configuration

Feature	Minimum Management Center	Minimum Threat Defense	Details
Access control performance improvements (object optimization).	7.2.4 7.4.0	Any	**Upgrade impact. First deployment after 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 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**. Version restrictions: Not supported with management center Version 7.3.x.
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. We recommend you use SAN if the certificate secures multiple domain names or IP addresses. For more information about SAN, see RFC 5280, section 4.2.1.6. New/modified screens: **System**(⚙) > **Configuration** > **HTTPS Certificate**

Feature	Minimum Management Center	Minimum 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 lifetime varies depending on the Firepower version being used when the certificate was generated. See Default HTTPS Server Certificates, on page 64 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. See Web Analytics, on page 112.
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 Management Center	Minimum Threat Defense	Details
Ability to disable Duplicate Address Detection (DAD) on management interfaces.	6.4	Any	When you enable IPv6, you can disable 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. 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. 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 device management interfaces for testing purposes. 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 can track the number of successful logins each FMC account has performed within a selected number of days. When this feature is enabled, on log in users see a message reporting how many times they have successfully logged in to the system in the past configured number of days. (Applies to web interface as well as 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

Feature	Minimum Management Center	Minimum Threat Defense	Details
HTTPS Certificates.	6.3	Any	The default HTTPS server certificate provided with the system now expires in three years. If your appliance uses a default server certificate that was generated before you upgraded to Version 6.3, the server certificate will expire 20 years from when it was first generated. If you are using the default HTTPS server certificate the system now provides the ability to renew it. New/modified screens: **System**(⚙) > **Configuration** > **HTTPS Certificate** > **Renew HTTPS Certificate** Supported platforms: FMC
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

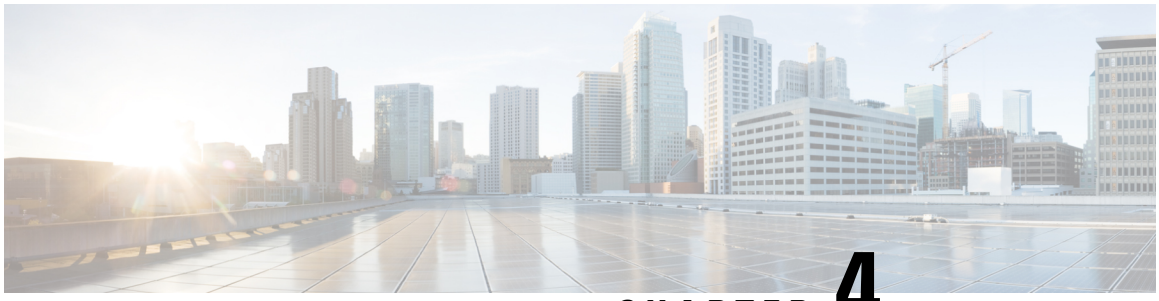

CHAPTER 4

Users

The management center includes default **admin** accounts for web and CLI access. This chapter discusses how to create custom user accounts. See Logging into the Management Center, on page 27 for detailed information about logging into the management center with a user account.

About Users

You can add custom user accounts on managed devices, either as internal users or as external users on a LDAP or RADIUS server. Each managed device maintains separate user accounts. For example, when you add a user to the management center, that user only has access to the 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 types of users:

- 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

The management center has a web interface, CLI (accessible from the console (either the serial port or the keyboard and monitor) or using SSH to the management interface), and Linux shell. For detailed information about the management UIs, see System User Interfaces, on page 29.

See the following information about management center user types, and which UI they can access:

- admin user—The management center supports two different internal **admin** users: one for the web interface, and another with CLI access. The system initialization process synchronizes the passwords for these two **admin** accounts so they start out the same, but they are tracked by different internal mechanisms and may diverge after initial configuration. See 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 management center CLI command **configure password**.)
- Internal users—Internal users added in the web interface have web interface access only.
- External users—External users have web interface access, and you can optionally configure CLI access.
- SSO users—SSO users have web interface access only.

⚠

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 explicit instructions in the 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; only use the procedures in this chapter.

User Roles

CLI User Role

CLI external users on the 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 as though you are going to make a custom role based on the predefined role. You can then see all of the privileges assigned.

Figure 17: View User Role Privileges

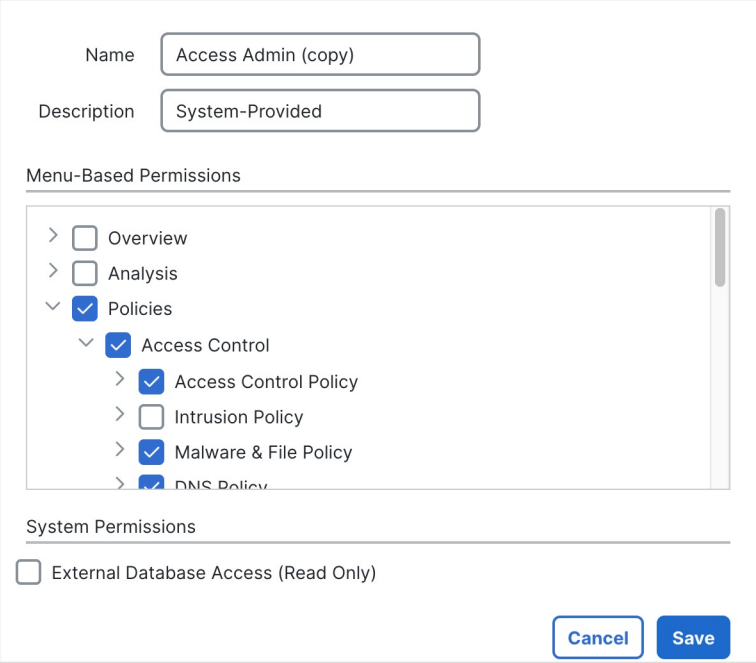

The management center includes the following predefined user roles:

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, so you cannot make them exempt from login session timeouts.

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

Discovery Admin

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

External Database User (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. Because this role's function does not involve the web interface, access is provided only for ease of support and password changes.

Intrusion Admin

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

Maintenance User

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

Network Admin

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

Security Analyst

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

Security Analyst (Read Only)

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

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

The following rules apply to passwords for internal user accounts on the management center, with Lights-Out Management (LOM) enabled or disabled. Different password requirements apply for externally authenticated accounts or in systems with security certifications compliance enabled. See Configure External Authentication for the Management Center, on page 127 and Security Certifications Compliance, on page 329 for more information.

During management center initial configuration, the system requires the **admin** user to set the account password to comply with strong password requirements described in the table below. For physical management centers, the strong password requirements with LOM enabled are used, and for virtual management centers, the strong password requirements with LOM not enabled are used. At this time the system synchronizes the passwords for the web interface **admin** and the CLI access **admin**. After initial configuration, the web interface **admin** can remove the strong password requirement, but the CLI access **admin** must always comply with strong password requirements with LOM not enabled.

	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 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 system checks passwords against a special dictionary containing not only many English dictionary words, but also 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 fourteen characters rather than twenty.) • 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 vary between different series of physical management centers. We recommend restricting your choice of special characters to those listed in the final bullet above. Do not include the user name in the password. The system checks passwords against a special dictionary containing not only many English dictionary words, but also other character strings that could 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. (See Add or Edit an Internal User, on page 125 for more information.)	Passwords must include:
		• Between eight and twenty characters (On MC 1000, MC 2500, and MC 4500 the upper limit is fourteen characters rather than twenty.)
		• Characters from at least three of the following four categories:
		  • Uppercase letters
		  • Lowercase letters
		  • Digits
		  • Special characters such as ! @ # * - _ +
		The rules for special characters vary between different series of physical management centers. We recommend restricting your choice of special characters to those listed in the final bullet above.
		Do not include the user name in the password.

Users and Domains

In a multidomain deployment, you can create user accounts in any domain in which you have been assigned Administrator access.

User Roles in Domains

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:

Figure 18: User Roles Per Domain

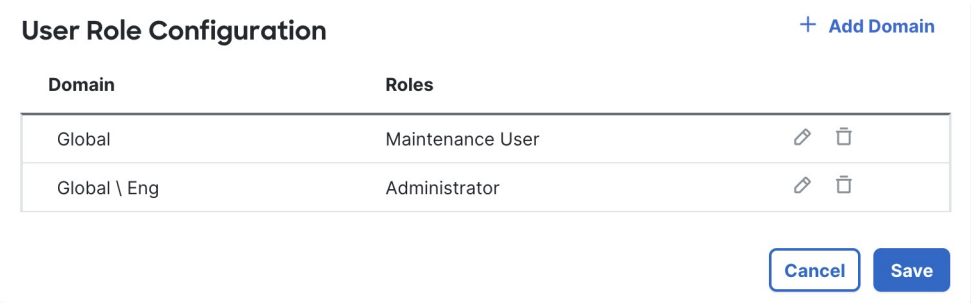

User Management

Users are only visible in the domain in which they are created.

If you add a user in the current domain but assign a user role in a subdomain, then that user will only show on the current domain's **Users** page, even though the user has a role in a subdomain. For example, from the Global domain, you add user **leaf** and assign a role for **Leaf1**, but because you were in the Global domain when you added the user, you can see it from the Global domain:

Figure 19: Global Domain Users

If you change domains to **Leaf1**, you cannot see the user **leaf**, but you can see the user **test**, which was added directly from the Leaf1 subdomain:

Figure 20: Subdomain Users

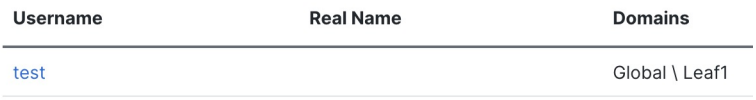

Logging In

Users added from the Global domain will log in with just their username, even if their roles are only in a subdomain. In this case, **leaf** only has a user role in the **Leaf1** subdomain, but because it was added from Global, do not include the subdomain in the login:

Figure 21: User Login When Added in Global

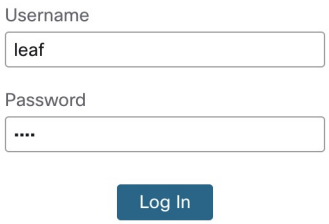

Users added directly in a subdomain need to log into the management center with the subdomain(s) as part of the login name, depending on which domain their user was added from: *subdomain1\subdomain2\username*. You do not need to enter the **Global** parent domain. For example, **test** was added from the **Leaf1** subdomain, so you need to include Leaf1 in the login name:

Figure 22: Global Domain User Login

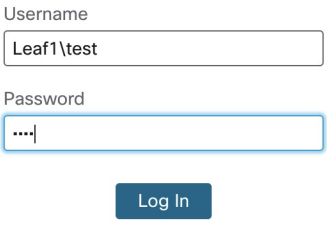

When you log in, you are placed in the domain where your username was added. For example, the admin user defaults to the Global domain:

Figure 23: User Domain at Login

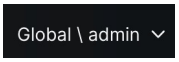

However, after login, you can change to a subdomain by clicking the down arrow:

Figure 24: Change to a Subdomain

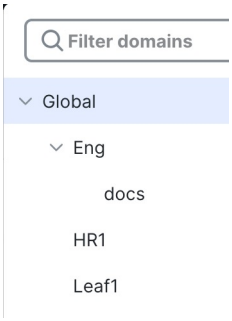

Guidelines and Limitations for User Accounts for Management Center

- The 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**; the system forces you to change this during the initialization process. See the *Getting Started Guide* for your model for more information about system initialization.
- By default, the following settings apply to all user accounts on the 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.
  - There are no user-defined limits on the number of read-only and read/write sessions that can be open at the same time.

  You can change these settings for all users as a system configuration. (**System** (⚙) > **Configuration** > **User Configuration**) See User Configuration, on page 107.
- Ensure that you follow the principles of least privilege when assigning default access roles to users at initial setup. When a user first logs in to the system with their credentials, their account will be assigned this default access role. We recommend that the default access role be the lowest possible privilege required for anyone to log in to the system. For example, common users can be given the Security Analyst (Read-Only) role as the default access role, and administrators can be added to a separate administrator's group to give them full administrator rights. 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. Note that this guideline applies to all users - internal, external, or CAC users.

If a user who has logged in with the default access role needs a temporary elevation of their privileges, a user with administrative privileges can temporarily provide that user the required higher level of access by assigning them a role with higher privilege. This privilege will be revoked after 24 hours of inactivity, and the user will return to their default access role.

If a user needs a permanent access role reassignment to a higher privilege level, such as System Admin, use the Group Controlled Access Roles method to provide admin access to the user. 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 on configuring Group Controlled Access Roles, see the Add an LDAP External Authentication Object for Management Center section.

Requirements and Prerequisites for User Accounts for Management Center

Model Support

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 143 also supports the Network Admin role.

Add or Edit an Internal User

This procedure describes how to add custom internal user accounts for the management center.

The **System** > **Users** > **Users** shows both internal users that you added manually and external users that are added automatically when a user logged 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, see Security Certifications Compliance, on page 329.

Note Avoid having multiple Admin users simultaneously creating new users on the 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 the following 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.

Step 4 **Real Name**: Enter descriptive information to identify the user or department to whom the account belongs.

Step 5 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.

Step 6 Enter values in the **Password** and **Confirm Password** fields.

The values must conform to the password options you set for this user.

Step 7 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.

Step 8 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.

Step 9 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.

Step 10 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.

Step 11 Set the following **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 120.
- **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.

Step 12 In the **User Role Configuration** area, assign the user roles. For more information about user roles, see Customize User Roles for the Web Interface, on page 197.

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 see 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, see Users and Domains, on page 122. See the following steps:
  - **a.** Click **Add Domain**.
  - **b.** Choose a domain from the **Domain** drop-down list.
  - **c.** Check the user roles that you want to assign the user.
  - **d.** Click **Save**.

Step 13 (Optional, for physical 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. See Lights-Out Management Overview, on page 95 for more information about Lights-Out Management.

Step 14 Click **Save**.

Configure External Authentication for the Management Center

To enable external authentication, you need to add one or more external authentication objects.

About External Authentication for the Management Center

When you enable external authentication, the management center verifies the user credentials with an LDAP or RADIUS server as specified in an *external authentication object*.

You can configure multiple external authentication objects for web interface access. For example, if you have 5 external authentication objects, users from any of them can be authenticated to access the web interface. You can use only one external authentication object for CLI access. If you have more than one external authentication object enabled, then users can authenticate using only the first object in the list.

External authentication objects can be used by the management center and threat defense devices. You can share the same object between the different appliance/device types, or create separate objects.

Note The timeout range is different for the threat defense and the management center, so if you share an object, be sure not to exceed the threat defense's smaller timeout range (1-30 seconds for LDAP, and 1-300 seconds for RADIUS). If you set the timeout to a higher value, the threat defense external authentication configuration will not work.

For the management center, enable the external authentication objects directly on the **System** > **Users** > **External Authentication** tab; this setting only affects management center usage, and it does not need to be enabled on this tab for managed device usage. For threat defense devices, you must enable the external authentication object in the platform settings that you deploy to the devices.

Web interface users are defined separately from CLI users in the external authentication object. For CLI users on RADIUS, you must pre-configure the list of RADIUS usernames in the external authentication object. For LDAP, you can specify a filter to match CLI users on the LDAP server.

You cannot use an LDAP object for CLI access that 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, which can present a security risk. Make sure that you:

- Restrict the list of users with CLI or Linux shell access.
- Do not create Linux shell users.

About LDAP

The Lightweight Directory Access Protocol (LDAP) 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.

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

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

About RADIUS

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.

Add an LDAP External Authentication Object for the 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 and not a hostname for the LDAP server on this procedure, the LDAP server may return a URI for authentication that can include a hostname. A DNS lookup is required to resolve the hostname. See Modify Management Center Management Interfaces, on page 81 to add DNS servers.
- If you are configuring an LDAP authentication object for use with CAC authentication, do not remove the CAC inserted in your computer. You must have a CAC inserted at all times after enabling user certificates.

Procedure

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

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

Step 3 Click **Add icon** (+) **Add External Authentication Object**.

Step 4 Set the **Authentication Method** to **LDAP**.

Step 5 (Optional) 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 143 to fully configure CAC authentication and authorization. You cannot use this object for CLI users.

Step 6 Enter a **Name** and optional **Description**.

Step 7 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 8 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.

Step 9 (Optional) Change the **Port** from the default.

Step 10 (Optional) Enter the **Backup Server** parameters.

Step 11 Enter **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 the following 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. If you choose SSL encryption, the port resets to 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 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 man-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 threat defense and the management center, so if you share an object, be sure not to exceed the threat defense's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the threat defense LDAP configuration will not work.

Step 12 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 13 (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 is used exactly as it is configured on the LDAP server. For this reason, the device limits the number of recursions of a search to 4 to prevent search syntax errors from causing infinite loops.

Example:

Enter the following in the **Administrator** field to authenticate names in the information technology organization at the Example company:

```
cn=itgroup,ou=groups, dc=example,dc=com
```

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 14 (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 the following 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 can present 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 management center user should be **admin**; do not include an **admin** user in the **CLI Access Filter**.

Step 15 (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, see Troubleshooting LDAP Authentication Connections, on page 202.

Step 16 (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**.

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.

Step 17 Click **Save**.

Step 18 Enable use of this server. See Enable External Authentication for Users on the Management Center, on page 142.

Examples

Basic Example

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

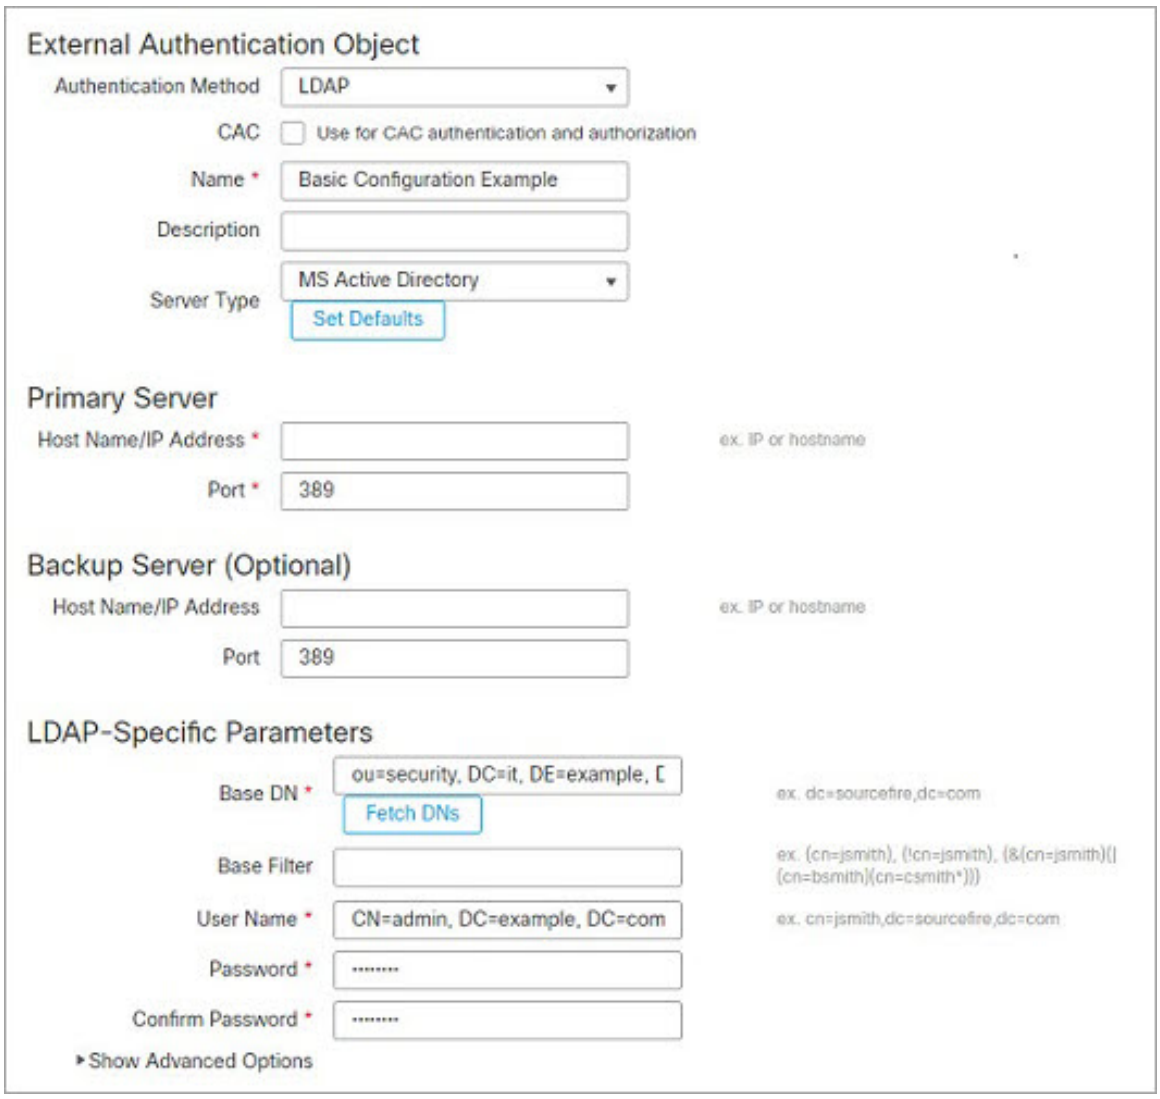

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company.

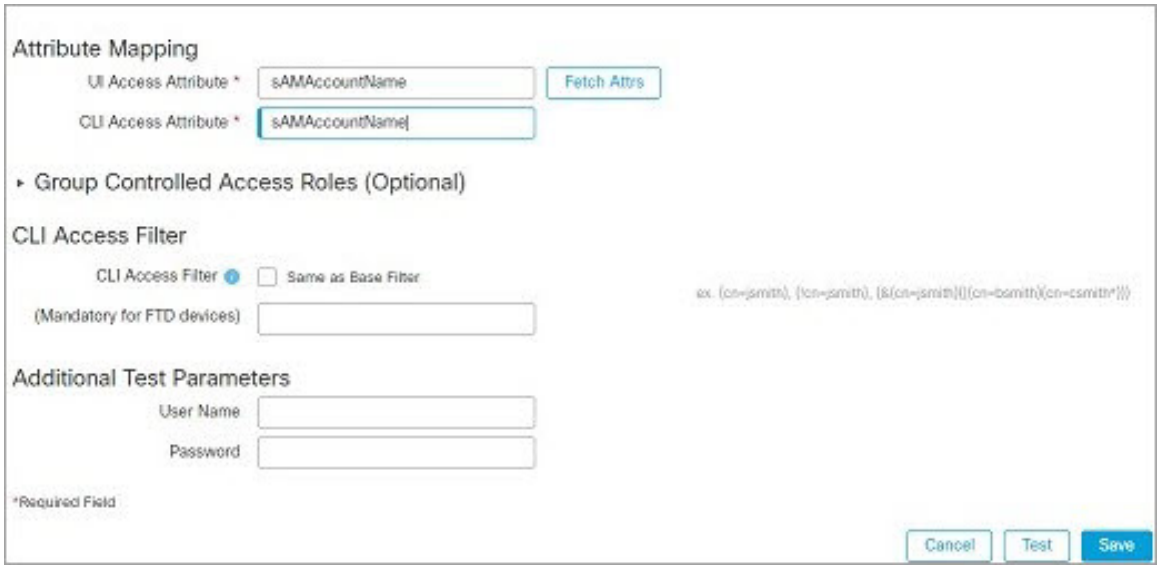

However, because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute. Choosing the MS Active Directory server type and clicking **Set Defaults** sets the UI Access Attribute to `sAMAccountName`. As a result, the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

In addition, a **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.

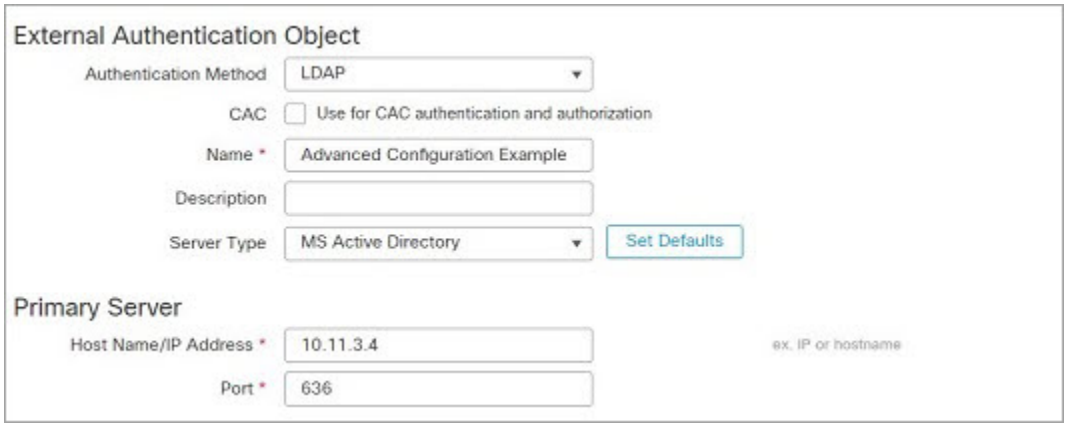

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company. However, note that this server has a base filter of `(cn=*smith)`. The filter restricts the users retrieved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

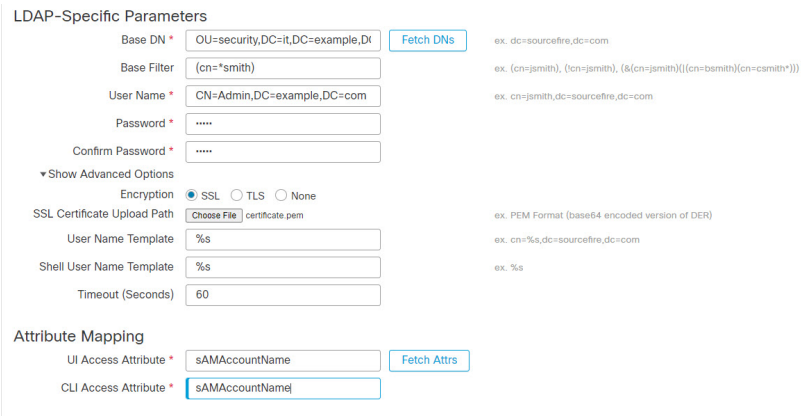

The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout (Seconds)** setting.

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`.

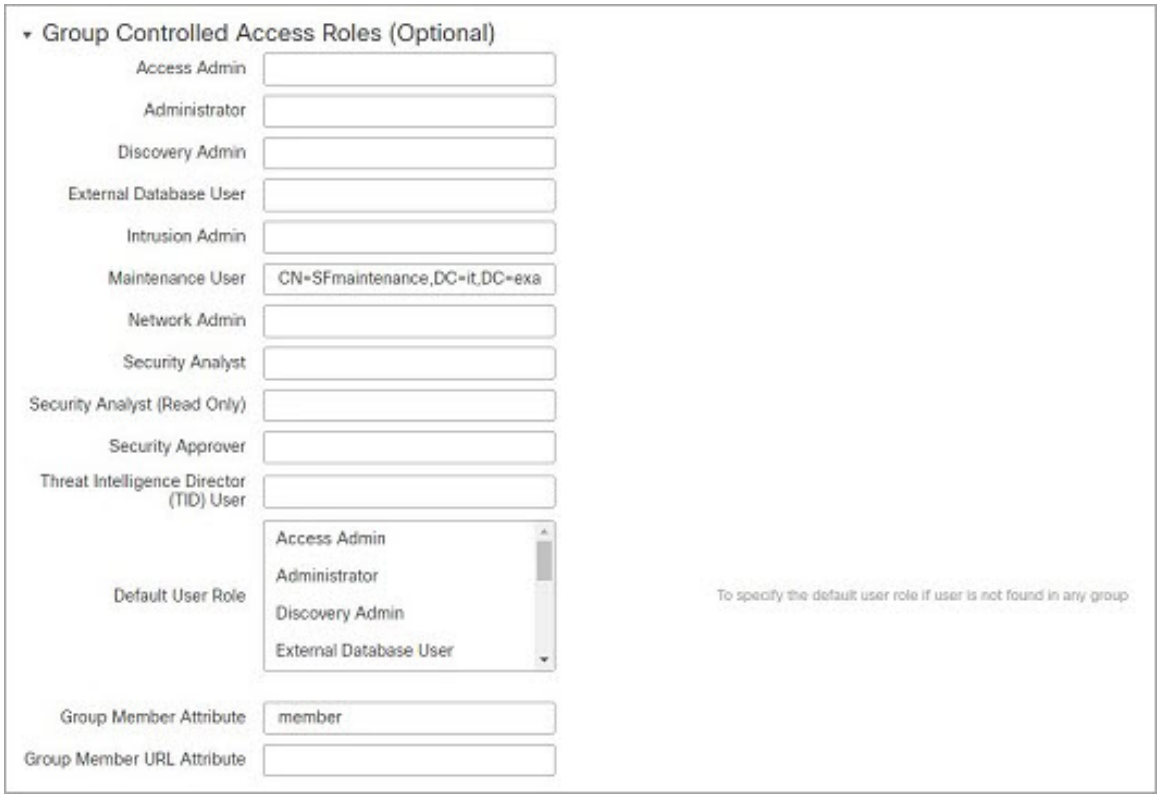

The **CLI Access Filter** is set to be the same as the base filter, so the same users can access the appliance through the CLI as through the web interface.

Add a RADIUS External Authentication Object for 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.

Procedure

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

Step 2 Click **External Authentication**.

Step 3 Click **Add icon** (+) **Add External Authentication Object**.

Step 4 Set the **Authentication Method** to **RADIUS**.

Step 5 Enter a **Name** and optional **Description**.

Step 6 Check the **RADIUS Server-Enabled Message Authenticator** check box to require the Message-Authenticator attribute in all RADIUS responses, ensuring that every response from the RADIUS server is securely verified by the Threat Defense.

This feature is enabled by default for new RADIUS servers. We recommend you enable it for existing servers after the upgrade. Disabling message authenticators may expose your firewalls to potential attacks. Ensure that your RADIUS server has the Message-Authenticator configuration.

Step 7 For the **Primary Server**, enter a **Host Name/IP Address**.

Step 8 (Optional) Change the **Port** from the default.

Step 9 Enter the **RADIUS Secret Key**.

Step 10 (Optional) Enter the **Backup Server** parameters.

Step 11 (Optional) Enter **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 threat defense and the management center, so if you share an object, be sure not to exceed the threat defense's smaller timeout range (1-300 seconds). If you set the timeout to a higher value, the 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.

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 12 (Optional) **Define Custom RADIUS Attributes**.

If your RADIUS server returns values for attributes not included in the `dictionary` file in `/etc/radiusclient/`, and you plan to use those attributes to set roles for users with those attributes, you need to define those attributes. You can locate the attributes returned for a user by looking at the user's profile on your RADIUS server.

a) Enter an **Attribute Name**.

When you define an attribute, you provide the name of the attribute, which consists of alphanumeric characters. Note that words in an attribute name should be separated by dashes rather than spaces.

b) Enter the **Attribute ID** as an integer.

The attribute ID should be an integer and should not conflict with any existing attribute 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 13 (Optional) In the **CLI Access Filter** area **Administrator CLI Access User List** field, enter the user names that should have CLI access, separated by commas.

Make sure that these usernames match usernames on the RADIUS server. The names must be Linux-valid usernames:

- 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 (/)

To prevent RADIUS authentication of CLI access, leave the field blank.

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 management center, the only internal CLI user is **admin**, so do not also create an **admin** external user.

Step 14 (Optional) Click **Test** to test management center connectivity to the RADIUS server.

Step 15 (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**.

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.

Step 16 Click **Save**.

Step 17 Enable use of this server. See Enable External Authentication for Users on the Management Center, on page 142.

Examples

Simple User Role Assignments

The following 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.

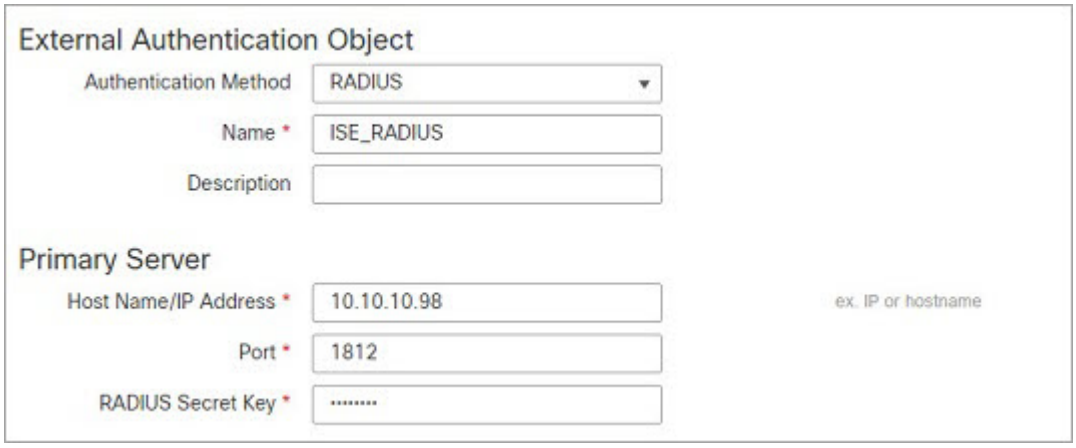

The following example shows RADIUS-specific parameters, including the timeout (30 seconds) and number of failed retries before the Secure Firewall System attempts to contact the backup server, if any.

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.

The following graphic depicts the role configuration for the example:

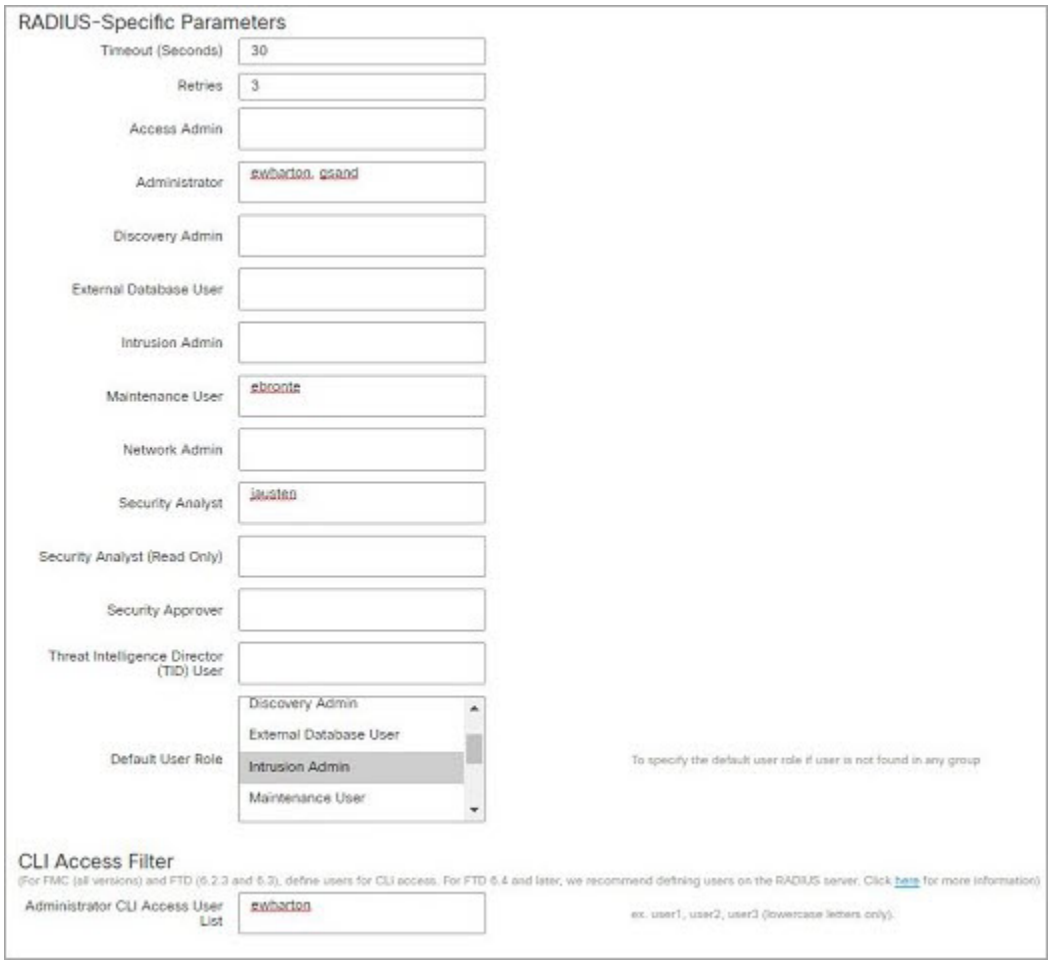

Roles for Users Matching an Attribute-Value Pair

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.

The following 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 v. 5.00 remote access server should receive the Security Analyst (Read Only) role, so you enter the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.

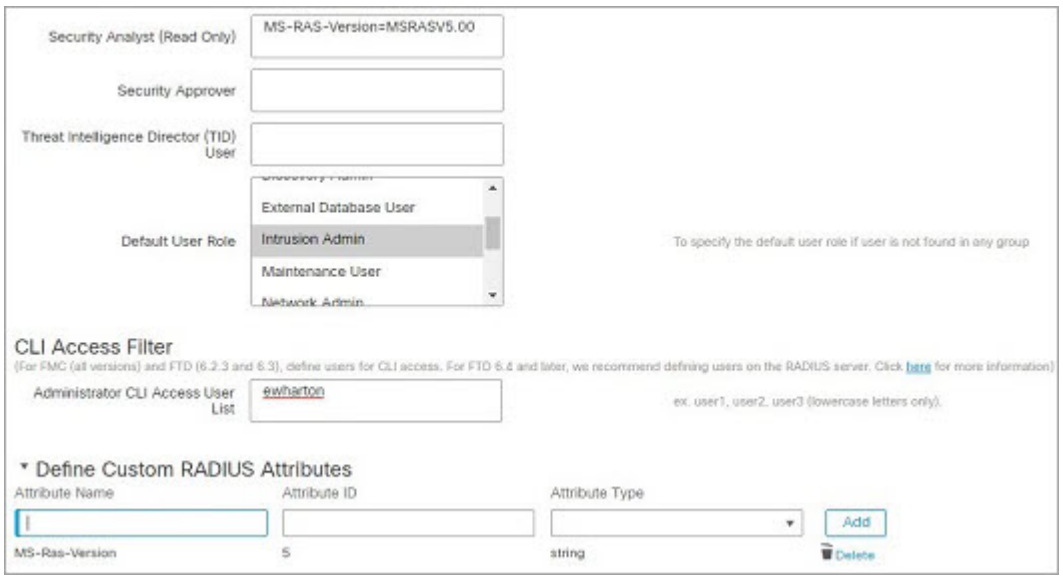

Enable External Authentication for Users on the Management Center

When you enable external authentication for management users, the 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 Management Center, on page 129 and Add a RADIUS External Authentication Object for Management Center, on page 137.

Procedure

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

Step 2 Click **External Authentication**.

Step 3 Set the default user role for external web interface users.

Users without a role cannot perform any actions. Any user roles defined in the external authentication object overrides this default user role.

a) Click the **Default User Role** value (by default, none selected).

a) In the **Default User Role Configuration** dialog box, check the role(s) that you want to use.

b) 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 1 object, then users are compared against servers in the order specified. See the next step to reorder servers.

If you enable shell authentication, you must enable an external authentication object that includes a **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 authentication they are accessed when an authentication request occurs.

Step 6 Choose **Shell Authentication** > **Enabled** if you want to allow CLI access for external users.

Note
The multidomain feature is not supported in CLI. Hence, the **Shell Authentication** option is available only in Global domain and not in Sub domains.

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**.

Configure Common Access Card Authentication with LDAP

If your organization uses Common Access Cards (CACs), you can configure LDAP authentication to authenticate management center users logging 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.

CAC-authenticated users are identified 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 users are re-added after each subsequent login, but you must reconfigure any manual changes to their user roles.

Caution When configuring CAC authentication with LDAP, ensure that you follow the principles of least privilege while assigning a default access role to the users. 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 has logged in with the default access role needs a temporary elevation of their privileges, a user with administrative privileges can temporarily provide that user the required higher level of access by assigning them a role with higher privilege. This privilege will be revoked after 24 hours of inactivity, and the user will return to their default access role.

If a user needs a permanent access role reassignment to a higher privilege level, such as System Admin, use the **Group Controlled Access Roles** method to provide admin access to the user. 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 on configuring Group Controlled Access Roles, see 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 terminates the session and the system logs you out of the web interface.

Procedure

Step 1 Insert a CAC as directed by your organization.

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 associated with the CAC you inserted in step 1.

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

Step 5 On the Login page, in the **Username** and **Password** fields, log in as a user with Administrator privileges. You **cannot** yet log in using your CAC credentials.

Step 6 Choose **System** > **Users** > **External Authentication**.

Step 7 Create an LDAP authentication object exclusively for CAC, following the procedure in Add an LDAP External Authentication Object for the Management Center, on page 129. You must configure the following:

- **CAC** check box.
- **LDAP-Specific Parameters** > **Show Advanced Options** > **User Name Template**.
- **Attribute Mapping** > **UI Access Attribute**.

Step 8 Click **Save**.

Step 9 Enable external authentication and CAC authentication as described in Enable External Authentication for Users on the Management Center, on page 142.

Step 10 Choose **System** (⚙) > **Configuration**, and click **HTTPS Certificate**.

Step 11 Import a HTTPS server certificate, if necessary, following the procedure outlined in Importing HTTPS Server Certificates, on page 68.

The same certificate authority (CA) must issue the HTTPS server certificate and the user certificates on the CACs you plan to use.

Step 12 Under **HTTPS Client Certificate Settings**, choose **Enable Client Certificates**. For more information, see Requiring Valid HTTPS Client Certificates, on page 70.

Step 13 Log in to the device according to Logging Into the Secure Firewall Management Center with CAC Credentials, on page 33.

Configure SAML Single Sign-On

You can configure your management center to use Single Sign-On, a system by which a central identity provider (IdP) provides authentication and authorization for users logging into the management center as well as other applications within an organization. The applications configured to take part in such an SSO arrangement are said to be federated service provider applications. SSO users can log in once to gain access to all service provider applications that are members of the same federation.

About SAML Single Sign-On

A management center configured for SSO presents a link for single sign-on on the Login page. Users configured for SSO access click on this link and are redirected to the IdP for authentication and authorization, rather than supplying a username and password on the management center Login page. Once successfully authenticated by the IdP, SSO users are redirected back to the management center web interface and logged in. All the communication between the management center and the IdP to accomplish this takes place using the browser as an intermediary; as a result, the management center does not require a network connection to directly access the identity provider.

The management center supports SSO using any SSO provider conforming to the Security Assertion Markup Language (SAML) 2.0 open standard for authentication and authorization.

Note The management center cannot sign SAML authentication request messages. Hence, if the IdP requires service provider's signature on the authentication requests, the SSO on the management center would fail.

The management center web interface offers configuration options for the following SSO providers:

- Okta
- OneLogin
- Azure
- PingID's PingOne for Customers cloud solution
- Other

Note The Cisco Secure Sign On SSO product does not recognize the management center as a pre-integrated service provider.

SSO Guidelines for the Management Center

Keep the following in mind when you configure a management center to be a member of an SSO federation:

- The management center can support SSO with only one SSO provider at a time—you cannot configure the management center to use, for instance, both Okta and OneLogin for SSO.
- Management Center management centers in a high availability configuration can support SSO, but you must keep the following considerations in mind:
  - SSO configuration is not synchronized between the members of the high availability pair; you must configure SSO separately on each member of the pair.
  - Both management centers in a high availability pair must use the same IdP for SSO. You must configure a service provider application at the IdP for each management center configured for SSO.
  - In a high availability pair of management centers where both are configured to support SSO, before a user can use SSO to access the secondary management center for the first time, that user must first use SSO to log into the primary management center at least once.
  - When configuring SSO for management centers in a high availability pair:

- If you configure SSO on the primary management center, you are not required to configure SSO on the secondary management center.
- If you configure SSO on the secondary management center, you are required to configure SSO on the primary management center as well. (This is because SSO users must login into the primary management center at least once before logging into the secondary management center.)

- In a management center that uses multi-tenancy, the SSO configuration can be applied only at the global domain level, and applies to the global domain and all subdomains.
- Only users with the Admin role authenticated internally or by LDAP or RADIUS can configure SSO.
- The management center does not support SSO initiated from the IdP.
- The management center does not support logging in with CAC credentials for SSO accounts.
- Do not configure SSO in deployments using CC mode.
- SSO activities are logged in the management center audit log with Login or Logout specified in the Subsystem field.

Related Topics

High Availability, on page 301
Domains, on page 213
Logging Into the Secure Firewall Management Center with CAC Credentials, on page 33
Security Certifications Compliance, on page 329
Audit Records, on page 407

SSO User Accounts

Identity providers can support user and group configuration directly, or they often can import users and groups from other user management applications such as Active Directory, RADIUS, or LDAP. This documentation focuses on configuring the management center to work with the IdP to support SSO assuming 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 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 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 management center.

The following account characteristics for SSO users can be configured from the management center web interface under **System** (⚙) > **Users** > **Edit User**:

- Real Name
- Exempt from Browser Session Timeout

User Role Mapping for SSO Users

By default, all the users who are given SSO access to a 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. After you have established and successfully tested the management center SSO configuration, you can configure user role mapping.

User role mapping requires coordinating configuration settings at the management center with settings at the SSO IdP application. User roles can be assigned to users or to groups defined at the IdP application. Users may or may not be members of groups, and user or group definitions may or may not be imported to the IdP from other user management systems within your organization, such as Active Directory. For this reason, to effectively configure management center SSO user role mapping you must be familiar with how your SSO federation is organized and how users, groups and their roles are assigned at the SSO IdP application. This documentation focuses on configuring the management center to work with the IdP to support user role mapping; to create users or groups within the IdP, or import users or groups into the IdP from a user management application, consult the IdP vendor documentation.

In user role mapping, the IdP maintains a role attribute for the management center service provider application, and each user or group with access to that management center is configured with a string or expression for the role attribute. Requirements for the attribute value are different for each IdP. At the management center, the name of the role attribute, a list of expressions assigned to a list of management center user roles are part of the SSO configuration. When a user logs into the management center using SSO, the management center compares the value of the role attribute for that user (or that user's group, depending upon configuration) against the expressions for each management center user role. The management center assigns the user all the roles where the expression matches the attribute value the user has provided.

Note You can configure management center roles to be mapped based on individual user permissions or based on group permissions, but a single management center application cannot support role mapping for both groups and individual users.

Enable Single Sign-On at the Management Center

Before you begin

- At the SAML SSO management application, configure a service provider application for the management center and assign users or groups to the service provider application:
  - To configure a management center service provider application for Okta, see Configure the Management Center Service Provider Application for Okta, on page 150.
  - To configure a management center service provider application for OneLogin, see Configure the Management Center Service Provider Application for OneLogin, on page 162.
  - To configure a management center service provider application for Azure, see Configure the Management Center Service Provider Application for Azure, on page 175.
  - To configure a management center service provider application for PingID's PingOne for Customers cloud solution, see Configure the Management Center Service Provider Application for PingID PingOne for Customers, on page 188.

- To configure a management center service provider application for any SAML 2.0-compliant SSO provider, see Configure Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider, on page 192.

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**.

What to do next

Proceed with the instructions appropriate to your choice of SSO provider:

- Configure the management center for Okta SSO; see Configure the Management Center for Okta SSO, on page 151.
- Configure the management center for SSO using PingID's PingOne for Customers cloud solution; see Configure the Management Center for SSO with the PingID PingOne for Customers, on page 190.
- Configure the management center for Azure SSO; see Configure the Management Center for Azure SSO, on page 177.
- Configure the management center for OneLogin SSO; see Configure the Management Center for OneLogin SSO, on page 164.
- Configure the management center for SSO using any SAML 2.0-compliant provider; see Configure the Management Center for SSO Using Any SAML 2.0-Compliant SSO Provider, on page 194.

Configure Single Sign-On with Okta

See the following tasks to configure SSO using Okta:

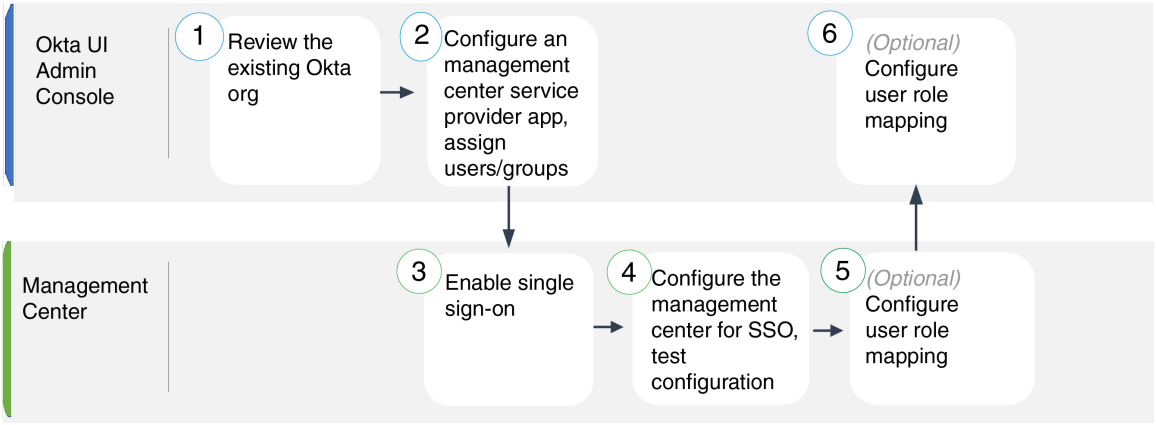

1	Okta UI Admin Console	Review the Okta Org, on page 149
2	Okta UI Admin Console	Configure the Management Center Service Provider Application for Okta, on page 150
3	management center	Enable Single Sign-On at the Management Center, on page 147
4	management center	Configure the Management Center for Okta SSO, on page 151
5	management center	Configure User Role Mapping for Okta in the Management Center, on page 152
6	Okta UI Admin Console	Configure User Role Mapping at the Okta IdP, on page 153

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 management center to an Okta org, be familiar with its configuration; consider the following questions:

- How many users will have access to the 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 management center?
- What kind of user role assignments do you want to make? (If you choose not to assign user roles, the 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 management center roles to be mapped based on individual user permissions or based on group permissions, but a single 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 Management Center Service Provider Application for Okta

Use these instructions at the Okta Classic UI Admin Console to create a 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 instance, to create users and groups, or to import user and group definitions from another user management application, see the Okta documentation.

Note If you plan to assign user groups to the management center application, do not also assign users within those groups as individuals.

Note The 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 management center.

Before you begin

- Familiarize yourself with the SSO federation and its user and groups; see Review the Okta Org, on page 149.
- Create user accounts and/or groups 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 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 management center.

- Confirm the login URL for the target management center (`https://ipaddress_or_hostname`).

Note If your 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 management center using the login URL that you configure in this task.

Procedure

Step 1 From the Okta Classic UI Admin Console, create a service provider application for the management center. Configure the management center application with the following selections:

- Select `Web` for the **Platform**.
- Select `SAML 2.0` for the **Sign on method**.
- Provide a **Single sign on URL**.

  This is the management center URL to which the browser sends information on behalf of the IdP.

  Append the string `saml/acs` to the management center login URL. For example: `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 management center), often formatted as a URL.

  Append the string `/saml/metadata` to the management center login URL. For example: `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 management center application. (If you plan to assign groups to the 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 management center application.

Step 4 (Optional) To make SSO setup at the management center easier, you can download the SAML XML metadata file for the management center service provider application from Okta to your local computer.

What to do next

Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Configure the Management Center for Okta SSO

Use these instructions at the management center web interface.

Before you begin

- Create a management center service provider application at the Okta Classic UI Admin Console; see Configure the Management Center Service Provider Application for Okta, on page 150.
- Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Procedure

Step 1 (This step continues directly from Enable Single Sign-On at the Management Center, on page 147.) 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 following values from the Okta SSO Service Provider application. (Retrieve these values from the Okta Classic UI Admin Console.)
     - **Identity Provider Single Sign-On (SSO) URL**
     - **Identity Provider Issuer**
     - **X.509 Certificate**
- If you saved the XML metadata file generated by Okta to your local computer (Step 4 in Configure the Management Center Service Provider Application for Okta, on page 150), you can upload the file to the 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 management center as well as the Okta service provider application configuration, correct any errors, and try again.

Step 5 When the system reports a successful configuration test, click **Apply**.

What to do next

You may optionally configure user role mapping for SSO users; see Configure User Role Mapping for Okta in the Management Center, on page 152. If you choose not to configure role mapping, by default all SSO users that log into the management center are assigned the user role you configure in Step 4 of Configure User Role Mapping for Okta in the Management Center, on page 152.

Configure User Role Mapping for Okta in the Management Center

The fields to configure for user role mapping in the 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; see Review the Okta Org, on page 149.
- Configure the management center as an SSO service provider application; see Configure the Management Center Service Provider Application for Okta, on page 150.
- Enable and configure single sign-on at the management center; see Enable Single Sign-On at the Management Center, on page 147, and Configure the Management Center for Okta SSO, on page 151.

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 management center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in Okta for management center user role mapping for users or groups. (See Step 1 of Configure a User Attribute for Role Mapping at the Okta IdP, on page 154 or Step 1 of Configure a Group Attribute for Role Mapping at the Okta IdP, on page 155.)

Step 5 Next to each management center user role you wish to assign to SSO users, enter a regular expression. (The management center uses a restricted version of Google's RE2 regular expression standard supported by Golang and Perl.) The management center compares these values against the user role mapping attribute value the IdP sends to the management center with SSO user information. The 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 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 User Role Mapping at the Okta IdP, on page 153.

Configure User Role Mapping at the Okta IdP

You can configure SSO user role mapping at the Okta Classic UI Admin Console based on individual user permissions or group permissions.

- To map based on individual user permissions, see Configure a User Attribute for Role Mapping at the Okta IdP, on page 154.
- To map based on group permissions, see Configure a Group Attribute for Role Mapping at the Okta IdP, on page 155.

When an SSO user logs in to the management center, Okta presents a user or group role attribute value configured at the Okta IdP, to the management center. The management center then compares that attribute value with the regular expressions assigned to each 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 management center

grants the user a configurable default user role.) The expression you assign to each management center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl. The management center treats the attribute value received from Okta as a regular expression using that same standard for purposes of comparison with the management center user role expressions.

Note A single management center cannot support role mapping for both groups and individual users; you must choose one mapping method for the management center service provider application and use it consistently. Furthermore, the management center can support group role mapping using only one group attribute statement per management center service provider application configured in Okta. Generally group-based role mapping is more efficient for a 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 management center you must configure a custom attribute in the profile to convey each user's role mapping expression to the 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 management center service provider application at the Okta IdP as described in Configure the Management Center Service Provider Application for Okta, on page 150.
- Configure SSO user role mapping at the management center as described in Configure User Role Mapping for Okta in the Management Center, on page 152.

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 management center, containing an expression to match for user role mapping. This variable name must match the string you entered at the management

center SSO configuration for **Group Member Attribute**. (See Step 5 in Configure User Role Mapping for Okta in the Management Center, on page 152.)

Step 2 For each user assigned to the 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 management center will assign to the user. The management center compares this string against the expressions you assigned to each management center user role in Step 6 of Configure User Role Mapping for Okta in the Management Center, on page 152. (For purposes of comparison with the management center user role expressions, the 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 management center service provider application. The management center can support group role mapping using only one group attribute statement per Okta 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 management center you must configure a custom attribute for the group to convey its role mapping expression to the 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. See the Okta documentation for details.

Before you begin

- Configure a management center service provider application at the Okta IdP; see Configure the Management Center Service Provider Application for Okta, on page 150.
- Configure user role mapping at the management center; Configure User Role Mapping for Okta in the Management Center, on page 152.

Procedure

Create a new SAML group attribute for the management center service provider application:

- For **Name**, use the same string you entered at the management center SSO configuration for **Group Member Attribute**. (See Step 4 in Configure User Role Mapping for Okta in the Management Center, on page 152.)

- For **Filter**, specify an expression to represent the role or roles the 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 management center the group names that match. The management center in turn compares those group names against the regular expressions you assigned to each management center user role in Step 5 of Configure User Role Mapping for Okta in the Management Center, on page 152.

Okta User Role Mapping Examples

As the following examples demonstrate, the SSO configurations at the 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 management center service provider application in Okta.

Note You can configure management center roles to be mapped based on individual user permissions or based on group permissions, but a single management center application cannot support role mapping for both groups and individual users. Furthermore, the management center can support group role mapping using only one group attribute statement per management center service provider application configured in Okta.

Okta Role Mapping Example for Individual User Accounts

In role mapping for individual users, the Okta management center service application has a custom attribute whose name matches the name of the Group Member Attribute on the 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 management center, it will use the custom attribute value assigned for the user in question.

The following diagrams illustrate how the relevant fields and values in the 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 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 management center.

- In this diagram sue@example.com uses the `FMCrole` value `FMCAdmin` and the management center assigns her the Administrator role.

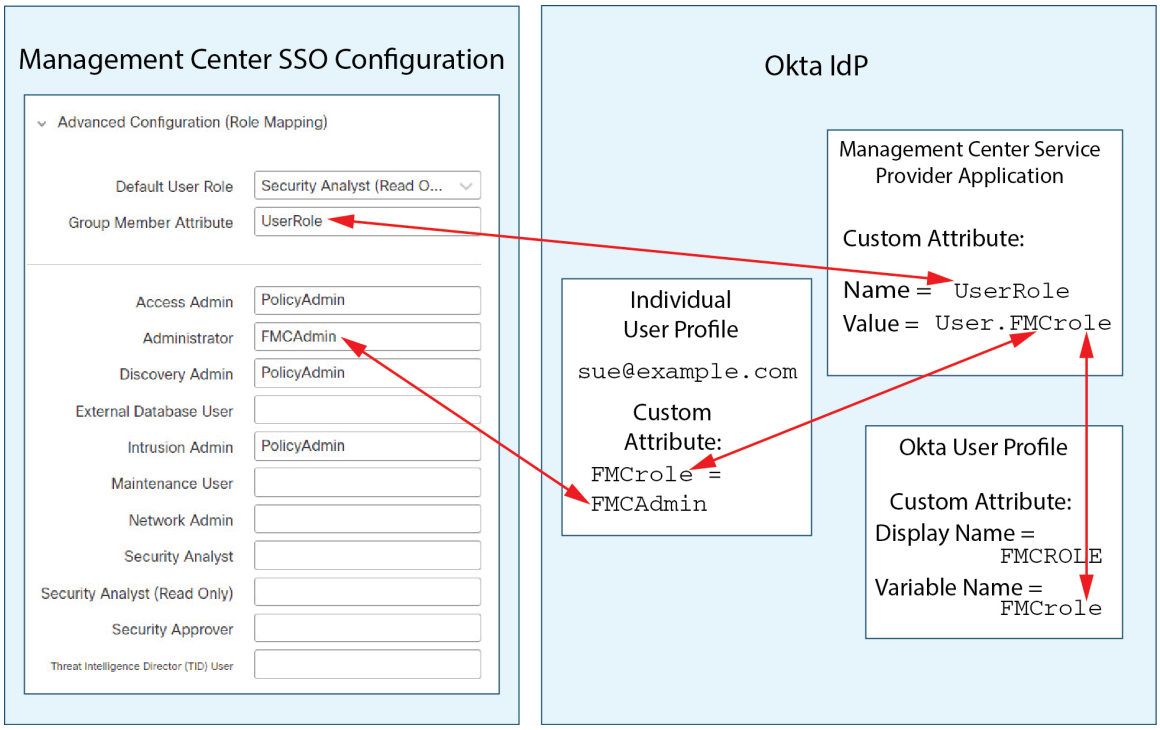

- In this diagram fred@example.com uses the `FMCrole` value `PolicyAdmin`, and the management center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.

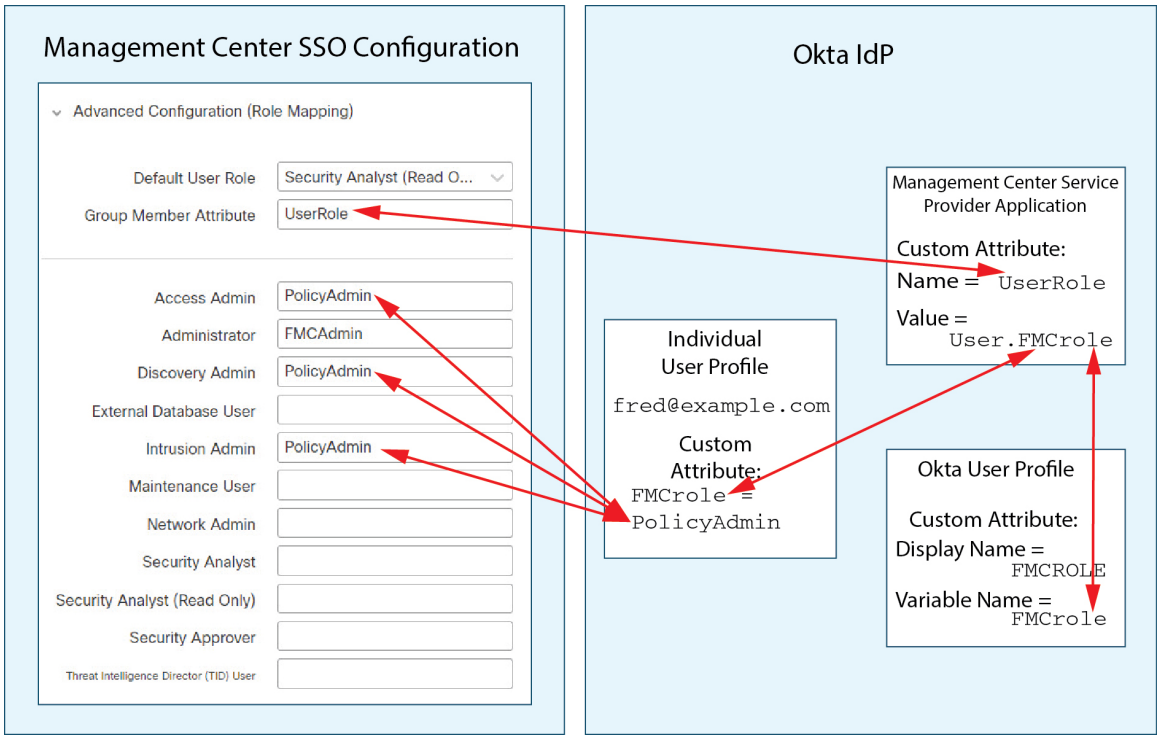

- Other users assigned to the Okta service application for this management center are assigned the default user role Security Analyst (Read Only) for one of the following reasons:
  - They have no value assigned to the `FMCrole` variable in their Okta user profile.
  - The value assigned to the `FMCrole` variable in their Okta user profile does not match any expression configured for a user role in the SSO configuration at the management center.

Okta Role Mapping Example for Groups

In role mapping for groups, the Okta management center service application has a custom group attribute whose name matches the name of the **Group Member Attribute** on the management center (in this example, `UserRole`). When Okta processes a request for management center SSO login, it compares the user's group membership against the expression assigned to the management center service application group attribute (in this case `^(.*)Admin$`). Okta sends to the management center the user's group membership(s) that match the group attribute. The management center compares the group names it receives against the regular expressions it has configured for each user role, and assigns user roles accordingly.

The following diagrams illustrate how the relevant fields and values in the management center and Okta configurations correspond to each other in user role mapping for groups. Each diagram uses the same SSO configurations at the 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 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 management center Fred's `Admin` group membership, and the management center assigns him the Administrator role.

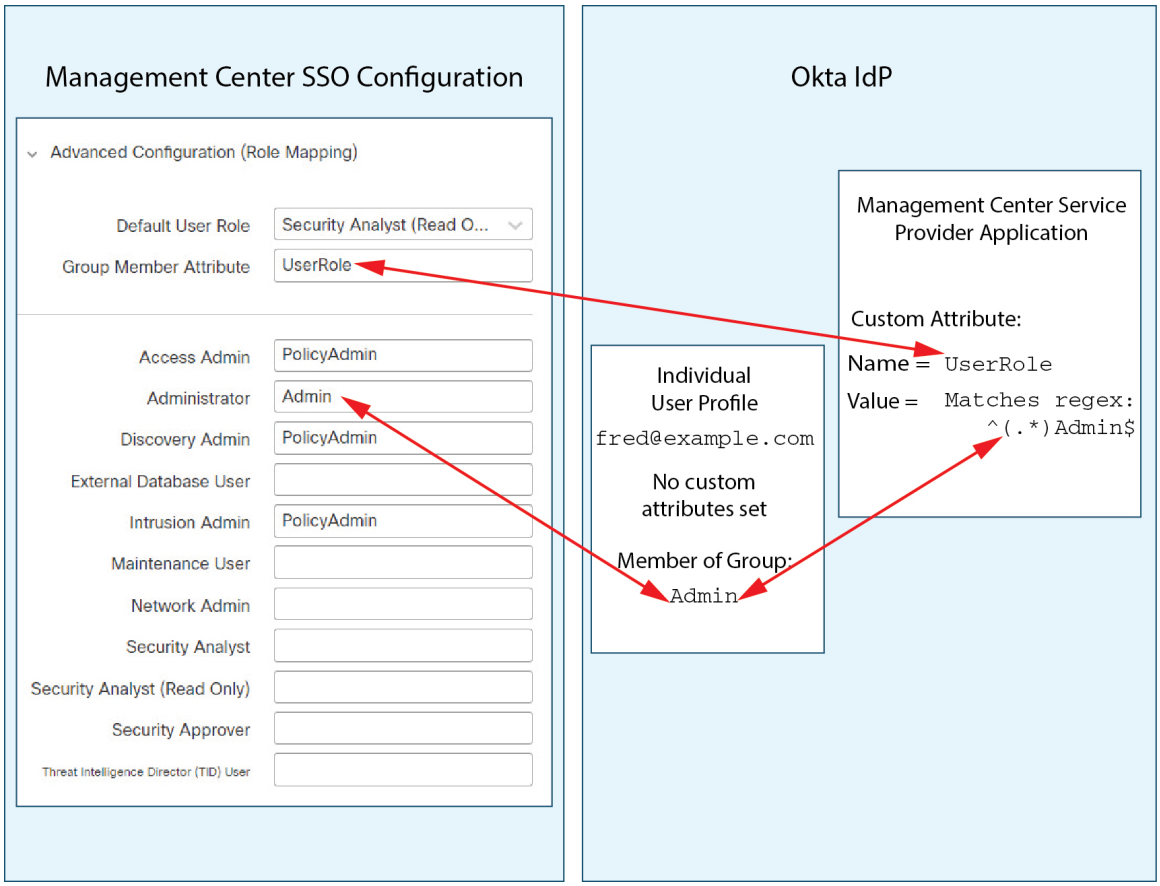

- In this diagram sue@example.com is a member of the Okta IdP group `PolicyAdmin`, which matches the expression `^(.*)Admin$`. Okta sends the management center Sue's `PolicyAdmin` group membership, and the management center assigns her the roles Access Admin, Discovery Admin, and Intrusion Admin.

  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 management center service application, Okta does not send information about Sue's `Maint` group membership to the management center, and her membership in the `Maint` group plays no part in the roles the management center assigns to her.

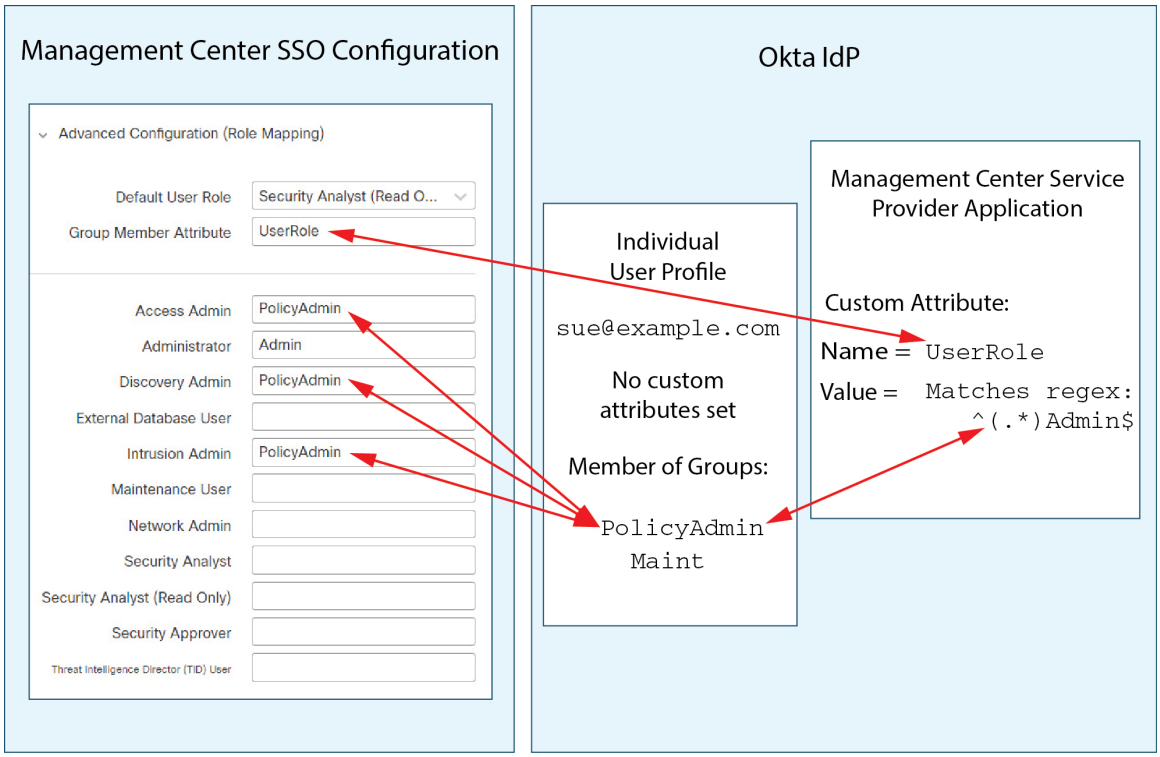

- In this diagram sean@example.com is a member of the Okta IdP group `Maint`. This group name does not match the expression `^(.*)Admin$`, so, when sean@example.com logs into the management center, Okta does not send information about Sean's `Maint` group membership to the management center and Sean is assigned the default user role (Security Analyst (Read Only)) rather than the Maintenance User role.

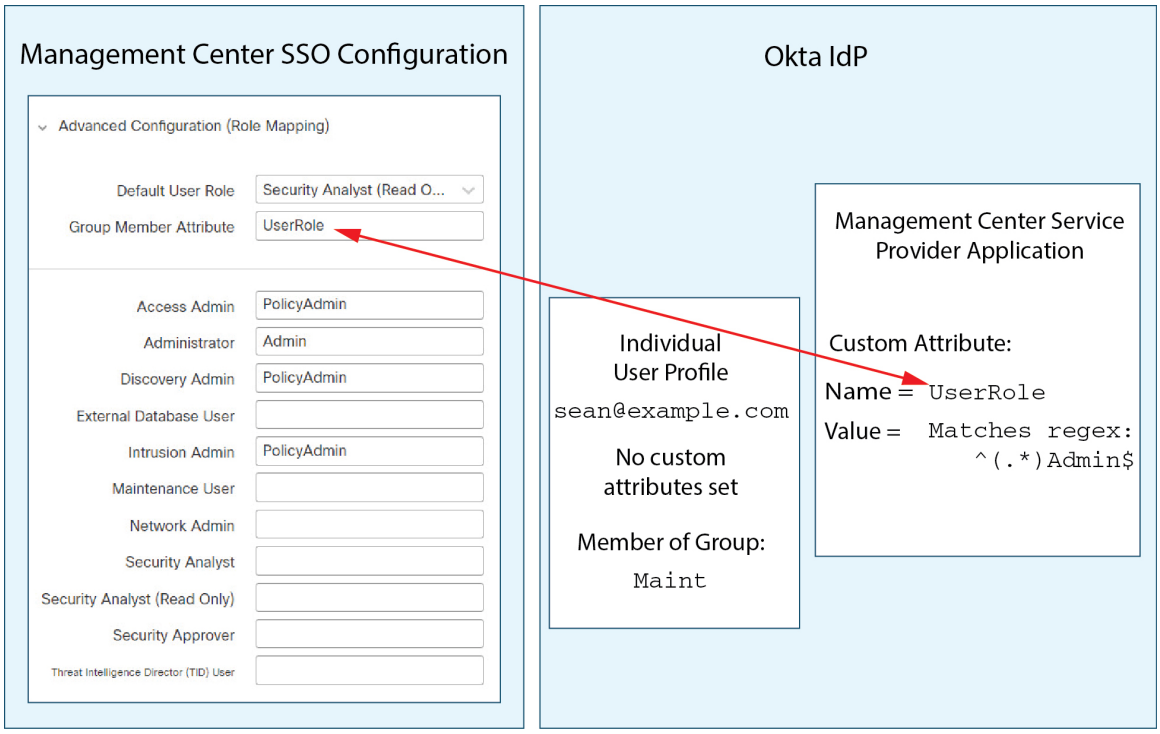

These diagrams illustrate the importance of advance planning when establishing a role mapping strategy. In this example, any Okta user with access to this management center who is a member of only the `Maint` group can be assigned only the default user role. The management center supports using only 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 management center SSO configuration. (The expression you assign to each management center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl.)

Configure Single Sign-On with OneLogin

See the following tasks to configure SSO using OneLogin:

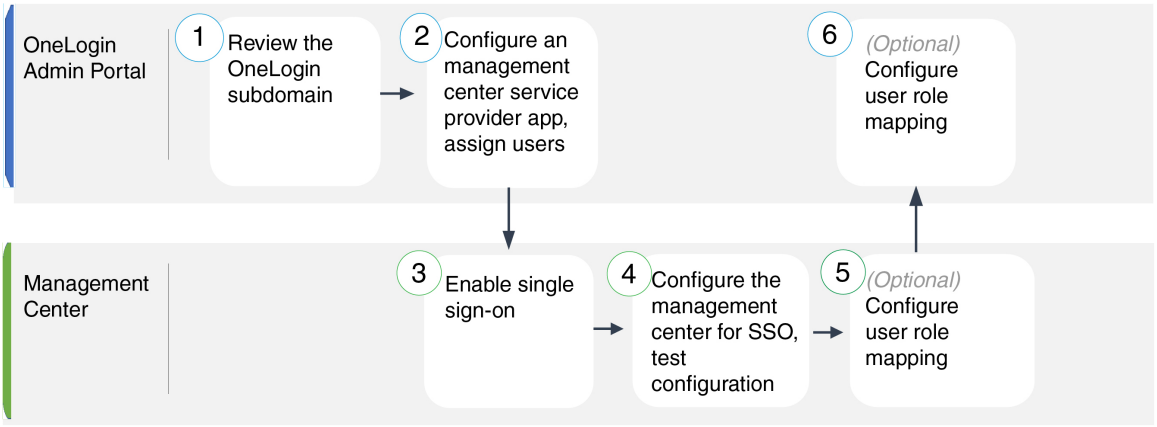

1	management center	Review the OneLogin Subdomain, on page 162
2	management center	Configure the Management Center Service Provider Application for OneLogin, on page 162
3	OneLogin Admin Portal	Enable Single Sign-On at the Management Center, on page 147
4	OneLogin Admin Portal	Configure the Management Center for OneLogin SSO, on page 164
5	OneLogin Admin Portal	Configure User Role Mapping for OneLogin in the Management Center, on page 165
6	management center	Configure User Role Mapping at the OneLogin IdP, on page 166

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 management center to a OneLogin subdomain, be familiar with its configuration; consider the following questions:

- How many users will have access to the 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 management center?
- What kind of management center user role assignments do you want to make? (If you choose not to assign user roles, the 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 management center roles to be mapped based on individual users or based on groups, but a single 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, see the OneLogin documentation available online.

Configure the Management Center Service Provider Application for OneLogin

Use these instructions at the OneLogin Admin Portal to create a 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

you need to establish a fully functional SSO org; for instance, to create users and groups, or to import user and group definitions from another user management application, see the OneLogin documentation.

Note If you plan to assign user groups to the management center application, do not also assign users within those groups as individuals.

Note The 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 management center.

Before you begin

- Familiarize yourself with the OneLogin subdomain and its users and groups; see Review the OneLogin Subdomain, on page 162.
- 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 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 management center.

- Confirm the login URL for the target management center (`https://ipaddress_or_hostname/`).

Note If your 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 management center using the login URL that you configure in this task.

Procedure

Step 1 Create the management center service provider application using the **SAML Test Connector (Advanced)** as its basis.

Step 2 Configure the application with the following settings:

- For the **Audience (Entity ID)**, append the string `/saml/metadata` to the management center login URL. For example: `https://ExampleFMC/saml/metadata`.

- For **Recipient**, append the string `/saml/acs` to the 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 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 management center login URL. For example: `https://ExampleFMC/saml/acs`.
- For **Login URL**, append the string `/saml/acs` to the management center login URL. For example: `https://ExampleFMC/saml/acs`.
- For the **SAML Initiator**, choose `Service Provider`.

Step 3 Assign OneLogin users to the management center service provider application.

Step 4 (Optional) To make SSO setup at the management center easier, you can download the SAML XML metadata for the management center service provider application from OneLogin to your local computer.

What to do next

Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Configure the Management Center for OneLogin SSO

Use these instructions at the management center web interface.

Before you begin

- Create a management center service provider application at the OneLogin Admin Portal; see Configure the Management Center Service Provider Application for OneLogin, on page 162.
- Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Procedure

Step 1 (This step continues directly from Enable Single Sign-On at the Management Center, on page 147.) 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 the following 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 saved the XML metadata file generated by OneLogin to your local computer (Step 4 in Configure the Management Center Service Provider Application for OneLogin, on page 162), you can upload the file to the 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 management center as well as the OneLogin service provider application configuration, correct any errors, and try again.

Step 5 When the system reports a successful configuration test, click **Apply**.

What to do next

You may optionally configure user role mapping for SSO users; see Configure User Role Mapping for OneLogin in the Management Center, on page 165. If you choose not to configure role mapping, by default all SSO users that log into the management center are assigned the user role you configure in Step 4 of Configure User Role Mapping for OneLogin in the Management Center, on page 165.

Configure User Role Mapping for OneLogin in the Management Center

The fields to configure for user role mapping at the 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, see Review the OneLogin Subdomain, on page 162.
- Configure an SSO service provider application for the management center; see Configure the Management Center Service Provider Application for OneLogin, on page 162.
- Enable and configure single sign-on at the management center; see Enable Single Sign-On at the Management Center, on page 147, and Configure the Management Center Service Provider Application for OneLogin, on page 162.

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 management center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in OneLogin for management center user role mapping for users or groups. See Step 1 of Configure User Role Mapping for Individual Users at the OneLogin IdP, on page 167 or Step 1 of Configure User Role Mapping for Groups at the OneLogin IdP, on page 168.

Step 5 Next to each management center user roll you wish to assign to SSO users, enter a regular expression. The management center compares these values against the user role mapping attribute the IdP sends to the management center with SSO user information. The 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 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 User Role Mapping at the OneLogin IdP, on page 166.

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 167.
- To map based on group permissions, see Configure User Role Mapping for Groups at the OneLogin IdP, on page 168.

When an SSO user logs into the management center, OneLogin presents to the management center a user or group role attribute value that gets its value from a custom user field configured at the OneLogin IdP. The management center compares that attribute value against the regular expressions assigned to each 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 management center grants the user a configurable default user role.) The expression you assign to each management center user role must comply with the restricted version of Google's RE2 regular expression standard supported by Golang and Perl. The management center treats the attribute value received from Okta as a regular expression using that same standard for purposes of comparison with the management center user role expressions.

Note A single management center cannot support role mapping for both groups and individual users; you must choose one mapping method for the management center service provider application and use it consistently. The 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 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 management center service provider application and a custom user field. These provide the means for OneLogin to pass user role information to the 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 162.
- Create and configure a management center service provider application in OneLogin; see Configure the Management Center Service Provider Application for OneLogin, on page 162.
- Configure SSO user role mapping as described in Configure User Role Mapping for OneLogin in the Management Center, on page 165.

Procedure

Step 1 Create a custom parameter for the management center service provider application.

- For the **Field Name**, use the same name you used for the **Group Member Attribute** in the management center SSO configuration. (See Step 4 in Configure User Role Mapping for OneLogin in the Management Center, on page 165.)
- 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 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 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 management center using SSO, the value you assign to this field for that user is the value the management center compares against the expressions you assigned to management center user roles

in the SSO configuration. (See Step 5 in Configure User Role Mapping for OneLogin in the Management Center, on page 165.)

What to do next

- Test your role mapping scheme by logging into the management center using SSO from various accounts and confirming that users are assigned 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 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 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 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 162.
- Create and configure a management center service provider application in OneLogin; see Configure the Management Center Service Provider Application for OneLogin, on page 162.
- Configure SSO user role mapping as described in Configure User Role Mapping for OneLogin in the Management Center, on page 165.

Procedure

Step 1 Create a custom parameter for the management center service provider application.

- For the **Field Name**, use the same name you used for the **Group Member Attribute** in the management center SSO configuration. (See Step 4 in Configure User Role Mapping for OneLogin in the Management Center, on page 165.)
- 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 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 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 management center compares this string against the expressions you assign to each management center user role in Step 5 of Configure User Role Mapping for OneLogin in the Management Center, on page 165.
- **Reapply All Mappings** when you have completed your changes.

What to do next

- Test your role mapping scheme by logging into the management center using SSO from various accounts and confirming that users are assigned management center user roles as you expect.

OneLogin User Role Mapping Examples

As the following examples demonstrate, the SSO configurations at the 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 management center service provider application in OneLogin.

Note A single management center cannot support role mapping for both groups and individual users; you must choose one mapping method for the management center service provider application and use it consistently. The 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 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 management center service application has a custom parameter whose name matches the name of the Group Member attribute on the 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 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 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 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 management center.

- In this diagram fred@example.com uses the `FMCUserRole` value `PolicyAdmin` and the management center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.

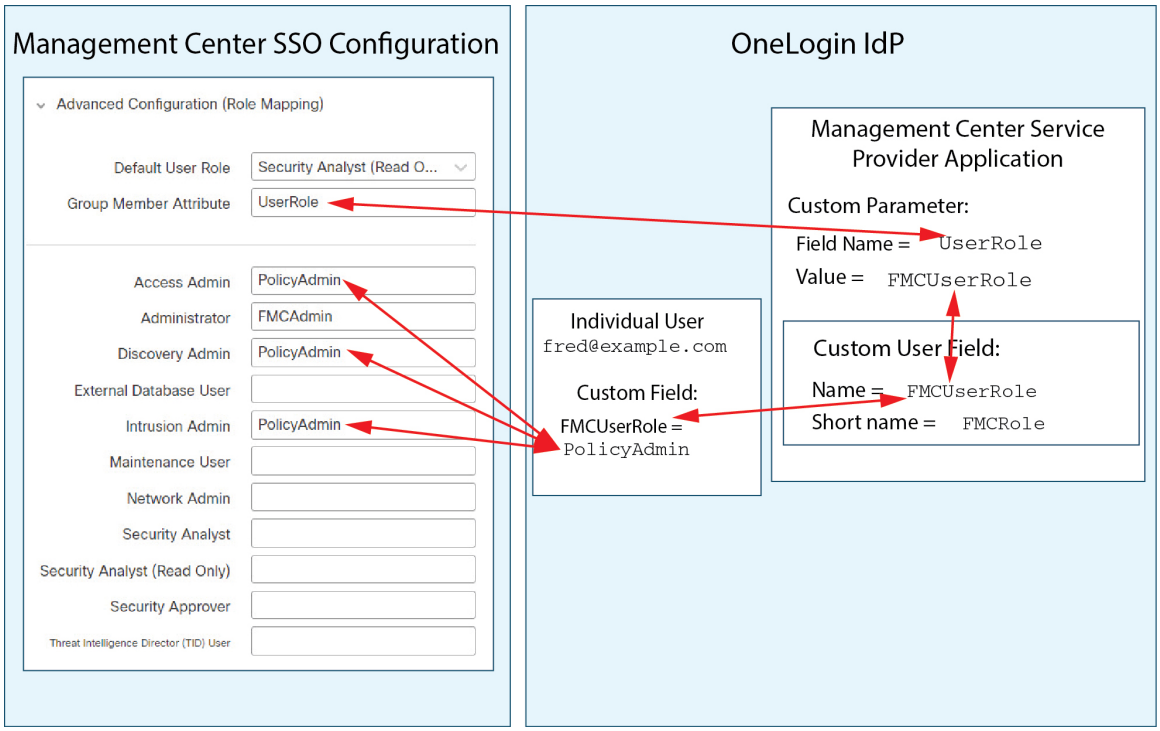

- In this diagram sue@example.com uses the `FMCUserRole` value `FMCAdmin`, and the management center assigns her the Administrator role.

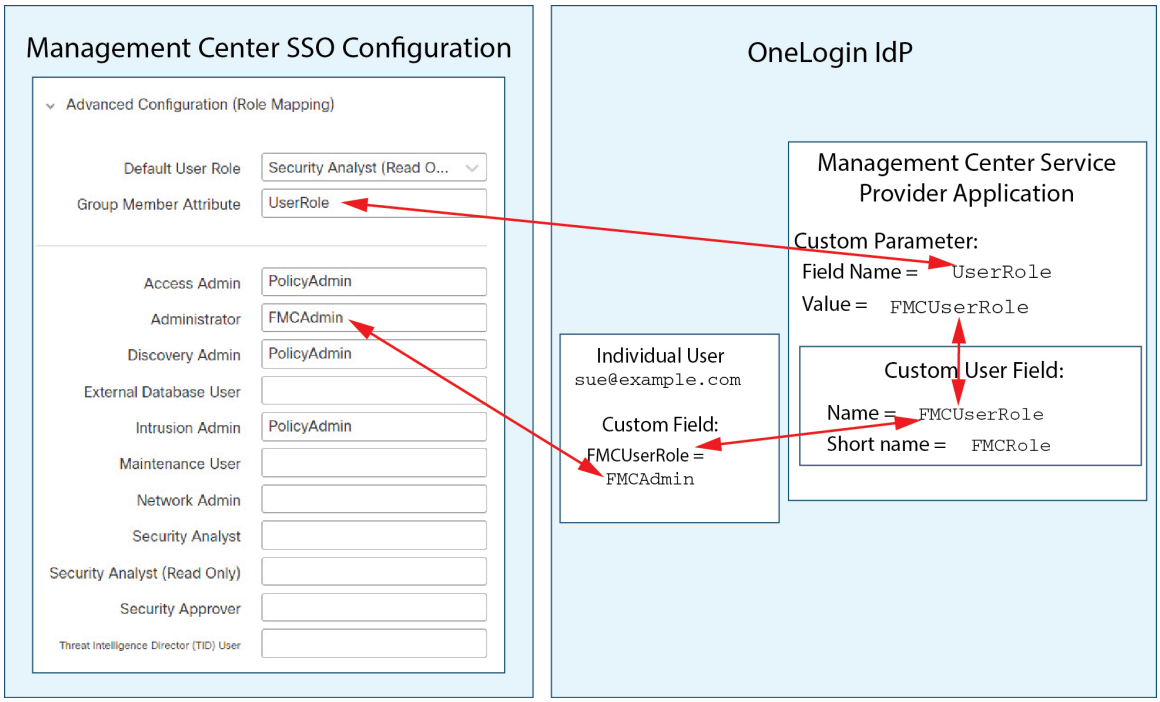

- Other users assigned to the OneLogin service application for this management center are assigned the default user role Security Analyst (Read Only) for one of the following reasons:
  - They have no value assigned to the `FMCUserRole` custom user field.
  - The value assigned to the `FMCUserRole` custom user field does not match any expression configured for a user role in the SSO configuration at the management center.

OneLogin Role Mapping Example for Groups

In role mapping for groups, the OneLogin management center service application has a has a custom parameter whose name matches the name of the Group Member attribute on the 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 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 management center and OneLogin configurations correspond to each other in user role mapping for groups. Each diagram uses the same SSO configurations at the 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 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 management center assigns Fred and other members of the `FMCPolicyAdminGroup` the roles Access Admin, Discovery Admin, and Intrusion Admin.

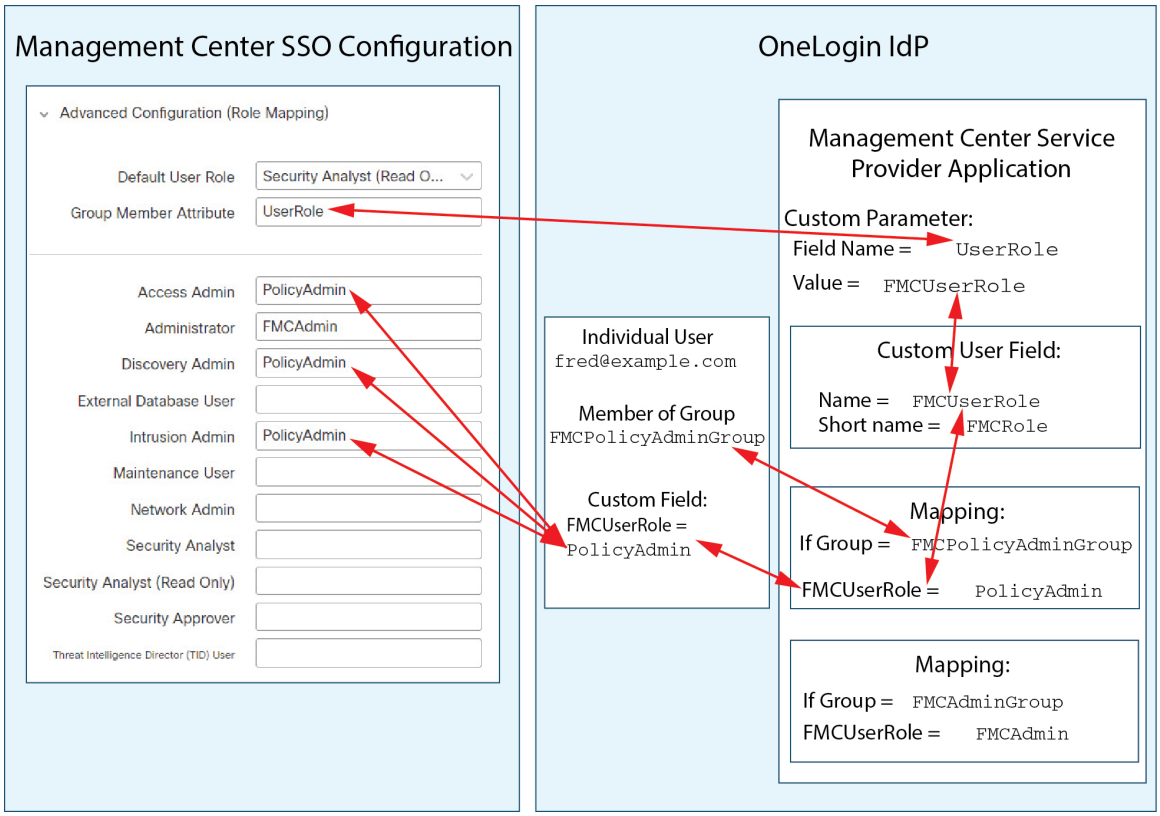

- In this diagram sue@example.com is a member of the OneLogin IdP group `FMCAdminGroup`. A OneLogin mapping assigns the value FMCAdmin to the custom user field `FMCUserRole` for members of the `FMCAdminGroup`. The management center assigns Sue and other members of the `FMCAdminGroup` the Administrator role.

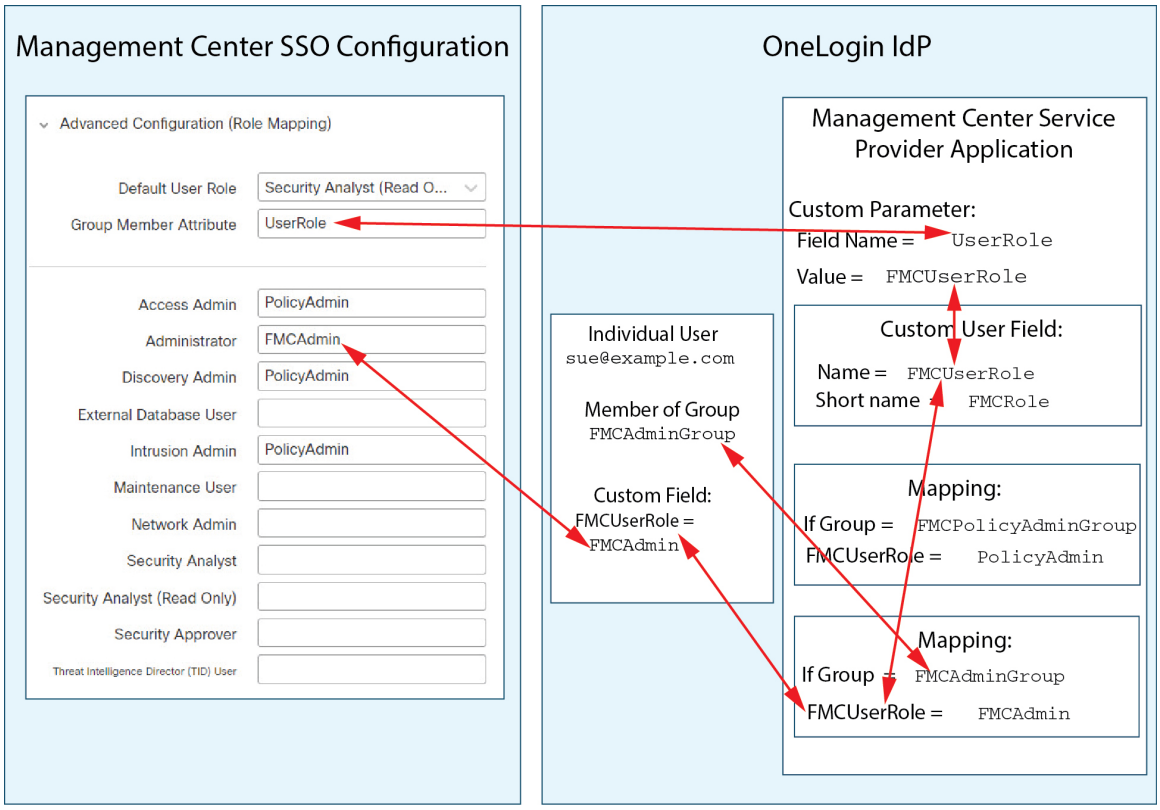

- In this diagram sean@example.com is a member of the Idp group `FMCMaintGroup`. There is no OneLogin mapping associated with this group, so OneLogin does not assign a value to the custom user field `FMCUserRole` for Sean. The management center assigns Sean the default user role (Security Analyst (Read Only)) rather than the Maintenance User role.

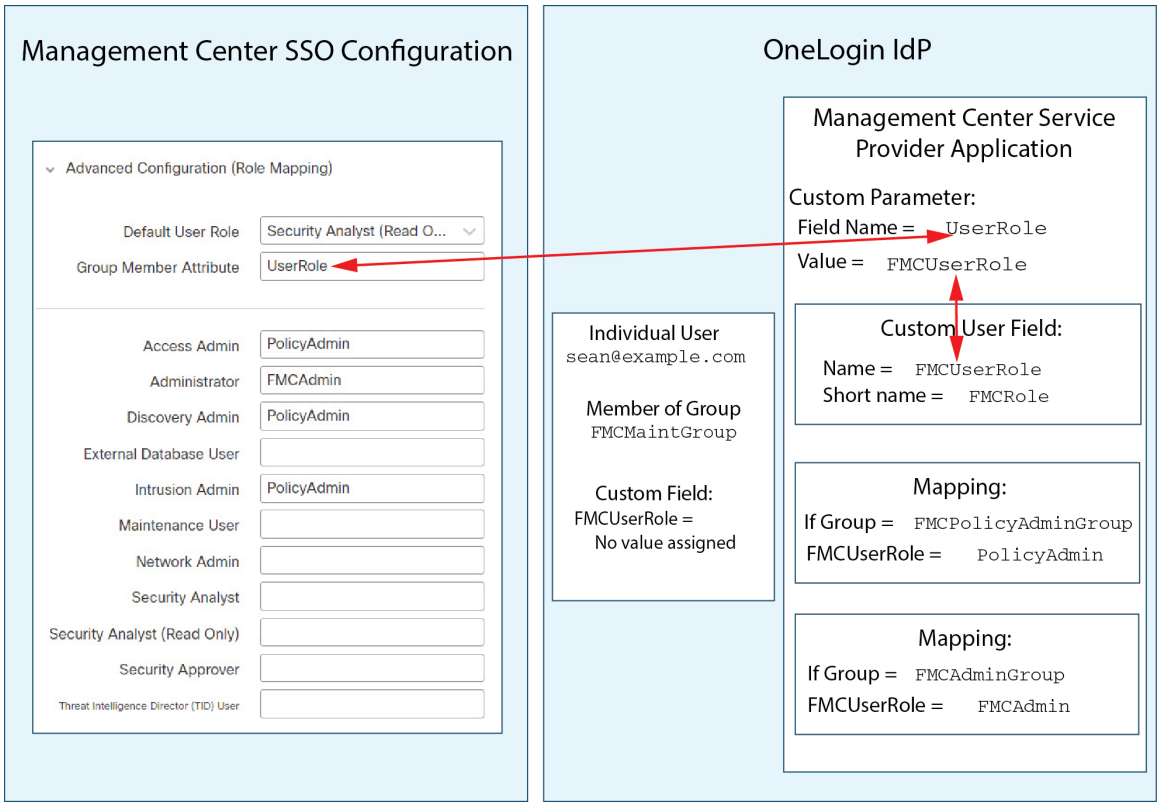

Configure Single Sign-On with Azure AD

See the following tasks to configure SSO using Azure:

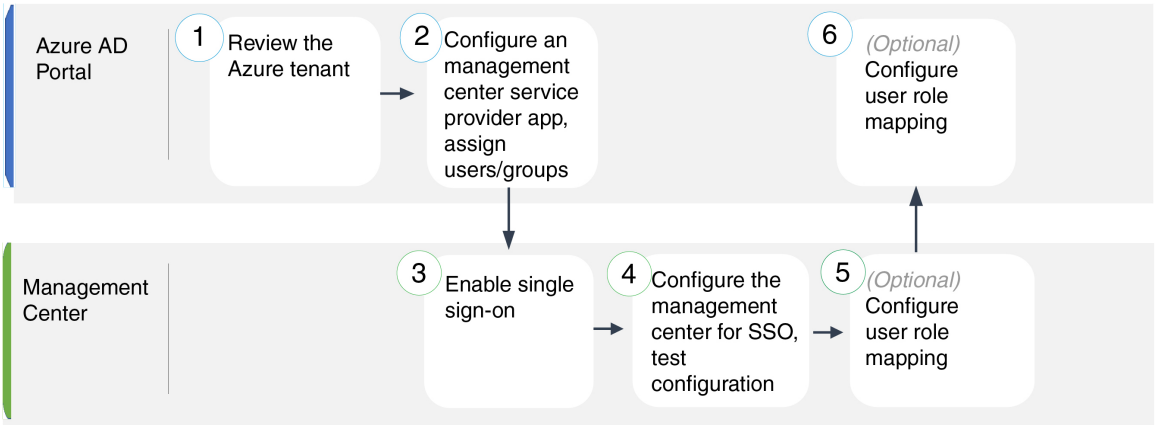

1	Azure AD Portal	Review the Azure Tenant, on page 175
2	Azure AD Portal	Configure the Management Center Service Provider Application for Azure, on page 175

3	management center	Enable Single Sign-On at the Management Center, on page 147
4	management center	Configure the Management Center for Azure SSO, on page 177
5	management center	Configure User Role Mapping for Azure in the Management Center, on page 178
6	Azure AD Portal	Configure User Role Mapping in the Azure IdP, on page 179

Review the Azure Tenant

Azure AD is Microsoft's multitenant cloud based identity and access management service. In Azure, the entity that encompasses all the federated devices that a user can access with the same SSO account is called a *tenant*. Before adding the management center to an Azure tenant, be familiar with its organization; consider the following questions:

- How many users will have access to the 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 management center?
- What kind of management center user role assignments do you want to make? (If you choose not to assign user roles, the 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?
- Keep in mind that you can configure management center roles to be mapped based on individual users or based on groups, but a single management center application cannot support role mapping for both groups and individual users.

This documentation assumes you are already familiar with the Azure Active Directory Portal and have an account with application admin privileges for the Azure AD tenant. Keep in mind that the management center supports Azure SSO only with tenant-specific single sign-on and single sign-out endpoints. You must have an Azure AD Premium P1 or above license and Global Administrator permissions; see Azure documentation for more information.

Configure the Management Center Service Provider Application for Azure

Use the Azure Active Directory Portal to create a management center service provider application within your Azure Active Directory tenant and establish basic configuration settings.

Note If you plan to assign user groups to the management center application, do not also assign users within those groups as individuals.

Note The 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 management center.

Before you begin

- Familiarize yourself with your Azure tenant and its users and groups; see Review the Azure Tenant, on page 175.
- Create user accounts and/or groups 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 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 management center.

- Confirm the login URL for the target management center (`https://ipaddress_or_hostname`)

Note If your 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 management center using the login URL that you configure in this task.

Procedure

Step 1 Create the management center service provider application using the Azure AD SAML Toolkit as its basis.

Step 2 Configure the application with the following settings for **Basic SAML Configuration**:

- For the **Identifier (Entity ID)** append the string `/saml/metadata` to the 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 management center login URL. For example: `https://ExampleFMC/saml/acs`.
- For the **Sign on URL** append the string `/saml/acs` to the 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 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 management center. Use the following 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 management center web interface

Step 6 In the SAML-based Sign-on information for the application, note the following values:

- **Login URL**
- **Azure AD Identifier**

You will need these values when you configure Azure SSO at the management center web interface.

Step 7 (Optional) to make SSO setup at the management center easier, you can download the SAML XML metadata file for the 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 management center service application.

Note
If you plan to assign user groups to the management center Application, do not also assign users within those groups as individuals.

Note
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 management center application cannot support role mapping for both groups and individual users.

What to do next

Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Configure the Management Center for Azure SSO

Use these instructions at the management center web interface.

Before you begin

- Create a management center service provider application at the Azure AD Portal; see Configure the Management Center Service Provider Application for Azure, on page 175.
- Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Procedure

Step 1 (This step continues directly from Enable Single Sign-On at the Management Center, on page 147.) 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 Management Center Service Provider Application for Azure, on page 175.
     - For **Identity Provider Issuer** enter the **Azure AD Identifier** you noted in Step 6 of Configure the Management Center Service Provider Application for Azure, on page 175.
     - For the **X.509 Certificate**, use the certificate you downloaded from Azure in Step 5 of Configure the Management Center Service Provider Application for Azure, on page 175. (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 Azure to your local computer (Step 7 of Configure the Management Center Service Provider Application for Azure, on page 175), you can upload the file the 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 management center as well as the Azure service provider application, correct any errors, and try again.

Step 5 When the system reports a successful configuration test, click **Apply**.

What to do next

You may optionally configure role mapping for SSO users; see Configure User Role Mapping for Azure in the Management Center, on page 178. If you choose not to configure role mapping, by default all SSO users that log into the management center are assigned the default user role you configure in Step 4 of Configure User Role Mapping for Azure in the Management Center, on page 178.

Configure User Role Mapping for Azure in the Management Center

The fields to configure for user role mapping at the 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 existing Azure users and groups; see Review the Azure Tenant, on page 175.
- Configure an SSO service provider application for the management center; see Configure the Management Center Service Provider Application for Azure, on page 175.
- Enable and configure single sign-on at the management center; see Enable Single Sign-On at the Management Center, on page 147, and Configure the Management Center for Azure 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 management center user role to assign users.

Step 4 In the **Group Member Attribute** field, enter an attribute configured in Azure for management center user role mapping for users or groups. See Step 1 of Configure User Role Mapping for Individual Users at the Azure IdP, on page 180 or Step 1 of Configure User Role Mapping for Groups at the Azure IdP, on page 181.

Step 5 Next to each management center user roll you wish to assign to SSO users, enter a regular expression. The management center compares these values against the user role mapping attribute the IdP sends to the management center with SSO user information. The 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 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 User Role Mapping in the Azure IdP, on page 179.

Configure User Role Mapping in the Azure IdP

You can configure SSO user role mapping at the Azure AD Portal based on individual user permissions or based on group permissions.

- To map based on individual user permissions, see Configure User Role Mapping for Individual Users at the Azure IdP.
- To map based on group permissions, see Configure User Role Mapping for Groups at the Azure IdP.

When an SSO user logs in to the management center, Azure presents a user or group role attribute value that gets its value from an application role configured at the Azure AD portal, to the management center. The management center then compares that attribute value with the regular expressions assigned to each 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 management center grants the user a configurable default user role.) The expression you assign to each management center user role must comply with the restricted version of Google's RE2 regular

expression standard supported by Golang and Perl. The management center treats the attribute value received from Okta as a regular expression using that same standard for purposes of comparison with the management center user role expressions.

Note A single management center cannot support role mapping for both groups and individual users; you must choose one mapping method for the management center service provider application and use it consistently. The management center can support role mapping using only one claim configured in Azure. Generally group-based role mapping is more efficient for a management center with many users. You should take into account user and group definitions established throughout your Azure tenant.

Configure User Role Mapping for Individual Users at the Azure IdP

To establish role mapping for individual users of the 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 Review the Azure Tenant, on page 175.
- Create and configure a management center service provider application in Azure; see Configure the Management Center Service Provider Application for Azure, on page 175.
- Configure SSO user role mapping as described in Configure User Role Mapping for Azure in the Management Center, on page 178.

Procedure

Step 1 Add a user claim to the SSO configuration for the management center service application with the following characteristics:

- **Name**: Use the same string you entered for the **Group Member Attribute** in the management center SSO configuration. (See Step 5 in Configure User Role Mapping for Azure in the Management Center, on page 178.)
- **Name identifier format**: Choose `Persistent`.
- **Source**: Choose `Attribute`.
- **Source attribute**: Choose `user.assignedroles`.

Step 2 Edit the manifest for the management center service application (in JSON format) and add application roles to represent 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 management center user roles. (Note: Azure does not permit spaces in this string.)

Step 3 For each user assigned to the 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 management center using SSO, the application role you assign to that user is the value Azure sends to the management center in the claim for the service application. The management center compares the claim against the expressions you assigned to management center user roles in the SSO configuration (See Step 6 of Configure User Role Mapping for Azure in the Management Center, on page 178.), and assigns the user all the management center user roles for which there is a match.

What to do next

- Test your role mapping scheme by logging into the management center using SSO from various accounts and confirming that users are assigned 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 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 Review the Azure Tenant, on page 175.
- Create and configure a management center service provider application in Azure; see Configure the Management Center Service Provider Application for Azure, on page 175.
- Configure SSO user role mapping as described in Configure User Role Mapping for Azure in the Management Center, on page 178.

Procedure

Step 1 Add a user claim to the SSO configuration for the management center service application with the following characteristics:

- **Name**: Use the same string you entered for the **Group Member Attribute** in the management center SSO configuration. (See Step 5 in Configure User Role Mapping for Azure in the Management Center, on page 178.)
- **Name identifier format**: Choose `Persistent`.
- **Source**: Choose `Attribute`.
- **Source attribute**: Choose `user.assignedroles`.

Step 2 Edit the manifest for the management center service application (in JSON format) and add application roles to represent 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 management center user roles. (Azure does not permit spaces in this string.)

Step 3 For each group assigned to the 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 management center using SSO, the application role you assign to that user's group is the value Azure sends to the management center in the claim for the service application. The management center compares the claim against the expressions you assigned to management center user roles in the SSO configuration (see Step 6 of Configure User Role Mapping for Azure in the Management Center, on page 178), and assigns the user all the management center user roles for which there is a match.

What to do next

Test your role mapping scheme by logging into the management center using SSO from various accounts and confirming that users are assigned management center user roles as you expect.

Azure User Role Mapping Examples

As the following examples demonstrate, the SSO configurations at the 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 management center service provider application in Azure.

Note You can configure management center roles to be mapped based on individual permissions or based on group permissions, but a single management center application cannot support role mapping for both groups and individual users. The 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 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 management center using SSO. Azure passes the claim value to the management center during the SSO login process, and the management center compares the claim value against strings assigned to each management center user role in the management center SSO configuration.

The following diagrams illustrate how the relevant fields and values in the 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 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 management center.

- In this diagram sue@ example.com uses the assignedroles attribute value `FMCAdmin`, and the management center assigns her the management center Administrator role.

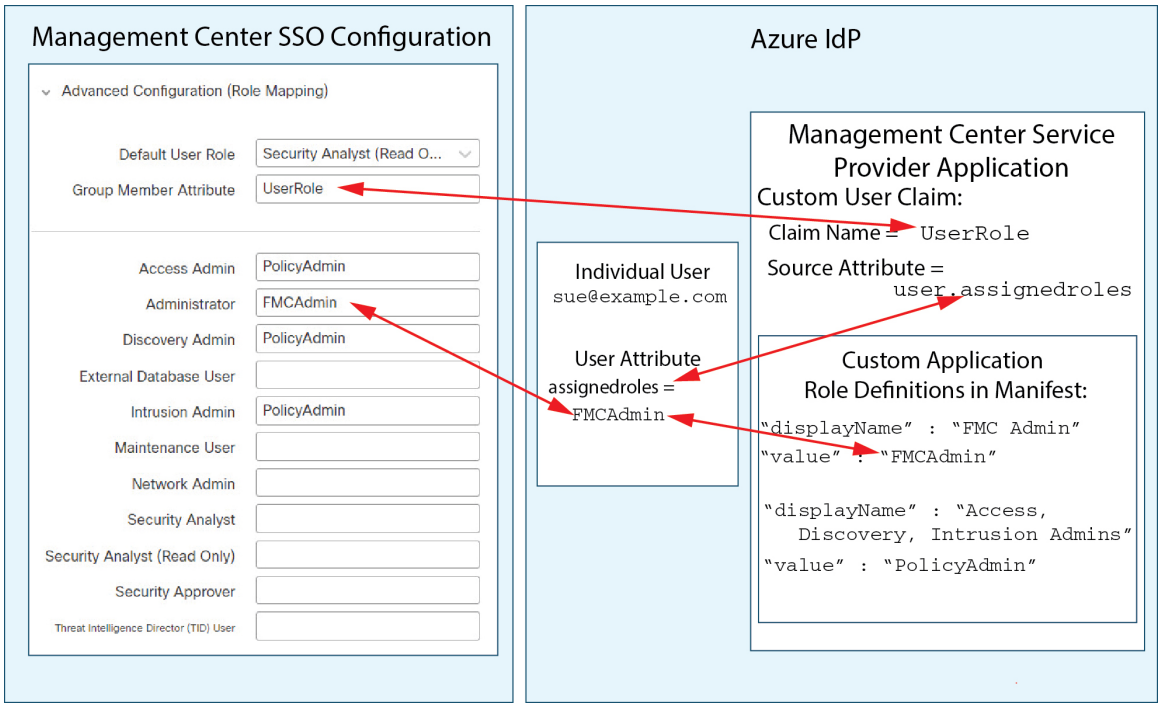

- In this diagram fred @ example .com uses the assignedroles attribute value `PolicyAdmin`, and the management center assigns him the roles Access Admin, Discovery Admin, and Intrusion Admin.

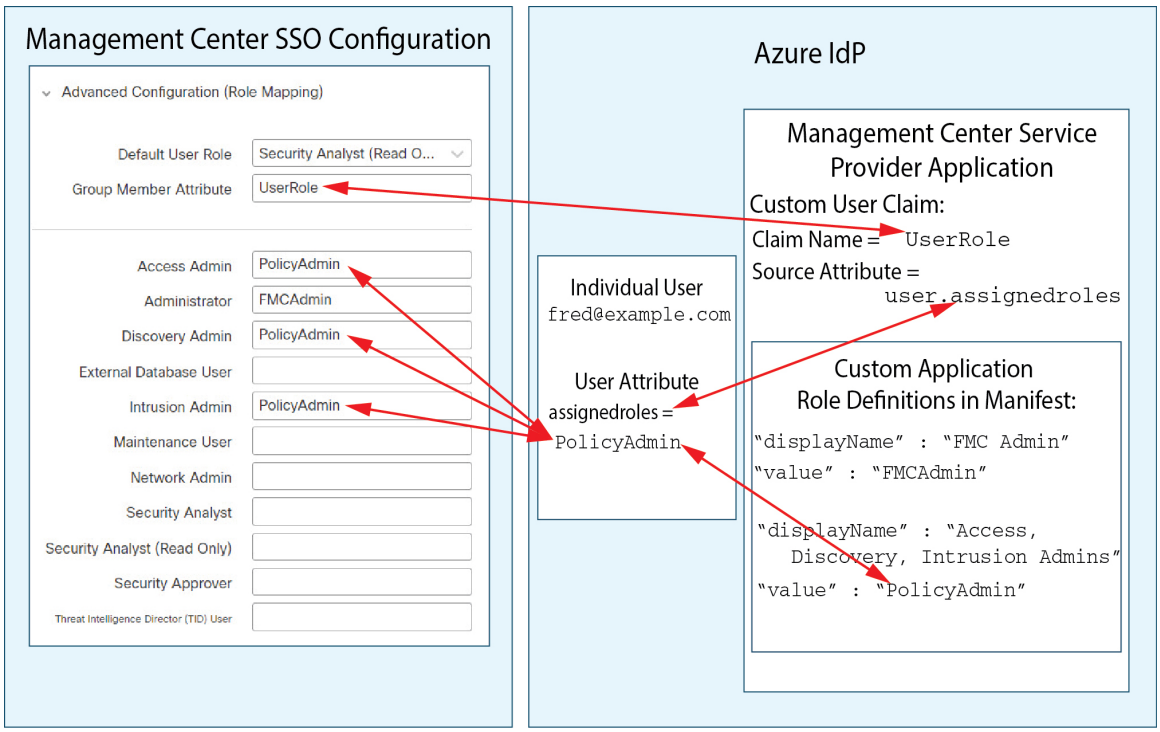

- Other users assigned to the Azure service application for this management center are assigned the default user role Security Analyst (Read Only) for one of the following reasons:

- 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 management center.

Azure Role Mapping Example for Groups

In role mapping for groups, the Azure management center service application has custom roles defined within its manifest. (In this case, FMCAdmin, AccessAdmin, Discovery Admin, and Maint.) These roles can be assigned to groups; Azure passes role assignments for each group to group members' 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 management center using SSO. Azure passes the claim value to the management center during the SSO login process, and the management center compares the claim value against strings assigned to each management center user role in the management center SSO configuration.

The following diagrams illustrate how the relevant fields and values in the management center and Azure configurations correspond to each other in user role mapping for groups. Each diagram uses the same SSO configurations at the 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 management center.

- In this diagram sue@example.com is a member of the groups `FMCAccessAdmins` and `FMCDiscoveryAdmins`. From these groups she inherits the custom roles `AccessAdmin` and `DiscoveryAdmin`. When Sue logs into the management center using SSO the management center assigns her the roles Access Admin and Discovery Admin.

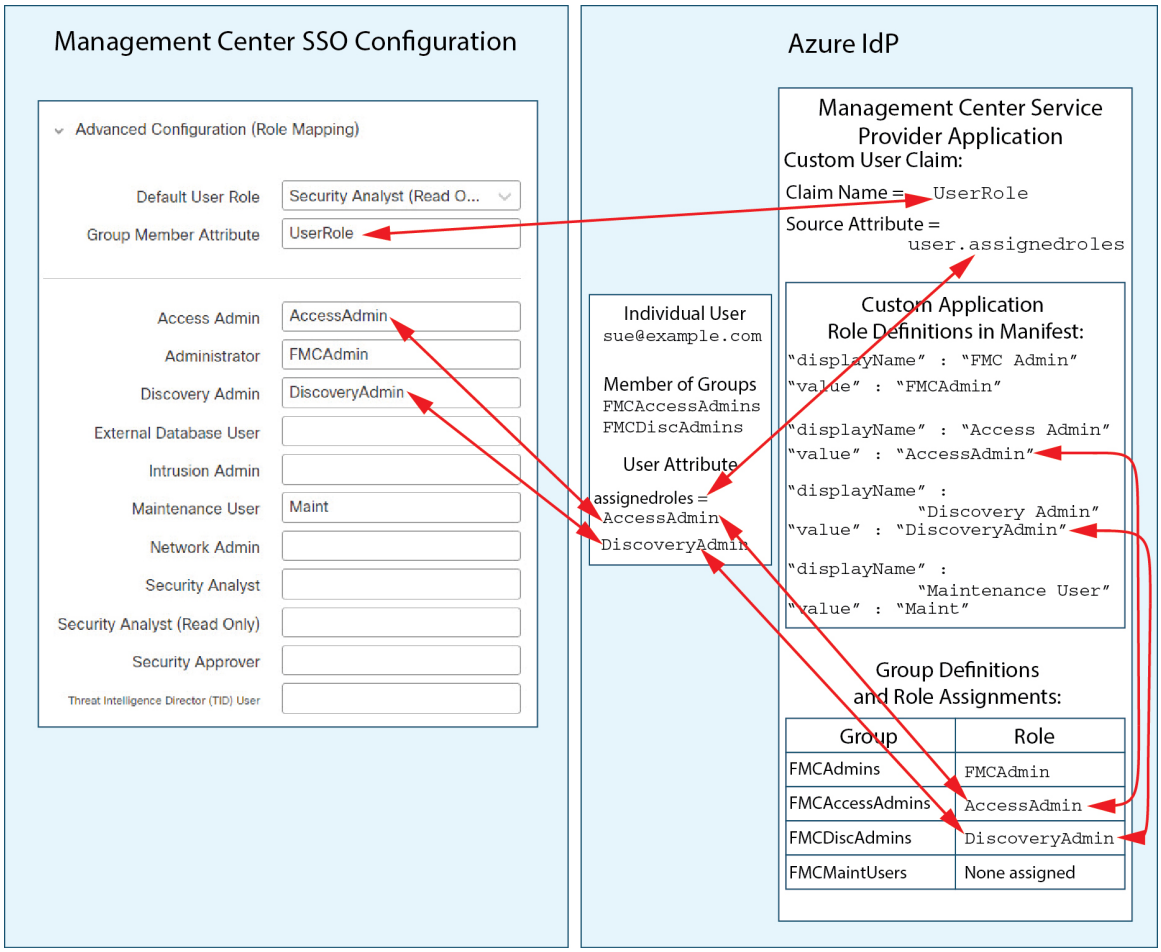

- In this diagram fred@example.com is a member of the `FMCAdmins` group, from which he inherits the custom role `FMCAdmin`. When Fred logs into the management center using SSO the management center assigns him the Administrator role.

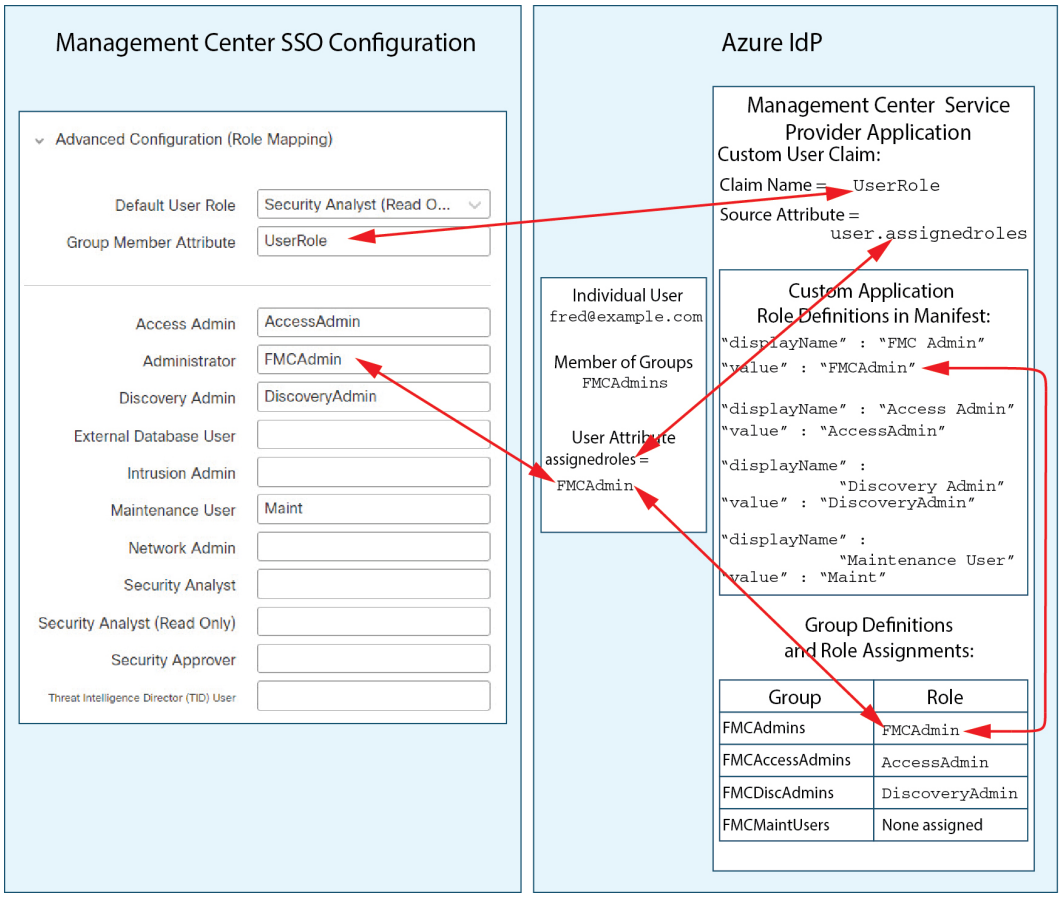

- In this diagram sean@example.com is a member of the `FMCMaintUsers` group, but because no custom role has been assigned to `FMCMaintUsers` within the Azure management center service provider application, Sean has no roles assigned to him, and when he logs into the management center using SSO, the management center assigns him the default role Security Analyst (Read Only).

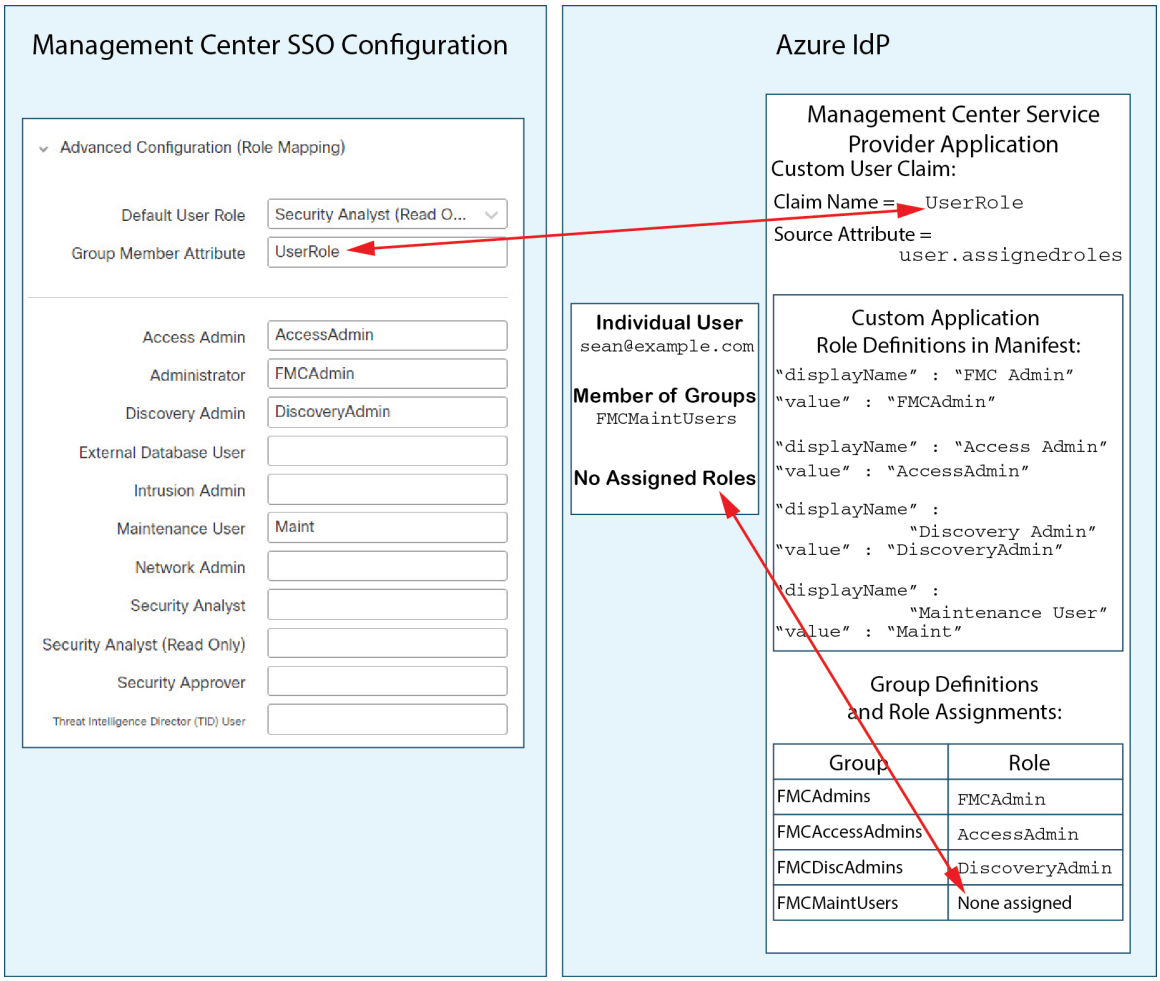

Configure Single Sign-On with PingID

See the following tasks to configure SSO using PingID's PingOne for Customers product:

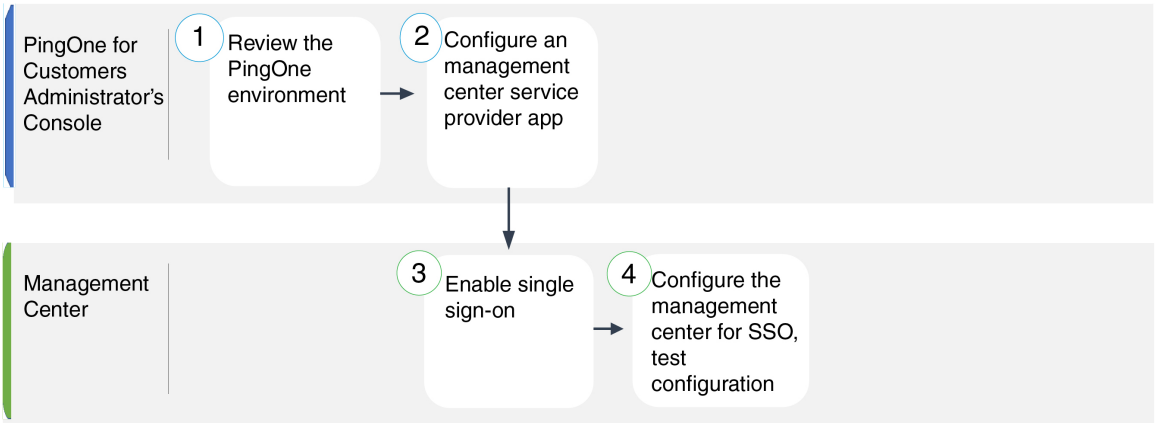

1	PingOne for Customers Administrator's Console	Review the PingID PingOne for Customers Environment, on page 188.
2	PingOne for Customers Administrator's Console	Configure the Management Center Service Provider Application for PingID PingOne for Customers, on page 188.
3	management center	Enable Single Sign-On at the Management Center, on page 147.
4	management center	Configure the Management Center for SSO with the PingID PingOne for Customers, on page 190.

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 management center to a PingOne environment, be familiar with its organization; consider the following questions:

- How many users will have access to the management center?
- Do you need to add more users to support SSO access to the 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 the Management Center Service Provider Application for PingID PingOne for Customers

Use the PingOne for Customers Administrator Console to create a management center service provider application within your PingOne for Customers environment and establish basic configuration settings. This documentation does not describe all the PingOne for Customers functions you need to establish a fully functional SSO environment; for instance, to create users see 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 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 management center.

- Confirm the login URL for the target management center (`https://ipaddress_or_hostname`)

Note If your 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 management center using the login URL that you configure in this task.

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 the following settings for the SAML Connection:

- For the **ACS URL**, append the string `/sam/acs` to the 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 the string `/saml/metadata` to the 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 the following 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 management center web interface.

Step 4 For **SAML ATTRIBUTES**, make the following 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 management center easier, you can download the SAML XML metadata file for the management center service provider application to your local computer.

Step 7 Enable the application.

What to do next

Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Configure the Management Center for SSO with the PingID PingOne for Customers

Before you begin

- Create a management center service provider application at the PingOne for Customers Administrator Console; see Configure the Management Center Service Provider Application for PingID PingOne for Customers, on page 188.
- Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Procedure

Step 1 (This step continues directly from Enable Single Sign-On at the Management Center, on page 147.) 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 the Management Center Service Provider Application for PingID PingOne for Customers, on page 188.
     - For **Identity Provider Issuer** enter the **Issuer ID** you noted in Step 3 of Configure the Management Center Service Provider Application for PingID PingOne for Customers, on page 188.
     - For the **X.509 Certificate**, use the certificate you downloaded from PingOne for Customers in Step 5 of Configure the Management Center Service Provider Application for PingID PingOne for Customers, on page 188. (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 the Management Center Service Provider Application for PingID PingOne for Customers, on page 188), click the **Upload XML File** radio button to upload the file to the management center:

Step 2 Click **Next**.

Step 3 At the **Verify Metadata** dialog, review the configuration parameters and click **Save**.

Step 4 Expand **Advanced Configuration (Role Mapping)**.

Step 5 From the **Default User Role** drop-down list, choose a default management center user role to assign users.

Step 6 In the **Group Member Attribute** field, enter an attribute configured in PingID PingOne for management center user role mapping for users or groups.

Step 7 Next to each management center user role you wish to assign to SSO users, enter a regular expression. The management center compares these values against the user role mapping attribute the IdP sends to the management center with SSO user information. The management center grants users a union of all the roles for which a match is found.

Step 8 Click **Test Configuration**. If the System displays an error message, review the SSO configuration for the management center as well as the PingOne for Customers service provider application, correct any errors, and try again.

Step 9 When the system reports a successful configuration test, click **Apply**.

Configure Single Sign-On with Any SAML 2.0-Compliant SSO Provider

The 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 management center for single sign-on using any SAML 2.0-compliant SSO provider:

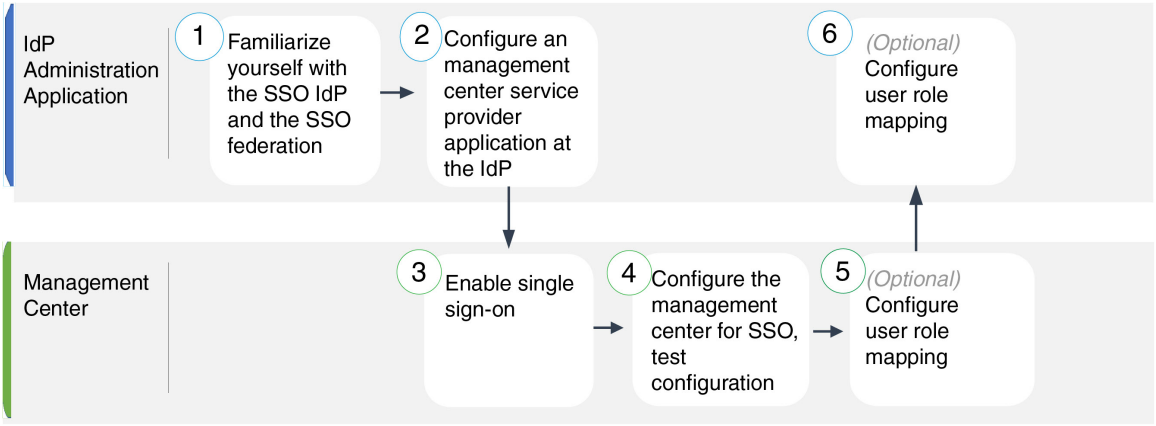

1	IdP Administration Application	Familiarize Yourself with the SSO Identity Provider and the SSO Federation, on page 192.
2	IdP Administration Application	Configure Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider, on page 192.
3	management center	Enable Single Sign-On at the Management Center, on page 147.
4	management center	Configure the Management Center for SSO Using Any SAML 2.0-Compliant SSO Provider, on page 194.

5	management center	Configure User Role Mapping in the Management Center for SAML 2.0-Compliant SSO Providers, on page 195.
6	IdP Administration Application	Configure Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 197.

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 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 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 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 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 Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider

Generally SSO providers require that you 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, but some IdP's automatically generate some configuration settings for you, while others require that you configure all settings yourself.

Note If you plan to assign user groups to the management center Application, do not also assign users within those groups as individuals.

Note The 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 management center.

Before you begin

- Familiarize yourself with the SSO federation and its users and groups; see Familiarize Yourself with the SSO Identity Provider and the SSO Federation, on page 192.
- Confirm your IdP account has the necessary permissions to perform this task.
- Create user accounts and/or groups 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 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 management center.

- Confirm the login URL for the target management center (`https://ipaddress_or_hostname`)

Note If your management center web interface can be reached with multiple URLs. (for instance, a full-qualified domain name as well as an IP address), SSO users must consistently access the management center using the login URL that you configure in this task.

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 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 management center), formatted as a URL. To create this, append the string `/saml/metadata` to the management center login URL, such as `https://ExampleFMC/saml/metadata`.
- Single Sign on URL, Recipient URL, Assertion Consumer Service URL: The service provider (management center) address to which the browser sends information on behalf of the IdP. To create this, append the string `saml/acs` to the management center login URL, such as `https://ExampleFMC/saml/acs`.
- X.509 Certificate: Certificate to secure communications between the management center and the IdP. Some IdP's 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 management center application. (If you plan to assign groups to the 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 management center application.

Step 5 (Optional) Some IdP's 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 management center.

What to do next

Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Configure the Management Center for SSO Using Any SAML 2.0-Compliant SSO Provider

Use these instructions at the management center web interface. To configure the management center for SSO using any SAML 2.0-compliant SSO provider, you need information from the IdP.

Before you begin

- Review the organization of your SSO federation, and its users and groups.
- Configure a management center service provider application at the IdP; see Configure the Management Center for SSO Using Any SAML 2.0-Compliant SSO Provider, on page 194.
- 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 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 management center and the IdP.
- Enable single sign-on; see Enable Single Sign-On at the Management Center, on page 147.

Procedure

Step 1 (This step continues directly from Enable Single Sign-On at the Management Center, on page 147.) 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 an the XML metadata file generated at the IdP (Step 5 in Configure Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider, on page 192), you can upload the file to the 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 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**.

What to do next

You may optionally configure user role mapping for SSO users; see Configure User Role Mapping in the Management Center for SAML 2.0-Compliant SSO Providers, on page 195. If you choose not to configure role mapping, by default all SSO users that log into the management center are assigned the default user role you configure in Step 4 of Configure User Role Mapping in the Management Center for SAML 2.0-Compliant SSO Providers, on page 195.

Configure User Role Mapping in the 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 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 management center once it has authenticated and authorized an SSO user.

- At the management center, associate values with each of the 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 management center, the management center compares the attribute value with the values associated with each management center user role, assigns the user all the roles that match, . The 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 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 management center; see Configure Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider, on page 192.
- Enable and configure single sign-on at the management center, see Enable Single Sign-On at the Management Center, on page 147, and Configure the Management Center for SSO Using Any SAML 2.0-Compliant SSO Provider, on page 194.

Procedure

Step 1 Choose **System** (⚙) > **Users** > **Single Sign-On**.

Step 2 Expand **Advanced Configuration (Role Mapping)**.

Step 3 Select a 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 management center service provider application for user role mapping using either users or groups. (See Step 1 of Configure Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 197.)

Step 5 Next to each management center user roll you wish to assign to SSO users, enter a regular expression. The management center compares these values against the user role mapping attribute the IdP sends to the management center with SSO user information. The 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 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 Management Center User Role Mapping at the IdP for SAML 2.0-Compliant SSO Providers, on page 197.

Configure 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 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 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 management center.
- Group role mapping is generally more efficient for a management center with many users.
- If you assign user groups to management center applications, do not also assign users within those groups as individuals.
- For the purpose of determining a match with management center user roles, the 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 management center service provider application at the IdP (see Configure Management Center Service Provider Application for Any SAML 2.0-Compliant SSO Provider, on page 192).

Procedure

Step 1 At the IdP, create or designate an attribute to be sent to the 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 management center SSO configuration.

Customize User Roles for the Web Interface

Each user account must be defined with a user role. This section describes how to manage user roles and how to configure a custom user role for web interface access. For default user roles, see User Roles, on page 118.

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 management center.

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 management center:
  1. **a.** On the other management center, click the **Export**() to save the role to your computer.
  2. **b.** On the new management center, choose **System** (⚙) > **Tools** > **Import/Export**.
  3. **c.** Click **Upload Package**, then follow the instructions to import the saved user role to the new 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 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 200.

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 User Configuration, on page 107.

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 5: 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

You can give custom user roles the permission, with a password, to temporarily gain the privileges of another, targeted user role in addition to those of the base role. This feature allows you to easily substitute one user for another during an absence, or to more closely track the use of advanced user privileges. Default user roles do not support escalation.

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

To configure user role escalation, see the following workflow.

Procedure

Step 1 Set the Escalation Target Role, on page 200. Only one user role at a time can be the escalation target role.

Step 2 Configure a Custom User Role for Escalation, on page 201.

Step 3 (For the logged in user) Escalate Your User Role, on page 201.

Set the Escalation Target Role

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

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.

Configure a Custom User Role for Escalation

Users for whom you want to enable escalation must belong to a custom user role with escalation enabled. This procedure describes how to enable escalation for a custom user role.

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

Before you begin

Set a target user role according to Set the Escalation Target Role, on page 200.

Procedure

Step 1 Begin configuring your custom user role as described in Create Custom User Roles, on page 198.

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**.

Escalate Your User Role

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.

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.

Troubleshooting LDAP Authentication Connections

If you create an LDAP authentication object and it either does not succeed in connecting to the server you select or does not retrieve the list of users you want, you can tune the settings in the object.

If the connection fails when you test it, try the following suggestions to troubleshoot your configuration:

- Use the messages displayed at the top of the web interface screen and in the test output to determine which areas of the object are causing the issue.
- Check that the user name and password you used for the object are valid:
  - Check that you have the rights to browse to the directory indicated in your base-distinguished name by connecting to the LDAP server using a third-party LDAP browser.
  - Check that the user name is unique to the directory information tree for the LDAP server.
  - If you see an LDAP bind error 49 in the test output, the user binding for the user failed. Try authenticating to the server through a third-party application to see if the binding fails through that connection as well.
- Check that you have correctly identified the server:
  - Check that the server IP address or host name is correct.
  - Check that you have TCP/IP access from your local appliance to the authentication server where you want to connect.
  - Check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.
  - If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used for the server.
  - Check that you have not used an IPv6 address for the server connection if you are authenticating CLI access.
  - If you used server type defaults, check that you have the correct server type and click **Set Defaults** again to reset the default values.

- If you typed in your base-distinguished name, click **Fetch DNs** to retrieve all the available base distinguished names on the server, and select the name from the list.
- If you are using any filters, access attributes, or advanced settings, check that each is valid and typed correctly.
- If you are using any filters, access attributes, or advanced settings, try removing each setting and testing the object without it.
- If you are using a base filter or a CLI access filter, make sure that the filter is enclosed in parentheses and that you are using a valid comparison operator (maximum 450 characters, including the enclosing parentheses).
- To test a more restricted base filter, try setting it to the base distinguished name for the user to retrieve just that user.
- If you are using an encrypted connection:
  - Check that the name of the LDAP server in the certificate matches the host name that you use to connect.
  - Check that you have not used an IPv6 address with an encrypted server connection.
- If you are using a test user, make sure that the user name and password are typed correctly.
- If you are using a test user, remove the user credentials and test the object.
- Test the query that you are using by connecting to the LDAP server and using this syntax:

```
ldapsearch -x -b 'base_distinguished_name'
-h LDAPserver_ip_address -p port -v -D
'user_distinguished_name' -W 'base_filter'
```

  For example, if you are trying to connect to the security domain on `myrtle.example.com` using the `domainadmin@myrtle.example.com` user and a base filter of (`cn=*`), you could test the connection using this statement:

```
ldapsearch -x -b 'CN=security,DC=myrtle,DC=example,DC=com'
-h myrtle.example.com -p 389 -v -D
'domainadmin@myrtle.example.com' -W '(cn=*)'
```

If you can test your connection successfully but authentication does not work after you deploy a platform settings policy, check that authentication and the object you want to use are both enabled in the platform settings policy that is applied to the device.

If you connect successfully but want to adjust the list of users retrieved by your connection, you can add or change a base filter or CLI access filter or use a more restrictive or less restrictive base DN.

While authenticating a connection to Active Directory (AD) server, rarely the connection event log indicates blocked LDAP traffic although the connection to AD server is successful. This incorrect connection log occurs when the AD server sends a duplicate reset packet. The threat defense device identifies the second reset packet as part of a new connection request and logs the connection with Block action.

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.

Changing Your Password

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

When password strength checking is enabled, passwords must comply with the strong password requirements described in Guidelines and Limitations for User Accounts for Management Center, on page 124.

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

Procedure

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

Step 2 Click **Change Password**.

Step 3 Optionally, check the **Show password** check box to see the password while using this 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 have the system create a password for you which complies with the listed criteria. (Generated passwords are non-mnemonic; take careful note of the password if you choose this option.)

Step 6 Click **Apply**.

Changing an Expired Password

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

Procedure

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

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

Tip
When password strength checking is enabled, passwords must comply with the strong password requirements described in Guidelines and Limitations for User Accounts for Management Center, on page 124.

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

Change the Web Interface Appearance

You can change the way the web interface appears.

Procedure

From the drop-down list under your user name, choose a theme:

- **Light**
- **Dusk**
- **Classic** (the look and feel before Version 6.6)

Specifying Your Home Page

You can specify the page within the web interface to use as your home page for the appliance. The default home page is the default dashboard (**Overview** > **Dashboards** > **Dashboard**), except for user accounts with no dashboard access, such as External Database users. (See Specifying Your Default Dashboard, on page 210 to set the default dashboard.)

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

Procedure

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

Step 2 Click **Home Page**.

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

The options in the drop-down list are based on the access privileges for your user account. For more information, see User Roles, on page 118.

Step 4 Click **Save**.

Configuring Event View Settings

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

Procedure

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

Step 2 Click **Event View Settings**.

Step 3 In the **Event Preferences** section, configure the basic characteristics of event views; see Event View Preferences, on page 206.

Step 4 In the **File Preferences** section, configure file download preferences; see File Download Preferences, on page 207.

Step 5 In the **Default Time Windows** section, configure the default time window or windows; see Default Time Windows, on page 208.

Step 6 In the **Default Workflow** sections, configure default workflows; see Default Workflows, on page 209.

Step 7 Click **Save**.

Event View Preferences

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

The following fields appear in the Event Preferences section:

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

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

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

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

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

- **All** - expand all sections

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

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

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

File Download Preferences

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

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

The following fields appear in the File Preferences section:

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

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

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

- When you download a captured file, the system creates a password-protected .zip archive containing the file. The **Zip File Password** field defines the password you want to use to restrict access to the .zip file. If you leave this field blank, the system creates archive files without passwords.

- The **Show Zip File Password** check box toggles displaying plain text or obfuscated characters in the **Zip File Password** field. When this field is cleared, the **Zip File Password** displays obfuscated characters.

Default Time Windows

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

User role access to this section is as follows:

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

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

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

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

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

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

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

There are three types of time window:

- *static*, which displays all the events generated from a specific start time to a specific end time
- *expanding*, which displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view
- *sliding*, which displays all the events generated from a specific start time (for example, one day ago) to the present; as time moves forward, the time window “slides” so that you see only the events for the range you configured (in this example, for the last day)

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

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

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

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

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

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

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

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

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

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

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

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

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

Default Workflows

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

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

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

Setting Your Default Time Zone

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

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

Note 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**.

Procedure

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

Step 2 Click **Time Zone** drop-down.

Step 3 Choose the continent or country and the state name that corresponds with the time zone you want to use.

Specifying Your Default Dashboard

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

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

Procedure

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

Step 2 Click **Dashboard Settings**.

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

Step 4 Click **Save**.

Configure How-To Settings

How To is a widget that provides walkthroughs to navigate through tasks on the management center. The walkthroughs guide you to perform the steps required to achieve a task by taking you through each step, one after the other irrespective of the various UI screens that you may have to navigate, to complete the task. The **How To** widget is enabled by default.

For a list of feature walkthroughs supported in the management center, see Feature Walkthroughs Supported in Secure Firewall Management Center.

Note

- The walkthroughs are generally available for all UI pages, and are not user role sensitive. However, depending on the privileges of the user, some of the menu items will not appear on the management center interface. Thereby, the walkthroughs will not execute on such pages.
- This feature is not available in the Classic theme.

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**.

What to do next

To open the How To widget, choose **Help** > **How-Tos**. You can search for How-To walkthroughs that address tasks of interest. For more information, see Search for How To Walkthroughs, on page 20.

History for Management Center User Accounts

Feature	Minimum Management Center	Minimum Threat Defense	Details
Added new field for assigning Shell user name template.	7.0.0	Any	Provision to specify a template for CLI access attributes for LDAP external authentication—**Shell User Name Template** was introduced. Thus, CLI attribute would have its own template to identify the LDAP CLI users. New/modified screens: **System**(⚙) > **Users** > **External Authentication**
Added support of Single Sign-On using 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 includes the ability to map user or group roles from the IdP to management center user roles. Only users with the Admin role authenticated internally or by LDAP or RADIUS can configure SSO. New/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. Choose the Light or Dusk theme, or use the Classic theme that appeared in previous releases. New/modified screens: User Name > **User Preferences** > **General** > **UI Theme**
Added a new field for name in user accounts.	6.6.0	Any	Added a field that can identify the user or department responsible for an internal user account. New/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 management center and Cisco Security Manager is no longer supported as of Firepower 6.5. New/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. New fields in the change password interface added: **Show Password** and **Generate Password**. New/modified screens: User Name > **User Preferences** > **General** > **Change Password**

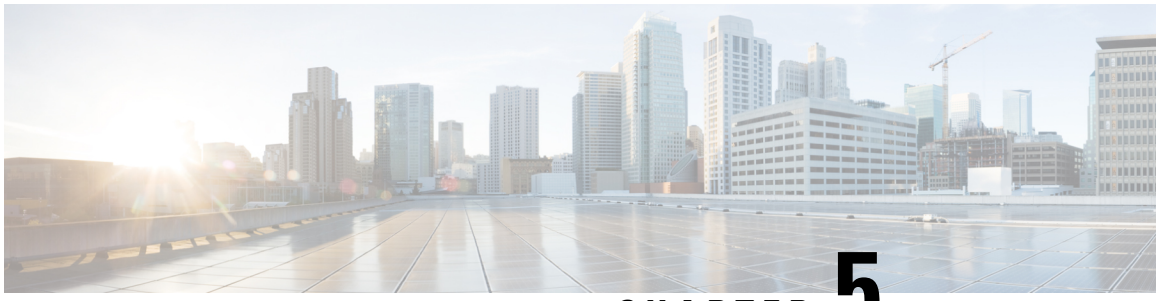

CHAPTER 5

Domains

The following topics describe how to manage multitenancy using domains:

Introduction to Multitenancy Using Domains

The management center allows you to implement multitenancy using *domains*. Domains segment user access to managed devices, configurations, and events. You can create up to 100 subdomains under a top-level Global domain, in two or three levels.

When you log into the management center, you log into a single domain, called the *current domain*. Depending on your user account, you may be able to switch to other domains.

In addition to any restrictions imposed by your user role, your current domain level can also limit your ability to modify various configurations. The management center limits most management tasks, like system software updates, to the Global domain.

The 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.

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 management center to manage network security for multiple customers:

- Administrators at the MSSP logging into the Global domain, cannot view or edit customers' deployments. They must log into respective second-level named subdomains to manage the customers' deployment.
- Administrators for each customer can log into second-level named subdomains to manage only the devices, configurations, and events applicable to their organizations. These local administrators cannot view or affect the deployments of other customers of the MSSP.

Three Domain Levels: Global, Second-Level, and Third-Level

In a three-level multidomain deployment, the Global domain has subdomains, at least one of which has its own subdomain. To extend the previous example, consider a scenario where an MSSP customer—already restricted to a subdomain—wants to further segment its deployment. This customer wants to separately manage two classes of device: devices placed on network edges and devices placed internally:

- Administrators for the customer logging into the second-level subdomain cannot view or edit the customer's edge network deployments. They must log into the respective leaf domain to manage the devices deployed on the network edge.
- Administrators for the customer's edge network can log into a third-level (leaf) domain to manage only the devices, configurations, and events applicable to devices deployed on the network edge. Similarly, administrators for the customer's internal network can log into a different third-level domain to manage internal devices, configurations, and events. Edge and internal administrators cannot view each other's deployment.

Note In the 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.

Related Topics

Configure SAML Single Sign-On, on page 144

Domains Terminology

This documentation uses the following terms when describing domains and multidomain deployments:

Global Domain

In a multidomain deployment, the top-level domain. If you do not configure multitenancy, all devices, configurations, and events belong to the Global domain. Administrators in the Global domain can manage the entire 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.

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 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 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 Limit
Network Discovery Data Storage Settings

Requirements and Prerequisites for Domains

Model Support

Any.

Supported Domains

Any

User Roles

- Admin

Managing Domains

To modify a domain's properties, you must have Administrator access in that domain's parent domain.

Procedure

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

Step 2 Manage your domains:

- Add — Click **Add Domain**, or click **Add Subdomain** next to the parent domain; see Creating New Domains, on page 217.
- Edit — Click **Edit** (✎) next to the domain you want to modify; see Domain Properties, on page 215.
- Delete — Click **Delete** (🗑) next to the empty domain you want to delete, then confirm your choice. Move devices from domains you want to delete by editing their destination domain.

Step 3 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 4 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Moving Data Between Domains, on page 218.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Moving Devices Between Domains, on page 219.

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.

Creating New Domains

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 215.

Step 7 Click **Save** to return to the domain management page.

The system warns you if any devices are assigned to non-leaf domains. Click **Create New Domain** to create a new domain for those devices. Click **Keep Unassigned** if you plan to move the devices to existing domains.

Step 8 When you are done making changes to the domain structure and all devices are associated with leaf domains, click **Save** to implement your changes.

Step 9 If prompted, make additional changes:

- If you changed a leaf domain to a parent domain, move or delete the old network map; see Moving Data Between Domains, on page 218.
- If you moved devices between domains and must assign new policies and security zones or interface groups, see Moving Devices Between Domains, on page 219.

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.

Moving Data Between Domains

Because events and network maps are associated with leaf domains, when you change a leaf domain to a parent domain, you have two choices:

- Move the network map and associated events to a new leaf domain.
- Delete the network map but retain the events. In this case, the events remain associated with the parent domain until the system prunes events as needed or as configured. Or, you can delete old events manually.

Before you begin

Implement a domain configuration where a former leaf domain is now a parent domain; see Managing Domains, on page 216.

Procedure

Step 1 For each former leaf domain that is now a parent domain:

- Choose a new **Leaf Domain** to inherit the **Parent Domain**'s events and network map.
- Choose **None** to delete the parent domain's network map, but retain old events.

Step 2 Click **Save**.

What to do next

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

Moving Devices Between Domains

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 site-to-site VPN configurations are in the target domain.
  - 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, see Creating New Domains, on page 217.
- Implement a domain configuration where you moved a device from domain to domain and now must assign new policies and security zones; see Managing Domains, on page 216.

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 the following:

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 the following:

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.

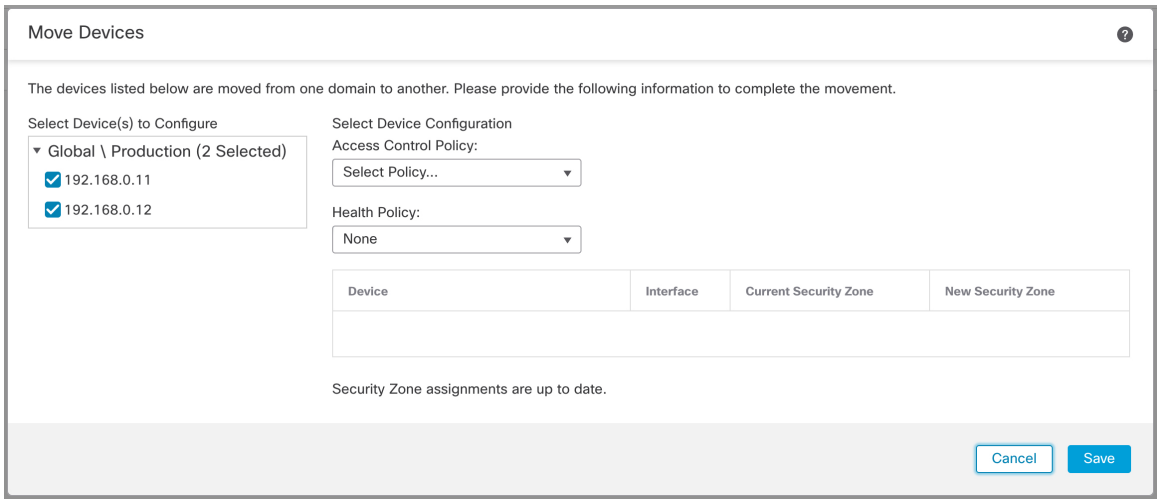

b. Choose an **Access Control Policy** to apply to the device, or choose **New Policy** to create a new policy.

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.

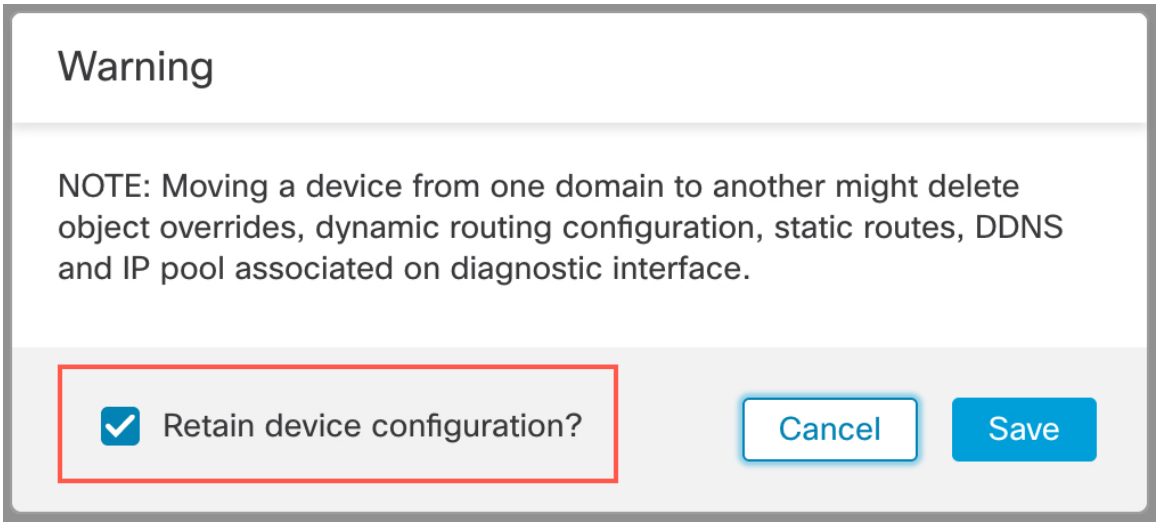

If you select this option, the system retains the device configurations after the devices are moved to the target domain. If you do not select this option, you must manually update the device configurations on the moved device that were affected by the move.

The following table shows how objects are handled in various scenarios.

Scenario	System Action
Objects exist in the target domain.	Reuse the objects.

Scenario	System Action
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**.

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, see *Export and Import the Device Configuration* in the Cisco Secure Firewall Management Center Device Configuration Guide.

History for Domain Management

Feature	Minimum Management Center	Minimum 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

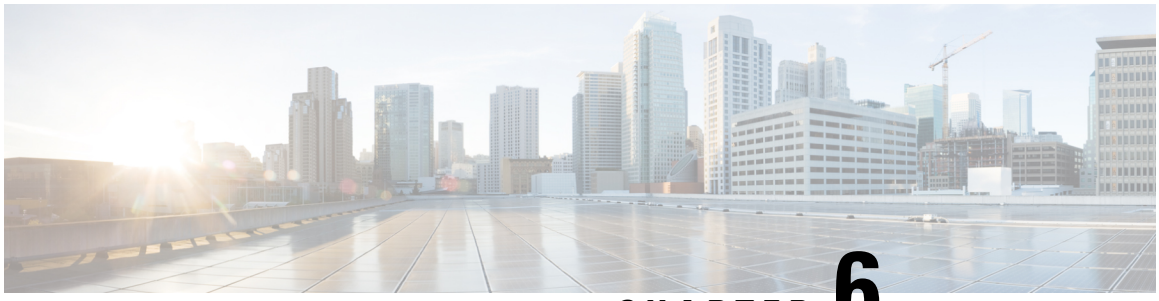

CHAPTER 6

Updates

This chapter explains how to perform content updates.

Important To upgrade the management center, or threat defense software or chassis, see the upgrade guide for the version that your *management center* is *currently* running: https://cisco.com/go/ftd-fmc-upgrade.

About System Updates

Use the management center to upgrade the system software for itself and the devices it manages. You can also update various databases and feeds that provide advanced services.

The system can obtain most updates from the internet. We recommend you schedule or enable automatic content updates whenever possible. Some updates are auto-enabled by the initial setup process or when you enable the related feature. Other updates you must schedule yourself. After initial setup, we recommend you review all auto-updates and adjust them if necessary.

Table 6: Upgrades and Updates

Component	Description	Details
System software	*Major* software releases contain new features, functionality, and enhancements. They may include infrastructure or architectural changes. *Maintenance* releases contain general bug and security related fixes. Behavior changes are rare, and are related to those fixes. *Patches* are on-demand updates limited to critical fixes with time urgency. *Hotfixes* can address specific customer issues.	**Direct Download:** Select patches and maintenance releases only, usually some time after the release is available for manual download. The length of the delay depends on release type, release adoption, and other factors. Both on-demand and scheduled downloads are supported. **Schedule Install:** Patches and maintenance releases only, as a scheduled task. **Uninstall:** Patches only. **Revert:** Major and maintenance releases for threat defense only. Revert is not supported for the management center or for Classic devices. **Reimage:** Major and maintenance releases only. **See:** Cisco Secure Firewall Threat Defense Upgrade Guide for Management Center
Vulnerability database (VDB)	The Cisco vulnerability database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.	**Direct Download:** Yes. **Schedule:** Yes, as a scheduled task. **Uninstall:** Starting with VDB 357, you can install any VDB as far back as the baseline VDB for the management center. **See:** Update the Vulnerability Database (VDB), on page 226
Geolocation database (GeoDB)	The Cisco geolocation database (GeoDB) maps IP addresses to countries/continents.	**Direct Download:** Yes. **Schedule:** Yes, from its own update page **Uninstall:** No. **See:** Update the Geolocation Database (GeoDB), on page 228
Intrusion rules (SRU/LSP)	Intrusion rule updates provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates may also delete rules, provide new rule categories and default variables, and modify default variable values.	**Direct Download:** Yes. **Schedule:** Yes, from its own update page. **Uninstall:** No. **See:** Update Intrusion Rules, on page 229
Security Intelligence feeds	Security Intelligence feeds are collections of IP addresses, domain names, and URLs that you can use to quickly filter traffic that matches an entry.	**Direct Download:** Yes. **Schedule:** Yes, from the object manager. **Uninstall:** No. **See:** Cisco Secure Firewall Management Center Device Configuration Guide

Component	Description	Details
URL categories and reputations	URL filtering allows you to control access to websites based on the URL's general classification (category) and risk level (reputation).	**Direct Download:** Yes. **Schedule:** Yes, when you configure integrations/cloud services, or as a scheduled task. **Uninstall:** No. **See:** Cisco Secure Firewall Management Center Device Configuration Guide

Requirements and Prerequisites for System Updates

Model Support

Any

Supported Domains

Global unless indicated otherwise.

User Roles

Admin

Guidelines and Limitations for System Updates

Before You Update

Before you update any component of your deployment (including intrusion rules, VDB, or GeoDB) read the release notes or advisory text that accompanies the update. These provide critical and release-specific information, including compatibility, prerequisites, new capabilities, behavior changes, and warnings.

Scheduled Updates

The system schedules tasks — including updates — in UTC. This means that when they occur locally depends on the date and your specific location. Also, because updates are scheduled in UTC, they do not adjust for Daylight Saving Time, summer time, or any such seasonal adjustments that you may observe in your location. If you are affected, scheduled updates occur one hour "later" in the summer than in the winter, according to local time.

Important We *strongly* recommend you review scheduled updates to be sure they occur when you intend.

Bandwidth Guidelines

To upgrade a the system software or perform a readiness check, the upgrade package must be on the appliance. Upgrade package sizes vary. Make sure you have the bandwidth to perform a large data transfer to your managed devices. See Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).

Update the Vulnerability Database (VDB)

The Cisco vulnerability database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The system uses the VDB to help determine whether a particular host increases your risk of compromise.

Cisco issues periodic updates to the VDB. The time it takes to update the VDB and its associated mappings on the management center depends on the number of hosts in your network map. As a rule of thumb, divide the number of hosts by 1000 to determine the approximate number of minutes to perform the update.

The initial setup on the management center automatically downloads and installs the latest VDB from Cisco as a one-time operation. It also schedules a weekly task to download the latest available software updates, which includes the latest VDB. We recommend you review this weekly task and adjust if necessary. Optionally, schedule a new weekly task to actually update the VDB and deploy configurations. For more information, see Vulnerability Database Update Automation, on page 499.

For VDB 343+, all application detector information is available through Cisco Secure Firewall Application Detectors. This site includes a searchable database of application detectors. The release notes provide information on changes for a particular VDB release.

Schedule VDB Updates

If your management center has internet access, we recommend you schedule regular VDB updates. See Vulnerability Database Update Automation, on page 499.

Manually Update the VDB

Use this procedure to manually update the VDB. Starting with VDB 357, you can install any VDB as far back as the baseline VDB for the management center.

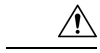

Caution Do not perform tasks related to mapped vulnerabilities while the VDB is updating. Even if the Message Center shows no progress for several minutes or indicates that the update has failed, do not restart the update. Instead, contact Cisco TAC.

In most cases, the first deploy after a VDB update restarts the Snort process, interrupting traffic inspection. The system warns you when this will happen (updated application detectors and operating system fingerprints require a restart; vulnerability information does not). Whether traffic drops or passes without further inspection during this interruption depends on how the targeted device handles traffic. For more information, see Snort Restart Traffic Behavior.

Before you begin

If the management center cannot access the internet or you are installing an older VDB, get the update yourself: https://www.cisco.com/go/firepower-software. Select or search for your model (or choose any model—you use the same VDB for all management centers), then browse to the *Coverage and Content Updates* page.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Content Updates**.

Step 2 Choose how you want to get the VDB onto the management center.

- Direct download: Click the **Download Updates** button.
- Manual upload: Click **Upload Update**, then **Choose File** and browse to the VDB. After you choose the file, 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) Next to the Vulnerability and Fingerprint Database update you want to install, click either the **Install** icon (for a newer VDB) or the **Rollback** icon (for an older VDB).
b) Choose the management center.
c) Click **Install**.

Monitor update progress in the Message Center. After the update completes, the system uses the new vulnerability information. However, you must deploy before updated application detectors and operating system fingerprints can take effect.

Step 4 Verify update success.

The VDB update page and **Help** (?) > **About** both show the current version.

What to do next

- Deploy configuration changes; see the Cisco Secure Firewall Management Center Device Configuration Guide.
- If you based configurations on vulnerabilities, application detectors, or fingerprints that are no longer available, examine those configurations to make sure you are handling traffic as expected. Also, keep in mind a scheduled task to update the VDB can undo a rollback. To avoid this, change the scheduled task or delete any newer VDB packages.

Update the Geolocation Database (GeoDB)

The geolocation database (GeoDB) is a database that you can leverage to view and filter traffic based on geographical location. We issue periodic updates to the GeoDB, and you must regularly update the GeoDB to have accurate geolocation information. As part of the initial configuration, the system schedules weekly GeoDB updates. We recommend you review this task and make changes if necessary, as described in Schedule GeoDB Updates, on page 228.

A GeoDB update overrides any previous versions. The management center automatically updates its managed devices, and unless the update adds new countries (this is rare) you do not need to redeploy. You can see your current version on **Help** (?) > **About**.

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.

Schedule GeoDB Updates

As part of the initial configuration, the system schedules weekly GeoDB updates. We recommend you review this task and make changes if necessary, as described in this procedure.

Note that the system does not automatically deploy after GeoDB updates because in most cases it is not necessary. However, after a scheduled GeoDB update adds a new country (this is rare), deploy as soon as you are able. 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.

Before you begin

Make sure the management center can access the internet.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Geolocation Updates**.

Step 2 Under **Recurring Geolocation Updates**, check **Enable Recurring Weekly Updates...**.

Step 3 Specify the **Update Start Time**.

Step 4 Click **Save**.

Manually Update the GeoDB

Use this procedure to perform an on-demand GeoDB update.

Before you begin

If the management center cannot access the internet, get the update yourself: https://www.cisco.com/go/firepower-software. Select or search for your model (or choose any model—you use the same GeoDB for all management centers), then browse to the *Coverage and Content Updates* page.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Geolocation Updates**.

Step 2 Under **One-Time Geolocation Update**, choose how you want to update the GeoDB.

- Direct download: Choose **Download and install...**.
- Manual upload: Choose **Upload and install...**, then click **Choose File** and browse to the package you downloaded earlier.

Step 3 Click **Import**.

Monitor update progress in the Message Center.

Step 4 Verify update success.

The GeoDB update page and **Help** (?) > **About** both show the current version.

What to do next

If the update adds a new country (this is rare), deploy now. 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.

Update Intrusion Rules

As new vulnerabilities become known, the Talos Intelligence Group releases intrusion rule updates. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules. Intrusion rule updates are cumulative and we recommend you keep the system up to date. You cannot import an intrusion rule update that either matches or predates the version of the currently installed rules.

An intrusion rule update may provide the following:

- **New and modified rules and rule states**—Rule updates provide new and updated intrusion and preprocessor rules. For new rules, the rule state may be different in each system-provided intrusion policy. For example, a new rule may be enabled in the Security over Connectivity intrusion policy and disabled in the Connectivity over Security intrusion policy. Rule updates may also change the default state of existing rules, or delete existing rules entirely.
- **New rule categories**—Rule updates may include new rule categories, which are always added.
- **Modified preprocessor and advanced settings**—Rule updates may change the advanced settings in the system-provided intrusion policies and the preprocessor settings in system-provided network analysis

policies. They can also update default values for the advanced preprocessing and performance options in your access control policies.

- **New and modified variables**—Rule updates may modify default values for existing default variables, but do not override your changes. New variables are always added.

In a multidomain deployment, you can import local intrusion rules in any domain, but you can import intrusion rule updates from Talos in the Global domain only.

Understanding When Intrusion Rule Updates Modify Policies

Intrusion rule updates can affect both system-provided and custom network analysis policies, as well as all access control policies:

- **System provided**—Changes to system-provided network analysis and intrusion policies, as well as any changes to advanced access control settings, automatically take effect when you re-deploy the policies after the update.
- **Custom**—Because every custom network analysis and intrusion policy uses a system-provided policy as its base, or as the eventual base in a policy chain, rule updates can affect custom network analysis and intrusion policies. However, you can prevent rule updates from automatically making those changes. This allows you to update system-provided base policies manually, on a schedule independent of rule update imports. Regardless of your choice (implemented on a per-custom-policy basis), updates to system-provided policies do not override any settings you customized.

Note that importing a rule update discards all cached changes to network analysis and intrusion policies. For your convenience, the Rule Updates page lists policies with cached changes and the users who made those changes.

Deploying Intrusion Rule Updates

For changes made by an intrusion rule update to take effect, you must redeploy configurations. When importing a rule update, you can configure the system to automatically redeploy to affected devices. This approach is especially useful if you allow the intrusion rule update to modify system-provided base intrusion policies.

Caution Although a rule update by itself does not restart the Snort process 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/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.

Recurring Intrusion Rule Updates

As part of the initial configuration, the system schedules daily intrusion rule updates. We recommend you review this task and make changes if necessary, as described in Schedule Intrusion Rule Updates, on page 231. You may want to change the frequency, or enable automatic deploy after rule imports. For high availability management centers, you only need to import the update on the active unit.

Importing Local Intrusion Rules

A local intrusion rule is a custom standard text rule that you import from a local machine as a plain text file with ASCII or UTF-8 encoding. You can create local rules using the instructions in the Snort users manual, which is available at http://www.snort.org.

In a multidomain deployment, you can import local intrusion rules in any domain. You can view local intrusion rules imported in the current domain and ancestor domains.

Schedule Intrusion Rule Updates

As part of the initial configuration, the system schedules daily intrusion rule updates. We recommend you review this task and make changes if necessary, as described in this procedure.

Before you begin

- Make sure your process for updating intrusion rules complies with your security policies.
- Consider the update's effect on traffic flow and inspection due to bandwidth constraints and Snort restarts. We recommend performing updates in a maintenance window.
- Make sure the management center can access the internet.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 Under **Recurring Rule Update Imports**, check **Enable Recurring Rule Update Imports**.

Step 3 Specify the **Import Frequency** and start time.

Step 4 (Optional) Check **Reapply all policies...** to deploy after each update.

Step 5 Click **Save**.

Manually Update Intrusion Rules

Use this procedure to perform an on-demand intrusion rule update.

Before you begin

- Make sure your process for updating intrusion rules complies with your security policies.
- Consider the update's effect on traffic flow and inspection due to bandwidth constraints and Snort restarts. We recommend performing updates in a maintenance window.
- If the management center cannot access the internet, get the update yourself: https://www.cisco.com/go/firepower-software. Select or search for your model (or choose any model—you use the same update for all management centers), then browse to the *Coverage and Content Updates* page.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 Under **One-Time Rule Update/Rules Import**, choose how you want to update intrusion rules.

- Direct download: Choose **Download new rule update...**.
- Manual upload: Choose **Rule update or text rule file...**, then click **Choose File** and browse to the intrusion rule update.

Step 3 (Optional) Check **Reapply all policies...** to deploy after the update.

Step 4 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 5 Verify update success.

The rule update page and **Help** (?) > **About** both show the current version.

What to do next

If you did not deploy as a part of the update, deploy now.

Import Local Intrusion Rules

Use this procedure to import local intrusion rules. Imported intrusion rules appear in the local rule category in a disabled state. You can perform this task in any domain.

Before you begin

- Make sure your local rule file follows the guidelines described in Best Practices for Importing Local Intrusion Rules, on page 233.
- Make sure your process for importing local intrusion rules complies with your security policies.
- Consider the import's effect on traffic flow and inspection due to bandwidth constraints and Snort restarts. We recommend scheduling rule updates during maintenance windows.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 (Optional) Delete existing local rules.

Click **Delete All Local Rules**, then confirm that you want to move all created and imported intrusion rules to the deleted folder.

Step 3 Under **One-Time Rule Update/Rules Import**, choose **Rule update or text rule file to upload and install**, then click **Choose File** and browse to your local rule file.

Step 4 Click **Import**.

You can monitor import 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 import. Instead, contact Cisco TAC.

What to do next

- Edit intrusion policies and enable the rules you imported.
- Deploy configuration changes; see the Cisco Secure Firewall Management Center Device Configuration Guide.

Best Practices for Importing Local Intrusion Rules

Observe the following guidelines when importing a local rule file:

- The rules importer requires that all custom rules are imported in a plain text file encoded in ASCII or UTF-8.
- The text file name can include alphanumeric characters, spaces, and no special characters other than underscore (_), period (.), and dash (-).
- The system imports local rules preceded with a single pound character (#), but they are flagged as deleted.
- The system imports local rules preceded with a single pound character (#), and does not import local rules preceded with two pound characters (##).
- Rules cannot contain any escape characters.
- In a multidomain deployment, the system assigns a GID of 1 to a rule imported into or created in the Global domain, and a domain-specific GID between 1000 and 2000 for all other domains.
- You do not have to specify a Generator ID (GID) when importing a local rule. If you do, specify only GID 1 for a standard text rule.
- When importing a rule for the first time, do *not* specify a Snort ID (SID) or revision number. This avoids collisions with SIDs of other rules, including deleted rules. The system will automatically assign the rule the next available custom rule SID of 1000000 or greater, and a revision number of 1.

  If you must import rules with SIDs, a SID can be any unique number 1,000,000 or greater.

  In a multidomain deployment, if multiple administrators are importing local rules at the same time, SIDs within an individual domain might appear to be non-sequential because the system assigned the intervening numbers in the sequence to another domain.
- When importing an updated version of a local rule you have previously imported, or when reinstating a local rule you have deleted, you *must* include the SID assigned by the system and a revision number greater than the current revision number. You can determine the revision number for a current or deleted rule by editing the rule.

Note The system automatically increments the revision number when you delete a local rule; this is a device that allows you to reinstate local rules. All deleted local rules are moved from the local rule category to the deleted rule category.

- Import local rules on the primary management center in a high availability pair to avoid SID numbering issues.
- The import fails if a rule contains any of the following: .
  - A SID greater than 2147483647.
  - A list of source or destination ports that is longer than 64 characters.
  - When importing into the Global domain in a multidomain deployment, a GID:SID combination uses GID 1 and a SID that already exists in another domain; this indicates that the combination existed before Version 6.2.1. You can reimport the rule using GID 1 and a unique SID.
- Policy validation fails if you enable an imported local rule that uses the deprecated `threshold` keyword in combination with the intrusion event thresholding feature in an intrusion policy.
- All imported local rules are automatically saved in the local rule category.
- The system always sets local rules that you import to the disabled rule state. You must manually set the state of local rules before you can use them in your intrusion policy.

View Intrusion Rule Update Logs

The system generates logs of rule updates/imports, listed by timestamp, user, and whether each update succeeded or failed. These logs contain detailed import information on all updated rules and components; see Intrusion Rule Update Log Details, on page 234. Use this procedure to view rule import logs. Note that deleting an import log does not delete the imported objects. In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **System** (⚙) > **Updates** > **Rule Updates**.

Step 2 Click **Rule Update Log**.

Step 3 (Optional) View details for any rule update by clicking **View** (👁) next to the log file.

Intrusion Rule Update Log Details

Tip You search the entire Rule Update Import Log database even when you initiate a search by clicking **Search** on the toolbar from the Rule Update Import Log detailed view with only the records for a single import file displayed. Make sure you set your time constraints to include all objects you want to include in the search.

Table 7: Intrusion Rule Update Log Details

Field	Description
Action	An indication that one of the following has occurred for the object type: • `new` (for a rule, this is the first time the rule has been stored on this appliance) • `changed` (for a rule update component or rule, the rule update component has been modified, or the rule has a higher revision number and the same GID and SID) • `collision` (for a rule update component or rule, import was skipped because its revision conflicts with an existing component or rule on the appliance) • `deleted` (for rules, the rule has been deleted from the rule update) • `enabled` (for a rule update edit, a preprocessor, rule, or other feature has been enabled in a default policy provided with the system) • `disabled` (for rules, the rule has been disabled in a default policy provided with the system) • `drop` (for rules, the rule has been set to Drop and Generate Events in a default policy provided with the system) • `error` (for a rule update or local rule file, the import failed) • `apply` (the **Reapply all policies after the rule update import completes** option was enabled for the import)
Default Action	The default action defined by the rule update. When the imported object type is `rule`, the default action is `Pass`, `Alert`, or `Drop`. For all other imported object types, there is no default action.
Details	A string unique to the component or rule. For rules, the GID, SID, and previous revision number for a changed rule, displayed as `previously (GID:SID:Rev)`. This field is blank for a rule that has not changed.
Domain	The domain whose intrusion policies can use the updated rule. Intrusion policies in descendant domains can also use the rule. This field is only present in a multidomain deployment.
GID	The generator ID for a rule. For example, `1` (standard text rule, Global domain or legacy GID) or `3` (shared object rule).
Name	The name of the imported object, which for rules corresponds to the rule Message field, and for rule update components is the component name.
Policy	For imported rules, this field displays `All`. This means that the rule was imported successfully, and can be enabled in all appropriate default intrusion policies. For other types of imported objects, this field is blank.
Rev	The revision number for a rule.
Rule Update	The rule update file name.
SID	The SID for a rule.
Time	The time and date the import began.

Field	Description
Type	The type of imported object, which can be one of the following: • `rule update component` (an imported component such as a rule pack or policy pack) • `rule` (for rules, a new or updated rule) • `policy apply` (the **Reapply all policies after the rule update import completes** option was enabled for the import)
Count	The count (`1`) for each record. The Count field appears in a table view when the table is constrained, and the Rule Update Log detailed view is constrained by default to rule update records. This field is not searchable.

Maintain Your Air-Gapped Deployment

If your management center is not connected to the internet, essential updates will not occur automatically. You must manually obtain and install these updates.

For more information, see:

- Software upgrade guides: https://cisco.com/go/ftd-fmc-upgrade
- Manually Update the VDB, on page 226
- Manually Update Intrusion Rules, on page 231
- Manually Update the GeoDB, on page 228

History for System Updates

Table 8: Version 7.3.0 Features

Feature	Minimum Management Center	Minimum Threat Defense	Details
Threat Defense Upgrade			
Choose and direct-download upgrade packages to the management center from Cisco.	7.3.0	Any	You can now choose which threat defense upgrade packages you want to direct download to the management center. Use the new **Download Updates** sub-tab on > **Updates** > **Product Updates**. Version restrictions: this feature is replaced by an improved package management system in Version 7.2.6/7.4.1. See: Download Upgrade Packages with the Management Center

Feature	Minimum Management Center	Minimum Threat Defense	Details
Upload upgrade packages to the management center from the threat defense wizard.	7.3.0	Any	You now use the wizard to upload threat defense upgrade packages or specify their location. Previously (depending on version), you used **System** (⚙) > **Updates** or **System** (⚙) > **Product Upgrades**. Version restrictions: this feature is replaced by an improved package management system in Version 7.2.6/7.4.1. See: Upgrade Threat Defense
Auto-upgrade to Snort 3 after successful threat defense upgrade is no longer optional.	7.3.0	Any	**Upgrade impact. All eligible devices upgrade to Snort 3 when you deploy.** When you upgrade threat defense to Version 7.3+, you can no longer disable the **Upgrade Snort 2 to Snort 3** option. After the software upgrade, all eligible devices will upgrade from Snort 2 to Snort 3 when you deploy configurations. Although you can switch individual devices back, Snort 2 is not supported on threat defense 7.7+. You should stop using it now. For devices that are ineligible for auto-upgrade because they use custom intrusion or network analysis policies, manually upgrade to Snort 3 for improved detection and performance. For migration assistance, see the Cisco Secure Firewall Management Center Snort 3 Configuration Guide for your version.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Combined upgrade and install package for Secure Firewall 3100.	7.3.0	7.3.0	**Reimage Impact.** In Version 7.3, we combined the threat defense install and upgrade package for the Secure Firewall 3100, as follows: • Version 7.1–7.2 install package: cisco-ftd-fp3k.*version*.SPA • Version 7.1–7.2 upgrade package: Cisco_FTD_SSP_FP3K_Upgrade-*version-build*.sh.REL.tar • Version 7.3+ combined package: Cisco_FTD_SSP_FP3K_Upgrade-*version-build*.sh.REL.tar Although you can upgrade threat defense without issue, you cannot reimage from older threat defense and ASA versions directly to threat defense Version 7.3+. This is due to a ROMMON update required by the new image type. To reimage from those older versions, you must "go through" ASA 9.19+, which is supported with the old ROMMON but also updates to the new ROMMON. There is no separate ROMMON updater. To get to threat defense Version 7.3+, your options are: • Upgrade from threat defense Version 7.1 or 7.2 — use the normal upgrade process. See the appropriate Upgrade Guide. • Reimage from threat defense Version 7.1 or 7.2 — reimage to ASA 9.19+ first, then reimage to threat defense Version 7.3+. See *Threat Defense→ASA: Firepower 1000, 2100; Secure Firewall 3100* and then *ASA→Threat Defense: Firepower 1000, 2100 Appliance Mode; Secure Firewall 3100* in the Cisco Secure Firewall ASA and Secure Firewall Threat Defense Reimage Guide. • Reimage from ASA 9.17 or 9.18 — upgrade to ASA 9.19+ first, then reimage to threat defense Version 7.3+. See the Cisco Secure Firewall ASA Upgrade Guide and then *ASA→Threat Defense: Firepower 1000, 2100 Appliance Mode; Secure Firewall 3100* in the Cisco Secure Firewall ASA and Secure Firewall Threat Defense Reimage Guide. • Reimage from threat defense Version 7.3+ — use the normal reimage process. See *Reimage the System with a New Software Version* in the Cisco FXOS Troubleshooting Guide for the Firepower 1000/2100 and Secure Firewall 3100/4200 with Firepower Threat Defense.
Content Updates			

Feature	Minimum Management Center	Minimum Threat Defense	Details
Automatic VDB downloads.	7.3.0	Any	The initial setup on the management center schedules a weekly task to download the latest available software updates, which now includes the latest vulnerability database (VDB). We recommend you review this weekly task and adjust if necessary. Optionally, schedule a new weekly task to actually update the VDB and deploy configurations. New/modified screens: The **Vulnerability Database** check box is now enabled by default in the system-created **Weekly Software Download** scheduled task.
Install any VDB.	7.3.0	Any	Starting with VDB 357, you can now install any VDB as far back as the baseline VDB for that management center. After you update the VDB, deploy configuration changes. If you based configurations on vulnerabilities, application detectors, or fingerprints that are no longer available, examine those configurations to make sure you are handling traffic as expected. Also, keep in mind a scheduled task to update the VDB can undo a rollback. To avoid this, change the scheduled task or delete any newer VDB packages. New/modified screens: On **System**(⚙) > **Updates** > **Product Updates** > **Available Updates**, if you upload an older VDB, a new **Rollback** icon appears instead of the **Install** icon.

Table 9: Version 7.2.10 Features

Feature	Minimum Management Center	Minimum Threat Defense	Details
Threat Defense Upgrade			
Threat defense and chassis upgrade wizards optimized for lower resolution screens.	7.2.10 7.6.0	Any	We optimized the threat defense and chassis upgrade wizards for lower resolution screens (and smaller browser windows). Text appears smaller and certain screen elements are hidden. If you change your resolution or window size mid-session, you may need to refresh the page for the web interface to adjust. Note that the minimum screen resolution to use the management center is 1280 x 720. New/modified screens: • **Devices** > **Threat Defense Upgrade** • **Devices** > **Chassis Upgrade** Version restrictions: Not supported with Version 7.2.0–7.2.9, 7.3.x, 7.4.0–7.4.2.

Table 10: Version 7.2.6 Features

Feature	Minimum Management Center	Minimum Threat Defense	Details
Upgrade			
Improved upgrade starting page and package management.	7.2.6 7.4.1	Any	A new upgrade page makes it easier to choose, download, manage, and apply upgrades to your entire deployment. This includes the management center, threat defense devices, and any older NGIPSv/ASA FirePOWER devices. The page lists all upgrade packages that apply to your current deployment, with suggested releases specially marked. You can easily choose and direct-download packages from Cisco, as well as manually upload and delete packages. Internet access is required to retrieve the list/direct download upgrade packages. Otherwise, you are limited to manual management. Patches are not listed unless you have at least one appliance at the appropriate maintenance release (or you manually uploaded the patch). You must manually upload hotfixes. New/modified screens: • **System**(⚙) > **Product Upgrades** is now where you upgrade the management center and all managed devices, as well as manage upgrade packages. • **System**(⚙) > **Content Updates** is now where you update intrusion rules, the VDB, and the GeoDB. • **Devices** > **Threat Defense Upgrade** takes you directly to the threat defense upgrade wizard. • **System**(⚙) > **Users** > **User Role** > **Create User Role** > **Menu-Based Permissions** allows you to grant access to **Content Updates** (VDB, GeoDB, intrusion rules) without allowing access to **Product Upgrades** (system software). Deprecated screens/options: • **System**(⚙) > **Updates** is deprecated. All threat defense upgrades now use the wizard. • The **Add Upgrade Package** button on the threat defense upgrade wizard has been replaced by a **Manage Upgrade Packages** link to the new upgrade page. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Suggested release notifications.	7.2.6 7.4.1	Any	The management center now notifies you when a new suggested release is available. If you don't want to upgrade right now, you can have the system remind you later, or defer reminders until the next suggested release. The new upgrade page also indicates suggested releases. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0. See: Cisco Secure Firewall Management Center New Features by Release
Updated internet access requirements for direct-downloading software upgrades.	7.2.6 7.4.1	Any	**Upgrade impact. The system connects to new resources.** The management center has changed its direct-download location for software upgrade packages from sourcefire.com to amazonaws.com. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Threat Defense Upgrade			
Enable revert from the threat defense upgrade wizard.	7.2.6 7.4.1	Any, if upgrading to 7.1+	You can now enable revert from the threat defense upgrade wizard. Other version restrictions: You must be upgrading threat defense to Version 7.1+. Not supported with management center Version 7.3.x or 7.4.0.
Select devices to upgrade from the threat defense upgrade wizard.	7.2.6	Any	Use the wizard to select devices to upgrade. You can now use the threat defense upgrade wizard to select or refine the devices to upgrade. On the wizard, you can toggle the view between selected devices, remaining upgrade candidates, ineligible devices (with reasons why), devices that need the upgrade package, and so on. Previously, you could only use the Device Management page and the process was much less flexible.
View detailed upgrade status from the threat defense upgrade wizard.	7.2.6 7.4.1	Any	The final page of the threat defense upgrade wizard now allows you to monitor upgrade progress. This is in addition to the existing monitoring capability on the Upgrade tab on the Device Management page, and on the Message Center. Note that as long as you have not started a new upgrade flow, **Devices** > **Threat Defense Upgrade** brings you back to this final wizard page, where you can view the detailed status for the current (or most recently complete) device upgrade. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Unattended threat defense upgrades.	7.2.6	Any	The threat defense upgrade wizard now supports unattended upgrades, using a new **Unattended Mode** menu. You just need to select the target version and the devices you want to upgrade, specify a few upgrade options, and step away. You can even log out or close the browser.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Simultaneous threat defense upgrade workflows by different users.	7.2.6	Any	We now allow simultaneous upgrade workflows by different users, as long as you are upgrading different devices. The system prevents you from upgrading devices already in someone else's workflow. Previously, only one upgrade workflow was allowed at a time across all users.
Skip pre-upgrade troubleshoot generation for threat defense devices.	7.2.6	Any	You can now skip the automatic generating of troubleshooting files before major and maintenance upgrades by disabling the new **Generate troubleshooting files before upgrade begins** option. This saves time and disk space. To manually generate troubleshooting files for a threat defense device, choose **System**(⚙) > **Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, then **Generate Troubleshooting Files**.
Management Center Upgrade			
New upgrade wizard for the management center.	7.2.6 7.4.1	Any	A new upgrade starting page and wizard make it easier to perform management center upgrades. After you use **System**(⚙) > **Product Upgrades** to get the appropriate upgrade package onto the management center, click **Upgrade** to begin. Other version restrictions: Only supported for management center upgrades from Version 7.2.6+/7.4.1+. Not supported for upgrades from Version 7.3.x or 7.4.0. See: Cisco Secure Firewall Threat Defense Upgrade Guide for Management Center
Hotfix high availability management centers without pausing synchronization.	7.2.6 7.4.1	Any	Unless otherwise indicated by the hotfix release notes or Cisco TAC, you do not have to pause synchronization to install a hotfix on high availability management centers. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Content Updates			
Scheduled tasks download patches and VDB updates only.	7.2.6 7.4.1	Any	**Upgrade impact. Scheduled download tasks stop retrieving maintenance releases.** The **Download Latest Update** scheduled task no longer downloads maintenance releases; now it only downloads the latest applicable patches and VDB updates. To direct-download maintenance (and major) releases to the management center, use **System**(⚙) > **Product Upgrades**. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.

Table 11: Version 7.2.0 Features

Feature	Details
Threat Defense Upgrade	
Copy upgrade packages ("peer-to-peer sync") from device to device.	Instead of copying upgrade packages to each device from the management center or internal web server, you can use the threat defense CLI to copy upgrade packages between devices ("peer to peer sync"). This secure and reliable resource-sharing goes over the management network but does not rely on the management center. Each device can accommodate 5 package concurrent transfers. This feature is supported for Version 7.2.x–7.4.x standalone devices managed by the same Version 7.2.x–7.4.x standalone management center. It is not supported for: • Container instances. • Device high availability pairs and clusters. These devices get the package from each other as part of their normal sync process. Copying the upgrade package to one group member automatically syncs it to all group members. • Devices managed by high availability management centers. • Devices in different domains, or devices separated by a NAT gateway. • Devices upgrading from Version 7.1 or earlier, regardless of management center version. • Devices running Version 7.6+. New/modified CLI commands: **configure p2psync enable**, **configure p2psync disable**, **show peers**, **show peer details**, **sync-from-peer**, **show p2p-sync-status**
Auto-upgrade to Snort 3 after successful threat defense upgrade.	When you use a Version 7.2+ management center to upgrade threat defense to Version 7.2+, you can now choose whether to **Upgrade Snort 2 to Snort 3**. After the software upgrade, eligible devices upgrade from Snort 2 to Snort 3 when you deploy configurations. For devices that are ineligible because they use custom intrusion or network analysis policies, we strongly recommend you manually upgrade to Snort 3 for improved detection and performance. For help, see the Cisco Secure Firewall Management Center Snort 3 Configuration Guide for your version. Version restrictions: Not supported for threat defense upgrades to Version 7.0.x or 7.1.x.
Upgrade for single-node clusters.	You can now use the device upgrade page (**Devices** > **Device Upgrade**) to upgrade clusters with only one active node. Any deactivated nodes are also upgraded. Previously, this type of upgrade would fail. This feature is not supported from the system updates page (**System**(⚙)**Updates**). Hitless upgrades are also not supported in this case. Interruptions to traffic flow and inspection depend on the interface configurations of the lone active unit, just as with standalone devices. Supported platforms: Firepower 4100/9300, Secure Firewall 3100

Feature	Details
Revert threat defense upgrades from the CLI.	You can now revert threat defense upgrades from the device CLI if communications between the management center and device are disrupted. Note that in high availability/scalability deployments, revert is more successful when all units are reverted simultaneously. When reverting with the CLI, open sessions with all units, verify that revert is possible on each, then start the processes at the same time. **Caution** Reverting from the CLI can cause configurations between the device and the management center to go out of sync, depending on what you changed post-upgrade. This can cause further communication and deployment issues. New/modified CLI commands: **upgrade revert**, **show upgrade revert-info**.
Management Center Upgrade	
Management center upgrade does not automatically generate troubleshooting files.	To save time and disk space, the management center upgrade process no longer automatically generates troubleshooting files before the upgrade begins. Note that device upgrades are unaffected and continue to generate troubleshooting files. To manually generate troubleshooting files for the management center, choose **System**(⚙) > **Health** > **Monitor**, click **Firewall Management Center** in the left panel, then **View System & Troubleshoot Details**, then **Generate Troubleshooting Files**.

Table 12: Version 7.1.0 Features

Feature	Details
Threat Defense Upgrade	
Revert a successful device upgrade.	You can now revert major and maintenance upgrades to FTD. Reverting returns the software to its state just before the last upgrade, also called a *snapshot*. If you revert an upgrade after installing a patch, you revert the patch as well as the major and/or maintenance upgrade. **Important** If you think you might need to revert, you must use **System**(⚙) > **Updates** to upgrade FTD. The System Updates page is the only place you can enable the **Enable revert after successful upgrade** option, which configures the system to save a revert snapshot when you initiate the upgrade. This is in contrast to our usual recommendation to use the wizard on the **Devices** > **Device Upgrade** page. This feature is not supported for container instances. Minimum FTD: 7.1

Feature	Details
Improvements to the upgrade workflow for clustered and high availability devices.	We made the following improvements to the upgrade workflow for clustered and high availability devices: • The upgrade wizard now correctly displays clustered and high availability units as groups, rather than as individual devices. The system can identify, report, and preemptively require fixes for group-related issues you might have. For example, you cannot upgrade a cluster on the Firepower 4100/9300 if you have made unsynced changes on Firepower Chassis Manager. • We improved the speed and efficiency of copying upgrade packages to clusters and high availability pairs. Previously, the FMC copied the package to each group member sequentially. Now, group members can get the package from each other as part of their normal sync process. • You can now specify the upgrade order of data units in a cluster. The control unit always upgrades last.

Table 13: Version 7.0.0 Features

Feature	Details
Threat Defense Upgrade	
Improved FTD upgrade performance and status reporting.	FTD upgrades are now easier faster, more reliable, and take up less disk space. A new **Upgrades** tab in the Message Center provides further enhancements to upgrade status and error reporting.

Feature	Details
Easy-to-follow upgrade workflow for FTD devices.	A new device upgrade page (**Devices > Device Upgrade**) on the FMC provides an easy-to-follow wizard for upgrading Version 6.4+ FTD devices. It walks you through important pre-upgrade stages, including selecting devices to upgrade, copying the upgrade package to the devices, and compatibility and readiness checks. To begin, use the new **Upgrade Firepower Software** action on the Device Management page (**Devices > Device Management > Select Action**). As you proceed, the system displays basic information about your selected devices, as well as the current upgrade-related status. This includes any reasons why you cannot upgrade. If a device does not "pass" a stage in the wizard, it does not appear in the next stage. If you navigate away from wizard, your progress is preserved, although other users with Administrator access can reset, modify, or continue the wizard. **Note** You must still use **System**(⚙) > **Updates** to upload or specify the location of FTD upgrade packages. You must also use the System Updates page to upgrade the FMC itself, as well as all non-FTD managed devices. **Note** In Version 7.0, the wizard does not correctly display devices in clusters or high availability pairs. Even though you must select and upgrade these devices as a unit, the wizard displays them as standalone devices. Device status and upgrade readiness are evaluated and reported on an individual basis. This means it is possible for one unit to appear to "pass" to the next stage while the other unit or units do not. However, these devices are still grouped. Running a readiness check on one, runs it on all. Starting the upgrade on one, starts it on all. To avoid possible time-consuming upgrade failures, *manually* ensure all group members are ready to move on to the next step of the wizard before you click **Next**.
Upgrade more FTD devices at once.	The number of devices you can upgrade at once is now limited by your management network bandwidth—not the system's ability to manage simultaneous upgrades. Previously, we recommended against upgrading more than five devices at a time. **Important** Only upgrades to FTD Version 6.7+ using the FTD upgrade wizard see this improvement. If you are upgrading devices to an older FTD release—even if you are using the new upgrade wizard—we still recommend you limit to five devices at a time.
Upgrade different device models together.	You can now use the FTD upgrade wizard to queue and invoke upgrades for all FTD models at the same time, as long as the system has access to the appropriate upgrade packages. Previously, you would choose an upgrade package, then choose the devices to upgrade using that package. That meant that you could upgrade multiple devices at the same time *only* if they shared an upgrade package. For example, you could upgrade two Firepower 2100 series devices at the same time, but not a Firepower 2100 series and a Firepower 1000 series.

Table 14: Version 6.7.0 Features

Feature	Details
Threat Defense Upgrade	
Upgrades remove PCAP files to save disk space.	Upgrades now remove locally stored PCAP files. To upgrade, you must have enough free disk space or the upgrade fails.
Improved FTD upgrade status reporting and cancel/retry options.	You can now view the status of FTD device upgrades and readiness checks in progress on the Device Management page, as well as a 7-day history of upgrade success/failures. The Message Center also provides enhanced status and error messages. A new Upgrade Status pop-up, accessible from both Device Management and the Message Center with a single click, shows detailed upgrade information, including percentage/time remaining, specific upgrade stage, success/failure data, upgrade logs, and so on. Also on this pop-up, you can manually cancel failed or in-progress upgrades (**Cancel Upgrade**), or retry failed upgrades (**Retry Upgrade**). Canceling an upgrade reverts the device to its pre-upgrade state. **Note** To be able to manually cancel or retry a failed upgrade, you must disable the new auto-cancel option, which appears when you use the FMC to upgrade an FTD device: **Automatically cancel on upgrade failure and roll back to the previous version**. With the option enabled, the device automatically reverts to its pre-upgrade state upon upgrade failure. Auto-cancel is not supported for patches. In an HA or clustered deployment, auto-cancel applies to each device individually. That is, if the upgrade fails on one device, only that device is reverted. New/modified screens: • **System**(⚙) > **Updates** > **Product Updates** > **Available Updates** > **Install** icon for the FTD upgrade package • **Devices** > **Device Management** > **Upgrade** • **Message Center > Tasks** New/modified CLI commands: **show upgrade status detail**, **show upgrade status continuous**, **show upgrade status**, **upgrade cancel**, **upgrade retry**
Content Updates	

Feature	Details
Custom intrusion rule import warns when rules collide.	The FMC now warns you of rule collisions when you import custom (local) 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**.

Table 15: Version 6.6.0 Features

Feature	Details
Threat Defense Upgrade	
Get FTD upgrade packages from an internal web server.	FTD devices can now get upgrade packages from your own internal web server, rather than from the FMC. This is especially useful if you have limited bandwidth between the FMC and its devices. It also saves space on the FMC. **Note** This feature is supported only for FTD devices running Version 6.6+. It is not supported for upgrades *to* Version 6.6, nor is it supported for the FMC or Classic devices. New/modified screens: We added a **Specify software update source** option to the page where you upload upgrade packages.
Content Updates	
Automatic VDB update during initial setup.	When you set up a new or reimaged FMC, the system automatically attempts to update the vulnerability database (VDB). This is a one-time operation. If the FMC has internet access, we recommend you schedule tasks to perform automatic recurring VDB update downloads and installations.

Table 16: Version 6.5.0 Features

Feature	Details
Content Updates	

Feature	Details
Automatic software downloads and GeoDB updates.	When you set up a new or reimaged FMC, the system automatically schedules: • A weekly task to download software updates for the FMC and its managed devices. • Weekly updates for the GeoDB. The tasks are scheduled in UTC, which means that when they occur locally depends on the date and your specific location. Also, because tasks are scheduled in UTC, they do not adjust for Daylight Saving Time, summer time, or any such seasonal adjustments that you may observe in your location. If you are affected, scheduled tasks occur one hour "later" in the summer than in the winter, according to local time. We recommend you review the auto-scheduled configurations and adjust them if necessary.

Table 17: Version 6.4.0 Features

Feature	Details
Management Center Upgrade	
Upgrades postpone scheduled tasks.	The management center upgrade process now postpones scheduled tasks. Any task scheduled to begin during the upgrade will begin five minutes after the post-upgrade reboot. **Note** Before you begin any upgrade, you must still make sure running tasks are complete. Tasks running when the upgrade begins are stopped, become failed tasks, and cannot be resumed. Note that this feature is supported for all upgrades *from* a supported version. This includes Version 6.4.0.10 and later patches, Version 6.6.3 and later maintenance releases, and Version 6.7.0+. This feature is not supported for upgrades *to* a supported version from an unsupported version.
Content Updates	

Feature	Details
Signed SRU, VDB, and GeoDB updates.	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 from the Cisco Support & Download site—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 FMC or ASA FirePOWER device, you must manually delete it. Leaving the package takes up disk space, and also may cause issues with future upgrades.

Table 18: Version 6.2.3 Features

Feature	Details
Device Upgrade	
Copy upgrade packages to managed devices before the upgrade.	You can now copy (or push) an upgrade package from the FMC to a managed device before you run the actual upgrade. This is useful because you can push during times of low bandwidth use, outside of the upgrade maintenance window. When you push to high availability, clustered, or stacked devices, the system sends the upgrade package to the active/control/primary first, then to the standby/data/secondary. New/modified screens: **System**(⚙) > **Updates**
Content Updates	
FMC warns of Snort restart before VDB updates.	The FMC now warns you that Vulnerability Database (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.

Feature	Details
Deprecated: Geolocation details	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.

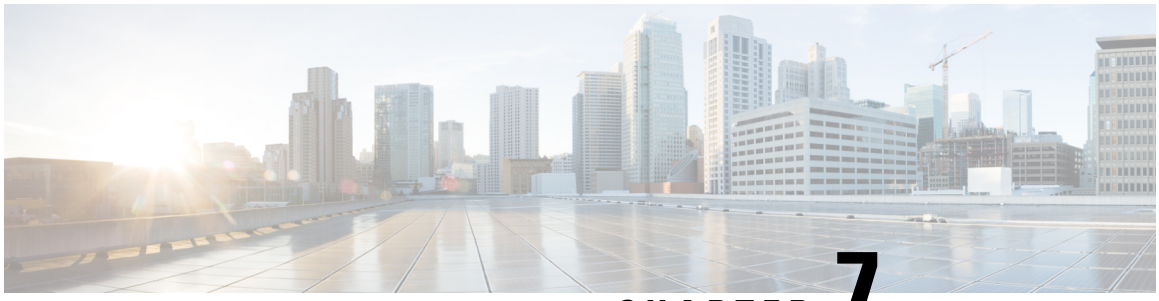

CHAPTER 7

Licenses

This chapter provides in-depth information about the different license types, service subscriptions, licensing requirements and more.

Note The Management Center supports either a Smart License or a legacy PAK (Product Activation Keys) license for its platform license. For more information about using the PAK license, see Configure Legacy Management Center PAK-Based Licenses, on page 298.

About 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:** My Cisco Entitlements (MCE) 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 Software Central (software.cisco.com).

For a more detailed overview on Cisco Licensing, go to cisco.com/go/licensingguide

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. The Smart Software Manager lets you create a primary account for your organization. For instructions, see 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

The following table compares the available licensing options for environments without internet access. Your sales representative may have additional advice for your specific situation.

Table 19: 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, see Register the Management Center with the Smart Software Manager On-Prem, on page 278	For more information, see Configure Specific License Reservation (SLR), on page 287

How Licensing Works for the Management Center and Devices

The 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 management center does not require a license for its own use. The management center virtual does require a platform license.

Periodic Communication with the Smart Software Manager

In order to maintain your product license entitlement, your product must communicate periodically with the Smart Software Manager.

You use a Product Instance Registration Token to register the management center with the Smart Software Manager. The Smart Software Manager issues an ID certificate for communication between the 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 management center may be removed from your account.

The 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 management center so the changes immediately take effect. You also can wait for the management center to communicate as scheduled.

Your 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 254. In non-airgapped deployments, normal license communication occurs every 30 days, but with the grace period, your management center will operate for up to 90 days without contacting the Smart Software Manager. Ensure that the management center contacts the Smart Software Manager before 90 days have passed, or else the management center will revert to an unregistered state.

Evaluation Mode

Before the management center registers with the Smart Software Manager, it operates for 90 days in evaluation mode. You can assign feature licenses to managed devices, and they will remain in compliance for the duration of evaluation mode. When this period ends, the management center becomes unregistered.

If you register the management center with the Smart Software Manager, the evaluation mode ends. If you later deregister the management center, you cannot resume evaluation mode, even if you did not initially use all 90 days.

For more information about the unregistered state, see Unregistered State, on page 256.

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 management center can become out of compliance in the following situations:

- Over-utilization—When the managed devices or the management center virtual uses unavailable licenses.

Note that even if one license is unavailable for a device, the 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.

- License expiration—When a managed device term-based license expires.

In an out-of-compliance state, see the following effects:

- Management Center Virtual platform license—Operation is not affected.
- All managed device licenses—Operation is not affected.

After you resolve the licensing problem, the 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 management center can become unregistered in the following situations:

- Evaluation mode expiration—Evaluation mode expires after 90 days.
- Manual deregistration of the management center
- Lack of communication with the Smart Software Manager—The management center does not communicate with the Smart Software Manager for 1 year. Note: After 90 days, the 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 management center is removed from your account so you will have to manually re-register the management center.

In an unregistered state, the 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.

Table 20: Smart Licenses

License You Assign	Duration	Granted Capabilities
Essentials	Perpetual or Subscription **Note** Essentials subscription licenses are supported only on Threat Defense Virtual.	Except for Specific License Reservation and the Secure Firewall 3100, Essentials perpetual licenses are automatically assigned with all threat defenses. User and application control Switching and routing NAT For details, see Essentials Licenses, on page 258.
IPS	Subscription	Intrusion detection and prevention File control Security Intelligence filtering For details, see IPS Licenses, on page 260
Malware defense	Subscription	Malware defense Secure Malware Analytics File storage (IPS license is a prerequisite for a Malware defense license.) For details, see Malware Defense Licenses, on page 259 and *License Requirements for File and Malware Policies* in the Cisco Secure Firewall Management Center Device Configuration Guide.
Carrier	Subscription for Firepower 4100/9300, Secure Firewall 3100, and Threat Defense Virtual	Diameter, GTP/GPRS, M3UA, and SCTP inspection For details, see Carrier License, on page 260.
URL Filtering	Subscription	Category and reputation-based URL filtering For details, see URL Filtering Licenses, on page 261. (IPS license is a prerequisite for a URL Filtering license.)
Management Center Virtual	• Regular Smart Licensing—Perpetual • Specific License Reservation—Subscription	The platform license determines the number of devices the management center virtual can manage. For details, see Management Center Virtual Licenses, on page 258.

License You Assign	Duration	Granted Capabilities
Export-Controlled Features	Perpetual	Features that are subject to national security, foreign policy, and anti-terrorism laws and regulations; see Licensing for Export-Controlled Functionality, on page 262.
Remote Access VPN: • Secure Client Premier • Secure Client Advantage • Secure Client VPN Only	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 threat defense can use any valid Secure Client license. The available features do not differ based on license type. For more information, see Secure Client Licenses, on page 262 and *VPN Licensing* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Note Subscription licenses are term-based licenses.

Management Center Virtual Licenses

The management center virtual requires a platform license that correlates with the number of devices it can manage.

The management center virtual supports Smart Licensing.

In regular Smart Licensing, these licenses are perpetual.

In Specific License Reservation (SLR), these licenses are subscription-based.

Note For the add-on license requirements of your new devices on FMCv, it is recommended to migrate to a higher management center virtual model that supports additional devices.

Essentials Licenses

The Essentials 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 IPS license is enabled.

Secure Firewall 3100

You obtain a Essentials license when you purchase the Secure Firewall 3100.

Other Models

Except in deployments using Specific License Reservation, a Essentials license is automatically added to your account when you register a device to the management center. For Specific License Reservation, you need to add the Essentials 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 IPS (TM) or IPS and URL Filtering (TMC) subscriptions. IPS 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 Defense 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 management center can receive Secure Endpoint malware events and indications of compromise (IOC) from the Secure Malware Analytics Cloud.

See also 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 management center moves to the out-of-compliance state.

IPS Licenses

A IPS license allows you to perform intrusion detection and prevention, file control, and Security Intelligence filtering:

- *Intrusion detection and prevention* allows you to analyze network traffic for intrusions and exploits and, optionally, drop offending packets.
- *File control* allows you to detect and, optionally, block users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols. *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.

You can purchase a IPS license as a stand-alone subscription (T) or in combination with URL Filtering (TC), Malware defense (TM), or both (TMC).

When you disable this license:

- The 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 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 IPS.

If the license expires, your entitlement for the above capabilities ceases and the management center moves to the out-of-compliance state.

Carrier License

The Carrier license enables the inspection of the following protocols:

- Diameter—Diameter is an Authentication, Authorization, and Accounting (AAA) protocol used in next-generation mobile and fixed telecom networks such as EPS (Evolved Packet System) for LTE (Long Term Evolution) and IMS (IP Multimedia Subsystem). It replaces RADIUS and TACACS in these networks.
- GTP/GPRS—GPRS Tunneling Protocol (GTP) is used in GSM, UMTS, and LTE networks for general packet radio service (GPRS) traffic. GTP provides a tunnel control and management protocol to provide GPRS network access for a mobile station by creating, modifying, and deleting tunnels. GTP also uses a tunneling mechanism for carrying user data packets.
- M3UA—MTP3 User Adaptation (M3UA) is a client/server protocol that provides a gateway to the Signaling System 7 (SS7) network for IP-based applications that interface with the SS7 Message Transfer Part 3 (MTP3) layer. M3UA makes it possible to run the SS7 User Parts (such as ISUP) over an IP network.

- SCTP—Stream Control Transmission Protocol (SCTP) is a transport-layer protocol that supports the SS7 protocol over IP networks. It supports the 4G LTE mobile network architecture. SCTP can handle multiple simultaneous streams, multiplexed streams, and provides more security features.

Note After you enable this license on a device, use a FlexConfig policy to enable the protocol inspection.

The Carrier license PIDs are available per family and not per device model. You can enable this license for each device either in the evaluation mode or with a Smart License.

The Carrier license for Firepower 4100/9300, Secure Firewall 3100, and Threat Defense Virtual is term-based. This license also supports Specific License Reservation.

Supported Devices

The devices that support the Carrier License are:

- Secure Firewall 3110
- Secure Firewall 3120
- Secure Firewall 3130
- Secure Firewall 3140
- Firepower 4112
- Firepower 4115
- Firepower 4125
- Firepower 4145
- Firepower 9300
- Threat Defense Virtual

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 IPS (TC) or Threat and Malware defense (TMC) subscriptions. IPS 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 management center will not download URL information. You cannot deploy the access control policy until you first add a URL Filtering license to the 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 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 management center moves to the out-of-compliance state.

Secure Client Licenses

You can configure remote access VPN using the Secure Client and standards-based IPSec/IKEv2.

To enable remote access VPN, you must purchase and enable one of the following licenses: Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only. You can select Secure Client Advantage and Secure Client Premier if you have both licenses and you want to use them both. The Secure Client VPN Onlylicense cannot be used with **Apex** or **Plus**. The Secure Client license must be shared with the Smart Account. For more instructions, see 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 Secure 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 threat defense requires strong encryption (which is higher than DES) for successfully establishing remote access VPN connections with Secure Clients.

You cannot deploy remote access VPN if the following are true:

- Smart Licensing on the management center is running in evaluation mode.
- Your Smart Account is not configured to use export-controlled features (strong encryption).

Licensing for Export-Controlled Functionality

Features that require export-controlled functionality

Certain software features are subject to national security, foreign policy, and anti-terrorism laws and regulations. These export-controlled features include:

- 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

How to determine whether export-controlled functionality is currently enabled for your system

To determine whether export-controlled functionality is currently enabled for your system: Go to **System** > **Licenses** > **Smart Licenses** and see if **Export-Controlled Features** displays **Enabled**.

About enabling export-controlled functionality

If **Export-Controlled Features** shows **Disabled** and you want to use features that require strong encryption, 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 280

- 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 management center, then you must deregister and then re-register the management center using a new Product Instance Registration Token with export-controlled functionality enabled.

If you registered devices to the management center in evaluation mode or before you enabled strong encryption on the 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.

Threat Defense Virtual Licenses

This section describes the performance-tiered license entitlements available for the threat defense virtual.

Any threat defense virtual license can be used on any supported threat defense virtual vCPU/memory configuration. This allows 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 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 Threat Defense Virtual Smart Licensing

Session limits for RA VPNs are determined by the installed 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 21: 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

Threat Defense Virtual Performance Tier Licensing Guidelines and Limitations

Please keep the following guidelines and limitations in mind when licensing your threat defense virtual device.

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

Note Make sure your Smart Licensing account contains the available licenses you need. It's important to choose the tier that matches the license you have in your account. If you are upgrading your threat defense virtual to Version 7.0, you can choose **FTDv - Variable** to maintain your current license compliance. Your 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 threat defense virtual device, or when provisioning the threat defense virtual using the REST API.
- Essentials licenses are subscription-based and mapped to performance tiers. Your virtual account needs to have the Essentials license entitlements for the threat defense virtual devices, as well as for IPS, Malware Defense, and URL Filtering licenses.
- Each HA peer consumes one entitlement, and the entitlements on each HA peer must match, including Essentials 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.

License PIDs

When you bought your device from Cisco or a reseller, your licenses should have been linked to your Smart Software License account. However, if you need to add licenses yourself, use the **Search All** field on the Cisco Commerce Workspace.

Figure 25: License Search

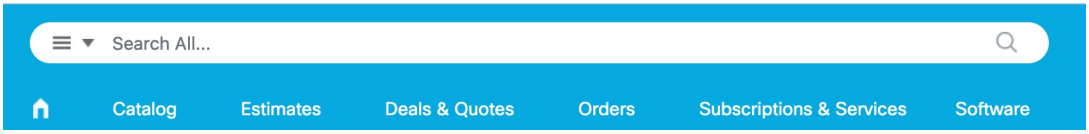

Choose **Products & Services** from the results.

Figure 26: Results

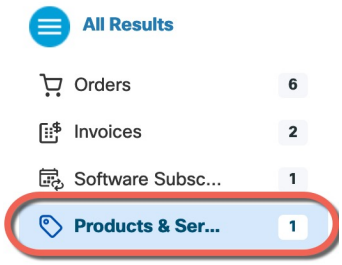

Management Center Virtual PIDs

- VMware:
  - SF-FMC-VMW-2-K9—2 devices
  - SF-FMC-VMW-10-K9—10 devices
  - SF-FMC-VMW-K9—25 devices
  - SF-FMC-VMW-300-K9—300 devices
- KVM:
  - SF-FMC-KVM-2-K9—2 devices
  - SF-FMC-KVM-10-K9—10 devices
  - SF-FMC-KVM-K9—25 devices
- PAK-based VMware:
  - FS-VMW-2-SW-K9—2 devices
  - FS-VMW-10-SW-K9—10 devices

- FS-VMW-SW-K9—25 devices

Threat Defense Virtual PIDs

When you order FTDV-SEC-SUB, you must choose a Essentials license and optional feature licenses (12 month term):

- Essentials license:
  - FTD-V-5S-BSE-K9
  - FTD-V-10S-BSE-K9
  - FTD-V-20S-BSE-K9
  - FTD-V-30S-BSE-K9
  - FTD-V-50S-BSE-K9
  - FTD-V-100S-BSE-K9
- IPS, Malware defense, and URL license combination:
  - FTD-V-5S-TMC
  - FTD-V-10S-TMC
  - FTD-V-20S-TMC
  - FTD-V-30S-TMC
  - FTD-V-50S-TMC
  - FTD-V-100S-TMC
- Carrier—FTDV_CARRIER
- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Firepower 1010 PIDs

- IPS, Malware defense, and URL license combination:
  - L-FPR1010T-TMC=

  When you add the above PID to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-FPR1010T-TMC-1Y
  - L-FPR1010T-TMC-3Y
  - L-FPR1010T-TMC-5Y
- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Firepower 1100 PIDs

- IPS, Malware defense, and URL license combination:
  - L-FPR1120T-TMC=
  - L-FPR1140T-TMC=
  - L-FPR1150T-TMC=

  When you add one of the above PIDs to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-FPR1120T-TMC-1Y
  - L-FPR1120T-TMC-3Y
  - L-FPR1120T-TMC-5Y
  - L-FPR1140T-TMC-1Y
  - L-FPR1140T-TMC-3Y
  - L-FPR1140T-TMC-5Y
  - L-FPR1150T-TMC-1Y
  - L-FPR1150T-TMC-3Y
  - L-FPR1150T-TMC-5Y
- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Firepower 2100 PIDs

- IPS, Malware defense, and URL license combination:
  - L-FPR2110T-TMC=
  - L-FPR2120T-TMC=
  - L-FPR2130T-TMC=
  - L-FPR2140T-TMC=

  When you add one of the above PIDs to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-FPR2110T-TMC-1Y
  - L-FPR2110T-TMC-3Y
  - L-FPR2110T-TMC-5Y
  - L-FPR2120T-TMC-1Y
  - L-FPR2120T-TMC-3Y
  - L-FPR2120T-TMC-5Y
  - L-FPR2130T-TMC-1Y

- L-FPR2130T-TMC-3Y
- L-FPR2130T-TMC-5Y
- L-FPR2140T-TMC-1Y
- L-FPR2140T-TMC-3Y
- L-FPR2140T-TMC-5Y

- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Secure Firewall 3100 PIDs

- Essentials license:
  - *Included automatically*
- IPS, Malware defense, and URL license combination:
  - L-FPR3110T-TMC=
  - L-FPR3120T-TMC=
  - L-FPR3130T-TMC=
  - L-FPR3140T-TMC=

  When you add one of the above PIDs to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-FPR3105T-TMC-1Y
  - L-FPR3105T-TMC-3Y
  - L-FPR3105T-TMC-5Y
  - L-FPR3110T-TMC-1Y
  - L-FPR3110T-TMC-3Y
  - L-FPR3110T-TMC-5Y
  - L-FPR3120T-TMC-1Y
  - L-FPR3120T-TMC-3Y
  - L-FPR3120T-TMC-5Y
  - L-FPR3130T-TMC-1Y
  - L-FPR3130T-TMC-3Y
  - L-FPR3130T-TMC-5Y
  - L-FPR3140T-TMC-1Y
  - L-FPR3140T-TMC-3Y
  - L-FPR3140T-TMC-5Y

- Carrier—L-FPR3K-FTD-CAR=
- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Firepower 4100 PIDs

- IPS, Malware defense, and URL license combination:
  - L-FPR4112T-TMC=
  - L-FPR4115T-TMC=
  - L-FPR4125T-TMC=
  - L-FPR4145T-TMC=

  When you add one of the above PIDs to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-FPR4112T-TMC-1Y
  - L-FPR4112T-TMC-3Y
  - L-FPR4112T-TMC-5Y
  - L-FPR4115T-TMC-1Y
  - L-FPR4115T-TMC-3Y
  - L-FPR4115T-TMC-5Y
  - L-FPR4125T-TMC-1Y
  - L-FPR4125T-TMC-3Y
  - L-FPR4125T-TMC-5Y
  - L-FPR4145T-TMC-1Y
  - L-FPR4145T-TMC-3Y
  - L-FPR4145T-TMC-5Y
- Carrier—L-FPR4K-FTD-CAR=
- Cisco Secure Client—See the Cisco Secure Client Ordering Guide.

Firepower 9300 PIDs

- IPS, Malware defense, and URL license combination:
  - L-FPR9K-40T-TMC=
  - L-FPR9K-48T-TMC=
  - L-FPR9K-56T-TMC=

  When you add one of the above PIDs to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:

- L-FPR9K-40T-TMC-1Y
- L-FPR9K-40T-TMC-3Y
- L-FPR9K-40T-TMC-5Y
- L-FPR9K-48T-TMC-1Y
- L-FPR9K-48T-TMC-3Y
- L-FPR9K-48T-TMC-5Y
- L-FPR9K-56T-TMC-1Y
- L-FPR9K-56T-TMC-3Y
- L-FPR9K-56T-TMC-5Y

- Carrier—L-FPR9K-FTD-CAR=
- Cisco Secure Client—See the Cisco AnyConnect Ordering Guide.

ISA 3000 PIDs

- IPS, Malware defense, and URL license combination:
  - L-ISA3000T-TMC=

  When you add the above PID to your order, you can then choose a term-based subscription corresponding with one of the following PIDs:
  - L-ISA3000T-TMC-1Y
  - L-ISA3000T-TMC-3Y
  - L-ISA3000T-TMC-5Y
- Cisco Secure Client—See the Cisco AnyConnect Ordering Guide.

Requirements and Prerequisites for Licensing

For Specific License Reservation requirements, see Requirements and Prerequisites for Specific License Reservation, on page 288.

General Prerequisites

- Make sure NTP is configured on the 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 management center.

Supported Domains

Global, except where indicated.

User Roles

- Admin

Requirements and Prerequisites for Licensing for High Availability, Clustering, and Multi-Instance

This section describes the licensing requirements for High Availability (for device High Availability and also management center virtual High Availability), clustering, and multi-instance deployments.

Licensing for Management Center High Availability

Each device requires the same licenses whether managed by a single management center or by management centers in a high availability pair (hardware or virtual).

Example: If you want to enable advanced malware protection for two devices managed by a management center pair, buy two Malware Defense licenses and two TM subscriptions, register the active management center with the Smart Software Manager, then assign the licenses to the two devices on the active management center.

Only the active management center is registered with the Smart Software Manager. When failover occurs, the system communicates with Smart Software Manager to release the license entitlements from the originally-active management center and assign them to the newly-active management center.

In Specific License Reservation deployments, only the primary management center requires a Specific License Reservation.

Hardware Management Center

No special license is required for hardware management centers in a high availability pair.

Management Center Virtual

You will need two identically licensed management center virtuals.

Example: For the management center virtual high availability pair managing 10 devices, you can use:

- Two (2) management center virtual 10 entitlements
- 10 device licenses

If you break the high availability pair, the management center virtual entitlements associated with the secondary management center virtual are released. (In the example, you would then have two standalone management center virtual 10s.)

Licensing for Device High-Availability

Both 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 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 Essentials license and a IPS license, and the standby unit has only a Essentials license, the management center communicates with the Smart Software Manager to obtain an available IPS 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.

Licensing for Device Clusters

Each 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.

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 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 in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the management center is licensed (and running in Evaluation mode), then when you license the 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.

Licensing for Multi-Instance Deployments

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 following details:

- Essentials 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 22: Sample License Usage for Container Instances on a Firepower 9300

Firepower 9300	Instance	Licenses
Security Module 1	Instance 1	Essentials, URL Filtering, Malware Defense
	Instance 2	Essentials, URL Filtering
	Instance 3	Essentials, URL Filtering

Firepower 9300	Instance	Licenses
Security Module 2	Instance 4	Essentials, IPS
	Instance 5	Essentials, URL Filtering, Malware Defense, IPS
Security Module 3	Instance 6	Essentials, Malware Defense, IPS
	Instance 7	Essentials, IPS

Table 23: Total Number of Licenses

Essentials	URL Filtering	Malware Defense	IPS
3	2	3	2

Create a Cisco Account

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.

The following figure shows an example.

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. See Create a Smart Account and Add Licenses, on page 274.

Create a Smart Account and Add Licenses

You should set up this account before you purchase licenses.

Before you begin

Your account representative or reseller may have set up a Smart Account on your behalf. If so, obtain the necessary information to access the account from that person instead of using this procedure, then verify that you can access the account.

You must create a new Cisco account if you don't already have one. For instructions, see Create a Cisco Account.

For general information about Smart Accounts, see 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**.

Step 4 In the **Your Company Details** section, click **Edit**.

Step 5 In the **Company or organization** field, type your organization name.

Step 6 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.

Step 7 If your company is not listed in our database, you can continue to enter your company information in the **Company or organization** field.

a) In the **Address** drop-down list, click the drop-down arrow and click **Add New Address**.

b) 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.

Step 8 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 9 Click **Update**.

A success message appears if your company details are validated.

Step 10 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 11 Click **Create Account**.

The Account Summary page displays your account details.

Step 12 Click **Done**.

Step 13 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 14 Make sure your Smart Licensing account contains the available licenses you need.

For license PIDs, see License PIDs, on page 265.

What to do next

To configure Smart License using the Smart Software Manager, see Configure Smart Licensing, on page 275.

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 Configure Specific License Reservation (SLR), on page 287.

Register the Management Center for Smart Licensing

You can register the 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 Management Center with the Smart Software Manager

Register the management center with the Smart Software Manager.

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. For license PIDs, see License PIDs, on page 265.

- Ensure that the management center can reach the Smart Software Manager at smartreceiver.cisco.com.
- Make sure you 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, you must configure NTP on the chassis using the same NTP server for the chassis as for the management center.
- If your organization has multiple management centers, make sure each management center has a unique name that clearly identifies and distinguishes it from other management centers 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**.

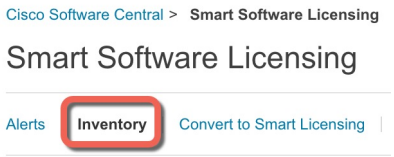

b) On the **General** tab, click **New Token**.

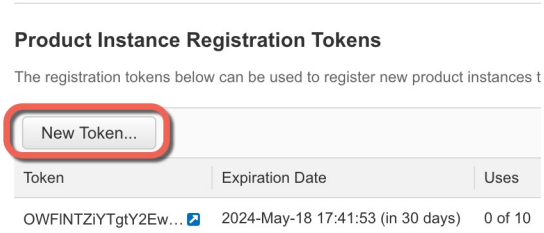

c) On the **Create Registration Token** dialog box enter the following settings, and then click **Create Token**:

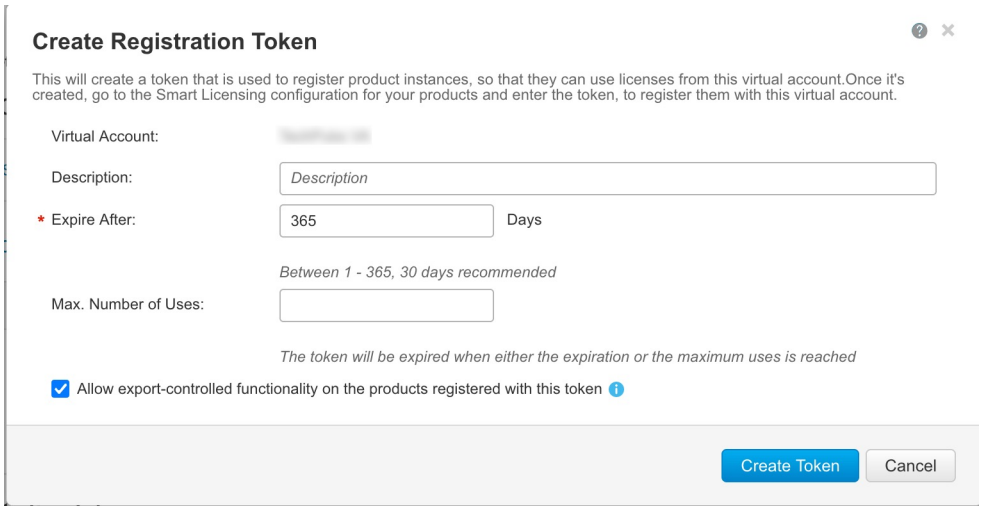

- **Description**
- **Expire After**—Cisco recommends 30 days.
- **Max. Number of Uses**
- **Allow export-controlled functionality on the products registered with this token**—Enables the export-compliance flag if you are in a country that allows for strong encryption. You must select this option now 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 threat defense.

Figure 27: View Token

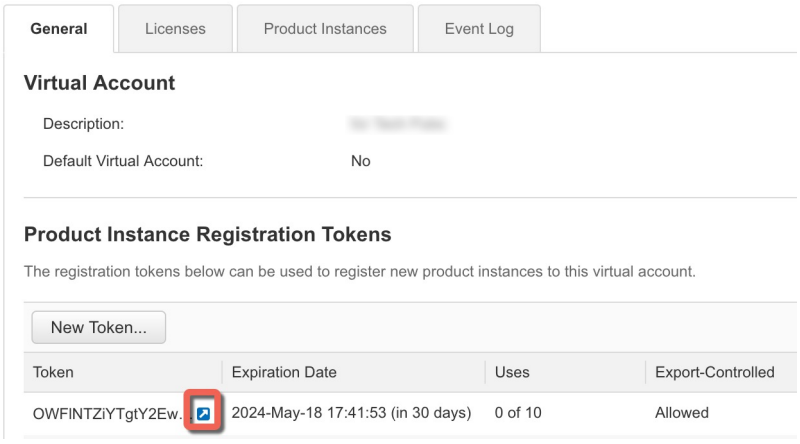

Figure 28: Copy Token

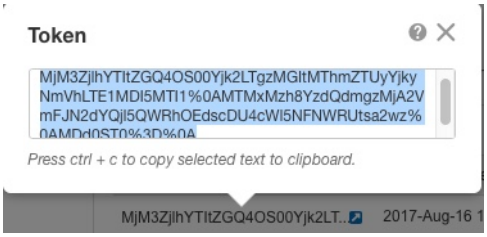

Step 2 In the management center, choose **System** (⚙) > **Licenses** > **Smart Licenses**.

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 Management Center to Share Usage Metrics and Statistics with Cisco , on page 44.
- **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 Management Center to Share Device Health Data with Cisco, on page 45.

Note

- When enabled, Cisco Support Diagnostics is enabled in the devices in the next sync cycle. The 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 management center.

Step 6 Click **Apply Changes**.

What to do next

- Add a Device to the management center; see *Add a Device to the 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 282.

Register the Management Center with the Smart Software Manager On-Prem

As described in Periodic Communication with the Smart Software Manager, on page 255, the 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 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 254.)

- Your 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, see 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 management center with the Smart Software Manager On-Prem.

a) Choose **Integration** > **Other Integrations**.
b) Click **Smart Software Satellite**.
c) Select **Connect to Cisco Smart Software Satellite Server**.
d) 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`*`/SmartTransport`

   The FQDN or hostname must match the CN value of the certificate presented by your Smart Software Manager On-Prem.

e) Add a new **SSL Certificate** and paste the certificate text that you copied earlier.
f) Click **Apply**.
g) Select **System > Licenses > Smart Licenses** and click **Register**.
h) Create a new token on Smart Software Manager On-Prem.
i) Copy the token.
j) Paste the token into the form on the management center page.
k) Click **Apply Changes**.

   The 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.

See 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, see 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, on the **Inventory** > **Licenses** page, verify that you have the license that corresponds to your management center:

Export Control License	Management Center Model
Cisco Virtual FMC Series Strong Encryption (3DES/AES)	All 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**

What to do next

You can now deploy configurations or policies that use the export-controlled features.

Remember The new export-controlled licenses and all features enabled by it do not take effect on the 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 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 management center. You can also assign licenses per device, or for multiple devices.

Assign Licenses to a Single Device

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 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: Go to **System** ( ) > **Licenses** > **Smart Licenses**, enter the hostname or IP address of the device into the filter at the top of the Smart Licenses table, and verify that only a green circle with a **Check Mark** ( ) appears for each device, for each license type. If you see any other icon, hover over the icon for more information.

Assign Licenses to Multiple Managed Devices

Devices managed by the management center obtain their licenses via the management center, not directly from the Smart Software Manager.

Use this procedure to enable licensing on multiple devices at once.

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 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.

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 Monitoring Smart Licenses, on page 284.

Manage Smart Licensing

This section describes how to manage Smart Licensing.

Deregister the Management Center

Deregister your 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. For example, deregister if you need to decommission the management center or reimage it.

See Unregistered State, on page 256 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 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**.

Monitoring Smart License Status

The **Smart License Status** section of the **System > Licenses > Smart Licenses** page provides an overview of license usage on the management center, as described below.

Usage Authorization

Possible status values are:

- **In-compliance** () — All licenses assigned to managed devices are in compliance and the 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 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 management center has not communicated with the Smart Software Manager in more than 90 days.

Product Registration

Specifies the last date when the 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 management center. If this deployment is not associated with a particular virtual account within your Smart Account, this information is not displayed.

Export-Controlled Features

If this option is enabled, you can deploy restricted features. For details, see Licensing for Export-Controlled Functionality, on page 262.

Cisco Success Network

Specifies whether you have enabled Cisco Success Network for the 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 Management Center to Share Usage Metrics and Statistics with Cisco , on page 44 for more information.

Monitoring Smart Licenses

To view the license status for the 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 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 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 353 and Dashboard Widget Availability by User Role, on page 341 and The Product Licensing Widget, on page 350.

- 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 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.

Troubleshooting Smart Licensing

Expected Licenses Do Not Appear in My Smart Account

If the licenses you expect to see are not in your Smart Account, try the following:

- Make sure they are not in a different Virtual Account. Your organization's license administrator may need to assist you with this.
- Check with the person who sold you the licenses to be sure that transfer to your account is complete.

Unable to Connect to Smart License Server

Check the obvious causes first. For example, make sure your management center has outside connectivity. See Internet Resources Accessed, on page 1031.

Unexpected Out-of-Compliance Notification or Other Error

- If a device is already registered to a different management center, you need to deregister the original management center before you can license the device under a new management center. See Deregister the Management Center, on page 283.

- Check if the term of the subscription license has expired.

Troubleshoot Other Issues

For solutions to other common issues, see https://www.cisco.com/c/en/us/support/docs/security/firepower-management-center/215838-fmc-and-ftd-smart-license-registration-a.html

Convert a Classic License for Use on the Threat Defense

You can convert licenses using either the License Registration Portal or the Smart Software Manager, and you can convert an unused Product Authorization Key (PAK) or a Classic license that has already been assigned to a device.

Note You cannot undo this process. You cannot convert a Smart License to a Classic license, even if the license was originally a Classic license.

In documentation on Cisco.com, Classic licenses may also be referred to as "traditional" licenses.

Before you begin

- It is easiest to convert a Classic license to a Smart License when it is still an unused PAK that has not yet been assigned to a product instance.
- Your hardware must be able to run threat defense. See the *Cisco Secure Firewall Threat Defense Compatibility Guide* at https://www.cisco.com/c/en/us/support/security/defense-center/products-device-support-tables-list.html.
- You must have a Smart Account. If you do not have one, create one. See Create a Smart Account and Add Licenses, on page 274.
- The PAKs or licenses that you want to convert must appear in your Smart Account.
- If you convert using the License Registration Portal instead of the Smart Software Manager, you must have your Smart Account credentials in order to initiate the conversion process.

Procedure

Step 1 The conversion process you follow depends on whether or not the license has been consumed:

- If the PAK that you want to convert has never been used, follow instructions for converting a PAK.
- If the PAK you want to convert has already been assigned to a device, follow instructions for converting a Classic license.

  Make sure your existing classic license is still registered to your device.

Step 2 See instructions for your type of conversion (PAK or installed Classic license) in the following documentation:

- To convert PAKs or licenses using the License Registration Portal:

- View a video that takes you through the License Registration Portal part of the conversion process at Cisco Classic Licensing Management with Smart Accounts.
- Sign in to the License Registration Portal at https://tools.cisco.com/SWIFT/LicensingUI/Home and follow the instructions in the documentation above.

- To convert PAKs or licenses using the Smart Software Manager:
  - *Converting Hybrid Licenses to Smart Software Licenses QRG*:

    https://community.cisco.com/t5/licensing-enterprise-agreements/converting-hybrid-licenses-to-smart-software-licenses-qrg/ta-p/3628609?attachment-id=134907
  - Sign in to the Smart Software Manager at https://software.cisco.com/#SmartLicensing-LicenseConversion and follow the instructions for your type of conversion (PAK or installed Classic license) in the documentation above.

Step 3 Freshly install threat defense on your hardware.

See the instructions for your hardware at Install and Upgrade Guides .

Step 4 If you will use the device manager to manage this device as a standalone device:

See information about licensing the device in the device manager configuration guide at Secure Firewall Device Manager Configuration Guides.

Skip the rest of this procedure.

Step 5 If you have already deployed Smart Licensing on your management center:

a) Set up Smart Licensing on your new threat defense.

See Assign Licenses to Multiple Managed Devices, on page 282.

b) Verify that the new Smart License has been successfully applied to the device.

See Monitoring Smart Licenses, on page 284.

Step 6 If you have not yet deployed Smart Licensing on your management center:

See Configure Smart Licensing, on page 275. (Skip any steps that do not apply or that you have already completed.)

Configure Specific License Reservation (SLR)

You can use the Specific License Reservation feature to deploy Smart Licensing in an air-gapped network.

Note Various names are used at Cisco for Specific License Reservation, including 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.

When Specific License Reservation is enabled, the management center reserves licenses from your virtual account for a specified duration without accessing the Smart Software Manager or using Smart Software Manager On-Prem.

Features that require access to the internet, such as URL Lookups or contextual cross-launch to public web sites, will not work.

Cisco does not collect web analytics or telemetry data for deployments that use Specific License Reservation.

Requirements and Prerequisites for Specific License Reservation

- If you are currently using regular Smart Licensing, de-register the management center before you implement Specific License Reservation. For information, see Deregister the Management Center, on page 283.

  All Smart Licenses that are currently deployed to the 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 uses the same licenses as regular Smart Licensing.
- (Recommended) If you deploy the management center pair in a high availability configuration, note the following:
  - Configure high availability before you assign licenses. If you already assigned licenses to devices on the secondary management center, be sure to unassign them.
  - If an SLR license is assigned to a primary management center, when the secondary management center becomes active after a failover, you cannot add the SLR license to the secondary management center. You must do one of the following:
    - Perform a failover to make the primary management center active.
    - Unassign and re-assign the license to the secondary 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 management center.

Before you begin

- Ensure that you have met the requirements described in Requirements and Prerequisites for Specific License Reservation, on page 288.
- 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 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 management center.

Procedure

Step 1 Access the management center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.

Step 2 Log into the 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 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 Management Center

Procedure

Step 1 Generate the reservation request code.

a) In the 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**.
e) Click **License Reservation**.
f) Enter the code that you generated from management center into the **Reservation Request Code** box.
g) Click **Next**.
h) Select **Reserve a specific license**.
i) Scroll down to display the entire License grid.
j) Under **Quantity To Reserve**, enter the number of each platform and feature license needed for your deployment.

Note

- You must explicitly include a Essentials license for each managed device, or, for multi-instance deployments, for each container.
- If you are using the 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-management center, you must select the appropriate license.

For the correct license name to choose for your management center, see the prerequisites in Enable the Export Control Feature for Accounts Without Global Permission, on page 280.

k) Click **Next**.
l) Click **Generate Authorization Code**.

At this point, the license is now in use according to the Smart Software Manager.

m) Download the Authorization Code in preparation for entering it into the management center.

Step 3 Enter the authorization code in the management center.

a) In the 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.
d)

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, see 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.

Important! 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 236.

To ensure that all functionality continues to work without interruption, monitor your license expiration dates (on the **Reserved Licenses** tab). If any of the licenses expire, the 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 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, the following actions are restricted:

- Device registration
- Policy deployment

Procedure

Step 1 In the management center, obtain the unique product instance identifier of this 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 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 management center is the correct management center. Click the name.

e) Look at the **UUID** and see if it is the UUID of the management center that you are trying to modify.

If not, you must repeat these steps until you find the correct management center.

Step 3 When you have located the correct 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, choose **Actions** > **Update Reserved Licenses**.

b) Update the reserved licenses as needed.

Note

- You must explicitly include a Essentials license for each managed device, or, for multi-instance deployments, for each container.

- If you are using the 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-management center, you must select the appropriate license.

    For the correct license name to choose for your management center, see the prerequisites in Enable the Export Control Feature for Accounts Without Global Permission, on page 280.

c) Click **Next** and verify the details.
d) Click **Generate Authorization Code**.
e) Download the Authorization Code in preparation for entering it into the 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 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 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) Choose **Actions** > **Enter Confirmation Code**:

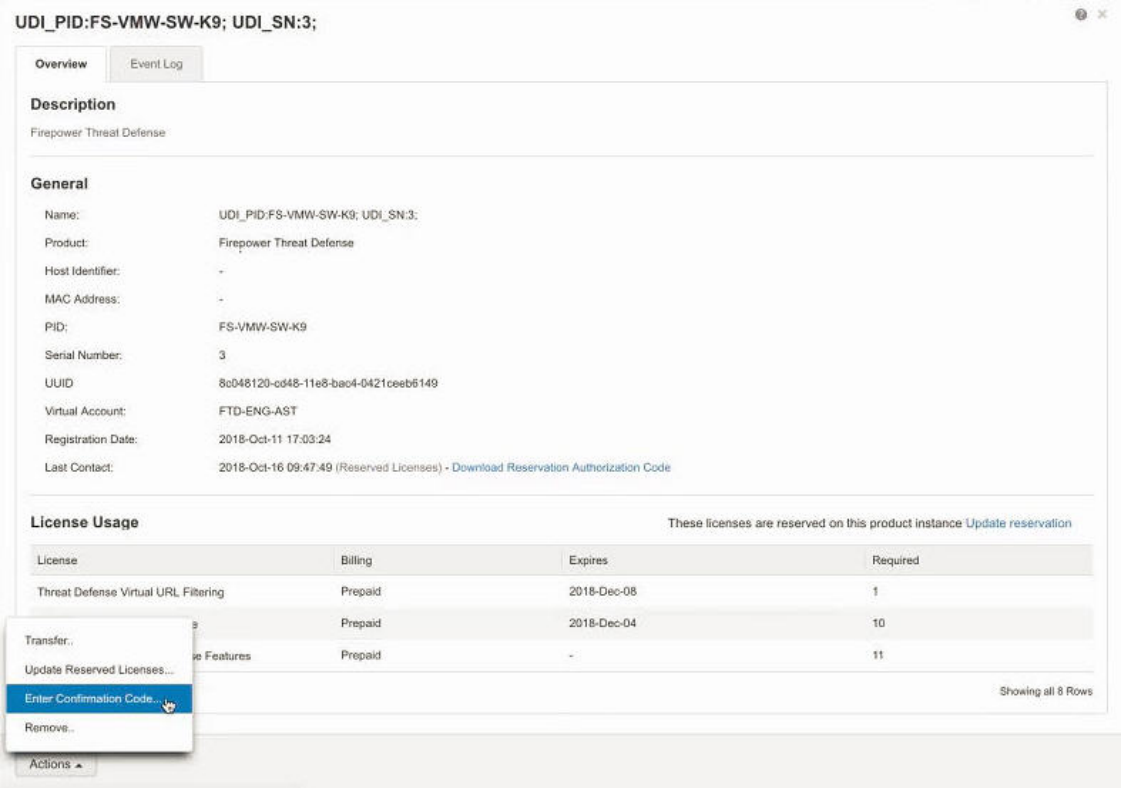

c) Enter the confirmation code that you generated from the management center.

Step 6 In the 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, see Monitoring Specific License Reservation Status, on page 296 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).

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.

This procedure releases all license entitlements associated with the management center back to your virtual account. After you de-register, no updates or changes on licensed features are allowed.

Procedure

Step 1 In the management center Web interface, select **System** > **Licenses** > **Specific Licenses**.

Step 2 Make a note of the **Product Instance** identifier for this management center.

Step 3 Generate a return code from the management center.

a) Click **Return SLR**.

The following figure shows Return SLR.

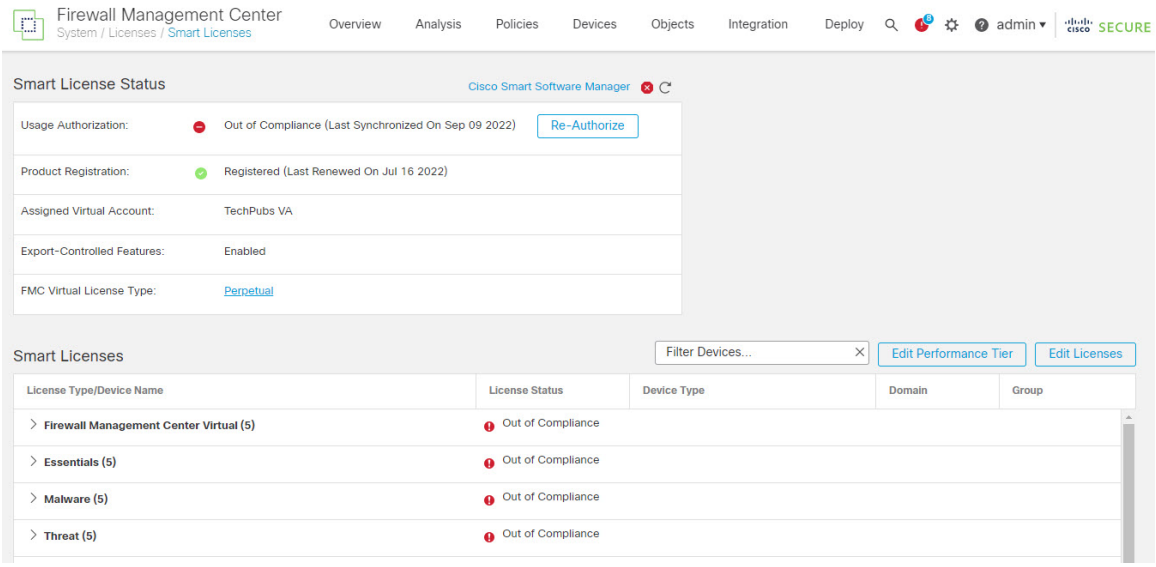

Devices become unlicensed and the management center moves to the de-registered state. It generates a return code and allows you to re-register the management center with SLR.

b) Make a note of the **Return Code**.

Step 4 In the Smart Software Manager, identify the 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 management center is the correct management center. Click the name.

e) Look at the **UUID** and see if it is the UUID of the management center that you are trying to modify.

If not, you must repeat these steps until you find the correct management center.

Step 5 When you have identified the correct management center, return the licenses to your Smart Account:

a) On the page that shows the correct UUID, choose **Actions** > **Remove**.

b) Enter the reservation return code that you generated from the 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 management center is removed from the Smart Software Manager Product Instances list.

Step 6 Disable the Specific License in the management center Linux shell:

a) Access the management center console using a USB keyboard and VGA monitor, or use SSH to access the management interface.
b) Log in to the 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 the following 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.

Monitoring Specific License Reservation Status

The **System** > **Licenses** > **Specific Licenses** page provides an overview of license usage on the management center, as described below.

Usage Authorization

Possible status values are:

- **Authorized** — The 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 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 management center.

Export-Controlled Features

Specifies whether you have enabled export-controlled functionality for the management center.

For more information about Export-Controlled Features, see Licensing for Export-Controlled Functionality, on page 262.

Product Instance

The Universally Unique Identifier (UUID) of this 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

How do I identify a particular management center in the Product Instance list in Smart Software Manager?

On 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 management center web interface, the UUID for the management center is the **Product Instance** value displayed on 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 the following:

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 294.
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 management center virtual

Make sure you have enough 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.

See Update a Specific License Reservation, on page 292.

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 see, Deactivate and Return the Specific License Reservation, on page 294.

What if I do not see a Specific License page in the 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 289.

I have disabled Specific Licensing, but forgot to copy the Return Code. What should I do?

The **Return Code** is saved in the management center virtual. You must re-enable the Specific License from the Linux shell (see Enable the Specific Licensing Menu Option, on page 289), then refresh the management center virtual web interface. Your **Return Code** will be displayed.

Configure Legacy Management Center PAK-Based Licenses

The 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 The license key uniquely identifies the 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 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 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.

Step 10 Click **Verify License**.

If the license is invalid, make sure that you correctly copied the license text.

Step 11 Click **Submit License**.

Additional Information about Licensing

For additional information to help resolve common licensing questions, see the following documents:

- FAQ—Licensing FAQ
- License Roadmap

History for Licenses

Feature	Minimum Management Center	Minimum Threat Defense	Details
Smart Licensing Standardization	7.3	Any	We changed the following license names in the management center GUI: • Base is now Essentials • Threat is now IPS • Malware is now Malware Defense • RA VPN/AnyConnect License is now Cisco Secure Client • AnyConnect Plus is now Secure Client Advantage • AnyConnect Apex is now Secure Client Premier • AnyConnect Apex and Plus is now Secure Client Premier and Advantage • AnyConnect VPN Only is now Secure Client VPN Only
Support for Carrier license	7.3	Any	The Carrier license enables inspection of Diameter, GTP/GPRS, SCTP, and M3UA protocols. New/Modified screens: **System** > **Smart Licenses**
Performance tier licensing for the 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 threat defense virtual models.
Licensing for multi-instance capability for the threat defense on the Firepower 4100/9300	6.3	Any	You can now deploy multiple 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: 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: management center, 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: management center, threat defense

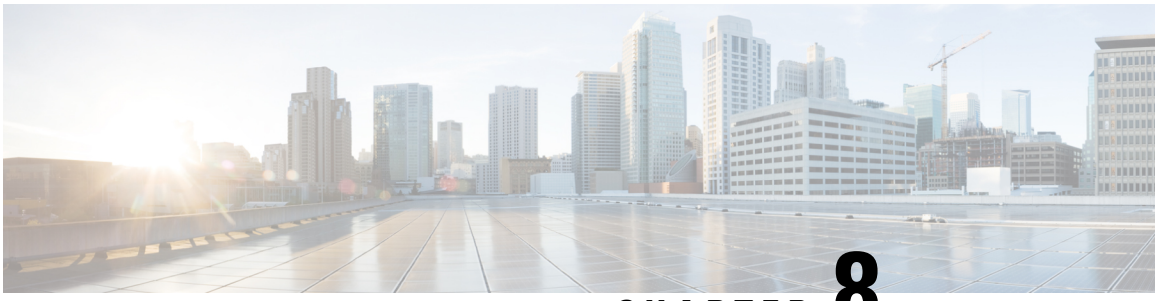

CHAPTER 8

High Availability

The following topics describe how to configure Active/Standby high availability of Cisco Secure Firewall Management Centers:

About Management Center High Availability

To ensure the continuity of operations, the high availability feature allows you to designate redundant management centers to manage devices. The management centers support Active/Standby high availability where one appliance is the active unit and manages devices. The standby unit does not actively manage devices. The active unit writes configuration data into a data store and replicates data for both units, using synchronization where necessary to share some information with the standby unit.

Active/Standby high availability lets you configure a secondary management center to take over the functionality of a primary management center if the primary fails. When the primary management center fails, you must promote the secondary management center to become the active unit.

Event data streams from managed devices to both management centers in the high availability pair. If one management center fails, you can monitor your network without interruption using the other management center.

Note that 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 management center, if that appliance fails, you must promote the standby management center to active.

Note Triggering a switchover on management center immediately after a successful change deployment can lead to preview configuration not working on the new active management center. This does not impact policy deploy functionality. It is recommended to trigger a switchover on the management center after the necessary sync is completed.

Similarly, when management center HA synchronization is in degraded state, triggering a switchover or changing roles could make management center HA to damage the database and it can become catastrophic. 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 Replacing Management Centers in a High Availability Pair, on page 320 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.

About 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 to have the same CertEnrollment object enrolled on the two devices before the high availability formation.

SNMP Behavior in Management Center High Availability

In an SNMP-configured HA pair, when you deploy an alert policy, the active management center sends the SNMP traps. When the primary management center fails, the secondary management center which becomes the active unit starts sending the SNMP traps without the need for any additional configuration.

Roles v. Status in Management Center High Availability

Primary/Secondary Roles

When setting up Secure Firewall Management Centers in a high availability pair, you configure one Secure 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 Secure Firewall Management Center becomes the active peer, while the secondary Secure Firewall Management Center becomes the standby peer, and the two units act as a single appliance for managed device and policy configuration.

Active/Standby Status

The main differences between the two Secure Firewall Management Centers in a high availability pair are related to which peer is active and which peer is standby. The active Secure Firewall Management Center remains fully functional, where you can manage devices and policies. On the standby Secure Firewall Management Center, functionality is hidden; you cannot make any configuration changes.

Event Processing on Management Center High Availability Pairs

Since both management centers in a high availability pair receive events from managed devices, the management IP addresses for the appliances are not shared. This means that you do not need to intervene to ensure continuous processing of events if one of the management center fails.

AMP Cloud Connections and Malware Information

Although they share file policies and related configurations, management centers in a high availability pair share neither Cisco AMP cloud connections nor malware dispositions. To ensure continuity of operations, and to ensure that detected files' malware dispositions are the same on both management centers, both primary and secondary management centers must have access to the AMP cloud.

URL Filtering and Security Intelligence

URL filtering and Security Intelligence configurations and information are synchronized between Secure Firewall Management Centers in a high availability deployment. However, only the primary Secure Firewall Management Center downloads URL category and reputation data for updates to Security Intelligence feeds.

If the primary Secure Firewall Management Center fails, not only must you make sure that the secondary Secure Firewall Management Center can access the internet to update threat intelligence data, but you must also use the web interface on the secondary Secure Firewall Management Center to promote it to active.

User Data Processing During Management Center Failover

If the primary management center fails, the Secondary management center propagates to managed devices user-to-IP mappings from the TS Agent identity source; and propagates SGT mappings from the ISE/ISE-PIC identity source. Users not yet seen by identity sources are identified as Unknown.

After the downtime, the Unknown users are re identified and processed according to the rules in your identity policy.

Configuration Management on Management Center High Availability Pairs

In a high availability deployment, only the active management center can manage devices and apply policies. Both management centers remain in a state of continuous synchronization.

If the active management center fails, the high availability pair enters a degraded state until you manually promote the standby appliance to the active state. Once the promotion is complete, the appliances leave maintenance mode.

Management Center High Availability Disaster Recovery

In case of a disaster recovery situation, a manual switchover must be performed. When the primary management center - FMC1 fails, access the web interface of the secondary management center - FMC2 and switch peers.

This is applicable conversely also in case the secondary (FMC2) fails. For more information, see Switching Peers in the Management Center High Availability Pair, on page 318.

For restoring a failed management center, refer to Replacing Management Centers in a High Availability Pair, on page 320.

Single Sign-On and High Availability Pairs

Management Centers in a high availability configuration can support Single Sign-On, but you must keep the following considerations in mind:

- SSO configuration is not synchronized between the members of the high availability pair; you must configure SSO separately on each member of the pair.
- Both management centers in a high availability pair must use the same identity provider (IdP) for SSO. You must configure a service provider application at the IdP for each management center configured for SSO.
- In a high availability pair of management centers where both are configured to support SSO, before a user can use SSO to access the secondary management center for the first time, that user must first use SSO to log into the primary management center at least once.
- When configuring SSO for management centers in a high availability pair:
  - If you configure SSO on the primary management center, you are not required to configure SSO on the secondary management center.
  - If you configure SSO on the secondary management center, you are required to configure SSO on the primary management center as well. (This is because SSO users must log in to the primary management center at least once before logging into the secondary management center.)

Related Topics

Configure SAML Single Sign-On, on page 144

Management Center High Availability Behavior During a Backup

When you perform a Backup on a management center high availability pair, the Backup operation pauses synchronization between the peers. During this operation, you may continue using the active management center, but not the standby peer.

After Backup is completed, synchronization resumes, which briefly disables processes on the active peer. During this pause, the High Availability page briefly displays a holding page until all processes resume.

Management Center High Availability Split-Brain

If the active management center in a high-availability pair goes down (due to power issues, network/connectivity issues), you can promote the standby management center to an active state. When the original active peer comes up, both peers can assume they are active. This state is defined as 'split-brain'. When this situation occurs, the system prompts you to choose an active appliance, which demotes the other appliance to standby.

If the active management center goes down (or disconnects due to a network failure), you may either break high availability or switch roles. The standby management center enters a degraded state.

Note Whichever appliance you use as the intended standby loses all of its device registrations and policy configurations when you resolve split-brain. For example, you would lose modifications to any policies that existed on the intended standby but not on the intended active. If the 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 management center before re-establishing high availability. You may then register the managed devices and import the policies to the intended active management center.

Troubleshooting Management Center High Availability

This section lists troubleshooting information for some common management center high availability operation errors.

Error	Description	Solution
You must reset your password on the active management center before you can log in to the standby.	You attempted to log into the standby management center when a force password reset is enabled for your account.	As the database is read-only for a standby management center, reset the password on the login page of the active management center.
500 Internal	May appear when attempting to access the web interface while performing critical management center high availability operations, including 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.	May appear when the management center reboots (manually or while recovering from a power down) during a high availability or data synchronization operation.	1. Access the management center shell and use the `manage_hadc.pl` command to access the 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 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	During the initial configuration of a threat defense, if the management center IP address and NAT ID are specified, the **Host** field can be left blank. However, in an HA environment with both the management centers behind a NAT, this error occurs when you add the threat defense on the secondary management center.	1. Delete the threat defense from primary management center. See *Delete a Device from the* Management Center in Cisco Secure Firewall Management Center Device Configuration Guide. 2. Remove managers from threat defense using the **configure manager delete** command. See Cisco Secure Firewall Threat Defense Command Reference. 3. Add threat defense to the management center with the IP address or name of the threat defense device in the **Host** field. See *Add a Device to the* Management Center in Cisco Secure Firewall Management Center Device Configuration Guide.
Device Registration Status:Host <*string*> is not reachable	The error occurs when adding threat defense device to the secondary management center center in a high-availability deployment where both the secondary management center and the threat defense device are behind NAT.	On the standby management center web interface, click **Integration** > **Other Integrations** > **High Availability**. Under the pending device registration table, click the IP address of the pending device, and then change the IP address to the public IP address of the threat defense. OR 1. Access the threat defense shell and use the `show managers` command to get the standby management center entry identifier value. 2. In the threat defense shell, edit the standby 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 Using CLI to Resolve Device Registration in Management Center High Availability, on page 317.

Requirements for Management Center High Availability

Model Support

See Hardware Requirements, on page 307.

Virtual Model Support

See Virtual Platform Requirements, on page 307.

Supported Domains

Global

User Roles

Admin

Hardware Requirements

- All management center hardware supports high availability. The peers must be the same model.
- The peers may be physically and geographically separated from each other in different data centers.
- 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 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 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 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.
- See also License Requirements for Management Center High Availability Configurations, on page 308.

Virtual Platform Requirements

High availability is supported for the following public cloud platforms:

- Amazon Web Services (AWS)
- Oracle Cloud Infrastructure (OCI)

And these on-prem/private cloud platforms:

- Cisco HyperFlex
- Kernel-based virtual machine (KVM)
- VMware vSphere/VMware ESXi

The management centers must have the same device management capacity (not supported on FMCv2) and be identically licensed. You also need one threat defense entitlement for each managed device. For more information, see License Requirements for Management Center High Availability Configurations, on page 308.

Note If you are managing Version 7.0.x Classic devices only (NGIPSv or ASA FirePOWER), you do not need FMCv entitlements.

Software Requirements

Access the **Appliance Information** widget to verify the software version, the intrusion rule update version and the vulnerability database update. By default, the widget appears on the **Status** tab of the **Detailed Dashboard** and the **Summary Dashboard**. For more information, see The Appliance Information Widget, on page 343

- The two 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 management centers in a high availability configuration must have the same version of the intrusion rule update installed.
- The two management centers in a high availability configuration must have the same version of the vulnerability database update installed.
- The two 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.

Warning If the software versions, intrusion rule update versions and vulnerability database update versions are not identical on both management centers, you cannot establish high availability.

License Requirements for Management Center High Availability Configurations

Each device requires the same licenses whether managed by a single management center or by management centers in a high availability pair (hardware or virtual).

Example: If you want to enable advanced malware protection for two devices managed by a management center pair, buy two Malware Defense licenses and two TM subscriptions, register the active management center with the Smart Software Manager, then assign the licenses to the two devices on the active management center.

Only the active management center is registered with the Smart Software Manager. When failover occurs, the system communicates with Smart Software Manager to release the license entitlements from the originally-active management center and assign them to the newly-active management center.

In Specific License Reservation deployments, only the primary management center requires a Specific License Reservation.

Hardware Management Center

No special license is required for hardware management centers in a high availability pair.

Management Center Virtual

You will need two identically licensed management center virtuals.

Example: For the management center virtual high availability pair managing 10 devices, you can use:

- Two (2) management center virtual 10 entitlements
- 10 device licenses

If you break the high availability pair, the management center virtual entitlements associated with the secondary management center virtual are released. (In the example, you would then have two standalone management center virtual 10s.)

Prerequisites for Management Center High Availability

Before establishing the management center high availability pair:

- Export required policies from the intended secondary management center to the intended primary management center. For more information, see Exporting Configurations, on page 509.
- Make sure that the intended secondary management center does not have any devices added to it. Delete devices from the intended secondary management center and register these devices to the intended primary management center. For more information see *Delete a Device from the Management Center* and *Add a Device to the Management Center* in the Cisco Secure Firewall Management Center Device Configuration Guide.
- Import the policies into the intended primary management center. For more information, see Importing Configurations, on page 510.
- On the intended primary 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 management center, associate the appropriate licenses to the newly added devices. For more information see Assign Licenses to a Single Device, on page 281.

You can now proceed to establish high availability. For more information, see Establishing Management Center High Availability, on page 310.

Establishing Management Center High Availability

Establishing high availability can take a significant amount of time, even several hours, depending on the bandwidth between the peers and the number of policies. It also depends on the number of devices registered to the active management center, which need to be synced to the standby 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 management centers adhere to the high availability system requirements. For more information , see Requirements for Management Center High Availability, on page 307.
- Confirm that you completed the prerequisites for establishing high availability. For more information, see Prerequisites for Management Center High Availability, on page 309.
- In a multidomain deployment, you must be in the Global domain to perform this task.

Procedure

Step 1 Log into the management center that you want to designate as the secondary.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Under Role for this management center, choose **Secondary**.

Step 5 Enter the hostname or IP address of the primary management center in the **Primary Firewall Management Center Host** text box.

You can leave this empty if the primary management center does not have an IP address reachable from the peer management center (which can be public or private IP address). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You need to specify the IP address of at least one management center to enable HA connection.

Step 6 Enter a one-time-use registration key in the **Registration Key** text box.

The registration key is any user-defined alphanumeric value up to 37 characters in length. This registration key will be used to register both -the secondary and the primary management centers.

Step 7 If you did not specify the primary IP address, or if you do not plan to specify the secondary IP address on the primary management center, then in the **Unique NAT ID** field, enter a unique alphanumeric ID. See NAT Environments, on page 77 for more information.

Step 8 Click **Register**.

Step 9 Using an account with Admin access, log into the management center that you want to designate as the primary.

Step 10 Choose **Integration** > **Other Integrations**.

Step 11 Choose **High Availability**.

Step 12 Under Role for this management center, choose **Primary**.

Step 13 Enter the hostname or IP address of the secondary management center in the **Secondary Firewall Management Center Host** text box.

You can leave this empty if the secondary management center does not have an IP address reachable from the peer management center (which can be public or private IP address). In this case, use both the **Registration Key** and the **Unique NAT ID** fields. You need to specify the IP address of at least one management center to enable HA connection.

Step 14 Enter the same one-time-use registration key in the **Registration Key** text box you used in step 6.

Step 15 If required, enter the same NAT ID that you used in step 7 in the **Unique NAT ID** text box.

Step 16 Click **Register**.

What to do next

After establishing the management center high availability pair, devices registered to the active management center are automatically registered to the standby management center.

Note When a registered device has a NAT IP address, automatic device registration fails and the secondary 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 management center High Availability page. If automatic registration otherwise fails on the standby management center, but the device appears to be registered to the active Secure Firewall Management Center, see Using CLI to Resolve Device Registration in Management Center High Availability, on page 317.

High Availability for Management Centers Hosted on Public Cloud

While establishing high availability between management centers hosted on public clouds, the combinations of IP addresses or hostnames for the primary and secondary management centers described below can successfully form high availability and get the devices registered on both the peers. In the **High Availability** page (**Integration** > **Other Integrations** > **High Availability**), perform one of the following configurations to successfully form high availability between management centers hosted in public cloud.

Using the Public IP Addresses or Hostnames for Both the Primary and Secondary Management Centers

1. On the secondary management center, do the following:
   a. Choose **Secondary** as the **Role for this Firewall Management Center**.
   b. Enter the public IP address or hostname for the secondary management center in the **Primary Firewall Management Center Host** field.
   c. Enter the registration key.
   d. Enter the same NAT ID that you used in the primary management center.

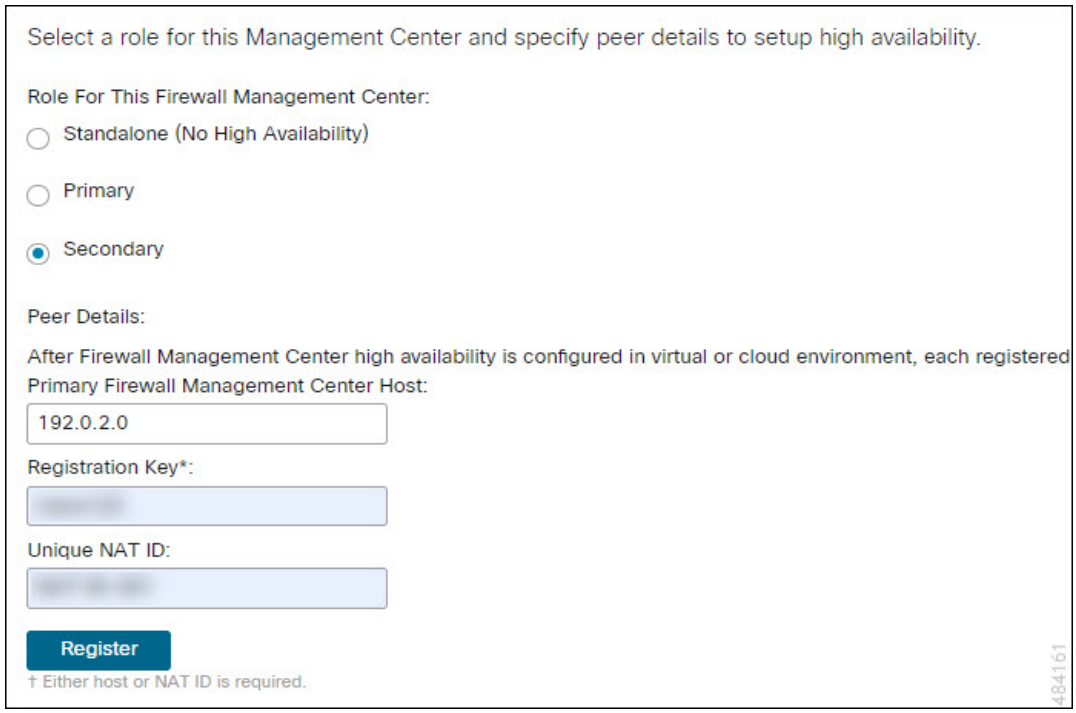

2. On the primary management center, do the following:
   a. Choose **Primary** as the **Role for this Firewall Management Center**.
   b. Enter the public IP address or hostname for the secondary management center in the **Secondary Firewall Management Center Host** field.
   c. Enter the registration key.
   d. Enter the unique NAT ID.

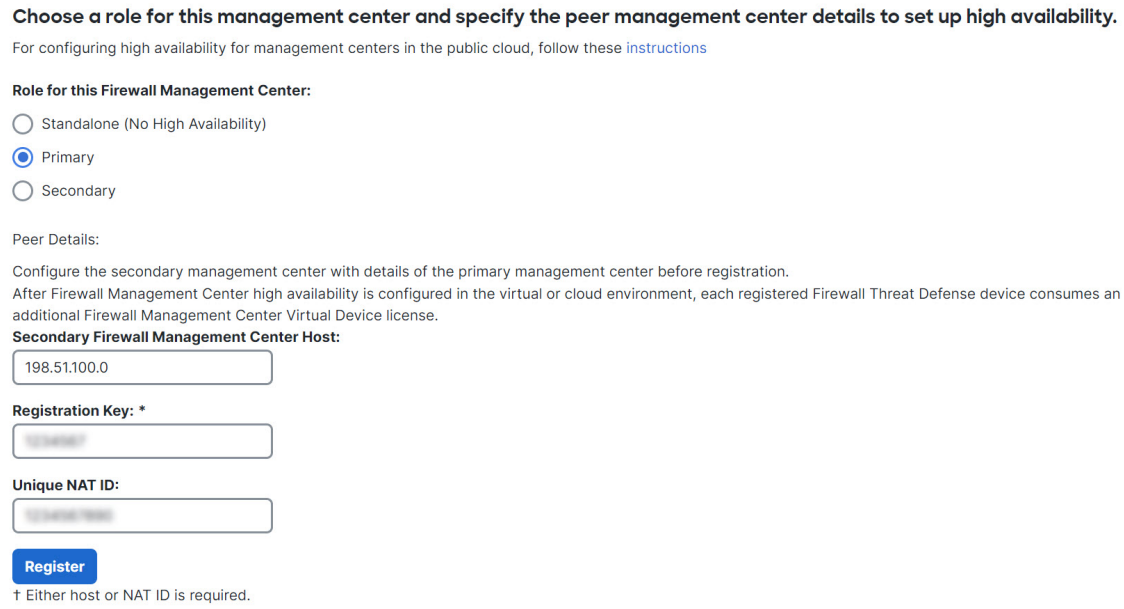

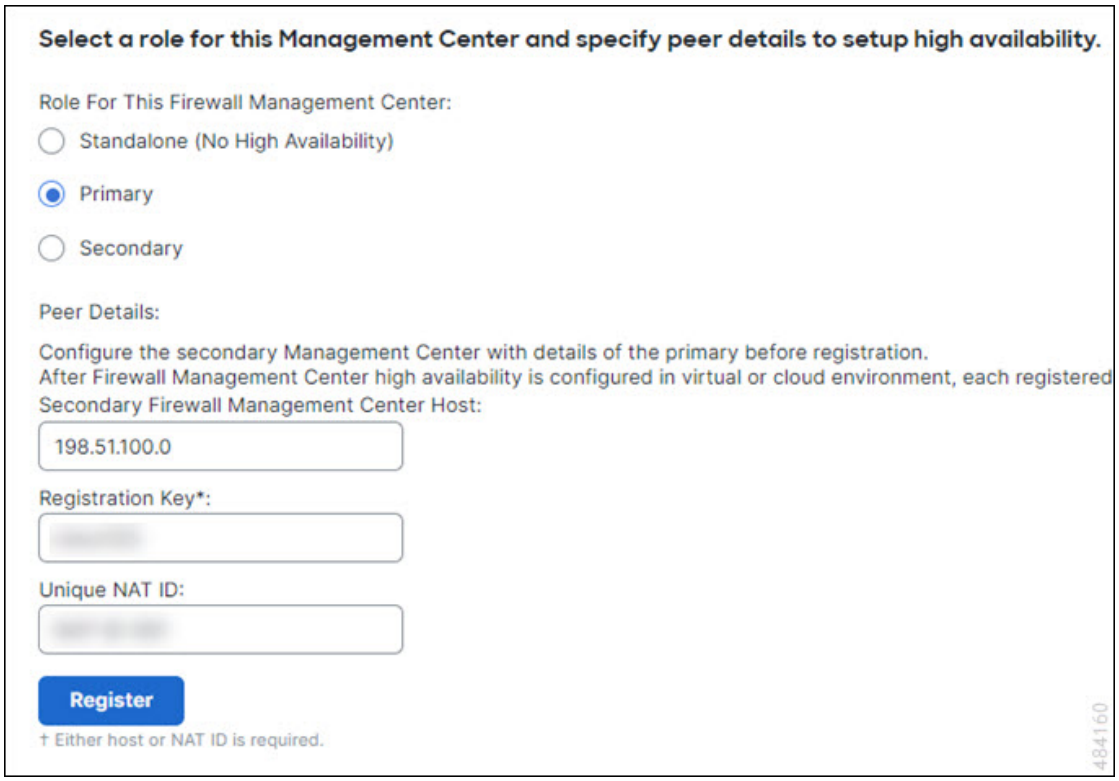

Using the Public IP Address or Hostname for the Secondary Management Center

1. On the secondary management center, do the following:
   a. Choose **Secondary** as the **Role for this Firewall Management Center**.

b. Enter **DONTRESOLVE** in the **Primary Firewall Management Center Host** field.

c. Enter the registration key.

d. Enter the same NAT ID that you used in the primary management center.

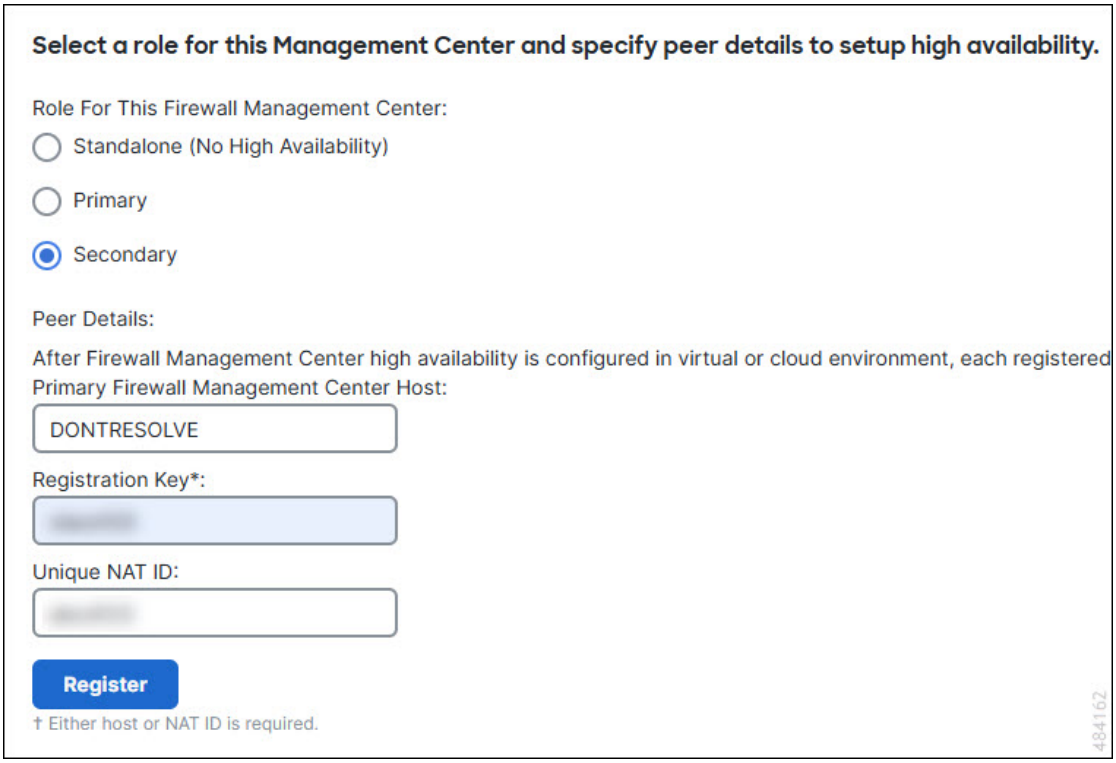

2. On the primary management center, do the following:

a. Choose **Primary** as the **Role for this Firewall Management Center**.

b. Enter the public IP address or hostname for the secondary management center in the **Secondary Firewall Management Center Host** field.

c. Enter the registration key.

d. Enter the unique NAT ID.

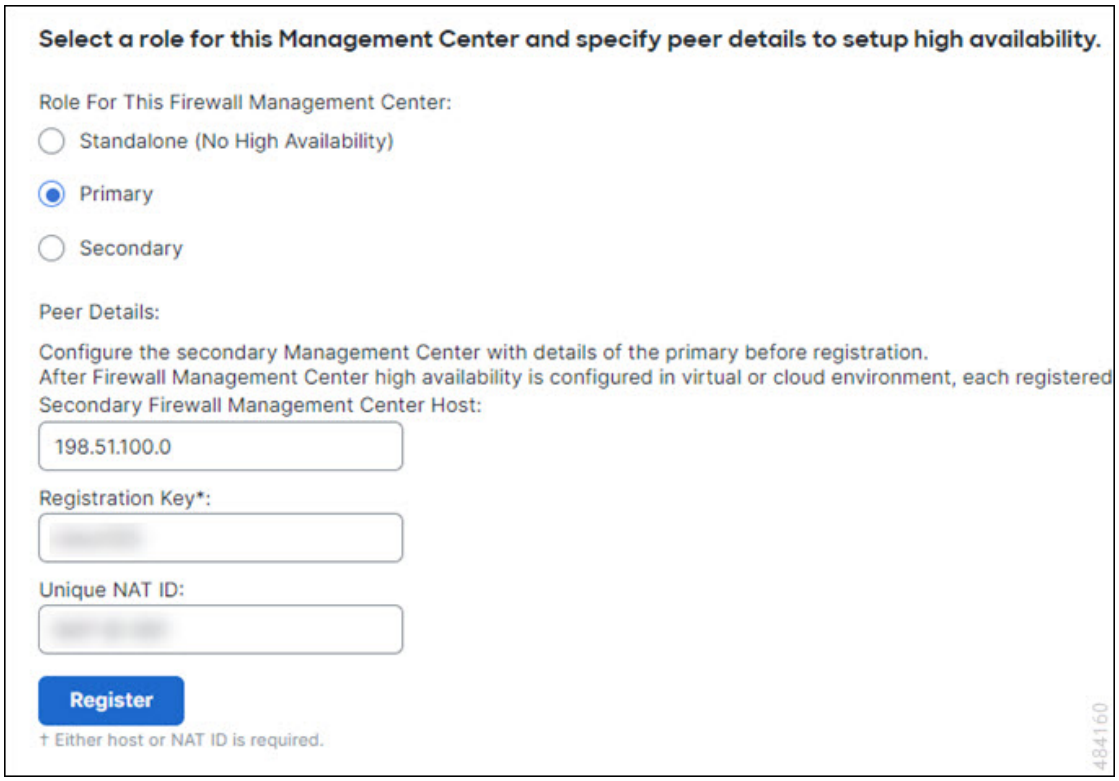

Viewing Management Center High Availability Status

After you identify your active and standby management centers, you can view information about the local 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 management centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

You can view:

Summary Information

- The health status of the high availability pair. The status of a correctly functioning system will oscillate between "Healthy" and "Synchronization task is in progress" as the standby unit receives configuration changes from the active unit.

- The current synchronization status of the high availability pair
- The IP address of the active peer and the last time it was synchronized
- The IP address of the standby peer and the last time it was synchronized

System Status

- The configured IP addresses for both peers
- The operating system for both peers
- The software version for both peers
- The appliance model of both peers

Note
You can view export control and compliance status only on the active management center.

Remote and Local Device Registration
You can view the list of devices that are pending or failed registration on management center.

Configurations Synced on Management Center High Availability Pairs

When you establish high availability between two management centers, the following configuration data is synced between them:

- License entitlements
- Access control policies
- Intrusion rules
- Malware and file policies
- DNS policies
- Identity policies
- SSL policies
- Prefilter policies
- Network discovery rules
- Application detectors
- Correlation policy rules
- Alerts
- Scanners
- Response groups

- Contextual cross-launch of external resources for investigating events
- Remediation settings, although you must install custom modules on both management centers. For more information on remediation settings, see Managing Remediation Modules, on page 1013.

Configuring External Access to the Management Center Database in a High Availability Pair

In a high availability setup, we recommend you to use only the active peer to configure the external access to the database. When you configure the standby peer for external database access, it leads to frequent disconnections. To restore the connectivity, you must pause and resume the synchronization of the standby peer. For information on how to enable external database access to management centers, see Enabling External Access to the Database, on page 63 .

Using CLI to Resolve Device Registration in Management Center High Availability

If automatic device registration fails on the standby management center, but appears to be registered to the active management center, complete the following steps:

Warning If you do an RMA of secondary management center or add a secondary 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 management center. See *Delete (Unregister) a Device from the management center* in Cisco Secure Firewall Management Center Device Configuration Guide.

Step 2 Complete the following steps to trigger automatic registration of the device on the standby 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 management center.

c. Run the CLI command: **configure manager add**.

This command configures the device to initiate a connection to a management center.

Tip
Configure remote management on the device, only for the active management center. When you establish high availability, the devices are automatically registered to the standby management center.

d. Log in to the active management center and register the device.

Step 3 If the standby management center is behind NAT, complete the following steps to edit the hostname of the standby management center:

a. Access the threat defense shell and use the `show managers` command to get the standby management center entry identifier value.

b. In the threat defense shell, edit the standby 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.

Switching Peers in the Management Center High Availability Pair

Because the system restricts some functionality to the active management center, if that appliance fails, you must promote the standby management center to active:

Procedure

Step 1 Log into one of the management centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Choose **Switch Peer Roles** to change the local role from Active to Standby, or Standby to Active. With the Primary or Secondary designation unchanged, the roles are switched between the two peers.

Pausing Communication Between Paired Management Centers

If you want to temporarily disable high availability, you can disable the communications channel between the management centers. You can pause synchronization from an active or standby peer.

Procedure

Step 1 Log into one of the management centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Choose **Pause Synchronization**.

Restarting Communication Between Paired Management Centers

If you temporarily disable high availability, you can restart high availability by enabling the communications channel between the management centers. You can resume synchronization from an active or standby peer.

Procedure

Step 1 Log into one of the management centers that you paired using high availability.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Choose **Resume Synchronization**.

Change the IP Address of the Management Center in a High Availability Pair

If the IP address for one of the high availability peers is changed, this change will not be automatically updated on the other peer, even after performing a high availability synchronization. To ensure that the remote peer management center is also updated, you must manually change the IP address.

Procedure

Step 1 Log in to the peer 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 or the name that resolves through the local DNS to a valid IP address (that is, the host name), or the host IP address.

Step 8 Click **Save**.

Disabling Management Center High Availability

Procedure

Step 1 Log into one of the management centers in the high availability pair.

Step 2 Choose **Integration** > **Other Integrations**.

Step 3 Choose **High Availability**.

Step 4 Choose **Break High Availability**.

Step 5 Choose one of the following options for handling managed devices:

- To control all managed devices with this 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 management center, choose **Manage registered devices from peer console**. All devices will be unregistered from this management center.
- To stop managing devices altogether, choose **Stop managing registered devices from both consoles**. All devices will be unregistered from both management centers.

Note
If you choose to manage the registered devices from the secondary management center, the devices will be unregistered from the primary management center. The devices are now registered to be managed by the secondary management center. However the licenses that were applied to these devices are deregistered on account of the high availability break operation. You must now proceed to re-register (enable) the licenses on the devices from the secondary management center. For more information see Assign Licenses to Devices, on page 281.

Step 6 Click **OK**.

Replacing Management Centers in a High Availability Pair

If you need to replace a failed unit in the management center high availability pair, you must follow one of the procedures listed below. The table lists four possible failure scenarios and their corresponding replacement procedures.

Failure Status	Data Backup Status	Replacement Procedure
Primary management center failed	Data backup successful	Replace a Failed Primary Management Center (Successful Backup), on page 321
	Data backup not successful	Replace a Failed Primary Management Center (Unsuccessful Backup), on page 322

Failure Status	Data Backup Status	Replacement Procedure
Secondary management center failed	Data backup successful	Replace a Failed Secondary Management Center (Successful Backup), on page 323
	Data backup not successful	Replace a Failed Secondary Management Center (Unsuccessful Backup), on page 323

Replace a Failed Primary Management Center (Successful Backup)

Two management centers, *FMC1* and *FMC2*, are part of a high availability pair. *FMC1* is the primary and *FMC2* is the secondary. This task describes the steps to replace a failed primary management center, *FMC1*, when data backup from the primary is successful.

Before you begin

Verify that the data backup from the failed primary management center is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed management center - *FMC1*.

Step 2 When the primary management center - *FMC1* fails, access the web interface of the secondary management center - *FMC2* and switch peers. For more information, see Switching Peers in the Management Center High Availability Pair, on page 318.

This promotes the secondary management center - *FMC2* to active.

You can use *FMC2* as the active management center until the primary management center - *FMC1* is replaced.

Caution
Do not break 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 management center with the same software version as *FMC1*.

Step 4 Restore the data backup retrieved from *FMC1* to the new management center.

Step 5 Install required management center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match *FMC2*.

The new management center and *FMC2* will now both be active peers, resulting in a high availability split-brain.

Step 6 When the 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 management center - *FMC1*.

Step 7 When the configuration syncs successfully, access the web interface of the secondary management center - *FMC2* and switch roles to make the primary management center - *FMC1* active. For more information, see Switching Peers in the Management Center High Availability Pair, on page 318.

What to do next

High availability has now been re-established and the primary and the secondary management centers will now work as expected.

Replace a Failed Primary Management Center (Unsuccessful Backup)

Two management centers - *FMC1* and *FMC2* are part of a high availability pair. *FMC1* is the primary and *FMC2* is the secondary. This task describes the steps to replace a failed primary management center -*FMC1* when data backup from the primary is unsuccessful.

Procedure

Step 1 Contact Support to request a replacement for a failed management center - *FMC1*.

Step 2 When the primary management center - *FMC1* fails, access the web interface of the secondary management center - *FMC2* and switch peers. For more information, see Switching Peers in the Management Center High Availability Pair, on page 318.

This promotes the secondary management center - *FMC2* to active.

You can use *FMC2* as the active management center until the primary management center - *FMC1* is replaced.

Caution
Do not break 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 management center with the same software version as *FMC1*.

Step 4 Install required management center patches, geolocation database (GeoDB) updates, vulnerability database (VDB) updates and system software updates to match *FMC2*.

Step 5 Deregister one of the management centers - *FMC2* from the Cisco Smart Software Manager. For more information, see Deregister the Management Center, on page 283.

Deregistering management center from the Cisco Smart Software Manager removes the Management Center from your virtual account. All license entitlements associated with the management center release back to your virtual account. After deregistration, the management center enters Enforcement mode where no update or changes on licensed features are allowed.

Step 6 Access the web interface of the secondary management center - *FMC2* and break management center high availability. For more information, see Disabling Management Center High Availability, on page 320. 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 management center- *FMC2*, will be removed and you cannot perform deployment activities from *FMC2*.

Step 7 Re-establish management center high availability, by setting up the management center - *FMC2* as the primary and management center - *FMC1* as the secondary. For more information , see Establishing Management Center High Availability, on page 310.

Step 8 Register a Smart License to the primary management center - *FMC2*. For more information see Register the Management Center with the Smart Software Manager, on page 275.

What to do next

High availability has now been re-established and the primary and the secondary management centers will now work as expected.

Replace a Failed Secondary Management Center (Successful Backup)

Two management centers - *FMC1* and *FMC2* are part of a high availability pair. *FMC1* is the primary and *FMC2* is the secondary. This task describes the steps to replace a failed secondary management center -*FMC2* when data backup from the secondary is successful.

Before you begin

Verify that the data backup from the failed secondary management center is successful.

Procedure

Step 1 Contact Support to request a replacement for a failed management center - *FMC2*.

Step 2 Continue to use the primary management center - *FMC1* as the active management center.

Step 3 Reimage the replacement management center with the same software version as *FMC2*.

Step 4 Restore the data backup from *FMC2* to the new management center.

Step 5 Install required 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 management center - *FMC2*, to synchronize the latest configuration from the primary management center - *FMC1*. For more information, see Restarting Communication Between Paired Management Centers, on page 319.

Classic and Smart Licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary management centers will now work as expected.

Replace a Failed Secondary Management Center (Unsuccessful Backup)

Two management centers - *FMC1* and *FMC2* are part of a high availability pair. *FMC1* is the primary and *FMC2* is the secondary. This task describes the steps to replace a failed secondary management center -*FMC2* when data backup from the secondary is unsuccessful.

Procedure

Step 1 Contact Support to request a replacement for a failed management center - *FMC2*.

Step 2 Continue to use the primary management center - *FMC1* as the active management center.

Step 3 Reimage the replacement management center with the same software version as *FMC2*.

Step 4 Install required 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 management center - *FMC1* and break management center high availability. For more information, see Disabling Management Center High Availability, on page 320. When prompted to select an option for handling managed devices, choose **Manage registered devices from this console**.

Step 6 Re-establish management center high availability, by setting up the management center - *FMC1* as the primary and management center - *FMC2* as the secondary. For more information , see Establishing Management Center High Availability, on page 310.

- When high availability is successfully established, the latest configuration from the primary management center - *FMC1* is synchronized to the secondary management center - *FMC2*.
- Classic and Smart Licenses work seamlessly.

What to do next

High availability has now been re-established and the primary and the secondary management centers will now work as expected.

Management Center High Availability Disaster Recovery

In case of a disaster recovery situation, a manual switchover must be performed. When the primary management center - FMC1 fails, access the web interface of the secondary management center - FMC2 and switch peers. This is applicable conversely also in case the secondary (FMC2) fails. For more information, see Switching Peers in the Management Center High Availability Pair, on page 318.

For restoring a failed management center, refer to Replacing Management Centers in a High Availability Pair, on page 320.

Restoring Management Center in a High Availability Pair (No Hardware Failure)

To restore a management center high availability pair when there is no hardware failure, follow these procedures:

- Restore Backup on the Primary Management Center , on page 325
- Restore Backup on the Secondary Management Center, on page 325

Restore Backup on the Primary Management Center

Before you begin

- There is no hardware failure and replacement of the management center.
- You are familiar with the backup and restore process. See Backup/Restore, on page 453.

Procedure

Step 1 Verify if backup of the primary management center is available—either a local storage in /var/sf/backup/, or a remote network volume.

Step 2 Pause synchronization on the primary management center. Choose **Integration** > **Other Integrations**, and then go to the **High Availability** tab to pause synchronization.

Step 3 Restore the backup on the primary management center. The management center reboots when the restoration is complete.

Step 4 Once the primary management center is active and its user interface is reachable, resume synchronization on the secondary management center. Choose **Integration** > **Other Integrations**, and then go to the **High Availability** tab to resume synchronization.

Restore Backup on the Secondary Management Center

Before you begin

- There is no hardware failure and replacement of the management center.
- You are familiar with the backup and restore process. See Backup/Restore, on page 453.

Procedure

Step 1 Verify if backup of the secondary management center is available—either a local storage in /var/sf/backup/, or a remote network volume.

Step 2 Pause synchronization on the primary management center. Choose **Integration** > **Other Integrations**, and then go to the **High Availability** tab to pause synchronization.

Step 3 Restore the backup on the secondary management center. The management center reboots when the restoration is complete.

Step 4 Once the secondary management center is active and its user interface is reachable, resume synchronization on the primary management center. Choose **Integration** > **Other Integrations**, and then go to the **High Availability** tab to resume synchronization.

Unified Backup of Management Centers in High Availability

You can perform a unified backup on an active management center, where a single backup file is created for both the active and standby management centers. The unified backup is applicable only for configuration-only backup. If eventing or TID backup is required, you must take separate backup for active and standby management centers. When you select configuration-only backup, by default, unified backup is applied. In a unified backup, if the active management center is unable to get a backup tar file from the standby management center, the normal backup file is generated for the active unit that can be used for restoration. There are several benefits of unified backup over the normal backup:

- Unified backup does not require you to take separate backups on active and standby management centers.
- Redundant data in backups and storage constraints are removed in a unified backup.
- In a normal backup, when the primary unit fails, and if a secondary unit backup is not available, you had to break the high availability pairing for the secondary RMA. This situation is eradicated in a unified backup.
- Typically, the backup of a standby unit cannot be scheduled. In an unified backup that is scheduled, both active and standby units' backup are taken.
- While executing unified backup, you do not have to pause the HA synchronization to perform backup on the standby unit.

You can use the unified backup to recover a new RMA device if an unanticipated incident occurs. 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 management center to restore and also select its State (Active/Standby).

Ensure that you select the appropriate state of the restored management center to prevent Split-Brain conflict.

Restore Management Center from Unified Backup

Use this procedure to restore management center from the unified backup(configuration-only).

Procedure

Step 1 Log into the 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**; see Manage Backups and Remote Storage, on page 481.

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. Because the unified backup stores the backup configuration of both primary and secondary management centers, you need to choose which unit you want to restore.

Step 5 To select the state of the restored management center, click the **Active** or **Standby** radio button. You must verify the role and state of your working management center to avoid having both peers with same role and

state configuration. Choosing the incorrect role and state for your management center when restoring can cause HA failure.

Step 6 Click **Restore**, and then **Confirm Restore** to begin the restoration.

History for Management Center High Availability

Feature	Minimum Management Center	Minimum Threat Defense	Details
Single backup file for high availability management centers.	7.4.1 7.2.6	Any	When performing a configuration-only backup of the active management center in a high availability pair, the system now creates a single backup file which you can use to restore either unit. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Support for high availability on KVM.	7.3.0	Any	We now support high availability on management center virtual for KVM.
Support for high availability on AWS and OCI.	7.1.0	Any	We now support high availability on management center virtual for AWS and OCI.
Support for high availability on HyperFlex.	7.0.0	Any	We now support high availability on management center virtual for HyperFlex.
Support for high availability on VMware.	6.7.0	Any	We now support high availability on 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.

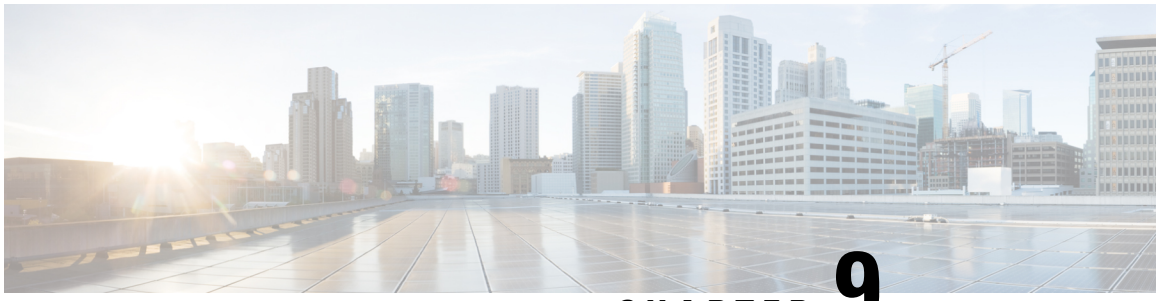

CHAPTER 9

Security Certifications Compliance

The following topics describe how to configure your system to comply with security certifications standards:

Security Certifications Compliance Modes

Your organization might be required to use only equipment and software complying with security standards established by the U.S. Department of Defense and global certification organizations. Secure Firewall supports compliance with the following 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

You can enable security certifications compliance in CC mode or UCAPL mode. Enabling security certifications compliance does not guarantee strict compliance with all requirements of the security mode selected. For more information on hardening procedures, refer to the guidelines for this product provided by the certifying entity.

Caution After you enable this setting, you cannot disable it. If you need to take an appliance out of CC or UCAPL mode, you must reimage.

Security Certifications Compliance Characteristics

The following table describes behavior changes when you enable CC or UCAPL mode. (Restrictions on login accounts refers to command line access, not 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
The `admin` user can be locked out after a maximum number of failed login attempts configurable through the web interface.	Yes	Yes	Yes	Yes	—	—

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 local appliance CLI.	No	No	Yes, regardless of security certifications compliance enablement.	Yes, regardless of security certifications compliance enablement.	Yes	Yes
The system automtically 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:

- To enable security certifications compliance in your deployment, enable it first on the Secure Firewall Management Center, then enable it in the same mode on all managed devices.

Caution The Secure 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 Secure 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 Secure 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 of the following features:
  - Email reports, alerts, or data pruning notifications.
  - Nmap Scan, Cisco IOS Null Route, Set Attribute Value, or ISE EPS remediations.
  - Remote storage for backups or reports.
  - Third-party client access to the system database.
  - External notifications or alerts transmitted via email (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.
- 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 Secure 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.

Note The system does not support CC or UCAPL mode for:

- Secure Firewall Threat Defense devices in clusters
- Secure Firewall Threat Defense container instances on the Firepower 4100/9300
- Exporting event data to an external client using eStreamer.

Appliance Hardening

For information about features you can use 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*, as well as the following topics within this document:

- Licenses, on page 253
- Users, on page 117
- Logging into the Management Center, on page 27
- Audit Log, on page 48
- Audit Log Certificate, on page 51
- Time Synchronization, on page 104

- *Configure NTP Time Synchronization for Threat Defense* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Creating an Email Alert Response, on page 557
- Configuring Email Alerting for Intrusion Events, on page 566
- *Configure SMTP* in the Cisco Secure Firewall Management Center Device Configuration Guide
- *About SNMP for the Firepower 1000/2100* in the Cisco Secure Firewall Management Center Device Configuration Guide
- *Configure SNMP* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Creating an SNMP Alert Response, on page 552
- *Configure Dynamic DNS* in the Cisco Secure Firewall Management Center Device Configuration Guide
- DNS Cache, on page 57
- Audit and Syslog, on page 405
- Access List, on page 46
- Security Certifications Compliance, on page 329
- Configure SSH for Remote Storage, on page 99
- Audit Log Certificate, on page 51
- HTTPS Certificates, on page 64
- Customize User Roles for the Web Interface, on page 197
- Add or Edit an Internal User, on page 125
- Session Timeout, on page 101
- *About Configuring Syslog* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Schedule Management Center Backups, on page 489
- *Site-to-Site VPNs for Threat Defense* 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

Protecting Your Network

See the following topics to learn about features you can configure to protect your network:

- Access Control Policies
- *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
- Update Intrusion Rules, on page 229
- *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 405
- Intrusion Events, on page 763
- Event Search, on page 685
- Workflows, on page 647
- *Device Management* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Login Banner, on page 74
- Updates, on page 223

Enable Security Certifications Compliance

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 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. If you need to take the appliance out of CC or UCAPL mode, you must reimage.

Before you begin

- We recommend you register all devices that you plan to be part of your deployment to the 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 be an Admin user to perform this task.

Procedure

Step 1 Depending on whether you are configuring a management center or a managed device:

- management center: Choose **System** (⚙) > **Configuration**.
- 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 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

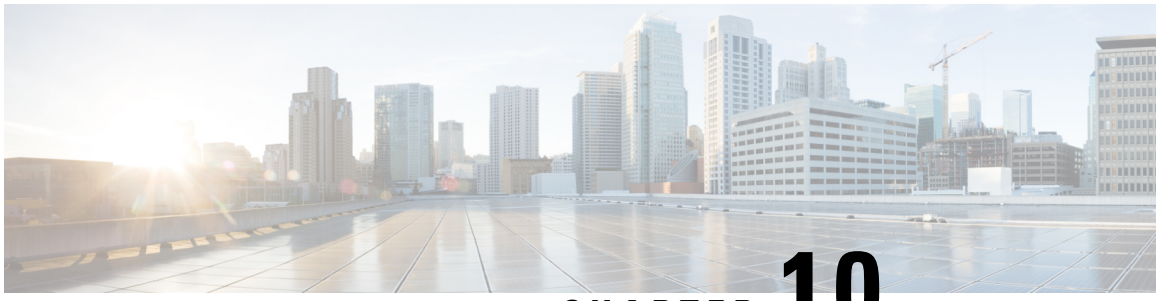

CHAPTER 10

Dashboards

The following topics describe how to use dashboards:

About Dashboards

Dashboards provide you with at-a-glance views of current system status, including data about the events collected and generated by the system. You can also use dashboards to see information about the status and overall health of the appliances in your deployment. Keep in mind that the information the dashboard provides depends on how you license, configure, and deploy the system.

Note Ensure that you have enabled REST API (**System** > **Configuration** > **REST API Preferences**) to view the correlated device metrics on the dashboard.

Tip The dashboard is a complex, highly customizable monitoring feature that provides exhaustive data. For a broad, brief, and colorful picture of your monitored network, use the Context Explorer.

A dashboard uses tabs to display widgets: small, self-contained components that provide insight into different aspects of the system. For example, the predefined Appliance Information widget tells you the appliance name, model, and currently running software version. The system constrains widgets by the dashboard time range, which you can change to reflect a period as short as the last hour or as long as the last year.

The system is delivered with several predefined dashboards, which you can use and modify. If your user role has access to dashboards (Administrator, Maintenance User, Security Analyst, Security Analyst [Read Only], and custom roles with the Dashboards permission), by default your home page is the predefined Summary Dashboard. However, you can configure a different default home page, including non-dashboards. You can also change the default dashboard. Note that if your user role cannot access dashboards, your default home page is relevant to the role; for example, a Discovery Admin sees the Network Discovery page.

You can also use predefined dashboards as the base for custom dashboards, which you can either share or restrict as private. Unless you have Administrator access, you cannot view or modify private dashboards created by other users.

Note Some drill-down pages and table views of events include a **Dashboard** toolbar link that you can click to view a relevant predefined dashboard. If you delete a predefined dashboard or tab, the associated toolbar links do not function.

In a multidomain deployment, you cannot view dashboards from ancestor domains; however, you can create new dashboards that are copies of the higher-level dashboards.

Dashboard Widgets

A dashboard has one or more tabs, each of which can display one or more widgets in a three-column layout. The system is delivered with many predefined dashboard widgets, each of which provides insight into a different aspect of the system. Widgets are grouped into three categories:

- *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.

The dashboard widgets that you can view depend on:

- the type of appliance you are using
- your user role
- your current domain (in a multidomain deployment)

In addition, each dashboard has a set of preferences that determines its behavior.

You can minimize and maximize widgets, add and remove widgets from tabs, as well as rearrange the widgets on a tab.

Note For widgets that display event counts over a time range, the total number of events may not reflect the number of events for which detailed data is available in the tables on pages under the Analysis menu. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.

Widget Availability

The dashboard widgets that you can view depend on the type of appliance you are using, your user role, and your current domain (in a multidomain deployment).

In a multidomain deployment, if you do not see a widget that you expect to see, switch to the Global domain. See Switching Domains on the Secure Firewall Management Center, on page 20.

Note that:

- An *invalid* widget is one that you cannot view because you are using the wrong type of appliance.
- An *unauthorized* widget is one that you cannot view because your user account does not have the necessary privileges.

For example, the Appliance Status widget is available only on the management center for users with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) account privileges.

Although you cannot add an unauthorized or invalid widget to a dashboard, an imported dashboard may contain unauthorized or invalid widgets. For example, such widgets can be present if the imported dashboard:

- Was created by a user with different access privileges, or
- Belongs to an ancestor domain.

Unavailable widgets are disabled and display error messages that indicate why you cannot view them.

Individual widgets also display error messages when those widgets have timed out or are otherwise experiencing problems.

Note You can delete or minimize unauthorized and invalid widgets, as well as widgets that display no data, keeping in mind that modifying a widget on a shared dashboard modifies it for all users of the appliance.

Dashboard Widget Availability by User Role

The following table lists the user account privileges required to view each widget. Only user accounts with Administrator, Maintenance User, Security Analyst, or Security Analyst (Read Only) access can use dashboards.

Users with custom roles may have access to any combination of widgets, or none at all, as their user roles permit.

Table 24: User Roles and Dashboard Widget Availability

Widget	Administrator	Maintenance User	Security Analyst	Security Analyst (R
Appliance Information	yes	yes	yes	yes
Appliance Status	yes	yes	yes	no
Correlation Events	yes	no	yes	yes
Current Interface Status	yes	yes	yes	yes
Current Sessions	yes	no	no	no
Custom Analysis	yes	no	yes	yes
Disk Usage	yes	yes	yes	yes
Interface Traffic	yes	yes	yes	yes
Intrusion Events	yes	no	yes	yes

Widget	Administrator	Maintenance User	Security Analyst	Security Analyst (RO)
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

The system is delivered with several predefined widgets that, when used on dashboards, can provide you with at-a-glance views of current system status. These views include:

- data about the events collected and generated by the system
- information about the status and overall health of the appliances in your deployment

Note The dashboard widgets you can view depend on the type of appliance you are using, your user role, and your current domain in a multidomain deployment.

The Allow List Events Widget

The Allow List Events widget shows the average events per second by priority, over the dashboard time range. It appears by default on the Correlation tab of the Default Dashboard.

You can configure the widget to display allow list events of different priorities by modifying the widget preferences.

In the widget preferences, you can:

- choose one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority
- choose **Show All** to display an additional graph for all allow list events, regardless of priority
- choose **Vertical Scale** to choose **Linear** (incremental) or **Logarithmic** (factor of ten) scale

The preferences also control how often the widget updates.

You can click a graph to view 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 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 management center virtual.
- For managed appliances, the name and status of the communications link with the managing appliance.
- For management centers in a high availability pair, the name, model, and system software and operating system versions of the peer management center, as well as how recently the 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.

The Appliance Status Widget

The Appliance Status widget indicates the health of the appliance and of any appliances it is managing. Note that because the management center does not automatically apply a health policy to managed devices, you must manually apply a health policy to devices or their status appears as `Disabled`. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to display appliance status as a pie chart or in a table by modifying the widget preferences.

The preferences also control how often the widget updates.

You can click a section on the pie chart or one of the numbers on the appliance status table to go to the Health Monitor page and view the compiled health status of the appliance and of any appliances it is managing.

The Correlation Events Widget

The Correlation Events widget shows the average number of correlation events per second, by priority, over the dashboard time range. It appears by default on the Correlation tab of the Detailed Dashboard.

You can configure the widget to display correlation events of different priorities by modifying the widget preferences, as well as to choose a linear (incremental) or logarithmic (factor of ten) scale.

Check one or more **Priorities** check boxes to display separate graphs for events of specific priorities, including events that do not have a priority. Choose **Show All** to display an additional graph for all correlation events, regardless of priority. The preferences also control how often the widget updates.

You can click a graph to view correlation events of a specific priority, or click the **All** graph to view all correlation events. In either case, the events are constrained by the dashboard time range; accessing correlation events via the dashboard changes the events (or global) time window for the appliance.

The Current Interface Status Widget

The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. On the 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.

The Current Sessions Widget

The Current Sessions widget shows which users are currently logged into the appliance, the IP address associated with the machine where the session originated, and the last time each user accessed a page on the appliance (based on the local time for the appliance). The user that represents you, that is, the user currently viewing the widget, is marked with a **User icon** and rendered in bold type. Sessions are pruned from this widget's data within one hour of logoff or inactivity. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

On the Current Sessions widget, you can:

- click any user name to manage user accounts on the User Management page.
- click the **Host icon** or **Compromised Host icon** next to any IP address to view the host profile for the associated machine.
- click any IP address or access time to view the audit log constrained by that IP address and by the time that the user associated with that IP address logged on to the web interface.

The widget preferences control how often the widget updates.

The Custom Analysis Widget

The Custom Analysis widget is a highly customizable widget that allows you to display detailed information on the events collected and generated by the 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** > **Add Widgets** to create a **Custom Analysis** widget.
3. Return to **Overview** > **Dashboards** and modify the new widget to customize with the scope of search.

Example: Configuration of Custom Analysis Widget

You can configure the Custom Analysis widget to display a list of recent intrusion events by configuring the widget to display data from the **Intrusion Events** table. Choosing the **Classification** field and aggregating this data by **Count** displays the number of events that were generated for each type.

On the other hand, aggregating by **Unique Events** displays the number of unique intrusion events of each type (for example, how many detections of network trojans, potential violations of corporate policy, attempted denial-of-service attacks, and so on).

You can further customize the widget using a saved search, either one of the predefined searches delivered with your appliance or a custom search that you created. For example, constraining the first example (intrusion events using the **Classification** field, aggregated by **Count**) using the **Dropped Events** search displays the number of intrusion events that were dropped for each type.

Related Topics

Modifying Dashboard Time Settings, on page 356

Custom Analysis Widget Preferences

The following table describes the preferences you can set in the Custom Analysis widget.

Different preferences appear depending on how you configure the widget. For example, a different set of preferences appears if you configure the widget to show relative occurrences of events (a bar graph) vs a graph over time (a line graph). Some preferences, such as Filter, only appear if you choose a specific table from which to display data.

Table 25: Custom Analysis Widget Preferences

Preference	Details
Title	If you do not specify a title for the widget, the system uses the configured event type as the title.
Preset	Custom Analysis presets provide quick access to information about your deployment. The predefined dashboards make extensive use of these presets. You can use these presets or you can create a custom configuration.
Table (required)	The table of events or assets that contains the data the widget displays.
Field (required)	The specific field of the event type you want to display. To show data over time (line graphs), choose **Time**. To show relative occurrences of events (bar graphs), choose another option.
Aggregate (required)	The aggregation method configures how the widget groups the data it displays. For most event types, the default option is **Count**.
Filter	You can use application filters to constrain data from the Application Statistics and Intrusion Event Statistics by Application tables.
Search	You can use a saved search to constrain the data that the widget displays. You do not have to specify a search, although some presets use predefined searches. Only you can access searches that you have saved as private. If you configure the widget on a shared dashboard and constrain its events using a private search, the widget resets to not using the search when another user logs in. This affects your view of the widget as well. If you want to make sure that this does not happen, save the dashboard as private. Only fields that constrain connection summaries can constrain Custom Analysis dashboard widgets based on connection events. Invalid saved searches are dimmed. If you constrain a Custom Analysis widget using a saved search, then edit the search, the widget does not reflect your changes until the next time it updates.
Show	Choose whether you want to display the most (**Top**) or the least (**Bottom**) frequently occurring events.
Results	Choose the number of result rows to display.
Show Movers	Choose whether you want to display the icons that indicate changes from the most recent results.
Time Zone	Choose the time zone you want to use to display results.
Color	You can change the color of the bars in the widget's bar graph.

Related Topics

Configuring Widget Preferences, on page 354

Viewing Associated Events from the Custom Analysis Widget

From a Custom Analysis widget, you can invoke an event view (workflow) that provides detailed information about the events displayed in the widget. The events appear in the default workflow for that event type, constrained by the dashboard time range. This also changes the appropriate time window on the management center, depending on how many time windows you configured and on the event type.

For example:

- If you configure multiple time windows, then access health events from a Custom Analysis widget, the events appear in the default health events workflow, and the health monitoring time window changes to the dashboard time range.
- If you configure a single time window and then access any type of event from the Custom Analysis widget, the events appear in the default workflow for that event type, and the global time window changes to the dashboard time range.

Procedure

You have the following choices:

- On any Custom Analysis widget, click **View** ( ) in the lower right corner of the widget to view all associated events, constrained by the widget preferences.
- On a Custom Analysis widget showing relative occurrences of events (bar graph), click any event to view associated events constrained by the widget preferences, as well as by that event.

The Disk Usage Widget

The Disk Usage widget displays the percentage of space used on the hard drive, based on disk usage category. It also indicates the percentage of space used on and capacity of each partition of the appliance's hard drive. The Disk Usage widget displays the same information for the malware storage pack if installed in the device, or if the management center manages a device containing a malware storage pack. This widget appears by default on the Status tabs of the Default Dashboard and the Summary Dashboard.

The By Category stacked bar displays each disk usage category as a proportion of the total available disk space used. The following table describes the available categories.

Table 26: Disk Usage Categories

Disk Usage Category	Description
Events	all events logged by the system
Files	all files stored by the system
Backups	all backup files
Updates	all files related to updates, such as rule updates and system updates
Other	system troubleshooting files and other miscellaneous files
Free	free space remaining on the appliance

You can hover your pointer over a disk usage category in the By Category stacked bar to view the percentage of available disk space used by that category, the actual storage space on the disk, and the total disk space available for that category. Note that if you have a malware storage pack installed, the total disk space available for the Files category is the available disk space on the malware storage pack.

You can configure the widget to display only the By Category stacked bar, or you can show the stacked bar plus the admin (`/`), `/Volume`, and `/boot` partition usage, as well as the `/var/storage` partition if the malware storage pack is installed, by modifying the widget preferences.

The widget preferences also control how often the widget updates, as well as whether it displays the current disk usage or collected disk usage statistics over the dashboard time range.

The Interface Traffic Widget

The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's management interface. 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 Defense 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.

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 763. The arrow (if any) that appears above the impact level numbers describes the inline result and is defined as follows:

  Table 27: 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.

On the widget, you can:

- Click a graph corresponding to dropped packets, to would have dropped packets, or to a specific impact to view intrusion events of that type.
- Click the graph corresponding to dropped events to view dropped events.
- Click the graph corresponding to would have dropped events to view would have dropped events.
- Click the **All** graph to view all intrusion events.

The resulting event view is constrained by the dashboard time range; accessing intrusion events via the dashboard changes the events (or global) time window for the appliance. Note that packets in a passive deployment are not dropped, regardless of intrusion rule state or the inline drop behavior of the intrusion policy.

The Network Compliance Widget

The Network Compliance widget summarizes your hosts' compliance with the 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.

You can also use the widget preferences to specify which of three different styles you want to use to display network compliance.

The **Network Compliance** style (the default) displays a pie chart that shows the number of hosts that are compliant, non-compliant, and that have not been evaluated. You can click the pie chart to view the host violation count, which lists the hosts that violate at least one 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 preferences control how often the widget updates. You can check the **Show Not Evaluated** box to hide events which have not been evaluated.

The Product Licensing Widget

The Product Licensing widget shows the device and feature licenses currently installed on the 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 management center, including temporary licenses, while the Expiring Licenses section displays only temporary and expired licenses.

The bars in the widget background show the percentage of each type of license that is being used; you should read the bars from right to left. Expired licenses are marked with a strikethrough.

You can configure the widget to display either the features that are currently licensed, or all the features that you can license, by modifying the widget preferences. The preferences also control how often the widget updates.

You can click any of the license types to go to the License page of the local configuration and add or delete feature licenses.

The Product Updates Widget

The Product Updates widget provides you with a summary of the software currently installed on the appliance as well as information on updates that you have downloaded, but not yet installed. This widget appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

Because the widget uses scheduled tasks to determine the latest version, it displays `Unknown` until you configure a scheduled task to download, push or install updates.

You can configure the widget to hide the latest versions by modifying the widget preferences. The preferences also control how often the widget updates.

The widget also provides you with links to pages where you can update the software. You can:

- Manually update an appliance by clicking the current version.
- Create a scheduled task to download an update by clicking the latest version.

The RSS Feed Widget

The RSS Feed widget adds an RSS feed to a dashboard. By default, the widget shows a feed of Cisco security news. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can also configure the widget to display a preconfigured feed of company news, the Snort.org blog, or the Cisco Threat Research blog, or you can create a custom connection to any other RSS feed by specifying its URL in the widget preferences. The management center can display encrypted RSS feeds only if they use trusted server certificates signed by a certificate authority (CA) that the management center recognizes. If you configure the RSS Feed widget to display an encrypted RSS feed that uses a CA the 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.

The System Time Widget

The System Time widget shows the local system time, uptime, and boot time for the appliance. It appears by default on the Status tabs of the Detailed Dashboard and the Summary Dashboard.

You can configure the widget to hide the boot time by modifying the widget preferences. The preferences also control how often the widget synchronizes with the appliance's clock.

Managing Dashboards

Procedure

Step 1 Choose **Overview** > **Dashboards** > **Dashboard**, and then choose the dashboard you want to modify from the menu.

Step 2 Manage your dashboards:

- Create Dashboards — Create a custom dashboard; see Creating Custom Dashboards, on page 354.
- Delete Dashboards — To delete a dashboard, click **Delete** ( ) next to the dashboard you want to delete. If you delete your default dashboard, you must define a new default or the appliance prompts you to choose a dashboard every time you attempt to view a dashboard.
- Edit Options — Edit custom dashboard options; see Editing Dashboards Options, on page 356.
- Modify Time Constraints — Modify the time display or pause/unpause the dashboard as described in Modifying Dashboard Time Settings, on page 356.

Step 3 Add (see Adding a Dashboard, on page 353), Delete (click **Close** ( )), and Rename (see Renaming a Dashboard, on page 358) dashboards.

Note
You cannot change the order of dashboards.

Step 4 Manage dashboard widgets:

- Add Widgets — Add widgets to a dashboard; see Adding Widgets to a Dashboard, on page 353.
- Configure Preferences — Configure widget preferences; see Configuring Widget Preferences, on page 354.
- Customize Display — Customize the widget display; see Customizing the Widget Display, on page 355.
- View Events — View associated events from the Custom Analysis Widget; see Viewing Associated Events from the Custom Analysis Widget, on page 347.

Tip
Every configuration of the Custom Analysis widget in the Cisco predefined dashboards corresponds to a system preset for that widget. If you change or delete one of these widgets, you can restore it by creating a new Custom Analysis widget based on the appropriate preset.

Adding a Dashboard

Procedure

Step 1 View the dashboard you want to modify; see Viewing Dashboards, on page 358.

Step 2 Click **Add** ( ).

Step 3 Enter a name.

Step 4 Click **OK**.

Adding Widgets to a Dashboard

Each tab can display one or more widgets in a three-column layout. When adding a widget to a dashboard, you choose the tab to which you want to add the widget. The system automatically adds it to the column with the fewest widgets. If all columns have an equal number of widgets, the new widget is added to the leftmost column. You can add a maximum of 15 widgets to a dashboard tab.

Tip After you add widgets, you can move them to any location on the tab. You cannot, however, move widgets from tab to tab.

The dashboard widgets you can view depend on the type of appliance you are using, your user role, and your current domain (in a multidomain deployment). Keep in mind that because not all user roles have access to all dashboard widgets, users with fewer permissions viewing a dashboard created by a user with more permissions may not be able to use all of the widgets on the dashboard. Although the unauthorized widgets still appear on the dashboard, they are disabled.

Procedure

Step 1 View the dashboard where you want to add a widget; see Viewing Dashboards, on page 358.

Step 2 Click the tab where you want to add the widget.

Step 3 Click **Add Widgets**. You can view the widgets in each category by clicking on the category name, or you can view all widgets by clicking **All Categories**.

Step 4 Click **Add** next to the widgets you want to add. The Add Widgets page indicates how many widgets of each type are on the tab, including the widget you want to add.

Tip
To add multiple widgets of the same type (for example, you may want to add multiple RSS Feed widgets, or multiple Custom Analysis widgets), click **Add** again.

Step 5 When you are finished adding widgets, click **Done** to return to the dashboard.

What to do next

- If you added a Custom Analysis widget, configure the widget preferences; see Configuring Widget Preferences, on page 354.

Related Topics

Widget Availability, on page 340

Configuring Widget Preferences

Each widget has a set of preferences that determines its behavior.

Procedure

Step 1 On the title bar of the widget whose preferences you want to change, click **Show Preferences** ( ).

Step 2 Make changes as needed.

Step 3 On the widget title bar, click **Hide Preferences** ( ) to hide the preferences section.

Creating Custom Dashboards

Tip Instead of creating a new dashboard, you can export a dashboard from another appliance, then import it onto your appliance. You can then edit the imported dashboard to suit your needs.

Procedure

Step 1 Choose **Overview** > **Dashboards** > **Management**.

Step 2 Click **Create Dashboard**.

Step 3 Modify the custom dashboard options as described in Custom Dashboard Options, on page 355.

Step 4 Click **Save**.

Custom Dashboard Options

The table below describes options you can use when creating or editing custom dashboards.

Table 28: Custom Dashboard Options

Option	Description
Copy Dashboard	When you create a custom dashboard, you can choose to base it on any existing dashboard, whether user-created or system-defined. This option makes a copy of the preexisting dashboard, which you can modify to suit your needs. Optionally, you can create a blank new dashboard by choosing **None**. This option is available only when you create a new dashboard. In a multidomain deployment, you can copy any non-private dashboards from ancestor domains.
Name	A unique name for the custom dashboard.
Description	A brief description of the custom dashboard.
Change Tabs Every	Specifies (in minutes) how often the dashboard should cycle through its tabs. Unless you pause the dashboard or your dashboard has only one tab, this setting advances your view to the next tab at the interval you specify. To disable tab cycling, enter `0` in the **Change Tabs Every** field.
Refresh Page Every	Determines how often the entire dashboard page automatically refreshes. Refreshing the entire dashboard allows you to see any preference or layout changes that were made to a shared dashboard by another user, or that you made to a private dashboard on another computer, since the last time the dashboard refreshed. A frequent refresh can be useful, for example, in a networks operations center (NOC) where a dashboard is displayed at all times. If you make changes to the dashboard at a local computer, the dashboard in the NOC automatically refreshes at the interval you specify, and no manual refresh is required. This refresh does not update the data, and you do not need to refresh the entire dashboard to see data updates; individual widgets update according to their preferences. This value must be greater than the **Change Tabs Every** setting. Unless you pause the dashboard, this setting will refresh the entire dashboard at the interval you specify. To disable the periodic page refresh, enter `0` in the **Refresh Page Every** field. **Note** This setting is separate from the update interval available on many individual widgets; although refreshing the dashboard page resets the update interval on individual widgets, widgets will update according to their individual preferences even if you disable the **Refresh Page Every** setting.
Save As Private	Determines whether the custom dashboard can be viewed and modified by all users of the appliance or is associated with your user account and reserved solely for your own use. Keep in mind that any user with dashboard access, regardless of role, can modify shared dashboards. If you want to make sure that only you can modify a particular dashboard, save it as private.

Customizing the Widget Display

You can minimize and maximize widgets, as well as rearrange the widgets on a tab.

Procedure

Step 1 View a dashboard; see Viewing Dashboards, on page 358.

Step 2 Customize the widget display:

- To rearrange a widget on a tab, click the title bar of the widget you want to move, then drag it to its new location.

  Note
  You cannot move widgets from tab to tab. If you want a widget to appear on a different tab, you must delete it from the existing tab and add it to the new tab.

- To minimize or maximize a widget on the dashboard, click **Minimize** ( ) or **Maximize** ( ) in a widget's title bar.
- To delete a widget if you no longer want to view it on a tab, click **Close** ( ) in the title bar of the widget.

Editing Dashboards Options

Procedure

Step 1 View the dashboard you want to edit; see Viewing Dashboards, on page 358.

Step 2 Click **Edit** ( ).

Step 3 Change the options as described in Custom Dashboard Options, on page 355.

Step 4 Click **Save**.

Modifying Dashboard Time Settings

You can change the time range to reflect a period as short as the last hour (the default) or as long as the last year. When you change the time range, the widgets that can be constrained by time automatically update to reflect the new time range.

The maximum number of data points in any graph is 300, and the time setting determines how much time is summarized within each data point. Following is the number of data points, and the time span covered, in the dashboards for each time range:

- 1 hour = 12 data points, 5 minutes each
- 6 hours = 72 data points, 5 minutes each
- 1 day = 288 data points, 5 minutes each
- 1 week = 300 data points, 33.6 minutes each

- 2 weeks = 300 data points, 67.2 minutes each
- 30 days = 300 data points, 144 minutes each
- 90 days = 300 data points, 432 minutes each
- 180 days = 300 data points, 864 minutes each
- 1 year = 300 data points, 1752 minutes each

Note that not all widgets can be constrained by time. For example, the dashboard time range has no effect on the Appliance Information widget, which provides information that includes the appliance name, model, and current version of the software.

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 the following effects:

- Individual widgets stop updating, regardless of any **Update Every** widget preference.
- Dashboard tabs stop cycling, regardless of the **Cycle Tabs Every** setting in the dashboard properties.
- Dashboard pages stop refreshing, regardless of the **Refresh Page Every** setting in the dashboard properties.
- Changing the time range has no effect.

When you are finished with your analysis, you can unpause the dashboard. Unpausing the dashboard causes all appropriate widgets on the page to update to reflect the current time range. In addition, dashboard tabs resume cycling and the dashboard page resumes refreshing according to the settings you specified in the dashboard properties.

If you experience connectivity problems or other issues that interrupt the flow of system information to the dashboard, the dashboard automatically pauses and an error notice appears until the problem is resolved.

Note Your session normally logs you out after 1 hour of inactivity (or another configured interval), regardless of whether the dashboard is paused. If you plan to passively monitor the dashboard for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings.

Procedure

Step 1 View the dashboard where you want to add a widget; see Viewing Dashboards, on page 358.

Step 2 Optionally, to change the dashboard time range, choose a time range from the **Show the Last** drop-down list.

Step 3 Optionally, pause or unpause the dashboard on the time range control, using **Pause** ( ) or **Play** ( ).

Renaming a Dashboard

Procedure

Step 1 View the dashboard you want to modify; see Viewing Dashboards, on page 358.

Step 2 Click the dashboard title you want to rename.

Step 3 Type a name.

Step 4 Click **OK**.

Viewing Dashboards

By default, the home page for your appliance displays the default dashboard. If you do not have a default dashboard defined, the home page shows the Dashboard Management page, where you can choose a dashboard to view.

Procedure

At any time, you can do one of the following:

- To view the default dashboard for your appliance, choose **Overview** > **Dashboards** > **Dashboard**.
- To view a specific dashboard, choose **Overview** > **Dashboards** > **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.

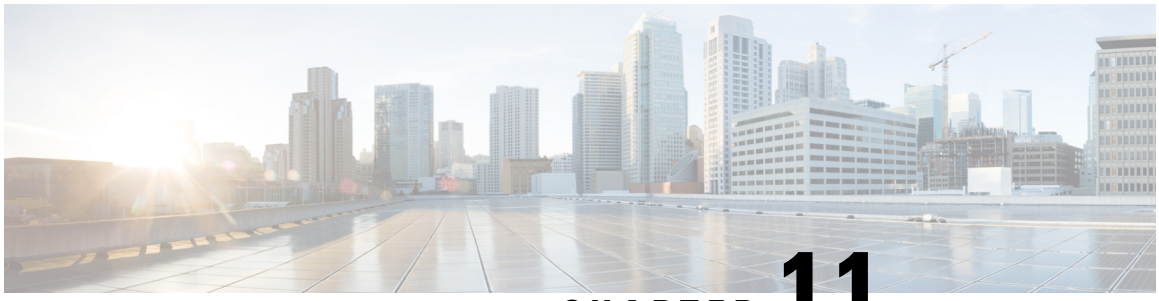

CHAPTER 11

Health

The following topics describe how to use health monitoring:

Requirements and Prerequisites for Health Monitoring

Model Support

Any

Supported Domains

Any

User Roles

Admin

Maintenance User

About Health Monitoring

The health monitor on the management center tracks various health indicators to ensure that the hardware and software in the system are working correctly. You can use the health monitor to check the status of critical functionality across your deployment.

You can configure the frequency for running the health modules for alerting. The management center also supports time series data collection. You can configure 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. The metric data is collected for analysis and hence no alerting is associated with it.

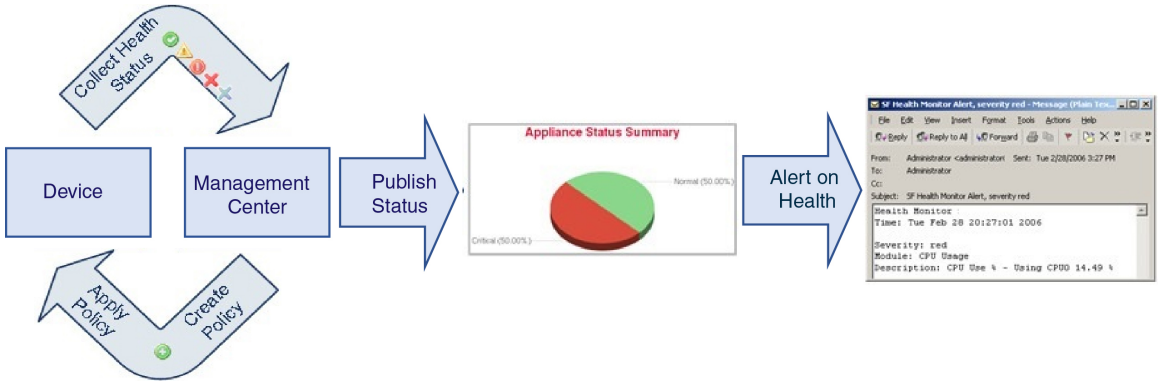

You can use the health monitor to create a collection of tests, referred to as a *health policy*, and apply the health policy to one or more appliances. The tests, referred to as *health modules*, are scripts that test for the criteria you specify. You can modify a health policy by enabling or disabling 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 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 management center and report the status of the 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.

If you are 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 management center. Individual appliance health monitors let you drill down into health details for a specific appliance.

Fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. These event views allow you to search and view event data and to access other information that may be related to the events you are investigating. For example, if you want to see all the occurrences of CPU usage with a certain percentage, you can search for the CPU usage module and enter the percentage value.

You can also configure email, SNMP, or syslog alerting in response to health events. A *health alert* is an association between a standard alert and a health status level. For example, if you 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 management center high-availability deployment running Version 6.7 or higher, the active management center creates a health monitor page that uses REST APIs to show detailed metric-based information. The standby management center 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 management center does not display the metric-based information.

Health Modules

Health modules, or *health tests,* test for the criteria you specify in a health policy.

There are two types of health module: alert and metrics. Alert modules (sometimes called *legacy* modules) monitor system infrastructure and report health status only. 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. You can create custom dashboards with your preferred health metrics, allowing you to monitor statistics or troubleshoot appliance health issues.

Table 29: 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.
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.

Module	Type	Description
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.
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.
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. See Disk Usage and Drain of Events Health Monitor Alerts, on page 434 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. See 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.

Module	Type	Description
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 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, see the *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.
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. **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 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 management center interface statistics are different. Due to the difference in sampling interval, throughput values in the 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** > **Dashboard** page) can be different as it displays the input and output rates from Snort.

Module	Type	Description
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 than1000, 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.
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, its 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. See Memory Usage Thresholds for Health Monitor Alerts, on page 433 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 management center directly using the **configure network management-data-interface** command. This module alerts when there is a conflict between the existingmanagement 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.

Module	Type	Description
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.
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 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 > AMP > Dynamic Analysis Connections** 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. See the **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 management center or on any devices, verify that the management center can reach the devices.
VPN Statistics	Metrics	Monitors site-to-site and remote access VPN tunnels between threat defense devices.
XTLS Counters	Metrics	Monitors XTLS/SSL flows, memory and cache effectiveness. Disabled by default.

Table 30: Management Center Health Modules

Module	Type	Description
Secure Endpoint Status	Alert	The module alerts if the management center cannot connect to the AMP cloud or Cisco AMP Private Cloud after an initial successful connection, or if the private cloud cannot contact the public AMP cloud. It also alerts if you deregister an AMP cloud connection using the Secure Endpoint management console.
AMP for Firepower Status	Alert	Alerts if: • The management center cannot contact the AMP cloud (public or private) or the Secure Malware Analytics Cloud or Appliance, or the AMP private cloud cannot contact the public AMP cloud. • The encryption keys used for the connection are invalid. • A device cannot contact the Secure Malware Analytics Cloud or Secure Malware Analytics Appliance to submit files for dynamic analysis. • An excessive number of files are detected in network traffic based on the file policy configuration. If your management center loses connectivity to the Internet, the system may take up to 30 minutes to generate a health alert.
Appliance Heartbeat	Alert	This module determines if an appliance heartbeat is being heard from the appliance and alerts based on the appliance heartbeat status.
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	Alert	Alerts if the configuration database size is too big. It also monitors the system for database schema or configuration data (sometimes called *EO*) integrity issues. If this module alerts, contact Cisco TAC. This will prevent upgrade.
Discovery Host Limit	Alert	This module determines if the number of hosts the management center can monitor is approaching the limit and alerts based on the warning level configured for the module. For more information, see Host Limit.
Disk Status	Alert	This module examines the performance of the hard disk and malware storage pack (if installed) on the appliance. This module generates a Warning (yellow) health alert when the hard disk and RAID controller (if installed) are in danger of failing, or if an additional hard drive is installed that is not a malware storage pack. This module generates an Alert (red) health alert when an installed malware storage pack cannot be detected.

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. See Disk Usage and Drain of Events Health Monitor Alerts, on page 434 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. See 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 management center.
Event Backlog Status	Alert	Alerts if the backlog of event data awaiting transmission from the device to the 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 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 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 management center hardware: fan speed, temperature, and power supply. Alerts when values exceed configurable thresholds.
Health Monitor Process	Alert	Monitors the health process itself, and alerts if there have been no health events in some number of minutes (configurable).

Module	Type	Description
ISE Connection Monitor	Alert	This module monitors the status of the server connections between the Cisco Identity Services Engine (ISE) and the management center. ISE provides additional user data, device type data, device location data, SGTs (Security Group Tags), and SXP (Security Exchange Protocol) services.
License Monitor	Alert	This module monitors expiration of Classic licenses.
Local Malware Analysis	Alert	This module monitors ClamAV updates for Local Malware Analysis.
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, its 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. See Memory Usage Thresholds for Health Monitor Alerts, on page 433 for additional information about how thresholds are calculated. Beginning with Version 6.6.0, the minimum required RAM for management center virtual upgrades to Version 6.6.0+ is 28 GB, and the recommended RAM for management center virtual deployments is 32 GB. We recommend you do not decrease the default settings: 32 GB RAM for most management center virtual instances, 64 GB for the management center virtual 300 (VMware only). **Attention** A critical alert is generated by the health monitor when insufficient RAM is allocated to a 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.
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.
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 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, see *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.
Security Intelligence	Alert	Alerts if Security Intelligence is in use and the management center cannot update a feed, or feed data is corrupt or contains no recognizable IP addresses. See also the Threat Data Updates on Devices module.
Smart License Monitor	Alert	Monitors Smart Licensing status and alerts if: • There is a communication error between the Smart Licensing Agent (Smart Agent) and the Smart Software Manager. • The Product Instance Registration Token has expired. • The Smart License usage is out of compliance. • The Smart License authorization or evaluation mode has expired.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the 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 > AMP > Dynamic Analysis Connections** 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. See the **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 management center or on any devices, verify that the management center can reach the devices.
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.

Module	Type	Description
Time Synchronization Status	Alert	This module tracks the synchronization of a device clock that obtains time using NTP with the clock on the NTP server and alerts if the difference in the clocks is more than ten seconds.
Unresolved Groups Monitor	Alert	Monitors Foreign Security Principals (FSPs) that are groups used in policies. Security principals are Active Directory objects, like authenticated user groups, to which security can be applied in access control policies. This module generates a warning alert for unresolved groups that exist but are not used in policies, and a critical alert for unresolved groups that are used in policies.
URL Filtering Monitor	Alert	Monitors connectivity with the Cisco cloud, which is required for downloading URL filtering data and doing URL filtering lookups.
VPN Tunnel Status	Alert	Alerts when VPN tunnels are down. Supported for both remote access and site-to-site VPN.

Configuring Health Monitoring

Procedure

Step 1 Determine which health modules you want to monitor as discussed in Health Modules, on page 361.

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 Creating Health Policies, on page 373.

Step 3 (Optional.) Configure health monitor alerts as discussed in Creating Health Monitor Alerts, on page 379.

You can set up email, syslog, or SNMP alerts that trigger when the health status level reaches a particular severity level for specific health modules.

Health Policies

A health policy contains configurable health test criteria for several modules. You can control which health modules run against each of your appliances and configure the specific limits used in the tests run by each module.

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

You can create one health policy that can be applied to every appliance in your system, customize each health policy to the specific appliance where you plan to apply it, or use the default health policy provided for you.

Note When you register an appliance, the 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 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 management center.

This initial health policy is based on a default health policy, which you can neither view nor edit, but which you can copy when you create a custom health policy.

Upgrades and the Default Health Policy

When you upgrade the 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.

Creating Health Policies

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you choose as a basis for the new policy. You can 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 374.
- Edit the policy to specify the module-level policy settings as described in Edit a Health Policy, on page 375.

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 management center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

If you apply a policy with all modules disabled to an appliance, it removes all applied health policies from the appliance, so no health policy is applied. However, you must have at least one health policy assigned to an appliance.

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 and apply it to the appliance. 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.

What to do next

- Optionally, monitor the task status; see View Task Messages, on page 431.

  Monitoring of the appliance starts when the policy is successfully applied.

Edit a Health Policy

You can edit a health policy that you want to modify.

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 (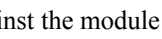) or turn off () to enable or disable testing of health status respectively.
- To execute a bulk enable or disable testing on the health modules, click the **Select All** toggle button

Note

- The modules and attributes are flagged with the supporting appliances—threat defense, management center, or both.
- You cannot choose to include or exclude the individual attributes of CPU and Memory modules.

For information on the modules, see Health Modules, on page 361.

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, see About Dashboards, on page 339. 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**.

Alternatively, you can apply the health policy to your appliance as described in Apply a Health Policy, on page 374

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 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.

Tip To stop health monitoring for an appliance, create a health policy with all modules disabled and apply it to the appliance.

Procedure

Step 1 Choose **System** ( ) > **Health** > **Policy**.

Step 2 Click **Delete** ( ) next to the policy that you want to delete, and then click **Delete health policy** to delete it.

Device Exclusion in Health Monitoring

In the course of normal network maintenance, you disable appliances or make them temporarily unavailable. Because those outages are deliberate, you do not want the health status from those appliances to affect the summary health status on your management center.

You can use the health monitor exclude feature to disable health monitoring status reporting on an appliance or module. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the management center from displaying a warning or critical state because of the lapsed connection to the device.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the appliance or module from the excluded list, the events that were generated during the exclusion continue to show a status of disabled.

To temporarily disable health events from an appliance, go to the exclusion configuration page and add an appliance to the device exclude list. After the setting takes effect, the system no longer considers the excluded appliance when calculating the overall health status. The Health Monitor Appliance Status Summary lists the appliance as disabled.

You can also disable an individual health module. For example, when you reach the host limit on the 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 appliances that are excluded if you expand to view the list of appliances with a particular status by clicking the arrow in that status row.

Note On 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 management center, the exclusion settings remain persistent. The newly re-registered device remains excluded.

Excluding Appliances from Health Monitoring

You can exclude appliances individually or by group, model, or associated health policy.

If you need to set the events and health status for an individual appliance to disabled, you can exclude the appliance. After the exclusion settings take effect, the appliance shows as disabled in the Health Monitor Appliance Module Summary, and health events for the appliance have a status of disabled.

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**.

What to do next

To exclude individual health policy modules on appliances, see Excluding Health Policy Modules, on page 377.

Excluding Health Policy Modules

You can exclude individual health policy modules on appliances. This allows you to prevent health events from the module from changing the status for the appliance to warning or critical.

Note Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

After the exclusion settings take effect, the appliance shows the number of modules being excluded in the device from health monitoring.

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 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 specific device only; for more information, see Health Modules, on page 361.

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 the following:

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**.

Expired Health Monitor Exclusions

When the exclusion period for a device or modules lapses, you can choose to clear or renew the exclusion.

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 appliance. 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 appliance, 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 appliance. 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

You can set up alerts to notify you through email, through SNMP, or through the syslog when the status changes for the modules in a health policy. You can associate an existing alert response with health event levels to trigger and alert when health events of a particular level occur.

For example, if you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

Health Monitor Alert Information

The alerts generated by the health monitor contain the following information:

- Severity, which indicates the severity level of the alert.
- Module, which specifies the health module whose test results triggered the alert.
- Description, which includes the health test results that triggered the alert.

The table below describes these severity levels.

Table 31: Alert Severities

Severity	Description
Critical	The health test results met the criteria to trigger a Critical alert status.
Warning	The health test results met the criteria to trigger a Warning alert status.
Normal	The health test results met the criteria to trigger a Normal alert status.
Error	The health test did not run.
Recovered	The health test results met the criteria to return to a normal alert status, following a Critical or Warning alert status.

Creating Health Monitor Alerts

You must be an Admin user to perform this procedure.

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

If you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

Before you begin

- Configure an alert response that governs the management center's communication with the SNMP, syslog, or email server where you send the health alert; see Secure Firewall Management Center Alert Responses, on page 551.

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.

Editing Health Monitor Alerts

You must be an Admin user to perform this procedure.

You can edit existing health monitor alerts to change the severity level, health module, or alert response associated with the health monitor alert.

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**.

Deleting Health Monitor Alerts

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor Alerts**.

Step 2 Click **Delete** (🗑) next to the health alert you want to delete, and then 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; see Secure Firewall Management Center Alert Responses, on page 551.

About the Health Monitor

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The health monitor provides the compiled health status for all devices managed by the management center, plus the 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 management center and all of the devices that the 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 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 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. When you use the Monitoring navigation pane:

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.

What to do next

- See Device Health Monitors, on page 386 for information about the compiled health status and metrics for any device managed by the management center.
- See Using Management Center Health Monitor, on page 382 for information about the health status of the management center.

To return to the Health Status landing page at any time, click **Home**.

Using Management Center Health Monitor

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The management center monitor provides a detailed view of the health status of the 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 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 management center.

- Process Health—The Process Health panel has an at-a-glace 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 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 management center data is stored.
- Hardware Statistics—The hardware statistics shows the fan speed, power supply, and temperature of the management center chassis. For more information, see Hardware Statistics on Management Center, on page 385.

Tip Your session normally logs you out after 1 hour of inactivity (or another configured interval). If you plan to passively monitor health status for long periods of time, consider exempting some users from session timeout, or changing the system timeout settings. See Add or Edit an Internal User, on page 125 and Configure Session Timeouts, on page 102 for more information.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Use the **Monitoring** navigation pane to access the management center and device-specific health monitors.

- A standalone management center is shown as a single node; a high-availability management center is shown as a pair of nodes.
- The health monitor is available to both the active and standby 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 management center, as access to various pages can be restricted in the standby management center.

Step 3 Explore the management center dashboard.

The management center dashboard includes a summary view of the HA state of the management center (if configured), as well as at-a-glance views of management center processes and device metrics such as CPU, memory, and disk usage.

Running All Modules for an Appliance

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run all health module tests on demand to collect up-to-date health information for the appliance.

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.

Running a Specific Health Module

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Health module tests run automatically at the policy run time interval you configure when you create a health policy. However, you can also run a health module test on demand to collect up-to-date health information for that module.

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.

Generating Health Module Alert Graphs

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

You can graph the results over a period of time of a particular health test for a specific appliance.

Procedure

Step 1 View the health monitor for the appliance.

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.

Hardware Statistics on Management Center

The hardware statistics on the management center appliance (only physical) includes information on its hardware entities, such as fan speed, power supply, and temperature. For SNMP to poll and send traps to monitor the health of a management center:

1. Enable SNMP on the management center for polling the MIBs. By default, the SNMP on the management center is disabled. See Configure SNMP Polling, on page 100.
2. Add an ACL entry for each of the required SNMP host to enable traps. Ensure to specify the host's IP address and select the port as SNMP. See Configure an Access List, on page 47.

To view the hardware statistics on the **Health** > **Monitor** page:

1. On the **Health** > **Policy** page, ensure that the Hardware Statistics module is enabled. You can change the default threshold values.
2. Add a portlet to the management center health monitoring dashboard—select Hardware Statistics metric group, and then select Fan Speed and Temperature metrics.

You can view the power supply status under the firewall management center in the **Health Monitoring** > **Home** page.

Note

- The fan speed is displayed in RPM.
- The temperature is displayed in °C (Celsius).
- When one slot of the power supply is active, the dashboard displays it as *Online* and the other slot as *No Power*.
- Each horizontal line in the graphs shows the status for each PSU and fan respectively.
- Hover over the graph to view the data of that individual statistics.

Device Health Monitors

The device health monitor provides the compiled health status for any device managed by the 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 the following 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.
  - 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.

See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Viewing System Details and Troubleshooting

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The System Details section provides a general system information for a selected device. You can also launch troubleshooting tasks for that device.

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 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 the following troubleshoot choices:

- Generate troubleshooting files; see Generate Troubleshooting Files for Specific System Functions, on page 438
- Generate and download advanced troubleshooting files; see Download Advanced Troubleshooting Files, on page 439.
- Create and modify health policies; see Creating Health Policies, on page 373.
- Create and modify health monitor alerts; see Creating Health Monitor Alerts, on page 379.

Viewing the Device Health Monitor

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The device health monitor provides a detailed view of the health status of a firewall device. The device health monitor compiles device metrics and provides health status and trends of the device in an array of dashboards.

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 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.

Step 4 You can 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 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
The **Process Health** widget in management center displays the CPU% usage for each process. The CPU% metric reflects usage relative to the number of cores in the 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. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device 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; see Correlating Device Metrics, on page 388.

Correlating Device Metrics

The device health monitor includes an array of key threat defense device metrics that serve to predict and respond to system events. The health of any 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. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Before you begin

- To view and correlate the time series data (device metrics) in the health monitor dashboard, enable REST API (**Settings** > **Configuration** > **REST API Preferences**).
- You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

Note Correlating device metrics is available only for threat defense 6.7 and later versions. Hence, for threat defense versions earlier than 6.7, the health monitor dashboard does not display these metrics even if you enable REST API.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the **Devices** list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed devices.

Step 3 Choose the device for which you want to modify the dashboard.

Step 4 Click the **Add Dashboard** (+) icon in the upper right corner of the device monitor to add a new dashboard.

Step 5 From the **Select Correlation Group** drop-down, choose a predefined correlation group or to create a custom group.

Step 6 To create a dashboard from a predefined correlation group, choose the group and click **Add**.

Step 7 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.

See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click **Add Metrics** to add and select metrics from another group.

Step 9 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.

Step 10 Click **Add** to add the dashboard to the health monitor.

Step 11 You can **Edit** or **Delete** custom correlation dashboards.

Cluster Health Monitor

When a threat defense is the control node of a cluster, the management center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the management center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.

- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing the Cluster Health Monitor

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the management center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⌄) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster 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 its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can 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 on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node 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 node. 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.

Step 7 (For node-specific health monitor) 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 that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Health Monitor Status Categories

Available status categories are listed by severity in the table below.

Table 32: 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

The Health Event View page allows you to view health events logged by the health monitor on the management center logs health events. The fully customizable event views allow you to quickly and easily analyze the health status events gathered by the health monitor. You can search event data to easily access other information that may be related to the events you are investigating. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events.

You can perform many of the standard event view functions on the health event view pages.

Viewing Health Events

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The Table View of Health Events page provides a list of all health events on the specified appliance.

When you access health events from the Health Monitor page on your 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.

Procedure

Choose **System** (⚙) > **Health** > **Events**.

Tip
If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the Select Workflow page, click **Health Events**.

Note
If no events appear, you may need to adjust the time range.

Viewing Health Events by Module and Appliance

Procedure

Step 1 View the health monitor for the appliance; see Viewing the Device Health Monitor, on page 387.

Step 2 In the **Module Status Summary** graph, click the color for the event status category you want to view.

The Alert Detail list toggles the display to show or hide events.

Step 3 In the **Alert Detail** row for the alert for which you want to view a list of events, click **Events**.

The Health Events page appears, containing results for a query with the name of the appliance and the name of the specified health alert module as constraints. If no events appear, you may need to adjust the time range.

Step 4 If you want to view all health events for the specified appliance, expand **Search Constraints**, and click the **Module Name** constraint to remove it.

Viewing the Health Events Table

You can view and modify the Health Events Table.

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.

The Health Events Table

The Health Monitor modules you choose to enable in your health policy run various tests to determine appliance health status. When the health status meets criteria that you specify, a health event is generated.

The table below describes the fields that can be viewed and searched in the health events table.

Table 33: Health Event Fields

Field	Description
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type `CPU`. The search should retrieve applicable CPU Usage and CPU temperature events.
Test Name (Search only)	The name of the health module that generated the event.
Time (Search only)	The timestamp for the health event.
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled `Unable to Execute`.

Field	Description
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the 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 management center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Critical, Yellow, Green, or Disabled) reported for the appliance.
Device	The appliance where the health event was reported.

History for Health Monitoring

Table 34:

Feature	Minimum Management Center	Minimum Threat Defense	Details
Health alerts for NTP server sync issues.		Any	Introduced the **Time Sever 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**.

Feature	Minimum Management Center	Minimum Threat Defense	Details
New cluster health monitor dashboard.	7.3	Any	A new dashboard to view the cluster health monitor metrics was introduced with the following components: • Overview—Displays information about the cluster topology, cluster statistics, and metric charts. • Load Distribution—Displays load distribution across the cluster nodes. • Member Performance—Displays current metrics of all the member nodes of the cluster. • CCL—Displays, graphically, the cluster control link data namely, the input, and output rate. **Note** These features are applicable only for a cluster. Hence, you must select the cluster under the **Devices** list on the **Monitoring** pane to view and use the cluster dashboard. New/modified screens: **System**(⚙) > **Health** > **Monitor**.
New hardware statistics module.	7.3	Any	The management center hardware and environment status statistics were added to the health monitor dashboard: • A new policy module, **Hardware Statistics**, was introduced to enable monitoring of hardware daemons on the management center hardware. The metrics included fan speed, temperature, and power supply. • A custom metric group, **Hardware Statistics**, was also added to view graphical representation of the hardware health metrics on the monitoring dashboard. • The power supply status is captured in **Health Alerts** of the management center. **Note** These features are applicable only for the management center. Hence, they are available only on the management center dashboard. New/modified screens: • **System**(⚙) > **Health** > **Monitor** • **System**(⚙) > **Health** > **Policy**

Feature	Minimum Management Center	Minimum Threat Defense	Details
New hardware and environment status metric group,	7.3	Any	The threat defense hardware and environment status statistics were added to the health monitor dashboard: • A custom metric group, **Hardware / Environment Status**, was introduced to view hardware-related statistics on the threat defense. The metrics included fan speed, chassis temperature, SSD status, and power supply. • The device **Health Alerts** was enhanced to include the power supply status of the threat defense hardware—*Critical* alert is displayed for abnormal thermal status, and *Normal* alert is displayed for normal thermal status. **Note** These features are applicable only for threat defense. Hence, you must select the appropriate device under the **Devices** list on the **Monitoring** pane. New/modified screens: **System**(⚙) > **Health** > **Monitor**.
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 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.
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.

Feature	Minimum Management Center	Minimum Threat Defense	Details
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 management center is running any of these versions, you may also continue to see alerts.

Feature	Minimum Management Center	Minimum Threat Defense	Details
New health modules.	7.0	Any	

Feature	Minimum Management Center	Minimum Threat Defense	Details
			We added the following health modules: • AMP Connection Status: Monitors AMP cloud connectivity from the threat defense. • AMP Threat Grid Status: Monitors AMP Threat Grid cloud connectivity from the threat defense. • ASP Drop: Monitors the connections dropped by the data plane accelerated security path. • Advanced Snort Statistics: Monitors Snort statistics related to packet performance, flow counters, and flow events. • Event Stream Status: Monitors connections to third-party client applications that use the Event Streamer. • FMC Access Configuration Changes: Monitors access configuration changes made directly on the management center. • FMC HA Status: Monitors the active and standby management center and the sync status between the devices. Replaces the HA Status module. • FTD HA Status: Monitors the active and standby threat defense HA pair and the sync status between the devices. • File System Integrity Check: Performs a file system integrity check if the system has CC mode or UCAPL mode enabled. • Flow Offload: Monitors hardware flow offload statistics on the Firepower 9300 and 4100 platforms. • Hit Count: Monitors the number of times a particular rule is hit on the access control policy. • MySQL Status: Monitors the status of the MySQL database. • NTP Status FTD: Monitors the NTP clock synchronization status of the managed device. • RabbitMQ Status: Monitors the status of the RabbitMQ messaging broker. • Routing Statistics: Monitors both IPv4 and IPv6 route information from the threat defense. • Security Services Exchange Connection Status: Monitors security services exchange cloud connectivity from the threat defense. • Sybase Status: Monitors the status of the Sybase database. • Unresolved Groups Monitor: Monitors the unresolved groups used in access control policies. • VPN Statistics: Monitors site-to-site and remote access VPN tunnel statistics.

Feature	Minimum Management Center	Minimum Threat Defense	Details
			• xTLS Counters: Monitors xTLS/SSL flows, memory and cache effectiveness.

Health monitor enhancements.	7.0	Any	The health monitor adds the following enhancements:
			• Enhanced management center dashboard with summary views of:
			  • High Availability
			  • Event Rate & Capacity
			  • Process Health
			  • CPU thresholds
			  • Memory
			  • Interface rates
			  • Disk Usage
			• Enhanced threat defense dashboard:
			  • Health alert for split brain scenario
			  • Additional health metrics available from new Health Modules

Feature	Minimum Management Center	Minimum 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 Management Center	Minimum 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 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, see 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 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.

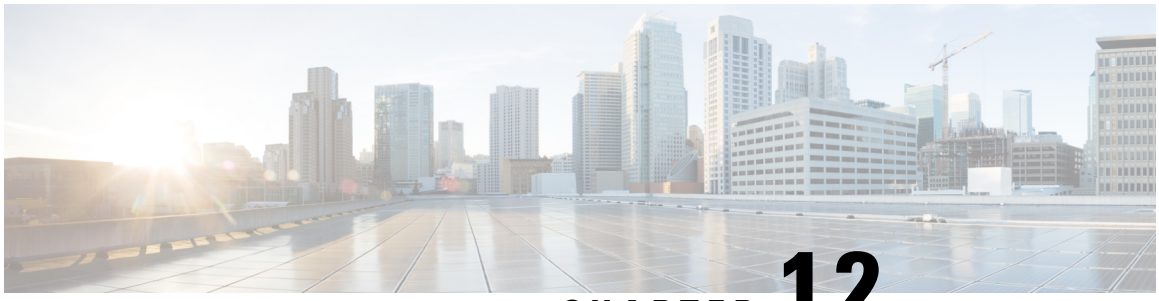

CHAPTER 12

Audit and Syslog

The following topics describe how to audit activity on your system:

The System Log

The System Log (syslog) page provides you with system log information for the appliance.

You can audit activity on your system in two ways. The appliances that are part of the system generate an audit record for each user interaction with the web interface, and also record system status messages in the system log.

The system log displays each message generated by the system. The following items are listed in order:

- the date that the message was generated
- the time that the message was generated
- the host that generated the message
- the message itself

Viewing the System Log

System log information is local. For example, you **cannot** use the management center to view system status messages in the system logs on your managed devices.

You can filter messages using most syntax accepted by the UNIX file search utility Grep. This includes using Grep-compatible regular expressions for pattern matching.

Before you begin

You must be an Admin or Maintenance user and be in the Global domain to view system statistics.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Syslog**.

Step 2 To search for specific message content in the system log:

a) Enter a word or query in the filter field as described in Syntax for System Log Filters, on page 406.

Only Grep-compatible search syntax is supported.

Examples:

To search for all log entries that contain the user name "Admin," use `Admin`.

To search for all log entries that are generated on November 27, use `Nov[[:space:]]*27` or `Nov.*27` (but not `Nov 27` or `Nov*27`).

To search for all log entries that contain authorization debugging information on November 5, use `Nov[[:space:]]*5.*AUTH.*DEBUG`.

b) To make your search case-sensitive, select **Case-sensitive**. (By default, filters are not case-sensitive.)

c) To search for all system log messages that do **not** meet the criteria you entered, select **Exclusion**.

d) Click **Go**.

Syntax for System Log Filters

The following table shows the regular expression syntax you can use in System Log filters:

Table 35: System Log Filter Syntax

Syntax Component	Description	Example
`.`	Matches any character or white space	`Admi.` matches `Admin`, `AdmiN`, `Admi1`, and `Admi`
`[[:alpha:]]`	Matches any alphabetic character	`[[:alpha:]]dmin` matches `Admin`, `bdmin`, and
`[[:upper:]]`	Matches any uppercase alphabetic character	`[[:upper:]]dmin` matches `Admin`, `Bdmin`, and
`[[:lower:]]`	Matches any lowercase alphabetic character	`[[:lower:]]dmin` matches `admin`, `bdmin`, and
`[[:digit:]]`	Matches any numeric character	`[[:digit:]]dmin` matches `0dmin`, `1dmin`, and
`[[:alnum:]]`	Matches any alphanumeric character	`[[:alnum:]]dmin` matches `1dmin`, `admin`, `2dmin`,
`[[:space:]]`	Matches any white space, including tabs	`Feb[[:space:]]29` matches logs from Februar
`*`	Matches zero or more instances of the character or expression it follows	`ab*` matches `a`, `ab`, `abb`, `ca`, `cab`, and `cabb` `[ab]*` matches anything
`?`	Matches zero or one instances	`ab?` matches `a` or `ab`
`\`	Allows you to search for a character typically interpreted as regular expression syntax	`alert\?` matches `alert?`

About System Auditing

The appliances that are part of the system generate an audit record for each user interaction with the web interface.

Related Topics

Standard Reports, on page 525

Audit Records

Secure Firewall Management Centers log read-only auditing information for user activity. Audit logs are presented in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete and report on audit information and can view detailed reports of the changes that users make.

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.

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

SSO Guidelines for the Management Center, on page 145

Viewing Audit Records

On the management center, you can view a table of audit records. The predefined audit workflow includes a single table view of events. You can manipulate the table view depending on the information you are looking for. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this procedure.

Procedure

Step 1 Access the audit log workflow using **System** (⚙) > **Monitoring** > **Audit**.

Step 2 If no events appear, you may need to adjust the time range. For more information, see Event Time Constraints, on page 672.

Note

Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

Step 3 You have the following choices:

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.

- To learn more about the contents of the columns in the table, see The System Log, on page 405.
- To sort and constrain events on the current workflow page, see Using Table View Pages, on page 664.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page. For more information, see Using Workflows, on page 656.
- To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 664.
- To constrain on a specific value, click a value within a row. If you click a value on a drill-down page, you move to the next page and constrain on the value. Note that clicking a value within a row in a table view constrains the table view and does **not** drill down to the next page. See Event View Constraints, on page 678 for more information.

  Tip
  Table views always include “Table View” in the page name.

- To delete audit records, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All** to delete all events in the current constrained view.
- To bookmark the current page so you can quickly return to it, click **Bookmark This Page**. For more information, see Bookmarks, on page 682.
- To navigate to the bookmark management page, click **View Bookmarks**. For more information, see Bookmarks, on page 682.
- To generate a report based on the data in the current view, click **Reporting**. For more information, see Creating a Report Template from an Event View, on page 529.
- To view a summary of system changes recorded in the audit log, click **Compare** next to applicable events in the **Message** column. For more information, see Using the Audit Log to Examine Changes, on page 410.

Related Topics

Event View Constraints, on page 678

Audit Log Workflow Fields

The following table describes the audit log fields that can be viewed and searched.

Table 36: 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.

Field	Description
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 management center for multitenancy.
Configuration Change (search only)	Specifies whether to view audit records of configuration changes in the search results. (`yes` or `no`)
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Related Topics

Event Searches, on page 685

The Audit Events Table View

You can change the layout of the event view or constrain the events in the view by a field value. When disabling columns, after you click the **Close** (✕) in the column heading that you want to hide, in the pop-up window that appears, click **Apply**. When you disable a column, it is disabled for the duration of your session (unless you add it back later). Note that when you disable the first column, the Count column is added.

To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page in the workflow.

Tip Table views always include “Table View” in the page name.

Related Topics

Using Workflows, on page 656

Using the Audit Log to Examine Changes

You can use the audit log to view detailed reports of some of the changes to your system. These reports compare the current configuration of your system to its most recent configuration before a supported change was made.

The Compare Configurations page displays the differences between the system configuration before changes and the running configuration in a side-by-side format. The audit event type, time of last modification, and name of the user who made the change are displayed in the title bar above each configuration.

Differences between the two configurations are highlighted:

- Blue indicates that the highlighted setting is different in the two configurations, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one configuration but not the other.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this procedure.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Audit**.

Step 2 Click **Compare** next to an applicable audit log event in the **Message** column.

Tip
You can navigate through changes individually by clicking **Previous** or **Next** above the title bar. If the change summary is more than one page long, you can also use the scroll bar on the right to view additional changes.

Suppressing Audit Records

If your auditing policy does not require that you audit specific types of user interactions with the 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 be an Admin user to perform this procedure.

Procedure

In the `/etc/sf` directory, create one or more `AuditBlock` files in the following form, where type is one of the types described in Audit Block Types, on page 411:

`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

The contents for each audit block type must be in a specific format, as described in the following table. Make sure you use the correct capitalization for the file names. Note also that the contents of the files are case sensitive.

Note that when you add an `AuditBlock` file, an audit record with a subsystem of `Audit` and a message of `Audit Filter type Changed` is added to the audit events. For security reasons, this audit record **cannot** be suppressed.

Table 37: 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 411 for a list of subsystems that are audited.
User	Create a file named `AuditBlock.user` and include, one per line, each user account that you want to suppress. You can use partial string matching provided that they map from the beginning of the username. For example, the partial username `IPSAnalyst` matches the user names `IPSAnalyst1` and `IPSAnalyst2`.

Audited Subsystems

The following table lists audited subsystems.

Table 38: 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
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 management centers in high availability pairs
IDS Impact Flag	Impact flag configuration for intrusion events
IDS Policy	Intrusion policies

Name	Includes user interactions with...
IDSRule sid:`sig_id` rev:`rev_num`	Intrusion rules by SID
Install	Installing updates
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 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 management center, see:

- Audit Log, on page 48
- Audit Log Certificate, on page 51

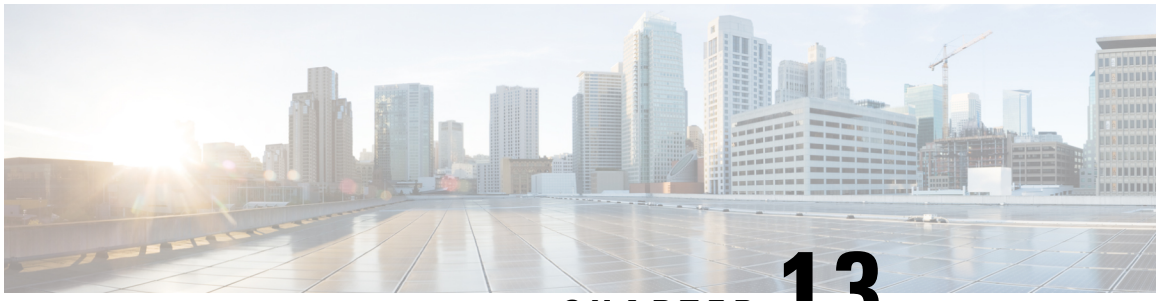

CHAPTER 13

Statistics

The following topics describe how to monitor the system:

About 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.

The Host Statistics Section

The following table describes the host statistics listed on the Statistics page.

Table 39: Host Statistics

Category	Description
Time	The current time on the system.
Uptime	The number of days (if applicable), hours, and minutes since the system was last started.
Memory Usage	The percentage of system memory that is being used.
Load Average	The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.
Disk Usage	The percentage of the disk that is being used. Click the arrow to view more detailed host statistics.
Processes	A summary of the processes running on the system.

The Disk Usage Section

The Disk Usage section of the Statistics page provides a quick synopsis of disk usage, both by category and by partition status. If you have a malware storage pack installed on a device, you can also check its partition status. You can monitor this page from time to time to ensure that enough disk space is available for system processes and the database.

Tip You can also use the Disk Usage health monitor to monitor disk usage and alert on low disk space conditions.

The Processes Section

The Processes section of the Statistics page allows you to see the processes that are currently running on an appliance. It provides general process information and specific information for each running process. You can use the management center's web interface to view the process status for any managed device.

Note that there are two different types of processes that run on an appliance: daemons and executable files. Daemons always run, and executable files are run when required.

Process Status Fields

When you expand the Processes section of the Statistics page, you can also view the following:

Cpu(s)

Lists the following CPU usage information:

- user process usage percentage
- system process usage percentage
- nice usage percentage (CPU usage of processes that have a negative nice value, indicating a higher priority). Nice values indicate the scheduled priority for system processes and can range between -20 (highest priority) and 19 (lowest priority).
- idle usage percentage

Mem

Lists the following memory usage information:

- total number of kilobytes in memory
- total number of used kilobytes in memory
- total number of free kilobytes in memory
- total number of buffered kilobytes in memory

Swap

Lists the following swap usage information:

- total number of kilobytes in swap
- total number of used kilobytes in swap
- total number of free kilobytes in swap
- total number of cached kilobytes in swap

The following table describes each column that appears in the Processes section.

Table 40: Process List Columns

Column	Description
Pid	The process ID number
Username	The name of the user or group running the process
Pri	The process priority
Nice	The *nice* value, which is a value that indicates the scheduling priority of a process. Values range between -20 (highest priority) and 19 (lowest priority)
Size	The memory size used by the process (in kilobytes unless the value is followed by `m`, which indicates megabytes)
Res	The amount of resident paging files in memory (in kilobytes unless the value is followed by `m`, which indicates megabytes)
State	The process state: • D — process is in uninterruptible sleep (usually Input/Output) • N — process has a positive nice value • R — process is runnable (on queue to run) • S — process is in sleep mode • T — process is being traced or stopped • W — process is paging • X — process is dead • Z — process is defunct • < — process has a negative nice value
Time	The amount of time (in hours:minutes:seconds) that the process has been running
Cpu	The percentage of CPU that the process is using
Command	The executable name of the process

Related Topics

System Daemons, on page 418
Executables and System Utilities, on page 419

System Daemons

Daemons continually run on an appliance. They ensure that services are available and spawn processes when required. The following table lists daemons that you may see on the Process Status page and provides a brief description of their functionality.

Note The table below is not an exhaustive list of all processes that may run on an appliance.

Table 41: System Daemons

Daemon	Description
crond	Manages the execution of scheduled commands (cron jobs)
dhclient	Manages dynamic host IP addressing
fpcollect	Manages the collection of client and server fingerprints
httpd	Manages the HTTP (Apache web server) process
httpsd	Manages the HTTPS (Apache web server with SSL) service, and checks for working SSL certificate authentication; runs in the background to provide secure web access to the appl
keventd	Manages Linux kernel event notification messages
klogd	Manages the interception and logging of Linux kernel messages
kswapd	Manages Linux kernel swap memory
kupdated	Manages the Linux kernel update process, which performs disk synchronization
mysqld	Manages database processes
ntpd	Manages the Network Time Protocol (NTP) process
pm	Manages all 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 occ logs runtime information to a file
SFDataCorrelator	Manages data transmission
sfestreamer (management center only)	Manages connections to third-party client applications that use the Event Streamer

Daemon	Description
sfmgr	Provides the RPC service for remotely managing and configuring an appliance using a connection to the appliance
SFRemediateD (management center only)	Manages remediation responses
sftimeserviced (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. Currently used only by health monitoring to send health ev from a managed device to the management center.
sftroughd	Listens for connections on incoming sockets and then invokes the correct executable (Cisco message broker, sfmb) to handle the request
sftunnel	Provides the secure communication channel for all processes requiring communication appliance
sshd	Manages the Secure Shell (SSH) process; runs in the background to provide SSH acce appliance
syslogd	Manages the system logging (syslog) process

Executables and System Utilities

There are a number of executables on the system that run when executed by other processes or through user action. The following table describes the executables that you may see on the Process Status page.

Table 42: System Executables and Utilities

Executable	Description
awk	Utility that executes programs written in the `awk` programming language
bash	GNU Bourne-Again Shell
cat	Utility that reads files and writes content to standard output
chown	Utility that changes user and group file permissions
chsh	Utility that changes the default login shell
SFDataCorrelator (management center only)	Analyzes binary files created by the system to generate events, connection data, and network maps
cp	Utility that copies files
df	Utility that lists the amount of free space on the appliance
echo	Utility that writes content to standard output

Executable	Description
egrep	Utility that searches files and folders for specified input; supports extended set of regular expressions not supported in standard grep
find	Utility that recursively searches directories for specified input
grep	Utility that searches files and directories for specified input
halt	Utility that stops the server
httpsdctl	Handles secure Apache Web processes
hwclock	Utility that allows access to the hardware clock
ifconfig	Indicates the network configuration executable. Ensures that the MAC address stays constant
iptables	Handles access restriction based on changes made to the Access Configuration page.
iptables-restore	Handles iptables file restoration
iptables-save	Handles saved changes to the iptables
kill	Utility that can be used to end a session and process
killall	Utility that can be used to end all sessions and processes
ksh	Public domain version of the Korn shell
logger	Utility that provides a way to access the syslog daemon from the command line
md5sum	Utility that prints checksums and block counts for specified files
mv	Utility that moves (renames) files
myisamchk	Indicates database table checking and repairing
mysql	Indicates a database process; multiple instances may appear
openssl	Indicates authentication certificate creation
perl	Indicates a perl process
ps	Utility that writes process information to standard output
sed	Utility used to edit one or more text files
sfheartbeat	Identifies a heartbeat broadcast, indicating that the appliance is active; heartbeat used to maintain contact between a device and management center.
sfmb	Indicates a message broker process; handles communication between management centers and device.
sh	Public domain version of the Korn shell

Executable	Description
shutdown	Utility that shuts down the appliance
sleep	Utility that suspends a process for a specified number of seconds
smtpclient	Mail client that handles email transmission when email event notification functionality is enabled
snmptrap	Forwards SNMP trap data to the SNMP trap server specified when SNMP notification functionality is enabled
snort	Indicates that Snort is running
ssh	Indicates a Secure Shell (SSH) connection to the appliance
sudo	Indicates a sudo process, which allows users other than admin to run executables
top	Utility that displays information about the top CPU processes **Note** The CPU usage output of this utility is a split-up of different types of usages of the CPU core. You must add both user and system processes usage to know the actual total CPU usage. For example, if the output of top command is: `%Cpu(s): 76.6 us, 22.1 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 1.3 si, 0.0 st` Here, 76.6% of CPU time 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 appear to be different from the Health Monitor dashboard. In addition, this utility uses a three seconds interval to calculate the CPU usage. Whereas, the management center health monitor uses one-second intervals.
touch	Utility that can be used to change the access and modification times of specified files
vim	Utility used to edit text files
wc	Utility that performs line, word, and byte counts on specified files

Related Topics

Configure an Access List, on page 47

The SFDataCorrelator Process Statistics Section

On the management center, you can view statistics about the Data Correlator and network discovery processes for the current day. As the managed devices perform data acquisition, decoding, and analysis, the network discovery process correlates the data with the fingerprint and vulnerability databases, then produces binary files that are processed by the Data Correlator running on the management center. The Data Correlator analyzes the information from the binary files, generates events, and creates network maps.

The statistics that appear for network discovery and the Data Correlator are averages for the current day, using statistics gathered between 12:00 AM and 11:59 PM for each device.

The following table describes the statistics displayed for the Data Correlator process.

Table 43: Data Correlator Process Statistics

Category	Description
Events/Sec	Number of discovery events that the Data Correlator receives and processes per second
Connections/Sec	Number of connections that the Data Correlator receives and processes per second
CPU Usage — User (%)	Average percentage of CPU time spent on user processes for the current day
CPU Usage — System (%)	Average percentage of CPU time spent on system processes for the current day
VmSize (KB)	Average size of memory allocated to the Data Correlator for the current day, in kilobytes
VmRSS (KB)	Average amount of memory used by the Data Correlator for the current day, in kilobytes

The Intrusion Event Information Section

On both the management center and managed devices, you can view summary information about intrusion events on the Statistics page. This information includes the date and time of the last intrusion event, the total number of events that have occurred in the past hour and the past day, and the total number of events in the database.

Note The information in the Intrusion Event Information section of the Statistics page is based on intrusion events stored on the managed device rather than those sent to the management center. No intrusion event information is listed on this page if the managed device cannot (or is configured not to) store intrusion events locally.

The following table describes the statistics displayed in the Intrusion Event Information section of the Statistics page.

Table 44: Intrusion Event Information

Statistic	Description
Last Alert Was	The date and time that the last event occurred
Total Events Last Hour	The total number of events that occurred in the past hour
Total Events Last Day	The total number of events that occurred in the past twenty-four hours
Total Events in Database	The total number of events in the events database

Viewing System Statistics

The display includes statistics for the management center and its managed devices.

Before you begin

You must be an Admin or Maintenance user and be in the Global domain to view system statistics.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Statistics**.

Step 2 Choose a device from the **Select Device(s)** list, and click **Select Devices**.

Step 3 View available statistics.

Step 4 In the Disk Usage section, you can:

- Hover your pointer over a disk usage category in the **By Category** stacked bar to view (in order):
  - the percentage of available disk space used by that category
  - the actual storage space on the disk
  - the total disk space available for that category
- Click the arrow next to **By Partition** to expand it. If you have a malware storage pack installed, the `/var/storage` partition usage is displayed.

Step 5 (Optional) Click the arrow next to **Processes** to view the information described in Viewing System Statistics, on page 423.

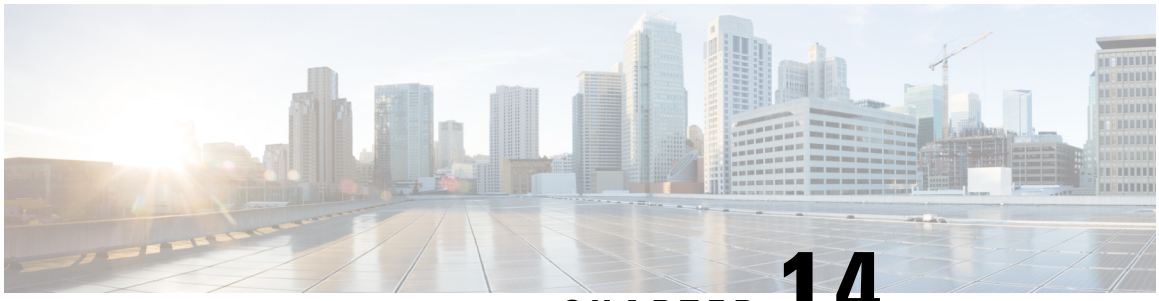

CHAPTER 14

Troubleshooting

The following topics describe ways to diagnose problems you may encounter:

Best Practices for Troubleshooting

- Before you make changes to try to fix a problem, generate a troubleshooting file to capture the original problem. See Health Monitor Reports for Troubleshooting, on page 438 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 426

- 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.

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.

To open the Message Center, click on the System Status icon, located next to the Deploy menu in the main menu. This icon can take one of the following forms, depending 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.

In addition to the Message Center, the web interface displays pop-up notifications in immediate response to your activities and ongoing system activities. Some pop-up notifications automatically disappear after five seconds, while others are "sticky," meaning they display until you explicitly dismiss them by clicking **Dismiss** ( ). Click the **Dismiss** link at the top of the notifications list to dismiss all notifications at once.

Tip Hovering your cursor over a non-sticky pop-up notification causes it to be sticky.

The system determines which messages it displays to users in pop-up notifications and the Message Center based on their licenses, domains, and access roles.

Message Types

The Message Center displays messages reporting system activities and status organized into three different tabs:

Deployments

This tab displays current status related to configuration deployment for each appliance in your system, grouped by domain. The system reports the following deployment status values on this tab. You can get additional detail about the deployment jobs by clicking **Show History**.

- Running (**Spinning**) — The configuration is in the process of deploying.
- **Success** — The configuration has successfully been deployed.
- **Warning** ( ) — Warning deployment statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Failure** — The configuration has failed to deploy; see Configuration Changes that Require Deployment. Failed deployments contribute to the message count displayed with the **Error System Status icon**.

Upgrades

This tab displays the current status related to software upgrade tasks for the managed devices. The system reports the following upgrade status values on this tab:

- **In progress**—Indicates that the upgrade task is in progress.
- **Completed**—Indicates that the software upgrade task is completed successful.
- **Failed**—Indicates that the software upgrade task has failed to complete.

Health

This tab displays current health status information for each appliance in your system, grouped by domain. Health status is generated by health modules as described in About Health Monitoring, on page 359. 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 a **Error icon**. Error statuses contribute to the message count displayed with the **Error System Status icon**.

You can click on links in the Health tab to view related detailed information on the Health Monitoring page. If there are no current health status conditions, the Health tab displays no messages.

Tasks

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. Cisco recommends you remove messages to reduce clutter in the Tasks tab as well as the message database.

Message Management

From the Message Center you can:

- 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 the display as well the database. When the number of messages in the database approaches 100,000, the system automatically deletes task status messages that you have removed.

View Basic System Information

The About page displays information about your appliance, including the model, serial number, and version information for various components of the 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

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 You have the following choices:

- Click **Deployments** to view messages related to configuration deployments. See View Deployment Messages, on page 429. 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 management center and the devices registered to it. See View Health Messages, on page 431. 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 431 or Manage Task Messages, on page 432. 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

You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.

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, the section is empty.
- **CLI Apply**—This section covers features that are configured using commands sent to the Lina process.
- **Infrastructure Messages**—This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands. Examining these errors can be particularly helpful if you are using FlexConfig policies to configure customized features. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note
There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that the 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 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.
```

View Upgrade Messages

You must be an Admin user or have the **Updates** permission to 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

You must be an Admin user or have the **Health** permission to view these messages.

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. The break-up of the severity, namely, warning, critical, and error is also displayed.
- Click on a status value to view only messages with that status.
- Hover your cursor over the relative time indicator for a message (for example, **3 day(s) ago**) to view the time of the most recent update for that message.
- To view detailed health status information for a particular message, click the message.
- To view complete health status on the Health Monitoring page, click **Health Monitor**.

Related Topics

About Health Monitoring, on page 359

View Task Messages

Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Tasks**.

Step 3 You have the following choices:

- Click **total** to view all current task statuses. To view the tasks based on the status, namely, waiting, running, retrying, success, and failures, click on them.
- Click a status value to view only messages for tasks with the that status.

Note
Messages for stopped tasks appear only in the total list of task status messages. You cannot filter on stopped tasks.

- Hover your cursor over the relative time indicator for a message (e.g., **3 day(s) ago**) to view the time of the most recent update for that message.
- Click any link within a message to view more information about the task.
- If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.

Manage Task Messages

Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click Tasks.

Step 3 You have the following choices:

- If more task status messages are available for display, click on **Fetch more messages** at the bottom of the message list to retrieve them.
- To remove a single message for a completed task (status stopped, success, or failure), click on **Remove** ( ) next to the message.
- To remove all messages for all tasks that have completed (status stopped, success, or failure), filter the messages on **total** and click on **Remove all completed tasks**.
- To remove all messages for all tasks that have completed successfully, filter the messages on **success**, and click on **Remove all successful tasks**.
- To remove all messages for all tasks that have failed, filter the messages on **failure**, and click on **Remove all failed tasks**.

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 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 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 45: Memory Usage Thresholds Based On Installed RAM

User-Input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
10%	10%	34%	72%	81%
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%

User-Input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
90%	90%	93%	97%	98%
100%	100%	100%	100%	100%

Disk Usage and Drain of Events Health Monitor Alerts

The Disk Usage health module compares disk usage on a managed device's hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds.

This topic describes the symptoms and troubleshooting guidelines for two health alerts generated by the Disk Usage health module:

- Frequent Drain of Events
- Drain of Unprocessed Events

The disk manager process manages the disk usage of a device. Each type of file monitored by the disk manager is assigned 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 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 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.

Important 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 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 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 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 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 721.

Communications Bottleneck – Sftunnel

Sftunnel is responsible for encrypted communications between the management center and the managed device. Events are sent over the tunnel to the management center. Connectivity issues and/or instability in the communication channel (sftunnel) between the managed device and the management center can be due to:

- Sftunnel is down or is unstable (flaps).

  Ensure that the 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 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 threat defense CLI using the **configure network management-interface** command.

Communications Bottleneck – SFDataCorrelator

The SFDataCorrelator manages data transmission between the management center and the managed device; on the 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 management center. This condition can happen when the managed device and the 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 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 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 436.

Disk Usage for Device Configuration History Files

The Disk Usage health module monitors the size of the device configuration history files on the 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 management center high-availability deployment, this health alert appears on the standby 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 management center. In a 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** > **Deployment History** > **Deployment Setting** > **Configuration Version Setting** and 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 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, see *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 that some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

Table 46: Selectable Troubleshoot Options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

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 Viewing the Device Health Monitor, on page 387.

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 management center web interface are stored on the 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.

Step 3 Choose **All Data** to generate all possible troubleshooting data, or check individual boxes as described in View Task Messages, on page 431.

Step 4 Click **Generate.**

Step 5 View task messages in the Message Center; see View Task Messages, on page 431.

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

You can download troubleshooting files.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 View the health monitor for the appliance; see , Viewing the Device Health Monitor, on page 387.

Step 2 Choose **System** (⚙) **> Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, and then click **Advanced Troubleshooting.**

Step 3 In **File Download**, enter the file name supplied by Support.

Step 4 Click **Download.**

Step 5 Follow your browser's prompts to download the file.

Note
For managed devices, the system renames the file by prepending the device name to the file name.

Step 6 Follow the directions from Support to send the troubleshooting files to Cisco.

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 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 440.

Troubleshoot a Connection

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 439.

Packet Capture Overview

The packet capture feature with trace option allows real packets that are captured on the ingress interface to be traced through the system. The trace information is displayed at a later stage. These packets are not dropped on the egress interface, as they are real data-path traffic. Packet capture for threat defense devices supports troubleshooting and analysis of data packets.

Once the packet is acquired, Snort detects the tracing flag that is enabled in the packet. Snort writes tracer elements, through which the packet traverses. Snort verdict as a result of capturing packets can be one of .the following:

Table 47: Snort Verdicts

Verdict	Description
Pass	Allow analyzed packet.
Block	Packet not forwarded.
Replace	Packet modified.
AllowFlow	Flow passed without inspection.
BlockFlow	Flow was blocked.
Ignore	Flow was blocked; occurs only for sessions with flows blocked on passive interfaces.

Verdict	Description
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 the following:

Table 48: Drop Reasons

Blocked or Flow Blocked by...	Cause
Snort	Snort is unable to process the packet, erg., snort can't decode packet since it is corrupted or has invalid format.
the App Id preprocessed	App Id module/preprocessed does not block packet by itself; but this may indicate that App Id detection causes other module (erg., firewall) to match a blocking rule.
the SSL preprocessed	There is a block/reset rule in SSL policy to match the traffic.
the firewall	There is a block/reset rule in firewall policy to match the traffic.
the captive portal preprocessed	There is a block/reset rule using the identity policy to match the traffic.
the safe search preprocessed	There is a block/reset rule using the safe-search feature in firewall policy to match the traffic.
the SI preprocessed	There is a block/reset rule a in Security Intelligence tab of AC Policy to block the traffic, erg., DNS or URL SI rule.
the filterer preprocessed	There is a block/reset rule in filterer tab of AC policy to match the traffic.
the stream preprocessed	There is an intrusion rule blocking/reset stream connection, erg., blocking when TCP normalization error.
the session preprocessed	This session was already blocked earlier by some other module, so session preprocessed is blocking further packets of the same session.

Blocked or Flow Blocked by...	Cause
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.
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 the 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 On the 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) (Optional) Check **SGT number** check box, and enter a Security Group Tag (SGT).

Step 7 Specify **Buffer** details:

a) (Optional) Enter a maximum **Packet Size**.
b) (Optional) Enter a minimum **Buffer Size**.
c) Select either **Continuous Capture** if you want the traffic captured without interruption, or **Stop when full** if you want the capture to stop when the maximum buffer size is reached.

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 128. 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 Overview

The Packet Tracer tool allows you to test policy configuration by modeling a packet with source and destination addresses, and protocol characteristics. The trace does a policy lookup to validate if the packet will be permitted or denied access based on the configured access rules, NAT, routing, access policies and rate-limiting policies. The packet flow is simulated based on interfaces, source address, destination address, ports, and protocols. This method of testing the packets allows you to verify the effectiveness of your policies and test whether the types of traffic you want to allow or deny are handled as required.

Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied access when they should be allowed. To simulate a packet fully, the packet tracer traces the data path—slow-path and fast-path modules. Initially, processing was transacted on per-session and per-packet basis. The Packet Tracer tool and Capture with Trace feature log the tracing data on per packet basis when the firewall processes packets per session or per packet.

PCAP File

You can initiate a packet tracer using a PCAP file, and that has a complete flow. Currently, only a PCAP with a single TCP/UDP-based flow and a maximum of 100 packets is supported. The packet tracer tool reads the PCAP file, and initializes the state for client and server replay entities. The tool starts replaying the packets in a synchronized manner by collecting and storing the trace output of each packet within the PCAP for subsequent processing and display.

PCAP Replay

Packet replay is executed by the sequence of the packet in the PCAP file, and interferences, if any, to the replay activity terminates it and concludes the replay. The trace output is generated for all the packets in the PCAP on the specified ingress interface and egress interface, thereby providing a complete context for flow evaluation.

PCAP replay is not supported for some features that dynamically modify the packet during replay, such as IPsec, VPN, SSL , HTTPs decryption, NAT, and so on.

Use the Packet Tracer

To use a packet tracer on Secure Firewall Threat Defense devices, you must be an Admin or Maintenance user.

Procedure

Step 1 On the management center, choose **Devices** > **Troubleshoot** > **Packet Tracer**.

Step 2 From the **Select Device** drop-down, choose the device on which 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 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 on 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 in the box to browse and upload the file. On selecting the file, the upload process starts automatically.

d) Go to this step.

Step 5 To define the trace parameters, from the **Protocol** drop-down menu, 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 SPI value for Source, 0-4294967295.
- **RAWIP**—Enter the port number, 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 has traveled through the system. Click on the individual packet to view the traces results for the packet. You can do the following:

- Copy (Copy) the trace results to clipboard.
- Expand or collapse (Expand or collapse) the displayed results.
- Maximize (Maximize) the trace result screen.

The time elapsed information that is useful to gauge the processing efforts are displayed for each phase. The total time that is taken for the entire flow of packets flowing from an ingress to an egress interface is also displayed in the results section.

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.

How to use the Threat Defense Diagnostic CLI from the Web Interface

You can execute the selected threat defense diagnostic CLI commands from the 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 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 management center high availability,diagnostic CLI is available only in the active 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

See the following table for feature-specific troubleshooting tips and techniques.

Table 49: 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	Troubleshooting LDAP Authentication Connections, on page 202
Licensing	Troubleshooting Smart Licensing, on page 285 Troubleshoot Specific License Reservation, on page 297

Feature	Relevant Troubleshooting Information
Management Center high availability	Troubleshooting Management Center High Availability, on page 305
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 Troubleshooting LDAP Authentication Connections, on page 202
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
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 440

PART IV

Tools

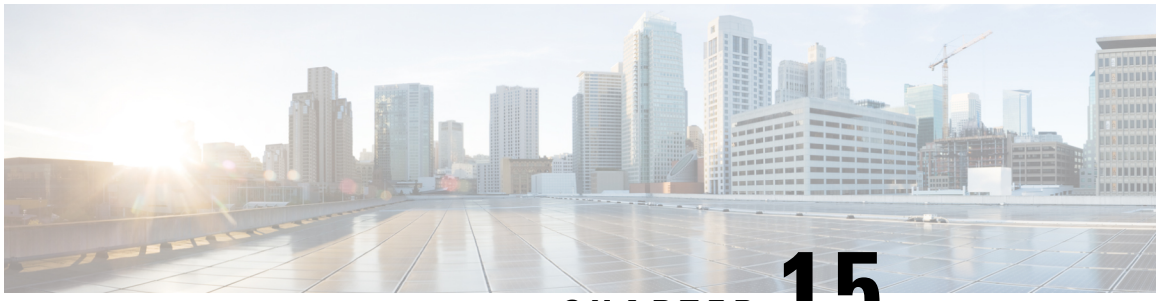

CHAPTER 15

Backup/Restore

About Backup and Restore

The ability to recover from a disaster is an essential part of any system maintenance plan. As part of your disaster recovery plan, we recommend that you perform periodic backups to a secure remote location.

What Is Backed Up?

Device backups are always configuration-only. Management center backups are as follows.

Table 50: 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 sertificate settings. • Public and private AMP cloud connections, in Version.
Events	All events in the management center database.	Intrusion event review status is not backed up. Restored intrusion events do not appear on Reviewed Events pages.

Backup Type	Backed Up	Not Backed Up
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.	

What Is Restored?

Restoring configurations overwrites *all* backed-up configurations, with very few exceptions. On the 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 the following:

- You cannot restore what is not backed up, as decribed above.
- Restoring fails VPN certificates.

  The threat defense restore process removes VPN certificates and all VPN configurations from threat defense devices, including certificates added after the backup was taken. After you restore the threat defense device, you must re-add/re-enroll all VPN certificates, and redeploy the device.
- Restoring to a configured management center — instead of factory-fresh or reimaged — merges intrusion events and file lists.

  The 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 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.

On-Demand Backups

You can perform on-demand backups for the management center and many threat defense devices from the management center.

For more information, see Backing Up Management Centers or Managed Devices, on page 462.

Scheduled Backups

You can use the scheduler on management center to automate backups. You can also schedule remote device backups from the management center.

The 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 Scheduled Backups, on page 489.

Storing Backup Files

You can store backups locally. However, we recommend you back up 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 management center to manage them.

For more information, see Remote Storage Device, on page 97 and Manage Backups and Remote Storage, on page 481.

Restoring from Backup

You restore the management center from the Backup Management page. You must use the threat defense CLI to restore 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 Management Centers and Managed Devices, on page 467.

Requirements for Backup and Restore

Backup and restore have the following requirements.

Platform Requirements: Backup

This table lists backup support by platform. Device backup is supported for both application and container instances.

Table 51: Backup Support by Platform

Platform	Backup Supported?		
	Standalone	**High Availability**	**Clusters**
Management center, hardware and virtual	**YES**	**YES**	—
Threat defense hardware	**YES**	**YES**	**YES**
Threat defense virtual, on-prem/private cloud	VMware HyperFlex Nutanix OpenStack	VMware	VMware
Threat defense virtual, public cloud	—	—	—

Platform Requirements: Restore

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 management centers, you can use backup and restore not only in an RMA scenario, but also to migrate configurations and events between management centers. For details, including supported target and destination models, see the Cisco Secure Firewall Management Center Model Migration Guide.

Version Requirements

As the first step in any backup, note the patch level. To restore a backup, the old and the new appliance must be running the same software version, including patches. To restore to a Firepower 4100/9300 chassis, you must be running a compatible FXOS.

For 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

To:

- Back up or restore the management center: Global only.
- Back up a device from the management center: Global only.
- 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 has the following guidelines and limitations.

Backup and Restore is for Disaster Recovery/RMA

Backup and restore is primarily intended for RMA scenarios. Before you begin the restore process of a faulty or failed physical appliance, contact Cisco TAC for replacement hardware.

You can also use backup and restore to migrate configurations and events between management centers. This makes it easier to replace 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 management center on a freshly reimaged management center. If you are restoring without reimaging and if you had registered threat defense on the management center after backing up, then when you register the device again on the restored 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 management center backup that is restored to.

Backup and Restore is not Configuration Import/Export

A backup file contains information that uniquely identifies an appliance, and cannot be shared. Do not use the backup and restore process to copy configurations between appliances or devices, or as a way to save configurations while testing new ones. Instead, use the import/export feature.

For example, threat defense device backups include the device's management IP address and all information the device needs to connect to its managing management center. Do not restore the threat defense backup to a device being managed by a different management center; the restored device will attempt to connect to the management center specified in the backup.

Restore is Individual and Local

You restore to management centers and managed devices individually and locally. This means:

- You cannot batch-restore to high availability or clustered management centers or devices.
- You cannot use the management center to restore a device. For the management center, you can use the web interface to restore. For threat defense devices, you must use the threat defense CLI, except for the ISA 3000 zero-touch restore, which uses an SD card and the reset button.
- While restoring the 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 management center user accounts to log into and restore one of its managed devices. The management center and devices maintain their own user accounts.

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.

Guidelines and Restrictions

- 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 Firepower 4100/9300 chassis, all existing configuration on the Firepower 4100/9300 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 Firepower 4100/9300 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 Firepower 4100/9300 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

Backup and restore has the following best practices.

When to Back Up

We recommend backing up during a maintenance window or other time of low use.

While the system collects backup data, there may be a temporary pause in data correlation (management center only), and you may be prevented from changing configurations related to the backup. If you include event data, event-related features such as eStreamer are not available.

You should back up in the following situations:

- Regular scheduled or on-demand backups.

  As part of your disaster recovery plan, we recommend that you perform periodic backups.

  The 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 Scheduled Backups, on page 489.

- After SLR changes.

  Back up the management center after you make changes to Specific Licensing Reservations (SLRs). If you make changes and then restore an older backup, you will have issues with your Specific Licensing return code and can accrue orphan entitlements.

- Before upgrade or reimage.

  If an upgrade fails catastrophically, you may have to reimage and restore. Reimaging returns most settings to factory defaults, including the system password. If you have a recent backup, you can return to normal operations more quickly.

Note Restoring from a backup does not reset the password that you had configured after the reimage or an RMA.

- After upgrade.

  Back up after you upgrade, so you have a snapshot of your freshly upgraded deployment. We recommend you back up the management center *after* you upgrade its managed devices, so your new management center backup file 'knows' that its devices have been upgraded.

Maintaining Backup File Security

Backups are stored as unencrypted archive (.tar) files.

Private keys in PKI objects—which represent the public key certificates and paired private keys required to support your deployment—are decrypted before they are backed up. The keys are reencrypted with a randomly generated key when you restore the backup.

Note We recommend you back up management centers and devices to a secure remote location and verify transfer success. Backups left locally may be deleted, either manually or by the upgrade process, which purges locally stored backups.

Especially because backup files are unencrypted, do *not* allow unauthorized access. If backup files are modified, the restore process will fail. Keep in mind that anyone with the Admin/Maint role can access the Backup Management page, where they can move and delete files from remote storage.

In the management center's system configuration, you can mount an NFS, SMB, or SSHFS network volume as remote storage. After you do this, all subsequent backups are copied to that volume, but you can still use the management center to manage them. For more information, see Remote Storage Device, on page 97 and Manage Backups and Remote Storage, on page 481.

Note that only the management center mounts the network volume. Managed device backup files are routed through the management center. Make sure you have the bandwidth to perform a large data transfer between the management center and its devices. For more information, see Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).

Backup and Restore in Management Center High Availability Deployments

In management center high availability deployments, backing up one management center does not back up the other. You should regularly back up both peers. Do not restore one HA peer with the backup file from the other. A backup file contains information that uniquely identifies an appliance, and cannot be shared.

Note that you can replace an HA management center without a successful backup. For more information on replacing HA management centers, both with and without successful backups, see Replacing Management Centers in a High Availability Pair, on page 320.

Backup and Restore in Threat Defense High Availability Deployments

In the threat defense high availability deployment, you should:

- Back up the device pair from the management center, but restore individually and locally from the threat defense CLI.

  The backup process produces unique backup files each peer. Do not restore one peer with the backup file from the other. A backup file contains information that uniquely identifies an appliance, and cannot be shared.

  The device's role is also noted in its backup file name. When you restore, make sure you choose the appropriate backup file: primary vs secondary.

- Do *not* suspend or break high availability before you restore.

  Maintaining the high availability configuration ensures replacement devices can easily reconnect after restore.

- Do *not* run the **restore** CLI command on both peers at the same time.

  Assuming you have successful backups, you can replace either or both peers. Any physical replacement tasks you can perform simultaneously: unracking, reracking, and so on. However, do *not* run the **restore** command on the second device until the restore process completes for the first device, including the reboot.

- When both peers fail, before the devices are decommissioned, ensure that unregister them both from the management center.

Note that you can replace a high availability device without a successful backup.

Backup and Restore in Threat Defense Clustering Deployments

In the threat defense clustering deployment, you should:

- Back up the entire cluster from the management center, but restore nodes individually and locally from the threat defense CLI.

  The backup process produces a bundled tar file that includes unique backup files for each cluster node. Do not restore one node with the backup file from another. A backup file contains information that uniquely identifies a device, and cannot be shared.

  The node's role is noted in its backup file name. When you restore, make sure you choose the appropriate backup file: control or data.

  You cannot back up individual nodes. If a data node fails to back up, the management center will still back up all other nodes. If the control node fails to back up, the backup is canceled.

- Do *not* suspend or break clustering before you restore.

  Maintaining the cluster configuration ensures replacement devices can easily reconnect after restore.

- Do *not* run the **restore** CLI command on multiple nodes at the same time. We recommend that you restore the control node first and wait until it rejoins the cluster before you restore any data nodes.

  Assuming you have successful backups, you can replace multiple nodes in the cluster. Any physical replacement tasks you can perform simultaneously: unracking, reracking, and so on. However, do *not* run the **restore** command on an additional node until the restore process completes for the previous node, including the reboot.

Backup and Restore for Firepower 4100/9300 Chassis

To restore 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. For additional best practices, see Configuration Import/Export Guidelines for Firepower 4100/9300 , on page 457.

Before Backup

Before you back up, you should:

- Update the VDB and SRU on the management center.

  We always recommend you use the latest vulnerability database (VDB) and intrusion rules (SRU). Before you back up the management center, check the Cisco Support & Download site for newer versions.

- Check disk space.

Before you begin a backup, make sure you have enough disk space on the appliance or on your remote storage server. The space available is displayed on the Backup Management page.

Backups can fail if there is not enough space. Especially if you schedule backups, make sure you regularly prune backup files or allocate more disk space to the remote storage location.

Before Restore

Before restore, you should:

- Revert licensing changes.

  Revert any licensing changes made since you took the backup.

  Otherwise, you may have license conflicts or orphan entitlements after the restore. However, do *not* unregister from Cisco Smart Software Manager (CSSM). If you unregister from CSSM, you must unregister again after you restore, then re-register.

  After the restore completes, reconfigure licensing. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.

- Disconnect faulty appliances.

  Disconnect the management interface, and for devices, the data interfaces.

  Restoring threat defense devices sets the management IP address of the replacement device to the management IP address of the old device. To avoid IP conflicts, disconnect the old device from the management network before you restore the backup on its replacement.

  Note that restoring the management center does *not* change the management IP address. You must set that manually on the replacement — just make sure you disconnect the old appliance from the network before you do.

- Do *not* unregister managed devices.

  Whether you are restoring the management center or managed device, do not unregister devices from the management center, even if you physically disconnect an appliance from the network.

  If you unregister, you will need to redo some device configurations, such as security zone to interface mappings. After you restore, the management center and devices should begin communicating normally.

- Reimage.

  In an RMA scenario, the replacement appliance will arrive configured with factory defaults. However, if the replacement appliance is already configured, we recommend you reimage. Reimaging returns most settings to factory defaults, including the system password. You can only reimage to major versions, so you may need to patch after you reimage.

  If you do not reimage, keep in mind that management center intrusion events and file lists are merged rather than overwritten.

After Restore

After restore, you should:

- Reconfigure anything that was not restored.

  This can include reconfiguring licensing, remote storage, and audit log server certificate settings. You also must re-add/re-enroll failed threat defense VPN certificates.

- Update the VDB and SRU on the management center.

  We always recommend you use the latest vulnerability database (VDB) and intrusion rules (SRU). This is especially important for the VDB, because the VDB in the backup will overwrite the VDB on the replacement management center.

- Deploy.

  Whether you are restoring the management center or device, you must deploy. For a restored device, you may need to force deploy: see *Redeploy Existing Configurations to a Device* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Backing Up Management Centers or Managed Devices

You can perform on-demand or scheduled backups for supported appliances.

You do not need a backup profile to back up devices from the management center. However, management center backups require backup profiles. The on-demand backup process allows you to create a new backup profile.

Back up the Management Center

Use this procedure to perform an on-demand 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

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.

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 management centers, if you choose to back up configuration only on an active management center, by default, both the active and standby management centers are backed up into a single unified backup file. For information of unified backup of management center in high availability, see Unified Backup of Management Centers in High Availability, on page 326.
- **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 About Backup and Restore, on page 453.

Step 4 Note the **Storage Location** for 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 481.

Step 5 (Optional) Enable **Copy when complete** to copy completed 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 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 management center to connect to a mail server: Configuring a Mail Relay Host and Notification Address, on page 62.

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 management center may write temporary files to the remote server. These files are cleaned up at the end of the backup process.

What to do next

If you configured remote storage or enabled **Copy when complete**, verify transfer success of the backup file.

Back up a Device from the Management Center

Use this procedure to perform an on-demand device 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

If you are backing up a Firepower 4100/9300 chassis, it is especially important that you also back up FXOS configurations: Exporting an FXOS Configuration File, on page 465.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**, then click **Managed Device Backup**.

Step 2 Select one or more **Managed Devices**.

For clustering, choose the cluster. You cannot perform backups on individual nodes.

Step 3 Note the **Storage Location** for device backup files.

This will either be local storage in `/var/sf/remote-backup/`, or a remote network volume. For 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 481.

Step 4 If you did not configure remote storage, choose whether you want to save the backup to the local storage in the management center or to the device using the **Retrieve to Management Center** check box.

- Enabled (default): Saves the backup to the management center in `/var/sf/remote-backup/`.

  For clusters, this option is always checked. The individual node backup files are copied to the management center and then bundled into a single compressed tar file before it is copied to any remote storage.

- Disabled: Saves the backup to the device in `/var/sf/backup`.

Step 5 Click **Start Backup** to start the on-demand backup.

Step 6 Monitor progress in the Message Center.

What to do next

If you configured remote storage, verify if the transfer of the backup file was successful.

Exporting an FXOS Configuration File

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 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 the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.

f) 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, chassis manager will display an error stating that `property pwd of org-root/cfg-exp-policy-default is out of range`.

g) In the **Location** field, enter the full path to where you want the configuration file exported including the filename.
h) Click **Export**.
The configuration file is created and exported to the specified location.

Create a Backup Profile

A backup profile is a saved set of preferences—what to back up, where to store the backup file, and so on.

Management Center backups require backup profiles. Backup profiles are not required to back up a device from the management center.

When you perform an on-demand 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.

The following procedure explains how to create a backup profile without performing an on-demand backup.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**, then click **Backup Profiles**.

Step 2 Click **Create Profile** and enter a **Name**.

Step 3 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 About Backup and Restore, on page 453.

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 481.

Step 5 (Optional) Enable **Copy when complete** to copy completed 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 management center to connect to a mail server: Configuring a Mail Relay Host and Notification Address, on page 62.

Step 7 Click **Save**.

Restoring Management Centers and Managed Devices

For the management center, you use the web interface to restore from backup. For threat defense devices, you must use the threat defense CLI. You cannot use the management center to restore a device.

The following sections explain how to restore management centers and managed devices.

Restore Management Center from Backup

When you restore 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 configurations overwrites *all* configurations, with very few exceptions. It also reboots the 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 management center from backup. For more information on backup and restore in management center HA deployments, see Replacing Management Centers in a High Availability Pair, on page 320.

Before you begin

You must read and understand the requirements, guidelines, limitations, and best practices. You do not want to skip any steps or ignore security concerns. Careful planning and preparation can help you avoid missteps.

- Requirements for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

Procedure

Step 1 Log into the 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 481.

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 management center reboots.

What to do next

- If necessary, reconfigure any licensing settings that you reverted before the restore. If you notice licensing conflicts or orphan entitlements, contact Cisco TAC.
- If necessary, reconfigure remote storage and audit log server certificate settings. These settings are not included in backups.
- 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 configuration changes; see the Cisco Secure Firewall Management Center Device Configuration Guide.

Restore Threat Defense from Backup: Firepower 1000/2100, Secure Firewall 3100, ISA 3000 (Non-Zero-Touch)

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 1000/2100, Secure Firewall 3100, or ISA 3000 threat defense device, standalone or in a High Availability pair or as a cluster. It assumes you have access to a successful backup of the device or devices you are replacing; see Back up a Device from the Management Center, on page 464. For zero-touch restore on the ISA 3000 using an SD card, see Zero-Touch Restore Threat Defense from Backup: ISA 3000, on page 471.

For high availability and clustered 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456

- Best Practices for Backup and Restore, on page 458

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 management center in `/var/sf/remote-backup`.
- In a remote storage location.

For threat defense high availability and clustered 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. For clusters, control and data node backup files are bundled together in a single compressed file. You must extract the files, which also indicate the device role.

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 481.

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 threat defense high availability deployments, this includes the failover link. For clustering, this includes the cluster control 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 threat defense high availability deployments, connect the failover link. For clustering, connect the cluster control 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 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 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 management center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing threat defense patch should have the same version. The threat defense CLI does not have an upgrade command. To patch:

a) From the 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 management center.

   If you do not unregister, you will have a ghost device registered to the 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`. For clustered devices, extract the appropriate backup file from the backup bundle.

Step 9 From the threat defense CLI, restore the backup.

Access the 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 threat defense high availability and clustering deployments, make sure you choose the appropriate backup file: primary vs secondary, or control vs. data. 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 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 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 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
  You do not require to manually execute the command for Threat Defense version 7.2.10 because the threat defense high availability automatically resumes after restoring from backup.

- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from 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.

Step 13 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Zero-Touch Restore Threat Defense from Backup: ISA 3000

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 an ISA 3000 threat defense device, 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 Management Center, on page 464.

For high availability and clustered 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

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 threat defense HA deployments, this includes the failover link.

Note
Do *not* unregister from the management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 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 (device manager) to install the patch.

The following procedure assumes you have a factory default configuration. If you already configured the device, you can log into 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 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 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.

Step 7 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 device manager to install a patch, you can reboot from the **Device** > **System Settings** > **Reboot/Shutdown** page. From the 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 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.

Step 8 Log into the management center and wait for the replacement device to connect.

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. Contact Cisco TAC.
- Resume high availability synchronization. From the 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
  You do not require to manually execute the command for Threat Defense version 7.2.10 because the threat defense high availability automatically resumes after restoring from backup.

- Re-add/re-enroll all VPN certificates. The restore process removes VPN certificates from 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.

Step 11 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Restore 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 or as a cluster. It assumes you have access to successful backups of:

- The logical device or devices you are replacing; see Back up a Device from the Management Center, on page 464.
- FXOS configurations; see Exporting an FXOS Configuration File, on page 465.

For high availability and clustered 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

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 management center in `/var/sf/remote-backup`.
- In a remote storage location.

For threat defense high availability and clustered 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. For clusters, control and data node backup files are bundled together in a single compressed file. You must extract the files, which also indicate the device role.

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 481.

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 threat defense high availability deployments, this includes the failover link. For clustering, this includes the cluster control 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 threat defense high availability deployments, connect the failover link. For clustering, connect the cluster control 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 chassis manager to import your backed-up FXOS configurations: Importing a Configuration File, on page 477.

Step 8 Use 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 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 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 management center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing threat defense patch should have the same version. The threat defense CLI does not have an upgrade command. To patch:

a) From the 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 management center.

   If you do not unregister, you will have a ghost device registered to the 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`. For clustered devices, extract the appropriate backup file from the backup bundle.

Step 11 From the threat defense CLI, restore the backup.

Access the 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 threat defense high availability and clustering deployments, make sure you choose the appropriate backup file: primary vs secondary, or control vs. data. 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 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 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 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.

Step 15 Connect the device's data interfaces.

See the hardware installation guide for your model: http://www.cisco.com/go/ftd-quick.

What to do next

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 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 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.
d) 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.

e) Enter the username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.
f) 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, chassis manager will display an error stating that `property pwd of org-root/cfg-exp-policy-default is out of range.`

g) In the **File Path** field, enter the full path to the configuration file including the file name.
h) 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.
i) 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 Threat Defense Virtual from Backup

Use this procedure to replace a faulty or failed threat defense virtual device.

For high availability and clustered 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 management center, even when disconnecting a device from the network. For threat defense high availability and clustered devices, do *not* suspend or break high availability or clustering . 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 for Backup and Restore, on page 455
- Guidelines and Limitations for Backup and Restore, on page 456
- Best Practices for Backup and Restore, on page 458

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 management center in `/var/sf/remote-backup`.
- In a remote storage location.

For threat defense high availability and clustered 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. For clusters, control and data node backup files are bundled together in a single compressed file. You must extract the files, which also indicate the device role.

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 481.

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 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 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 management center. The replacement device should be unmanaged from the physical network and the new hardware as well as the replacing threat defense patch should have the same version. The threat defense CLI does not have an upgrade command. To patch:

a) From the 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 management center.

If you do not unregister, you will have a ghost device registered to the 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`. For clustered devices, extract the appropriate backup file from the backup bundle.

Step 7 From the threat defense CLI, restore the backup.

Access the 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 threat defense high availability and clustering deployments, make sure you choose the appropriate backup file: primary vs secondary, or control vs. data. 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 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 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 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.

Step 11 Add and configure data interfaces.

See the getting started guide: https://www.cisco.com/go/ftdv-quick.

What to do next

Verify that the restore succeeded and the replacement device is passing traffic as expected.

Manage Backups and Remote Storage

Backups are stored as unencrypted archive (.tar) files. The file name includes identifying information that can include:

- The name of the backup profile or scheduled task associated with the backup.
- The display name or IP address of the backed-up appliance.
- The appliance's role, such as a member of an HA pair.

We recommend you back up appliances to a secure remote location and verify transfer success. Backups left on an appliance may be deleted, either manually or by the upgrade process; upgrades purge locally stored backups. For more information on your options, see Backup Storage Locations, on page 483.

Caution Especially because backup files are unencrypted, do *not* allow unauthorized access. If backup files are modified, the restore process will fail. Keep in mind that anyone with the Admin/Maint role can access the Backup Management page, where they can move and delete files from remote storage.

The following procedure describes how to manage backup files.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Backup/Restore**.

The Backup Management page lists available backups. It also lists how much disk space you have available to store backups. Backups can fail if there is not enough space.

Step 2 Do one of the following:

Table 52: Remote Storage and Backup File Management

To	Do This
Enable or disable remote storage for backups without having to edit the 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 Management Center, on page 464).
Move a file between the 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 management center, where it is stored on the management center depends on the kind of backup: • 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.

To	Do This
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 management center. Store your downloaded backup in a secure location.

Backup Storage Locations

The following table describes backup storage options for management centers and managed devices.

Table 53: 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 Management Center, on page 464). In the management center's system configuration, you can mount an NFS, SMB, or SSHFS network volume as remote storage for management center and device backups; see Remote Storage Device, on page 97.) After you do this, all subsequent management center backups *and management center-initiated device backups* are copied to that volume, but you can still use the management center to manage them (restore, download, upload, delete, move). Note that only the management center mounts the network volume. Managed device backup files are routed through the management center. Make sure you have the bandwidth to perform a large data transfer between the management center and its devices. For more information, see Guidelines for Downloading Data from the Firepower Management Center to Managed Devices (Troubleshooting TechNote).

Location	Details
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 Management Center, on page 464). For the 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 management center system configuration, do *not* copy backup files to the same directory using **Copy when complete**.
Local, on the management center.	If you do not configure remote storage by mounting a network volume, you can save backup files on the management center: • management center backups are saved to `/var/sf/backup`. • Device backups are saved to `/var/sf/remote-backup` on the 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

Table 54: History for Backup and Restore

Feature	Minimum Management Center	Minimum 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.
Single backup file for high availability management centers.	7.4.1 7.2.6	Any	When performing a configuration-only backup of the active management center in a high availability pair, the system now creates a single backup file which you can use to restore either unit. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Back up and restore device clusters.	7.3.0	Any	You can now use the management center to back up device clusters, except in the public cloud. To restore, use the device CLI. New/modified screens: **System** (⚙) > **Tools** > **Backup/Restore** > **Managed Device Backup** New/modified commands: **restore remote-manager-backup**
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.
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 for Backup and Restore, on page 455. New/modified screens: **System > Tools > Backup/Restore > Managed Device Backup** New/modified FTD CLI commands: **restore**

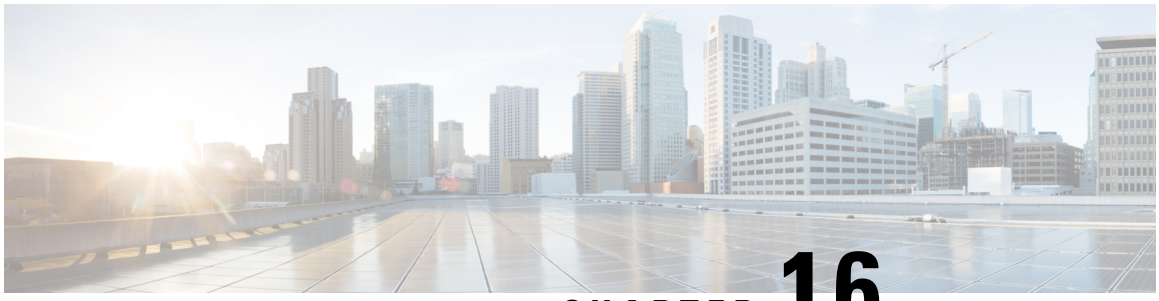

CHAPTER 16

Scheduling

The following topics explain how to schedule tasks:

About Task Scheduling

You can schedule various tasks to run at designated times, either once or on a recurring basis.

Tasks are scheduled in UTC on the back end, which means when they occur locally depends on the date and your specific location. Also, because tasks are scheduled in UTC, they do not adjust for Daylight Saving Time, summer time, or any such seasonal adjustments that you may observe in your location. If you are affected, scheduled tasks occur one hour "later" in the summer than in the winter, according to local time.

Some tasks are automatically scheduled or performed by the initial setup process:

- A one-time task to download and install the latest VDB.
- A weekly scheduled task to download the latest available software updates and VDB.
- A weekly scheduled task to perform a locally stored configuration-only backup of the management center.

You should review the weekly tasks and adjust if necessary. Optionally, schedule new recurring tasks to actually update the VDB and/or software, and deploy configurations.

Important We *strongly* recommend you review scheduled tasks to be sure they occur when you intend. Some tasks (such as those involving automated software updates or that require pushing updates to managed devices) may place a significant load on networks with low bandwidths. You should schedule tasks like these to run during periods of low network use. Other tasks, such as deploying configurations, can cause traffic interruptions. You should schedule tasks like these during maintenance windows.

Requirements and Prerequisites for Task Scheduling

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Maintenance User

Configuring a Recurring Task

You set the frequency for a recurring task using the same process for all types of tasks.

Note that the time displayed on most pages on the web interface is the local time, which is determined by using the time zone you specify in your local configuration. Further, the management center automatically adjusts its local time display for daylight saving time (DST), where appropriate. However, recurring tasks that span the transition dates from DST to standard time and back do not adjust for the transition. That is, if you create a task scheduled for 2:00 AM during standard time, it will run at 3:00 AM during DST. Similarly, if you create a task scheduled for 2:00 AM during DST, it will run at 1:00 AM during standard time.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** drop-down list, select the type of task that you want to schedule.

Step 4 Click **Recurring** next to the **Schedule task to run** option.

Step 5 In the **Start On** field, specify the date when you want to start your recurring task.

Step 6 In the **Repeat Every** field, specify how often you want the task to recur.

You can either type a number or click **Up** (▲) and **Down** (▼) to specify the interval. For example, type `2` and click **Days** to run the task every two days.

Step 7 In the **Run At** field, specify the time when you want to start your recurring task.

Step 8 For a task to be run on a weekly or monthly basis, select the days when you want to run the task in the **Repeat On** field.

Step 9 Give the job a name.

Step 10 Select the remaining options for the type of task you are creating:

- Backup - Schedule backup jobs as described in Schedule Management Center Backups, on page 489 .

- Download CRL - Schedule certificate revocation list downloads as described in Configuring Certificate Revocation List Downloads, on page 491.
- Deploy Policies - Schedule policy deployment as described in Automating Policy Deployment, on page 492.
- Nmap Scan - Schedule Nmap scans as described in Scheduling an Nmap Scan, on page 493.
- Report - Schedule report generation as described in Automating Report Generation, on page 494.
- Cisco Recommended Rules - Schedule automatic updates as described in Automating Cisco Recommendations, on page 496.
- Download Latest Update - Schedule software or VDB update downloads as described in Automating Software Downloads, on page 497 or Automating VDB Update Downloads, on page 500.
- Install Latest Update - Schedule installation of software or VDB updates on a management center or managed device as described in Automating Software Installs, on page 499 or Automating VDB Update Installs, on page 500.
- Push Latest Update - Schedule push of software updates to managed devices as described in Automating Software Pushes, on page 498.
- Update URL Filtering Database - Schedule automatic update of URL filtering data as described in Automating URL Filtering Updates Using a Scheduled Task, on page 501.

Step 11 Click **Save**.

Scheduled Backups

You can use the scheduler on a Secure Firewall Management Center to automate its own backups. You can also schedule remote device backups from the management center. For more information on backups, see Backup/Restore, on page 453.

Note that not all devices support remote backups.

Schedule Management Center Backups

You can use the scheduler on the management center to automate both management center and device backups. Note that not all devices support remote backups. For more information, see Backup/Restore, on page 453.

Note As part of the initial configuration, the system schedules weekly configuration-only management center backups (locally stored). We recommend you review this task and make changes if necessary, as described in this topic.

Before you begin

Create a backup profile that specifies your backup preferences. See Create a Backup Profile, on page 466.

You must be in the global domain to perform this task.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Scheduling**.

Step 2 From the **Job Type** list, select **Backup**.

Step 3 Specify whether you want to back up **Once** or **Recurring**.

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488.

Step 4 Enter a **Job Name**.

Step 5 For the **Backup Type**, click **Management Center**.

Step 6 Choose a **Backup Profile**.

Step 7 (Optional) Enter a **Comment**.

Keep comments brief. They will appear in the Task Details section of the schedule calendar page.

Step 8 (Optional) Enter an email address, or a comma-separated list of email addresses, in the **Email Status To:** field.

For information on setting up an email relay server to send task status messages, see Configuring a Mail Relay Host and Notification Address, on page 62.

Step 9 Click **Save**.

Schedule Remote Device Backups

You can use the scheduler on the management center to automate both management center and device backups. Note that not all devices support remote backups. For more information, see Backup/Restore, on page 453.

You must be in the global domain to perform this task.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Scheduling**.

Step 2 From the **Job Type** list, select **Backup**.

Step 3 Specify whether you want to back up **Once** or **Recurring**.

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488.

Step 4 Enter a **Job Name**.

Step 5 For the **Backup Type**, click **Device**.

Step 6 Select one or more devices.

If your device is not listed, it does not support remote backup.

Step 7 If you did not configure remote storage for backups, choose whether you want to **Retrieve to Management Center**.

- Enabled (default): Saves the backup to the management center in `/var/sf/remote-backup/`.
- Disabled: Saves the backup to the device in `/var/sf/backup/`.

If you configured remote backup storage, backup files are saved remotely and this option has no effect. For more information, see Manage Backups and Remote Storage, on page 481.

Step 8 (Optional) Enter a **Comment**.

Keep comments brief. They will appear in the Task Details section of the schedule calendar page.

Step 9 (Optional) Enter an email address, or a comma-separated list of email addresses, in the **Email Status To:** field.

For information on setting up an email relay server to send task status messages, see Configuring a Mail Relay Host and Notification Address, on page 62.

Step 10 Click **Save**.

Configuring Certificate Revocation List Downloads

You must perform this procedure using the local web interface for the management center.

The system automatically creates the Download CRL task when you enable downloading a certificate revocation list (CRL) in the local configuration on an appliance where you enable user certificates or audit log certificates for the appliance. You can use the scheduler to edit the task to set the frequency of the update.

Before you begin

- Enable and configure user certificates or audit log certificates and set one or more CRL download URLs. See Requiring Valid HTTPS Client Certificates, on page 70 and Require Valid Audit Log Server Certificates, on page 53 for more information.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Download CRL**.

Step 4 Specify how you want to schedule the CRL download, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 7 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured on the management center to send status messages.

Step 8 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Automating Policy Deployment

After modifying configuration settings in the management center, you must deploy those changes to the affected devices.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort Restart Traffic Behavior and Configurations that Restart the Snort Process When Deployed or Activated.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Deploy Policies**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Device** field, select a device where you want to deploy policies.

Step 7 Select or deselect the **Skip deployment for up-to-date devices** check box, as required.

By default, the **Skip deployment for up-to-date devices** option is enabled to improve performance during the policy deployment process.

Note
The system does not perform a scheduled policy deployment task if a policy deployment initiated from the management center web interface is in progress. Correspondingly, the system does not permit you to initiate a policy deployment from the web interface if a scheduled policy deployment task is in-progress.

Step 8 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page. The comment field must contain between 3 and 255 characters and cannot include any special characters other than hyphen (-), colon (:), and parentheses ().

Step 9 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 10 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62
Configuration Changes that Require Deployment

Nmap Scan Automation

You can schedule regular Nmap scans of targets on your network. Automated scans allow you to refresh information previously supplied by an Nmap scan. Because the system cannot update Nmap-supplied data, you need to rescan periodically to keep that data up to date. You can also schedule scans to automatically test for unidentified applications or servers on hosts in your network.

Note that a Discovery Administrator can also use an Nmap scan as a remediation. For example, when an operating system conflict occurs on a host, that conflict may trigger an Nmap scan. Running the scan obtains updated operating system information for the host, which resolves the conflict.

If you have not used the Nmap scanning capability before, you configure Nmap scanning before defining a scheduled scan.

Related Topics

Nmap Scanning

Scheduling an Nmap Scan

After Nmap replaces a host's operating system, applications, or servers detected by the system with the results from an Nmap scan, the system no longer updates the information replaced by Nmap for the host. Nmap-supplied service and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, you may want to set up regularly scheduled scans to keep Nmap-supplied operating systems, applications, or servers up to date. If the host is deleted from the network map and re-added, any Nmap scan results are discarded and the system resumes monitoring of all operating system and service data for the host.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, select **Nmap Scan**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.

- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Nmap Remediation** field, select an Nmap remediation.

Step 7 In the **Nmap Target** field, select the scan target.

Step 8 In the **Domain** field, select the domain whose network map you want to augment.

Step 9 If you want to comment on the task, type a comment in the **Comment** field.

Tip
The comment field appears in the Task Details section of the calendar schedule page; keep comments brief.

Step 10 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 11 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62
Nmap Scanning

Automating Report Generation

You can automate reports so that they run at regular intervals.

Before you begin

- For reports other than risk reports: Create a report template. See Report Templates, on page 526 for more information.
- If you want to distribute email reports using the scheduler, configure a mail relay host and specify report recipients and message information. See Configuring a Mail Relay Host and Notification Address, on page 62 and (for reports other than risk reports) Distributing Reports by Email at Generation Time, on page 546 or (for risk reports) Generating, Viewing, and Printing Risk Reports, on page 524.
- (Optional) Set or change the file name, output format, time window, or email distribution settings of the scheduled report. See Specify Report Generation Settings for a Scheduled Report, on page 495.
- If you will choose PDF as the report output format, look at the report template and verify that the number of results in each section of the template does not exceed the limit for PDFs. For information, see Report Template Fields, on page 526.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select a job.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 In the **Report Template** field, select a risk report or report template.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Tasks Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Note
Configuring this option does **not** distribute the reports.

Step 9 If you do not want to receive report email attachments when reports have no data (for example, when no events of a certain type occurred during the report period), select the **If report is empty, still attach to email** check box.

Step 10 Click **Save**.

Specify Report Generation Settings for a Scheduled Report

You must have Admin or Security Analyst privileges to perform this task.

To specify or change the file name, output format, time window, or email distribution settings of a scheduled report:

Procedure

Step 1 Select **Overview > Reporting > Report Templates**.

Step 2 Click **Edit** for the report template to change.

Step 3 If you will select PDF output:

a) Look to see whether any of the sections in the report shows a yellow triangle beside the number of results.
b) If you see any yellow triangles, mouse over the triangle to view the maximum number of results allowable for that section for PDF output.
c) For each section with a yellow triangle, reduce the number of results to a number below the limit.
d) When there are no more yellow triangles, click **Save**.

Step 4 Click **Generate**.

Note
If you want to change report generation settings without generating the report now, you must click **Generate** from the template configuration page. Changes will not be saved if you click **Generate** from the template list view unless you generate the report.

Step 5 Modify settings.

Step 6 To save the new settings without generating the report, click **Cancel**.

To save the new settings and generate the report, click **Generate** and skip the rest of the steps in this procedure.

Step 7 Click **Save**.

Step 8 If you see a prompt to save even though you haven't made changes, click **OK**.

Automating Cisco Recommendations

You can automatically generate rule state recommendations based on network discovery data for your network using the most recently saved configuration settings in a custom intrusion policy.

Note If the system automatically generates scheduled recommendations for an intrusion policy with unsaved changes, you must discard your changes in that policy and commit the policy if you want the policy to reflect the automatically generated recommendations.

When the task runs, the system automatically generates recommended rule states, and modifies the states of intrusion rules based on the configuration of your policy. Modified rule states take effect the next time you deploy your intrusion policy.

Before you begin

- Configure Cisco recommended rules in an intrusion policy as described in the Cisco Secure Firewall Management Center Device Configuration Guide.
- If you want to email task status messages, configure a valid email relay server.
- You must have the IPS Smart Licenseor Protection Classic License to generate recommendations.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From **Job Type**, choose **Cisco Recommended Rules**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Enter a name in the **Job Name** field.

Step 6 Next to **Policies**, choose one or more intrusion policies where you want to generate recommendations. Check **All Policies** check box to choose all intrusion policies.

Step 7 (Optional) Enter a comment in the **Comment** field.

Keep comments brief. Comments appear in the Task Details section of the schedule calendar page.

Step 8 (Optional) To email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field.

Step 9 Click **Save**.

Related Topics

Conflicts and Changes: Network Analysis and Intrusion Policies
About Cisco Recommended Rules
Configuring a Mail Relay Host and Notification Address, on page 62

Software Upgrade Automation

You can automatically download and apply maintenance releases and patches.

To upgrade the management center, schedule Download and Install tasks. To upgrade managed devices, schedule Download, Push, and Install tasks. Make sure you leave adequate time between the tasks; for example, installations scheduled to occur while a push is still running will fail.

This feature is not supported for major releases. Internet access is required to download upgrade packages. When scheduling upgrades to device groups, the upgrade will run on all grouped devices simultaneously.

Note As part of the initial configuration, the system schedules weekly downloads. We recommend you review this task and make changes if necessary, as described in Automating Software Downloads, on page 497. This task only downloads the updates. It is your responsibility to install any updates this task downloads.

Related Topics

Management Interfaces, on page 74
Updates, on page 223

Automating Software Downloads

Use this procedure to schedule the download of select patches and maintenance releases. You must be in the global domain.

Before you begin

Make sure the management center can access the internet.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Download Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 Next to **Update Items**, check **Software** check box.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Automating Software Pushes

If you want to automate the installation of software updates on managed devices, you must push the updates to the devices before installing.

When you create the task to push software updates to managed devices, make sure you allow enough time between the push task and a scheduled install task for the updates to be copied to the device.

You must be in the global domain to perform this task.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Push Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the device that you want to update.

Step 7 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Automating Software Installs

Make sure you allow enough time between the task that pushes the update to a managed device and the task that installs the update.

You must be in the global domain to perform this task.

Caution Depending on the update being installed, the appliance may reboot after the software is installed.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Install Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the appliance (including the management center) where you want to install the update.

Step 7 Next to **Update Items**, check the **Software** check box.

Step 8 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page; keep comments brief.

Step 9 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 10 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Vulnerability Database Update Automation

You can use the scheduling feature to update the Cisco vulnerability database (VDB), thereby ensuring that you are using the most up-to-date information to evaluate the hosts on your network. You must schedule the download, install, and subsequent deploy as separate tasks, allowing enough time between tasks.

Note The initial setup on the management center automatically downloads and installs the latest VDB from Cisco as a one-time operation. It also schedules a weekly task to download the latest available software updates, which includes the latest VDB. We recommend you review this weekly task and adjust if necessary. Optionally, schedule a new weekly task to actually update the VDB and deploy configurations.

Related Topics

Management Interfaces, on page 74

Automating VDB Update Downloads

You must be in the global domain to perform this task.

Before you begin

Make sure the management center has internet access.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Download Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 Next to **Update Items**, check the **Vulnerability Database** check box.

Step 7 (Optional) Type a brief comment in the **Comment** field.

Step 8 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 9 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Automating VDB Update Installs

Allow enough time between the task that downloads the VDB update and the task that installs the update.

You must be in the global domain to perform this task.

Caution In most cases, the first deploy after a VDB update restarts the Snort process, interrupting traffic inspection. The system warns you when this will happen (updated application detectors and operating system fingerprints require a restart; vulnerability information does not). Whether traffic drops or passes without further inspection during this interruption depends on how the targeted device handles traffic. For more information, see Snort Restart Traffic Behavior.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Install Latest Update**.

Step 4 Specify how you want to schedule the task, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488 for details.

Step 5 Type a name in the **Job Name** field.

Step 6 From the **Device** drop-down list, select the management center.

Step 7 Next to **Update Items**, check the **Vulnerability Database** check box.

Step 8 (Optional) Type a brief comment in the **Comment** field.

Step 9 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 10 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Automating URL Filtering Updates Using a Scheduled Task

In order to ensure that threat data for URL filtering is current, the system must obtain data updates from the Cisco Collective Security Intelligence (CSI) cloud.

By default, when you enable URL filtering, automatic updates are enabled. However, if you need to control when these updates occur, use the procedure described in this topic instead of the default update mechanism.

Although daily updates tend to be small, if it has been more than five days since your last update, new URL filtering data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

Before you begin

- Make sure the management center has internet access: Security, Internet Access, and Communication Ports, on page 1027.

- Make sure you have a URL Filtering license and enable URL filtering. For more information, see the *Enable URL Filtering Using Category and Reputation* in the Cisco Secure Firewall Management Center Device Configuration Guide.
- Verify that **Enable Automatic Updates** is not selected on the **Cloud Services** under the **Integration** > **Other Integrations** menu.
- You must be in the Global domain to perform this task.

Procedure

Step 1 Select **System** (⚙) > **Tools** > **Scheduling**.

Step 2 Click **Add Task**.

Step 3 From the **Job Type** list, select **Update URL Filtering Database**.

Step 4 Specify how you want to schedule the update, **Once** or **Recurring**:

- For one-time tasks, use the drop-down lists to specify the start date and time.
- For recurring tasks, see Configuring a Recurring Task, on page 488.

Step 5 Type a name in the **Job Name** field.

Step 6 If you want to comment on the task, type a comment in the **Comment** field.

The comment field appears in the Task Details section of the schedule calendar page. Keep comments brief.

Step 7 If you want to email task status messages, type an email address (or multiple email addresses separated by commas) in the **Email Status To:** field. You must have a valid email relay server configured to send status messages.

Step 8 Click **Save**.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Scheduled Task Review

After adding scheduled tasks, you can view them and evaluate their status. The View Options section of the page allows you to view scheduled tasks using a calendar and a list of scheduled tasks.

The Calendar view option allows you to view which scheduled tasks occur on which day.

The Task List shows a list of tasks along with their status. The task list appears below the calendar when you open the calendar. In addition, you can view it by selecting a date or task from the calendar.

You can edit a scheduled task that you previously created. This feature is especially useful if you want to test a scheduled task once to make sure that the parameters are correct. Later, after the task completes successfully, you can change it to a recurring task.

There are two types of deletions you can perform from the Schedule View page. You can delete a specific one-time task that has not yet run or you can delete every instance of a recurring task. If you delete an instance

of a recurring task, all instances of the task are deleted. If you delete a task that is scheduled to run once, only that task is deleted.

Task List Details

Table 55: Task List Columns

Column	Description
Name	Displays the name of the scheduled task and the comment associated with it.
Type	Displays the type of scheduled task.
Start Time	Displays the scheduled start date and time.
Frequency	Displays how often the task is run.
Last Run Time	Displays the actual start date and time. For a recurring task, this applies to the most recent execution.
Last Run Status	Describes the current status for a scheduled task: • A **Check mark**( ) indicates that the task ran successfully. • A question mark icon (**Question Mark** ( )) indicates that the task is in an unknown state. • An exclamation mark icon ( ) indicates that the task failed. For a recurring task, this applies to the most recent execution.
Next Run Time	Displays the next execution time for a recurring task. Displays N/A for a one-time task.
Creator	Displays the name of the user that created the scheduled task.
Edit	Edits the scheduled task.
Delete	Deletes the scheduled task.

Viewing Scheduled Tasks on the Calendar

You can view a scheduled task on the calendar.

Procedure

Step 1 Select **System** ( ) > **Tools** > **Scheduling**.

Step 2 You can perform the following tasks using the calendar view:

- Click **Double Left Arrow** ( ) to move back one year.
- Click **Single Left Arrow** ( ) to move back one month.
- Click **Single Right Arrow** ( ) to move forward one month.
- Click **Double Right Arrow** ( ) to move forward one year.
- Click **Today** to return to the current month and year.
- Click **Add Task** to schedule a new task.
- Click a date to view all scheduled tasks for the specific date in a task list table below the calendar.
- Click a specific task on a date to view the task in a task list table below the calendar.

Editing Scheduled Tasks

You can edit scheduled tasks.

Procedure

Step 1 Select **System** ( ) > **Tools** > **Scheduling**.
Step 2 On the calendar, click either the task that you want to edit or the day on which the task appears.
Step 3 In the **Task Details** table, click **Edit** ( ) next to the task you want to edit.
Step 4 Edit the task.
Step 5 Click **Save**.

Deleting Scheduled Tasks

You can delete a scheduled task.

Procedure

Step 1 Select **System** ( ) > **Tools** > **Scheduling**.
Step 2 In the calendar, click the task you want to delete. For a recurring task, click an instance of the task.
Step 3 In the **Task Details** table, click **Delete** ( ), then confirm your choice.

History for Scheduled Tasks

Feature	Minimum Management Center	Minimum Threat Defense	Details
Scheduled tasks download patches and VDB updates only.	7.2.6	Any	**Upgrade impact. Scheduled download tasks stop retrieving maintenance releases.** The **Download Latest Update** scheduled task no longer downloads maintenance releases; now it only downloads the latest applicable patches and VDB updates. To direct-download maintenance (and major) releases to the management center, use **System**(⚙) > **Product Upgrades**. Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.
Automatic VDB downloads.	7.3.0	Any	Initial setup schedules a weekly task to download the latest available software updates, which now includes the latest VDB. We recommend you review this weekly task and adjust if necessary, as well as schedule a new weekly task to actually update the VDB. You must deploy configurations for new application detectors and operating system fingerprints to take effect. New/modified screens: The **Vulnerability Database** check box is now enabled by default in the system-created **Weekly Software Download** scheduled task.
Automatic intrusion rule updates.	6.6	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 configuration backups.	6.5	Any	Initial setup schedules weekly tasks to: • Download the latest available software updates for the FMC 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**: management center vs device. Platform restrictions: Device must support on-demand backup; see Requirements for Backup and Restore, on page 455.

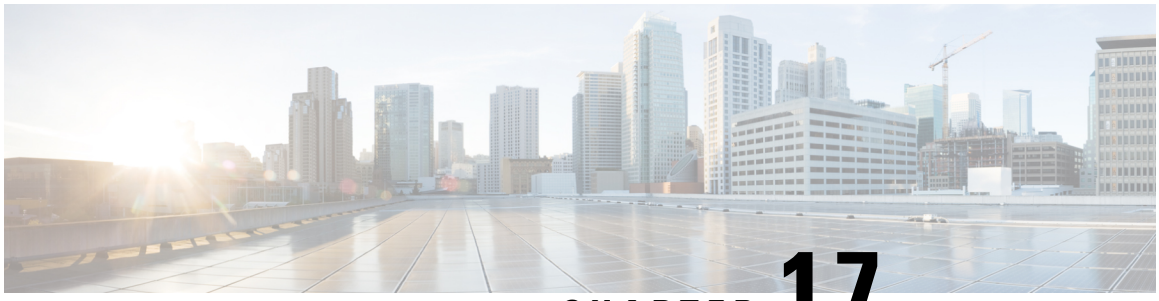

CHAPTER 17

Import/Export

The following topics explain how to use the Import/Export feature:

About Configuration Import/Export

You can use the Import/Export feature to copy configurations between appliances. Import/Export is not a backup tool, but can simplify the process of adding new appliances to your deployment.

You can export a single configuration, or you can export a set of configurations (of the same type or of different types) with a single action. When you later import the package onto another appliance, you can choose which configurations in the package to import.

An exported package contains revision information for that configuration, which determines whether you can import that configuration onto another appliance. When the appliances are compatible but the package includes a duplicate configuration, the system offers resolution options.

Note The importing and exporting appliances must be running the same software version. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match. If the versions do not match, the import fails. You cannot use the Import/Export feature to update intrusion rules. Instead, download and apply the latest rule update version.

Configurations that Support Import/Export

Import/Export is supported for the following configurations:

- Access control policies and the policies they invoke: prefilter, network analysis, intrusion, SSL, file, Threat Defense Service Policy
- Intrusion policies, independently of access control
- NAT policies (Secure Firewall Threat Defense only)

- 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

Special Considerations for Configuration Import/Export

When you export a configuration, the system also exports other required configurations. For example, exporting an access control policy also exports any subpolicies it invokes, objects and object groups it uses, ancestor policies, and so on. As another example, if you export a platform settings policy with external authentication enabled, the authentication object is exported as well. There are some exceptions, however:

- System-provided databases and feeds—The system does not export URL filtering category and reputation data, Cisco Intelligence Feed data, or the geolocation database (GeoDB). Make sure all the appliances in your deployment obtain up-to-date information from Cisco.
- Global Security Intelligence lists—The system exports Global Security Intelligence Block and Do Not Block lists associated with exported configurations. 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 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 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 management center, for values not set in the export package. Therefore, an imported intrusion policy may behave differently than expected if the importing management center has differently configured default variables.

- Custom user objects—If you have created custom user groups or objects in your management center and if such a custom user object is a part of any rule in your access control policy, note that the export file (.sfo) does not carry the user object information and therefore while importing such a policy, any reference to such custom user objects will be removed and will not be imported to the destination management center. To avoid detection issues due to the missing user group, add the customized user objects manually to the new management center and re-configure the access control policy after import.

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 system 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 management center, the system appends autogenerated numbers to the imported object and group names to make them unique.
- You must map any security zones and interface groups used in the imported configurations to matching-type zones and groups managed by the importing management center.
- If you export a configuration that uses PKI objects containing private keys, the system decrypts the private keys before export. On import, the system encrypts the keys with a randomly generated key.

Requirements and Prerequisites for Configuration Import/Export

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Exporting Configurations

Depending on the number of configurations being exported and the number of objects those configurations reference, the export process may take several minutes.

Tip Many list pages include an **YouTube EDU** () next to list items. Where this icon is present, you can use it as a quick alternative to the export procedure that follows.

Before you begin

- Confirm that the importing and exporting appliances are running the same software version. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Import/Export**.

Step 2 Click **Collapse** ( ) and **Expand** ( ) to collapse and expand the list of available configurations.

Step 3 Check the configurations you want to export and click **Export**.

Step 4 Follow your web browser's prompts to save the exported package to your computer.

Importing Configurations

Depending on the number of configurations being imported and the number of objects those configurations reference, the import process may take several minutes.

Note If you log out of the system, 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.

Before you begin

- Confirm that the importing and exporting appliances are running the same software version. For access control and its subpolicies (including intrusion policies), the intrusion rule update version must also match.

Procedure

Step 1 On the importing appliance, choose **System** (⚙) > **Tools** > **Import/Export**.

Step 2 Click **Upload Package**.

Step 3 Enter the path to the exported package or browse to its location, then click **Upload**.

Step 4 If there are no version mismatches or other issues, choose the configurations you want to import, then click **Import**.
If you do not need to perform any conflict resolution or interface object mapping, the import completes and a success message appears. Skip the rest of this procedure.

Step 5 If prompted, on the Import Conflict Resolution page, map interface objects used in the imported configurations to zones and groups with matching interface types managed by the importing management center.

Interface object type (security zone or interface group) and interface type (passive, inline, routed, and so on) of source and destination objects must match. For information, see Interface.

If the configurations you are importing reference security zones or interface groups that do not already exist, you can map them to existing interface objects, or create new ones.

Note
For individual access control policies, you have the option of replacing an existing policy with the imported ones. However, for nested access control policies, you can only import them as new policies.

Step 6 Click **Import**.

Step 7 If prompted, on the Import Resolution page, expand each configuration and choose the appropriate option as described in Import Conflict Resolution, on page 511.

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 decryption policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms**, then click **Download Now**( ) (**Download Now**).

- Optionally, view a report summarizing the imported configurations; see View Task Messages, on page 431.

Import Conflict Resolution

When you attempt to import a configuration, the system determines whether a configuration of the same name and type already exists on the appliance. When an import includes a duplicate configuration, the system offers resolution options suitable to your deployment from among the following:

- **Keep existing**

  The system does not import that configuration.

- **Replace existing**

  The system overwrites the current configuration with the configuration selected for import.

- **Keep newest**

  The system imports the selected configuration only if its timestamp is more recent than the timestamp on the current configuration on the appliance.

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 decryption policies, access control policies, and possibly other policies. (**Integration** > **Other Integrations** > **Realms**, then click **Download Now**( ) (**Download Now**).

- **Import as new**

  The system imports the selected duplicate configuration, appending a system-generated number to the name to make it unique. (You can change this name before completing the import process.) The original configuration on the appliance remains unchanged.

The resolution options the system offers depends on whether your deployment uses domains, and whether the imported configuration is a duplicate of a configuration defined in the current domain, or a configuration defined in an ancestor or descendant of the current domain. The following table lists when the system does or does not present a resolution option.

Resolution Option	Secure Firewall Management Center		Managed Device
	Duplicate in current domain	Duplicate in ancestor or descendant domain	
Keep existing	Yes	Yes	Yes
Replace existing	Yes	No	Yes
Keep newest	Yes	No	Yes
Import as new	Yes	Yes	Yes

When you import an access control policy with a file policy that uses clean or custom detection file lists and a file list presents a duplicate name conflict, the system offers conflict resolution options as described in the table above, but the action the system performs on the policies and file lists varies as described in the table below:

Resolution Option	System Action	
	Access control policy and its associated file policy are imported as new and the file lists are merged	Existing access control policy and its associated file policy and file lists remain unchanged
Keep existing	No	Yes
Replace existing	Yes	No

Resolution Option	System Action	
	Access control policy and its associated file policy are imported as new and the file lists are merged	**Existing access control policy and its associated file policy and file lists remain unchanged**
Import as new	Yes	No
Keep newest and access control policy being imported is the newest	Yes	No
Keep newest and existing access control policy is the newest	No	Yes

If you modify an imported configuration on an appliance, and later re-import that configuration to the same appliance, you must choose which version of the configuration to keep.

CHAPTER 18

Data Purge and Storage

Data Stored on the Management Center

For	See
General information about data storage on the management center	The Disk Usage Widget, on page 348
Purging old data	Purging Data from the Management Center Database, on page 516
Allowing external access to the data on the management center (this is an advanced feature)	External Database Access, on page 62
Backups	Manage Backups and Remote Storage, on page 481 and subtopics
Reports	Configure Local Storage, on page 97
Events	Connection Logging, on page 713 Database, on page 58 and subtopics
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

For	See
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

Purging Data from the Management Center Database

You can use the database purge page to purge discovery, identity, connection, and security-related connection data files from the management center databases. Note that when you purge a database, the appropriate process is restarted.

Caution Purging a database removes the data you specify from the 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. To perform this action, you must be in the global domain.

Procedure

Step 1 Choose **System** (⚙) > **Tools** > **Data Purge**.

Step 2 Under **Discovery and Identity**, perform any or all of the following:

- Check the **Network Discovery Events** check box to remove all network discovery events from the database.
- Check the **Hosts** check box to remove all hosts and Host Indications of Compromise flags from the database.
- Check the **User Activity** check box to remove all user activity events from the database.
- Check the **User Identities** check box to remove all user login and user history data from the database, as well as User Indications of Compromise flags.

Step 3 Under **Connections**, perform any or all of the following:

- Check the **Connection Events** check box to remove all connection data from the database.
- Check the **Connection Summary Events** check box to remove all connection summary data from the database.

- Check the **Security-Related Connection Events** check box to remove all security-related connection data from the database.

Note
Checking the **Connection Events** check box does not remove Security Intelligence events. Connections with Security Intelligence data will still appear in the Security Intelligence event page (available under the Analysis > Connections menu). Correspondingly, checking the **Security-Related Connection Events** check box does not remove connection events with associated security-related connection data.

Step 4 Click **Purge Selected Events**.
The items are purged and the appropriate processes are restarted.

External Data Storage

You can optionally use remote data storage for store certain types of data.

For	See
Backups	Manage Backups and Remote Storage, on page 481 and subtopics Remote Storage Device, on page 97 and subtopics
Reports	Remote Storage Device, on page 97 and subtopics Moving Reports to Remote Storage, on page 548
Events	Information about syslog and other resources in Event Analysis Using External Tools, on page 617 Remote Data Storage in Cisco Secure Cloud Analytics, on page 518 Remote Data Storage on a Secure Network Analytics Appliance, on page 519 If you store connection events remotely, consider disabling storage of connection events on your management center. For information, see Database, on page 58 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

Similar but different options for storing event data externally to your 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 519.	For more information, see the links in Remote Data Storage in Cisco Secure Cloud Analytics, on page 518.

Remote Data Storage in Cisco Secure Cloud Analytics

Send select Secure Firewall event data to Secure Cloud Analytics using Security Analytics and Logging (SaaS). Supported events: Connection, Security Intelligence, intrusion, file, and malware.

For details, see the *Secure Firewall Management Center and Cisco Security Analytics and Logging (SaaS) Integration Guide*.

You can send events either directly or via syslog.

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

If you require more data storage than your Secure Firewall appliance can provide, you can use Security Analytics and Logging (On Premises) to store Secure Firewall data on a Secure Network Analytics appliance. For complete information, see the documentation available from Cisco Security Analytics and Logging.

You can view connection events in management center even if they are stored on a Secure Network Analytics appliance. See Work in Secure Firewall Management Center with Connection Events Stored on a Secure Network Analytics Appliance, on page 665.

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.

History for Data Storage

Feature	Minimum Management Center	Minimum 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 management center because you are storing them on a remote volume, those events do not count towards the flow rate limits for your 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 59. 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 management center, and to cross-launch from 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 519. New/Modified pages: The **System > Logging > Security Analytics & Logging** page now displays the wizard instead of the configuration for creating cross-launch options.

Feature	Minimum Management Center	Minimum Threat Defense	Details
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 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 519.
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

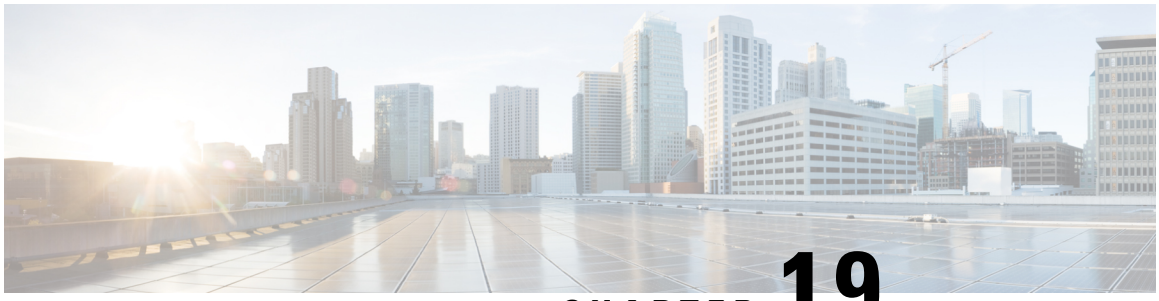

CHAPTER 19

Reports

The following topics describe how to work with reports:

Requirements and Prerequisites for Reports

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Maintenance User (risk reports only)
- Security Analyst

Introduction to Reports

The system offers two types of reports:

- Risk Reports, on page 524 — High-level summaries of risks found on your network.
- Standard Reports, on page 525 — Detailed, customizable reports about all aspects of your system.

Risk Reports

Risk reports are portable, high-level, easy-to-interpret summaries of risks found in your organization. You can use these reports to share information about areas of risk, and recommendations for addressing these risks, with people who do not have access to your system and who may not be network security experts. These reports are intended to facilitate discussion about areas for investment in the security of your network.

Risk Report Templates

- **Advanced Malware Risk Report**
- **Attacks Risk Report**. The following 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**

Generating, Viewing, and Printing Risk Reports

Templates for standard reports do not apply to risk reports.

Reports pertain to the current domain.

Each risk report generates as an HTML file.

To schedule risk report generation, see Automating Report Generation, on page 494.

Before you begin

- Make sure your system is configured to detect the risks that you want to summarize.
- If you want to email the report and you have not yet configured a Relay Host, you can do so now. For information, see Configuring a Mail Relay Host and Notification Address, on page 62.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Generate Report** for the desired report.

Step 4 Enter information.

- Information that you enter in the Input Parameters section will appear on the title page of the report. You can leave these fields blank.

Step 5 Click **Generate**.

Step 6 Click **OK**.

What to do next

- To view, download, move, or delete a risk report, see About Working with Generated Reports, on page 546.
- You can print to PDF any risk report from most supported browsers. For best results, enable background colors, images, and optionally headers and footers, in the print or print preview settings of your browser. Supported page sizes are A4 and US letter.

Standard Reports

The 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 management center. You can also design your own custom reports from scratch.

A report is a document file formatted in PDF, HTML, or CSV with the content you want to communicate. A report template specifies the data searches and formats for the report and its sections. The system includes a powerful report designer that automates the design of report templates. You can replicate the content of any event view table or dashboard graphic displayed in the web interface.

You can build as many report templates as you need. Each report template defines the individual sections in the report and specifies the database search that creates the report's content, as well as the presentation format (table, chart, detail view, and so on) and the time frame. Your template also specifies document attributes, such as the cover page and table of contents and whether the document pages have headers and footers (available only for reports in PDF format). You can export a report template in a single configuration package file and import it for reuse on another management center.

You can include input parameters in a template to expand its usefulness. Input parameters allow you to produce tailored variations of the same report. When you generate a report with input parameters, the generation process prompts you to enter a value for each input parameter. The values you type constrain the report contents on a one-time basis. For example, you can place an input parameter in the destination IP field of the search that produces an intrusion event report; at report generation time, you can specify a department's network segment when prompted for the destination IP address. The generated report then contains only information concerning that particular department.

About Designing Reports

Report Templates

You use report templates to define the content and format of the data in each of the report's sections, as well as the document attributes of the report file (cover page, table of contents, and page headers and footers). After you generate a report, the template stays available for reuse until you delete it.

Your reports contain one or more information sections. You choose the format (text, table, or chart) for each section individually. The format you select for a section may constrain the data that can be included. For example, you cannot show time-based information in certain tables using a pie chart format. You can change the data criteria or format of a section at any time to obtain optimum presentation.

You can base a report's initial design on a predefined event view, or you can start your design by importing content from any defined dashboard, workflow, or summary. You can also start with an empty template, adding sections and defining their attributes one by one.

Note In a multidomain deployment, you can view but not edit report templates belonging to ancestor domains. To generate reports from these templates, you must copy them to your current domain.

Report Template Fields

The following table describes the fields you can use to build a section in your report template. Not all fields are used in all types of sections; after you choose the section format, the system displays the appropriate fields.

Field Name	Section Types	Definition
Format	n/a	Choose the format of the section data: **Bar chart** (): Compares quantities of the selected variables. **Line chart** (): Shows trends/changes over time of a selected variable. Available only for time-based tables. **Pie chart** (): Shows each selected variable as a percentage of the whole. Variables with quantities of zero are dropped from the chart. Very small quantities are clustered into a category labeled **Other**. **Table view** (): Shows values of attributes for each record. Not available for summary or statistical data. **Detail view** (): Shows complex object data associated with certain events, such as packets (for intrusion events) and host profiles (for host events). This format is available only for certain event types that involve such objects. Output may degrade performance if large numbers are requested.
Table	All	Choose the table from which the section data is extracted.
Preset	All	Predefined searches. Select an appropriate preset to initialize the search criteria when you define a new search.

Field Name	Section Types	Definition
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.
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 management center if available. If data on the 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 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.

Field Name	Section Types	Definition
Color	Bar chart Line chart	Colors for graphed data in the section.

Report Template Creation

A report template is a framework of sections, each independently built from its own database query.

You can build a new report template by creating a new template, using an existing template, basing a template off an event view, or importing a dashboard or workflow.

If you do not want to copy an existing report template, you can create an entirely new template. The first step in creating a template is to generate the framework that allows you to add and format the sections. Then, in the order you prefer, you design the individual template sections and set attributes for the report document.

Each template section consists of a dataset generated by a search or filter, and has a format specification (table, pie chart, and so on) that determines the mode of presentation. You further determine section content by selecting the fields in the data records you want to include in the output, as well as the time frame and number of records to show.

Note Use the section preview utility to check the column selection and output characteristics such as pie chart colors. It is not a reliable indicator of the correctness of your configured search.

The report you generate from the template has several document attributes that span all sections and control features, such as the cover page, headers and footers, page numbering, and so on.

Note that if you selected CSV as your document format, you have no document attributes to set.

If you identify a good model among your existing templates, you can copy the template and edit its attributes to create a new report template. Cisco also provides a set of predefined report templates, visible on the **Reports Tab** in the list of templates.

From an event view, you can create a report template and modify it to meet your needs. You can add additional sections, modify automatically included sections, and delete sections.

You can quickly create a new report by importing dashboards, workflows, and statistics summaries. The import creates a section for each widget graphic in your dashboard and each event view in your workflow. You can delete any unnecessary sections to focus on the most important information.

Creating a Custom Report Template

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Create Report Template**.

Step 4 Enter a name for your new template in the **Report Title** field.

Step 5 To add an input parameter to the report title, place your cursor in the title where the parameter value should appear, then click insert **Input Parameter** ( ).

Step 6 Use the set of add under the Report Sections title bar to insert sections as necessary.

Step 7 Configure section content as described in Report Template Configuration, on page 531.

Tip
You can click **Preview** at the bottom of the section window to view the column layout or graphic format you chose.

Step 8 Click **Advanced** to set attributes for PDF and HTML reports as described in Document Attributes in a Report Template, on page 539.

Step 9 Click **Save**.

If you see an error, look for a yellow triangle beside the results value in each section. If you see any such triangles, do one of the following:

- For each field that displays a yellow triangle, mouse over the triangle and reduce the number of results to the number indicated.
- Click **Generate** and include an output format other than PDF.

Creating a Report Template from an Existing Template

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Copy** ( ) next to the report template you want to copy.

Step 4 In the **Report Title** field, enter a name.

Step 5 Make changes to the template as needed.

Step 6 Click **Save**.

Creating a Report Template from an Event View

Procedure

Step 1 Populate an event view with the events you want in the report:

- Use an event search to define the events you want to view.
- Drill down through a workflow until you have the appropriate events in your event view.

Step 2 From the event view page, click **Report Designer**.

The Report Sections page displays a section for each view in the captured workflow.

Step 3 Optionally, enter a new name in the **Report Title** field and click **Save**.

Step 4 You can:

- Add a cover page, table of contents, starting page number, or header and footer text — Click **Advanced** settings.
- Add page breaks — Click **Add Page Break** ( ), and drag the new page break object from the template bottom to the front of the section that should start the new page.
- Add text sections — Click **Add Text Section** ( ), and drag the new text section from the template bottom to the place where you want it to appear in the report template.
- Change the title of a section — Click the section title in the title bar, enter the section title, and click **OK**.
- Configure the report sections — Adjust the field settings in each section.

  Tip
  To view the current column layout or chart formatting for a section, click the section's **Preview** link.

- Exclude template sections from the report — Click **Delete** ( ) in the section's title bar, and confirm the deletion.

  Note
  The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the management center.

Step 5 Click **Save**.

Creating a Report Template by Importing a Dashboard or Workflow

Procedure

Step 1 Identify the dashboard, workflow, or summary you want to replicate in your report.

Step 2 Choose **Overview** > **Reporting**.

Step 3 Click **Report Templates**.

Step 4 Click **Create Report Template**.

Step 5 Enter a name for your new report template in the **Report Title** field.

Step 6 Click **Save**.

Step 7 Click **Import Section** ( ). You can choose any of the data sources described in Data Source Options on Import Report Sections, on page 531.

Step 8 Choose a dashboard, workflow, or summary from the drop-down menus.

Step 9 For the data sources you want to add, click **Import**.

For dashboards, each widget graphic will have its own section; for workflows, each event view will have its own section.

Step 10 Make changes to the content of your sections as needed.

Note

The last report section in some workflows contains detail views that show packets, host profiles, or vulnerabilities, depending on the workflow. Retrieving large numbers of events with these detail views when generating your report may affect performance of the management center.

Step 11 Click **Save**.

Data Source Options on Import Report Sections

Table 56: Data Source Options on Import Report Sections Window

Select this option...	To import...
Import Dashboard	any custom analysis widget on the selected dashboard.
Import Workflow	any predefined or custom workflow. Selections have the format: `Table - Workflow name` For example, `Connection Events - Traffic by Port` imports the views in the Traffic by Port workflow generated from the Connection Events table.
Import Summary Sections	any of the following generic summaries: • Intrusion Detailed Summary • Intrusion Short Summary • Discovery Detailed Summary • Discovery Short Summary

Report Template Configuration

You can modify and customize a report template after you create it. You can modify a variety of report section attributes to adjust the content of the section and its data presentation.

Each section in a report template queries a database table to generate content for that section. Changing the section's data format uses the same data query, but modifies the fields that appear in the section according to the analytical purpose of the format type. For example, the table view of intrusion events populates the section with a large number of data fields per event record, while a pie chart section shows the portion of all matching records that each selected attribute represents, with no details about individual events. Bar chart sections compare the total counts of matching records that have specific attributes. Line charts summarize changes in the matching records over time with respect to a single attribute. Line charts are available only for data that is time-based, not for information about hosts, users, third-party vulnerabilities, and so on.

The search or filter in a report section specifies the database query on which the section content is based. For most tables, you can constrain a report using a predefined or saved search, or you can create a new search on the fly:

- Predefined searches serve as examples for searching certain event tables and can provide quick access to important information about your network that you may want to include in reports.

- Saved event searches include all public event searches that you or others have created, plus all your saved private event searches.
- Saved searches for the current report template are accessible only in the report template itself. The search names of saved report template searches end with the string "Custom Search." Users create these searches while designing reports.

For the Application Statistics table, you use a user-defined application filter to constrain a report.

If you include table data in a section, you can choose which fields in the data record to show. All fields in the table are available for inclusion or exclusion. You select fields that accomplish the purpose of the report, then order and sort them accordingly.

You can add text sections to your templates to provide custom text, such as an introduction, for the whole report or for individual sections.

You can add page breaks before or after any section in the template. This feature is particularly helpful for multi-section reports with text pages that introduce the various sections.

A report template's time window defines the template's reporting period.

Note Security Analysts can edit only report templates they created. In multidomain deployments, you cannot edit report templates from ancestor domains, but you can copy them to create descendant versions.

Set the Table and Data Format for a Report Template Section

Procedure

Step 1 Click **Overview** > **Reporting** > **Report Templates** > **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 Fields, on page 526

Specify the Search or Filter for a Report Template Section

Procedure

Step 1 In the report template section, choose the database table to query from the **Table** drop-down menu:

- For most tables, the **Search** drop-down list appears.
- For the Application Statistics table, the **Filter** drop-down list appears.

Step 2 Choose the search or filter you want to use to constrain the report.

You can view the search criteria or create a new search by clicking **Edit** ( ).

Modify Fields in the Report Template Table Format Sections

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 can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images.

Tip Text sections are useful for introductions to your report or your report sections.

Procedure

Step 1 In the report template editor, click **Add Text Section** ( ).

Step 2 Drag the new text section to its intended position in the report template.

Step 3 If you want to position the text section first or last on a page, add page breaks before or after the text section.

Step 4 If you want to change the text section's generic name, click section's name in the title bar, and enter a new name.

Step 5 Add formatted text and images to the body of the text section.

You can include input parameters that dynamically update when you generate the report.

Step 6 Click **Save**.

Related Topics

Input Parameters, on page 536

Add a Page Break to a Report Template

Procedure

Step 1 In the report template editor, click **Add Page Break** ( ).

A page break appears at the bottom of the template.

Step 2 Drag the page break to its intended location, before or after a section.

Step 3 Click **Save**.

Global Time Windows and Report Template Sections

Report templates with time-based data (such as intrusion or discovery events) have a global time window, which the time-based sections in the template inherit by default when created. Changing the global time window changes the local time window for the sections that are configured to inherit the global time window. You can disable time window inheritance for an individual section by clearing its **Inherit Time Window** check box. You can then edit the local time window.

Note Global time window inheritance applies only to report sections with data from time-based tables, such as intrusion events and discovery events. For sections that report on network assets (hosts and devices) and related information (such as vulnerabilities), you must set each time window individually.

Setting the Global Time Window for a Report Template and Its Sections

Tip Your report can have different time ranges per section. For example, your first section could be a summary for the month, and the remaining sections could drill down into details at the week level. In such cases, you set the section-level time windows individually.

Procedure

Step 1 In the report template editor, click **Generate**.

Step 2 To modify the global time window, click **Time Window** ( ).

Step 3 Modify time settings in **Events Time Window**.

Step 4 Click **Apply**.

Step 5 Click **Generate** to generate the report and **Yes** to confirm.

Setting the Local Time Window for Report Template Sections

Procedure

Step 1 On the Report Sections page of a template, clear the **Inherit Time Window** check box for the section if it is present.

Step 2 To change the section's local time window, click **Time Window** ( ).

Note
Sections with data from statistics tables can have only sliding time windows.

Step 3 Click **Apply** on the Events Time Window.

Step 4 Click **Save**.

Rename a Report Template Section

Procedure

Step 1 In the report template editor, click the current section name in the section header.

Step 2 Enter a new name for the section.

Step 3 Click **OK**.

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**.

Searches in Report Template Sections

The key to generating successful reports is defining the searches that populate the report's sections. The system provides a search editor to view the searches available in your report templates and to define new custom searches.

Searching in Report Template Sections

Procedure

Step 1 From the relevant section in the report template, click **Edit** ( ) next to the **Search** field.

Step 2 If you want to base a custom search on a predefined search, you must choose a predefined search from the **Saved Searches** drop-down list.

This list includes all available predefined searches for this table, including system-wide and report-specific predefined searches.

Step 3 Edit the search criteria in the appropriate fields.

For certain fields, your constraints can include the same operators (<, <>, and so on) as event searches. If you enter multiple criteria, the search returns only the records that match all the criteria.

Step 4 If you want to insert an input parameter from the drop-down menu instead of entering a constraint value, you must click **Input Parameter** ( ).

Note
When you edit the constraints of a reporting search, the system saves your edited search under the following name: *section* `custom search`, where *section* is the name in the section title bar followed by the string `custom search`. To have meaningful names for your saved custom searches, be sure you change the section name before you save the edited search. You cannot rename a saved reporting search.

Step 5 Click **OK.**

Input Parameters

You can use input parameters in a report template that the report can dynamically update at generation time. The **Input Parameter** ( ) indicates the fields that can process them. There are two kinds of input parameters:

- *Predefined input parameters* are resolved by internal system functions or configuration information. For example, at report generation time, the system replaces the `$<Time>` parameter with the current date and time.
- *User-defined input parameters* supply constraints in section searches. Constraining a search with an input parameter instructs the system to collect a value at generation time from the person who requests the report. In this way, you can dynamically tailor a report at generation time to show a particular subset of data without changing the template. For example, you can provide an input parameter for the **Destination IP** field of a report section's search. Then, when you generate the report, you can enter the IP network segment for a particular department to get data for that department only.

  You can also define string-type input parameters to add dynamic text in certain fields of your report, such as in emails (subject or body), report file names, and text sections. You can personalize reports for different departments, with customized report file names, email addresses, and email messages, using the same template for all.

Predefined Input Parameters

Table 57: 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 management center
`$<Model Number>`	The model number of the management center
`$<Time Window>`	The time window currently applied to the report section
`$<Constraints>`	The search constraints currently applied to the report section

Table 58: 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
`$<Model Number>`	yes	yes	yes	yes	yes	yes
`$<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 59: 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 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

Creating User-Defined Input Parameters

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Add Input Parameter** ( ).

Step 3 Enter the parameter **Name**.

Step 4 Choose a value from the **Type** drop-down list.

Step 5 Click **OK** to add the parameter.

Step 6 Click **OK** to return to the editor.

Editing User-Defined Input Parameters

The **Input Parameters** section of the report template lists all available user-defined parameters for the template.

Procedure

Step 1 In the report template editor, click **Advanced**.

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

Step 3 Enter a new **Name**.

Step 4 Use the **Type** drop-down list to change the parameter type.

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

Step 6 If you want to delete an input parameter, click **Delete** ( ) next to the input parameter and confirm.

Step 7 Click **OK** to return to the report template editor.

Constraining a Search with User-Defined Input Parameters

Input parameters you define are available only for search fields that match their parameter type. For example, a parameter of type **Network/IP** is available only for fields that accept IP addresses or network segments in CIDR format.

Procedure

Step 1 In the report template editor, click **Edit** ( ) next to the **Search** field within the section.

Fields that can take an input parameter are marked with **Input Parameter** ( ).

Step 2 Click **Input Parameter** ( ) next to the field, then choose the input parameter from the drop-down menu.

User-defined input parameters are marked with ( ).

Step 3 Click **OK**.

Document Attributes in a Report Template

Before you generate your report, you can set document attributes that affect the report's appearance. These attributes include the optional cover page and table of contents. Support for some attributes depends on the selected report format: PDF, HTML, or CSV.

Table 60: Document Attribute Support

Attribute	PDF Support?	HTML Support?	CSV Support?
Cover page	yes, with optional logo and custom appearance	yes, with optional logo and custom appearance	no
Table of contents	yes	yes	no
Page headers and footers	yes, with optional text or logo in any field	no	no
Custom starting page number	yes	no	no
Option to suppress numbering of first page	yes	no	no

Editing Document Attributes in a Report Template

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 You have the following choices:

- Add cover page —To add a cover page, check the **Include Cover Page** check box.
- Customize cover page —To edit the cover page design, see Customizing a Cover Page, on page 540.
- Add table of contents — To add a table of contents, check the **Include Table of Contents** check box.
- Manage logos — To manage the logo image associated with the template, see Managing Report Template Logos, on page 540.
- Configure header and footer —To specify elements for the header and footer of the template, use the drop-down lists in the **Header** and **Footer** fields.
- Set first page number — To specify the page number of the report's first page, enter a **Page Number Start** value.
- Show first page number —To show the page number on the report's first page, check the **Number First Page?** check box. If you choose this option, the cover page is not numbered.

Step 3 Click **OK** to save your changes.

Customizing a Cover Page

You can customize a report template's cover page. Cover pages can have rich text with multiple font sizes and styles (bold, italic, and so on) as well as input parameters and imported images.

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** () next to **Cover Page Design**.

Step 3 Edit the cover page design within the rich text editor.

Step 4 Click **OK**.

Managing Report Template Logos

You can store multiple logos on the 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 management center, it is available for:

- all report templates on the 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 management center. For example, if you reuse a template, you can associate a logo for a different organization with the report.

You can delete any uploaded logos. Deleting a logo removes it from all templates where it is used. The deletion cannot be undone. Note that you cannot delete the predefined Cisco logo.

Procedure

Step 1 In the report template editor, click **Advanced**.

The logo currently associated with the template appears under **Logo** in **General Settings**.

Step 2 Click **Edit** ( ) next to the logo.

Step 3 You have the following choices:

- Add — Add a new logo as described in Adding a New Logo, on page 541.
- Change — Change a report template's logo as described in Changing the Logo for a Report Template, on page 541.
- Delete — Delete a logo as described in Deleting a Logo, on page 542.

Adding a New Logo

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 Click **Upload Logo**.

Step 4 Click **Browse**, browse to the file's location, and click **Open**.

Step 5 Click **Upload**.

Step 6 If you want to associate the new logo with the current template, choose it, and click **OK**.

Changing the Logo for a Report Template

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo to associate with the report template.

Step 4 Click **OK**.

Deleting a Logo

Procedure

Step 1 In the report template editor, click **Advanced**.

Step 2 Click **Edit** ( ) next to the **Logo** field.

Step 3 From the Select Logo dialog, choose the logo you want to delete.

Step 4 Click **Delete Logo**.

Step 5 Click **OK**.

Managing Report Templates

In a multidomain deployment, the system displays report templates created in the current domain, which you can edit. It also displays report templates created in ancestor domains, which you cannot edit. To view and edit report templates in a lower domain, switch to that domain. The system displays reports created in the current domain only.

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 You have the following choices:

- Delete — Next to the template you want to delete, click **Delete** ( ) and confirm.

  You cannot delete system-provided report templates. Security Analysts can delete only report templates they created. In a multidomain deployment, you can delete report templates belonging to the current domain only.

- Edit — To edit report templates; see Editing Report Templates, on page 542.
- Export — To export report templates, see Exporting Report Templates, on page 543.

  Tip
  You can also export report templates using the standard configuration export process; see Exporting Configurations, on page 509.

- Import — To import report templates, see Importing Configurations, on page 510.

Editing Report Templates

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

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Edit** (✎) for the template you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 You have the following choices:

- Add a page break; see Add a Page Break to a Report Template, on page 534.
- Add a text section; see Add a Text Section to a Report Template, on page 533.
- Configure section content as described in Report Template Configuration, on page 531.
- Create input parameters; see Creating User-Defined Input Parameters, on page 538.
- Edit input parameters; see Editing User-Defined Input Parameters, on page 538.
- Edit document attributes; see Editing Document Attributes in a Report Template, on page 540.
- Search template sections; see Searching in Report Template Sections, on page 536.
- Set document attributes described in Document Attributes in a Report Template, on page 539 by clicking **Advanced**.
- Set the global time window; see Setting the Global Time Window for a Report Template and Its Sections, on page 534.
- Set the local time window; see Setting the Local Time Window for Report Template Sections, on page 535.
- Set the search fields; see Modify Fields in the Report Template Table Format Sections, on page 533.
- Set the table and data format; see Set the Table and Data Format for a Report Template Section, on page 532.
- Specify searches and filters; see Specify the Search or Filter for a Report Template Section, on page 533.

Exporting Report Templates

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Report Templates**.

Step 3 For the template you want to export, click the **Export** icon.

About Generating Reports

Generating Reports

After you create and customize your report template, you are ready to generate the report. The generation process lets you select the report's format (HTML, PDF, or CSV). You can also adjust the report's global time window, which applies a consistent time frame to all sections except those you exempt.

For PDF reports:

- File names using Unicode (UTF-8) characters are not supported.
- Any report sections that include special Unicode file names (such as those appearing in file or malware events) display these file names in transliterated form.
- 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 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.

Step 10 You have the following choices:

- Click the report link to display the report in a new window.
- Click **OK** to return to the report template editor.

Report Generation Options

You can configure report generation options to:

- Schedule generation of future reports, either once or recurring. See Automating Report Generation, on page 494. You can customize the schedule on a full range of time frames such as daily, weekly, monthly, and so on.
- Distribute email reports using the scheduler. You must configure your report template and a mail relay host **before** scheduling the task.
- Automatically send the report as an email attachment to a list of recipients when you generate a report. You must have a properly configured mail relay host to deliver a report by email.
- Save newly generated report files to your configured remote storage location. To use remote storage, you must first configure a remote storage location.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Distributing Reports by Email at Generation Time

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Report Templates**.

Step 3 Click **Report** ( ) next to the template you want to use to generate a report.

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

Tip
To generate a report from an ancestor's template, copy the template into the current domain.

Step 4 Expand the **Email** section of the window.

Step 5 In the **Email Options** field, choose **Send Email**.

Step 6 In the **Recipient List**, **CC**, and **BCC** fields, enter recipients' email addresses in comma-separated lists.

Step 7 In the **Subject** field, enter an email subject.

Tip
You can provide input parameters in the **Subject** field and the message body to dynamically generate information in the email, such as a timestamp or the name of the management center.

Step 8 Enter a cover letter in the email body as necessary.

Step 9 Click **OK** and confirm.

Related Topics

Configuring a Mail Relay Host and Notification Address, on page 62

Schedule Future Reports

See Automating Report Generation, on page 494.

About Working with Generated Reports

Access and work with previously-generated reports on the Reports tab page.

Viewing Reports

The Reports lists all previously generated reports, with report name, date and time of generation, generating user, and whether the report is stored locally or remotely. A status column indicates whether the report is already generated, is in the generation queue (for example, for scheduled tasks), or failed to generate (for example, due to lack of disk space).

Note that users with Administrator access can view all reports; other users can view only the reports they generated.

In a multidomain deployment, you can view reports generated in the current domain only.

The Reports page shows all locally stored reports. It shows remotely stored reports as well, if remote storage is currently configured. The **Location** column data for remotely-stored reports is `Remote`.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Click the report you want to view.

Downloading Reports

You can download any report file to your local computer. From there, you can email it or distribute it electronically by other available means.

In a multidomain deployment, you can download reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 Check the check boxes next to the reports you want to download, then click **Download**.

Tip
Click the check box at the top left of the page to download all reports on the page. If you have multiple pages of reports, a second check box appears that you can click to download all reports on all pages.

Step 4 Follow your browser's prompts to download the reports. If you chose multiple reports, they are downloaded in a single `.zip` file.

Storing Reports Remotely

The location of your currently configured report storage appears at the bottom of the Overview> Reporting > Reports page, with disk usage for local, NFS, and SMB storage. If you access remote storage using SSH, disk usage data is not available.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

- Configure a remote storage location as described in Remote Storage Device, on page 97.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Check the **Enable Remote Storage of Reports** check box at the bottom of the page.

What to do next

- Move reports from local storage to remote storage; see Moving Reports to Remote Storage, on page 548.

Related Topics

Remote Storage Device, on page 97
Moving Reports to Remote Storage, on page 548

Moving Reports to Remote Storage

You can move your reports in local storage to a remote storage location in batch mode or singly.

Note If you store remotely and then switch back to local storage, the reports in remote storage do not appear on the Reports tab list. Similarly, if you switch from one remote storage location to another, the reports in the previous location do not appear in the list.

Before you begin

- Configure a remote storage location as described in Remote Storage Device, on page 97.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Choose **Reports**.

Step 3 Choose the check boxes next to the reports you want to move, then click **Move**.

Tip
Check the check box at the top left of the page to move all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to move all reports on all pages.

Step 4 Confirm that you want to move the reports.

Deleting Reports

You can delete your report files at any time. The procedure completely removes the files, and no recovery is possible. Although you still have the report template that generated the report, it may be difficult to regenerate a particular report file if the time window was expanding or sliding. Regeneration may also be difficult if your template uses input parameters.

In a multidomain deployment, you can delete reports generated in the current domain only.

Procedure

Step 1 Choose **Overview** > **Reporting**.

Step 2 Click **Reports**.

Step 3 You have the following choices:

- Delete selected — Check the check boxes next to the reports you want to delete, then click **Delete**.
- Delete all — Check the check box at the top left of the page to delete all reports on the page. If you have multiple pages of reports, a second check box appears that you can check to delete all reports on all pages.

Step 4 Confirm the deletion.

History for Reporting

Feature	Minimum Management Center	Minimum 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.

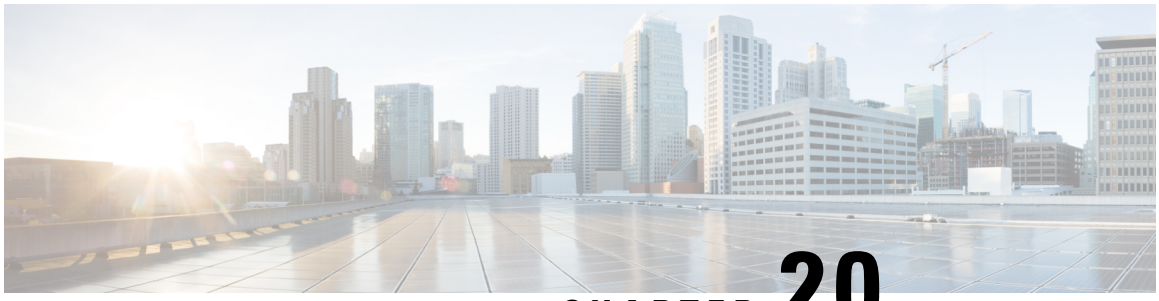

CHAPTER 20

External Alerting with Alert Responses

The following topics describe how to send external event alerts from the Secure Firewall Management Center using alert responses:

Secure Firewall Management Center Alert Responses

External event notification via SNMP, syslog, or email can help with critical-system monitoring. The Secure Firewall Management Center uses configurable *alert responses* to interact with external servers. An *alert response* is a configuration that represents a connection to an email, SNMP, or syslog server. They are called *responses* because you can use them to send alerts in response to events detected by Secure Firewall. You can configure multiple alert responses to send different types of alerts to different monitoring servers and/or people.

Note Depending on your device and Secure Firewall version, alert responses may not be the best way to send syslog messages. See the *About Syslog* chapter in the Cisco Secure Firewall Management Center Device Configuration Guide and Best Practices for Configuring Security Event Syslog Messaging, on page 625..

Note Alerts that use alert responses are sent by the Secure Firewall Management Center. Intrusion email alerts, which do not use alert responses, are also sent by the Secure Firewall Management Center. By contrast, SNMP and syslog alerts that are based on individual intrusion rules triggering are sent directly by managed devices. For more information, see External Alerting for Intrusion Events, on page 561.

In most cases, the information in an external alert is the same as the information in any associated event you logged to the database. However, for correlation event alerts where the correlation rule contains a connection

tracker, the information you receive is the same as for an alert on a traffic profile change, regardless of the base event type.

You create and manage alert responses on the Alerts page (**Policies** > **Actions** > **Alerts**). New alert responses are automatically enabled. To temporarily stop alert generation, you can disable alert responses rather than deleting them.

Changes to alert responses take effect immediately, except when sending connection logs to an SNMP trap or syslog server.

Configurations Supporting Alert Responses

After you create an alert response, you can use it to send the following external alerts from the Secure Firewall Management Center.

Alert/Event Type	For More Information
Intrusion events, by impact flag	Configuring Impact Flag Alerting, on page 557
Discovery events, by type	Configuring Discovery Event Alerting, on page 558
Malware and retrospective malware events detected by malware defense ("network-based")	Configuring Malware defense Alerting, on page 559
Correlation events, by correlation policy violation	Adding Responses to Rules and Allow Lists, on page 955
Connection events, by the logging rule or default action (email alerts not supported)	Other Connections You Can Log, on page 714
Health events, by health module and severity level	Creating Health Monitor Alerts, on page 379

Requirements and Prerequisites for Alert Responses

Model Support

Any.

Supported Domains

Any

User Roles

- Admin

Creating an SNMP Alert Response

You can create SNMP alert responses using SNMPv1, SNMPv2, or SNMPv3 for threat defense devices.

Note When selecting SNMP versions for the SNMP protocol, note that SNMPv2 only supports read-only communities and SNMPv3 only supports read-only users. SNMPv3 also supports encryption with AES128.

If you want to monitor 64-bit values with SNMP, you must use SNMPv2 or SNMPv3. SNMPv1 does not support 64-bit monitoring.

Before you begin

- If your network management system requires the management center's management information base (MIB) file, obtain it at `/etc/sf/DCEALERT.MIB`.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create SNMP Alert**.

Step 3 Edit the SNMP Alert Configuration fields:

a) **Name—**Enter a name to identify the SNMP response.

b) **Trap Server—**Enter the hostname or IP address of the SNMP trap server.

Note
The system does **not** warn you if you enter an invalid IPv4 address (such as 192.169.1.456) in this field. Instead, the invalid address is treated as a hostname.

c) **Version—**Choose the SNMP version you want to use from the drop-down list. SNMPv3 is the default.

Choose from:

- **SNMPv1** or **SNMPv2**: Enter a read-only SNMP community name in the **Community String** field, then skip to the end of the procedure.

  Note
  Do not include special characters (< > / % # & ? ', etc.) in the SNMP community string name.

- For **SNMPv3**: Enter the name of the user that you want to authenticate with the SNMP server in the **User Name** field and continue to the next step.

d) **Authentication Protocol—**Choose the protocol you want to use to encrypt authentication from the drop-down list.

Choose from:

- **MD5**—Message Digest 5 (MD5) hash function.
- **SHA**—Secure Hash Algorithm (SHA) hash function.

e) **Authentication Password—**Enter the password to enable authentication.

f) **Privacy Protocol—**Choose the protocol you want to use to encrypt a private password from the drop-down list.

Choose from:

- **DES**—Data Encryption Standard (DES) using 56-bit keys in a symmetric secret-key block algorithm.
- **AES**—Advanced Encryption Standard (AES) using 56-bit keys in a symmetric cipher algorithm.
- **AES128**—AES using 128-bit keys in a symmetric cipher algorithm. A longer key provides higher security but a reduction in performance.

g) **Privacy Password—**Enter the privacy password required by the SNMP server. If you specify a private password, privacy is enabled, and you must also specify an authentication password.

h) **Engine ID—**Enter an identifier for the SNMP engine, in hexadecimal notation, using an even number of digits.

When you use SNMPv3, the system uses an Engine ID value to encode the message. Your SNMP server requires this value to decode the message.

Cisco recommends that you use the hexadecimal version of the management center's IP address. For example, if the management center has an IP address of 10.1.1.77, use 0a01014D0.

Step 4 Click **Save**.

What to do next

Changes take effect immediately, except if you are using alert responses to send connection logs, you must deploy configuration changes after you edit those alert responses.

Creating a Syslog Alert Response

When configuring a syslog alert response, you can specify the severity and facility associated with the syslog messages to ensure that they are processed properly by the syslog server. The facility indicates the subsystem that creates the message and the severity defines the severity of the message. Facilities and severities are not displayed in the actual message that appears in the syslog, but are instead used to tell the system that receives the syslog message how to categorize it.

Tip For more detailed information about how syslog works and how to configure it, refer to the documentation for your system. On UNIX systems, the `man` pages for `syslog` and `syslog.conf` provide conceptual information and configuration instructions.

Although you can choose any type of facility when creating a syslog alert response, you should choose one that makes sense based on your syslog server; not all syslog servers support all facilities. For UNIX syslog servers, the `syslog.conf` file should indicate which facilities are saved to which log files on the server.

Syslog messages are transmitted over either UDP or TCP, depending on the configuration of the syslog server.

Before you begin

- This procedure is not the recommended way to send syslog messages in many cases.
- Confirm that the syslog server can accept remote messages.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Syslog Alert**.

Step 3 Enter a **Name** for the alert.

Step 4 In the **Host** field, enter the hostname or IP address of your syslog server.

Note
The system does **not** warn you if you enter an invalid IPv4 address (such as 192.168.1.456) in this field. Instead, the invalid address is treated as a hostname.

Step 5 In the **Port** field, enter the port the server uses for syslog messages. By default, this value is 514.

Step 6 From the **Facility** list, choose a facility described in Syslog Alert Facilities, on page 555.

Step 7 From the **Severity** list, choose a severity described in Syslog Severity Levels, on page 556.

Step 8 In the **Tag** field, enter the tag name that you want to appear with the syslog message.

For example, if you wanted all messages sent to the syslog to be preceded with `FromMC`, enter `FromMC` in the field.

Step 9 Click **Save**.

What to do next

Changes take effect immediately, EXCEPT:

If you are using alert responses to send connection logs to a syslog server, you must deploy configuration changes after you edit those alert responses.

If you will use this alert response for security events, you MUST specify the alert response in a policy. See Configuration Locations for Security Event Syslogs, on page 629.

Syslog Alert Facilities

The following table lists the syslog facilities you can select.

Table 61: Available 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. Note that syslog servers running a Linux operating system will use the `CRON` facility.

Facility	Description
DAEMON	A message generated by a system daemon.
FTP	A message generated by the FTP daemon.
KERN	A message generated by the kernel. On many systems, these messages are printed to the console when they appear.
LOCAL0-LOCAL7	A message generated by an internal process.
LPR	A message generated by the printing subsystem.
MAIL	A message generated by a mail system.
NEWS	A message generated by the network news subsystem.
NTP	A message generated by the NTP daemon.
SECURITY	A message generated by the audit subsystem.
SYSLOG	A message generated by the syslog daemon.
SOLARIS-CRON	A message generated by the clock daemon. Note that syslog servers running a Windows operating system will use the `CLOCK` facility.
USER	A message generated by a user-level process.
UUCP	A message generated by the UUCP subsystem.

Syslog Severity Levels

The following table lists the standard syslog severity levels you can select.

Table 62: Syslog Severity Levels

Level	Description
ALERT	A condition that should be corrected immediately.
CRIT	A critical condition.
DEBUG	Messages that contain debugging information.
EMERG	A panic condition broadcast to all users.
ERR	An error condition.
INFO	Informational messages.
NOTICE	Conditions that are not error conditions, but require attention.

Level	Description
WARNING	Warning messages.

Creating an Email Alert Response

Before you begin

- Ensure that the Secure Firewall Management Center can reverse-resolve its own IP address. Some mail servers may perform reverse DNS lookups to verify the sender's identity as a measure to prevent spam and unauthorized access.
- Configure your mail relay host as described in Configuring a Mail Relay Host and Notification Address, on page 62.

Note You cannot use email alerting to log connections.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 From the **Create Alert** drop-down menu, choose **Create Email Alert**.

Step 3 Enter a **Name** for the alert response.

Step 4 In the **To** field, enter the email addresses where you want to send alerts, separated by commas.

Step 5 In the **From** field, enter the email address that you want to appear as the sender of the alert.

Step 6 Next to **Relay Host**, verify the listed mail server is the one that you want to use to send the alert.

Tip
To change the email server, click **Edit** ( ).

Step 7 Click **Save**.

Configuring Impact Flag Alerting

You can configure the system to alert you whenever an intrusion event with a specific impact flag occurs. Impact flags help you evaluate the impact an intrusion has on your network by correlating intrusion data, network discovery data, and vulnerability information.

You must have the IPS Smart License or Protection Classic License to configure these alerts.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Impact Flag Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

Tip
To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Impact Configuration** section, check the appropriate check boxes to specify the alerts you want to receive for each impact flag.

For definitions of the impact flags, see Intrusion Event Impact Levels, on page 777.

Step 5 Click **Save**.

Configuring Discovery Event Alerting

You can configure the system to alert you whenever a specific type of discovery event occurs.

Before you begin

- Configure your network discovery policy to log the discovery event types you want to configure alerting for as described in 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 you want to use for each alert type.

Tip
To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Events Configuration** section, check the check boxes that correspond to the alerts you want to receive for each discovery event type.

Step 5 Click **Save**.

Configuring Malware defense Alerting

You can configure the system to alert you whenever any malware event, including a retrospective event, is generated by malware defense (that is, a "network-based malware event" is generated.) You cannot alert on malware events generated by Secure Endpoint ("endpoint-based malware events.")

Before you begin

- Configure a file policy to perform malware cloud lookups and associate that policy with an access control rule. See *Access Control Overview* in the Cisco Secure Firewall Management Center Device Configuration Guide for more information.
- You must have the Malware Defense license to configure these alerts.

Procedure

Step 1 Choose **Policies** > **Actions** > **Alerts**.

Step 2 Click **Advanced Malware Protections Alerts**.

Step 3 In the **Alerts** section, choose the alert response you want to use for each alert type.

Tip
To create a new alert response, choose **New** from any drop-down list.

Step 4 In the **Event Configuration** section, check the check boxes that correspond to the alerts you want to receive for each malware event type.

Keep in mind that **All network-based malware events** includes **Retrospective Events**.

(By definition, network-based malware events do not include events generated by Secure Endpoint.)

Step 5 Click **Save**.

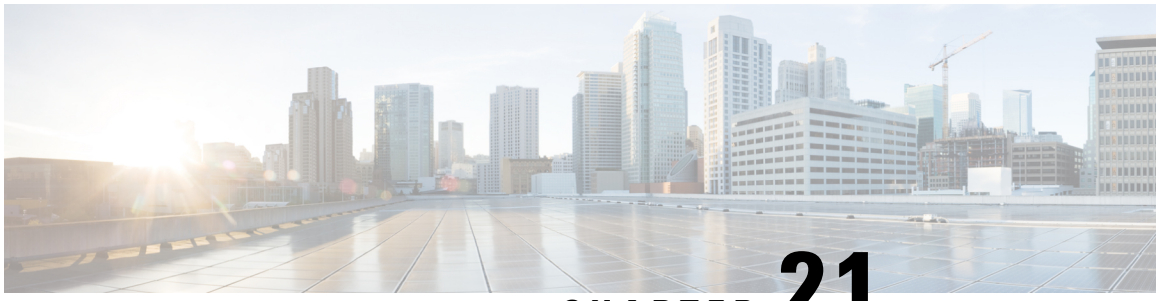

CHAPTER 21

External Alerting for Intrusion Events

The following topics describe how to configure external alerting for intrusion events:

About External Alerting for Intrusion Events

External intrusion event notification can help with critical-system monitoring:

- SNMP—Configured per intrusion policy and sent from managed devices. You can enable SNMP alerting per intrusion rule.
- Syslog—Configured per intrusion policy and sent from managed devices. When you enable syslog alerting in an intrusion policy, you turn it on for every rule in the policy.
- Email—Configured across all intrusion policies and sent from the Secure Firewall Management Center. You can enable email alerts per intrusion rule, as well as limit their length and frequency.

Keep in mind that if you configured intrusion event suppression or thresholding, the system may not generate intrusion events (and thus may not send alerts) every time a rule triggers.

Note The Secure Firewall Management Center also uses SNMP, syslog, and email *alert responses* to send different types of external alerts; see Secure Firewall Management Center Alert Responses, on page 551. The system does **not** use alert responses to send alerts based on individual intrusion events.

Related Topics

Intrusion Event Notification Filters in an Intrusion Policy

License Requirements for External Alerting for Intrusion Events

Threat Defense License

IPS

Classic License

Protection

Requirements and Prerequisites for External Alerting for Intrusion Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Configuring SNMP Alerting for Intrusion Events

After you enable external SNMP alerting in an intrusion policy, you can configure individual rules to send SNMP alerts when they trigger. These alerts are sent from the managed device.

Procedure

Step 1 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 2 Make sure **SNMP Alerting** is **Enabled**, then click **Edit**.
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration.

Step 3 Choose an **SNMP Version**, then specify configuration options as described in Intrusion SNMP Alert Options, on page 563.

Step 4 In the navigation pane, click **Rules**.

Step 5 In the rules pane, choose the rules where you want to set SNMP alerts, then choose **Alerting** > **Add SNMP Alert**.

Step 6 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

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

Intrusion SNMP Alert 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**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.
Community String	The community name.

SNMP v3 Options

Managed devices encode SNMPv3 alerts with an Engine ID value. To decode the alerts, your SNMP server requires this value, which is the hexadecimal version of the sending device's management interface IP address, appended with "01."

For example, if the device sending the SNMP alert has a management interface IP address of 172.16.1.50, the Engine ID value is 0xAC10013201.

Option	Description
Trap Type	The trap type to use for IP addresses that appear in the alerts. If your network management system correctly renders the INET_IPV4 address type, choose **as Binary**. Otherwise, choose **as String**. For example, HP OpenView requires **as String**.
Trap Server	The server that will receive SNMP traps notification. You can specify a single IP address or hostname.

Option	Description
Authentication Password	The password required for authentication. SNMP v3 uses either the Message Digest 5 (MD5) hash function or the Secure Hash Algorithm (SHA) hash function to encrypt this password, depending on configuration. If you specify an authentication password, authentication is enabled.
Private Password	The SNMP key for privacy. SNMP v3 uses the Data Encryption Standard (DES) block cipher to encrypt this password. When you enter an SNMP v3 password, the password displays in plain text during initial configuration but is saved in encrypted format. If you specify a private password, privacy is enabled, and you must also specify an authentication password.
User Name	Your SNMP user name.

Configuring Syslog Alerting for Intrusion Events

After you enable syslog alerting in an intrusion policy, the system sends all intrusion events to the syslog, either on the managed device itself or to an external host or hosts. If you specify an external host, syslog alerts are sent from the managed device.

Procedure

Step 1 In the intrusion policy editor's navigation pane, click **Advanced Settings**.

Step 2 Make sure **Syslog Alerting** is **Enabled**, then click **Edit**.
A message at the bottom of the page identifies the intrusion policy layer that contains the configuration. The **Syslog Alerting** page is added under **Advanced Settings**.

Step 3 Enter the IP addresses of the **Logging Hosts** where you want to send syslog alerts.

If you leave the **Logging Hosts** field blank, the logging hosts details are taken from Logging in the associated Access Control Policy.

Step 4 Choose **Facility** and **Severity** levels as described in Facilities and Severities for Intrusion Syslog Alerts, on page 565.

Step 5 To save changes you made in this policy since the last policy commit, choose **Policy Information**, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

What to do next

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

Facilities and Severities for Intrusion Syslog Alerts

Managed devices can send intrusion events as syslog alerts using a particular facility and **Severity**, so that the logging host can categorize the alerts. The *facility* specifies the subsystem that generated it. These facility and **Severity** values do not appear in the actual syslog messages.

Choose values that make sense based on your environment. Local configuration files (such as syslog.conf on UNIX-based logging hosts) may indicate which facilities are saved to which log files.

Syslog Alert Facilities

Facility	Description
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

Level	Description
WARNING	Warning messages
NOTICE	Conditions that are not error conditions, but require attention
INFO	Informational messages
DEBUG	Messages that contain debug information

Configuring Email Alerting for Intrusion Events

If you enable intrusion email alerting, the system can send email when it generates an intrusion event, regardless of which managed device or intrusion policy detected the intrusion. These alerts are sent from the Secure Firewall Management Center.

Before you begin

- Configure your mail host to receive email alerts; see Configuring a Mail Relay Host and Notification Address, on page 62.
- Ensure that the Secure Firewall Management Center can reverse-resolve is own IP address. Some mail servers may perform reverse DNS lookups to verify the sender's identity as a measure to prevent spam and unauthorized access.

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 566.

Step 4 Click **Save**.

Intrusion Email Alert Options

On/Off

Enables or disables intrusion email alerts.

Note Enabling it will enable alerting for all rules unless individual rules are selected.

From/To Addresses

The email sender and recipients. You can specify a comma-separated list of recipients.

Max Alerts and Frequency

The maximum number of email alerts (**Max Alerts**) that the Secure Firewall Management Center will send per time interval (**Frequency**).

Coalesce Alerts

Reduces the number of alerts sent by grouping alerts that have the same source IP and rule ID.

Summary Output

Enables brief alerts, suitable for text-limited devices. Brief alerts contain:

- Timestamp
- Protocol
- Source and destination IPs and ports
- Message
- The number of intrusion events generated against the same source IP

For example: `2011-05-18 10:35:10 10.1.1.100 icmp 10.10.10.1:8 -> 10.2.1.3:0 snort_decoder: Unknown Datagram decoding problem! (116:108)`

If you enable **Summary Output**, also consider enabling **Coalesce Alerts**. You may also want to lower **Max Alerts** to avoid exceeding text-message limits.

Time Zone

The time zone for alert timestamps.

Email Alerting on Specific Rules Configuration

Allows you to choose the rules where you want to set email alerts.

PART VI

Event and Asset Analysis Tools

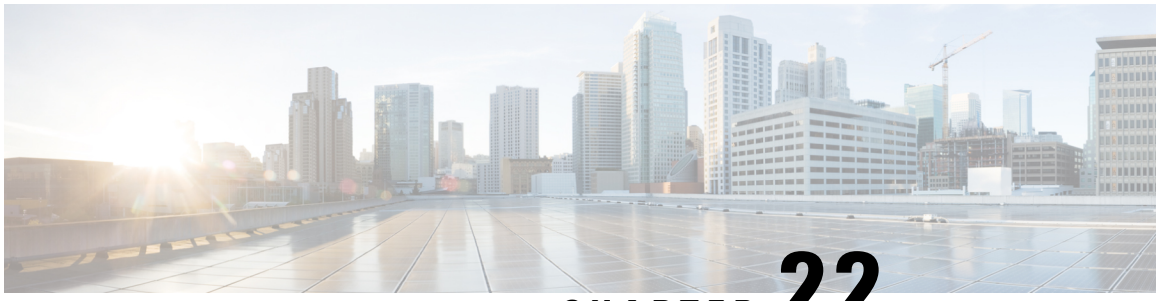

CHAPTER 22

Context Explorer

The following topics describe how to use the Context Explorer:

About the Context Explorer

The Context Explorer displays detailed, interactive graphical information in context about the status of your monitored network, including data on applications, application statistics, connections, geolocation, indications of compromise, intrusion events, hosts, servers, Security Intelligence, users, files (including malware files), and relevant URLs. Distinct sections present this data in the form of vivid line, bar, pie, and donut graphs, accompanied by detailed lists. The first section, a line chart of traffic and event counts over time, provides an at-a-glance picture of recent trends in your network's activity.

You can easily create and apply custom filters to fine-tune your analysis, and you can examine data sections in more detail by simply clicking or hovering your cursor over graph areas. You can also configure the explorer's time range to reflect a period as short as the last hour or as long as the last year. Only users with the Administrator, Security Analyst, or Security Analyst (Read Only) user roles have access to the Context Explorer.

The dashboard is highly customizable and compartmentalized and updates in real time. In contrast, the Context Explorer is manually updated, designed to provide broader context for its data, and has a single, consistent layout designed for active user exploration.

You use the dashboard to monitor real-time activity on your network and appliances according to your own specific needs. Conversely, you use the Context Explorer to investigate a predefined set of recent data in granular detail and clear context: for example, if you notice that only 15% of hosts on your network use Linux, but account for almost all YouTube traffic, you can quickly apply filters to view data only for Linux hosts, only for YouTube-associated application data, or both. Unlike the compact, narrowly focused dashboard widgets, the Context Explorer sections are designed to provide striking visual representations of system activity in a format useful to both expert and casual users.

The data displayed depends on such factors as how you license and deploy your managed devices, and whether you configure features that provide the data. You can also apply filters to constrain the data that appears in all Context Explorer sections.

In a multidomain deployment, the Context Explorer displays aggregated data from all subdomains when you view it in an ancestor domain. In a leaf domain, you can view data specific to that domain only.

Differences Between the Dashboard and the Context Explorer

The following table summarizes some of the key differences between the dashboard and the Context Explorer.

Table 63: 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

About Dashboards, on page 339

The Traffic and Intrusion Event Counts Time Graph

At the top of the Context Explorer is a line chart of traffic and intrusion events over time. The X-axis plots time intervals (which range from five minutes to one month, depending on the selected time window). The Y-axis plots traffic in kilobytes (blue line) and intrusion event count (red line).

Note that the smallest X-axis interval is five minutes. To accommodate this, the system will round the beginning and ending points in your selected time range down to the nearest five-minute interval.

By default, this section shows all network traffic and all generated intrusion events for the selected time range. If you apply filters, the chart changes to display only traffic and intrusion events associated with the criteria

specified in the filters. For example, filtering on the **OS Name** of `Windows` causes the time graph to display only traffic and events associated with hosts using Windows operating systems.

If you filter the Context Explorer on intrusion event data (such as a **Priority** of `High`), the blue Traffic line is hidden to allow greater focus on intrusion events alone.

You can hover your pointer over any point on the graph lines to view exact information about traffic and event counts. Hovering your pointer over one of the colored lines also brings that line to the forefront of the graph, providing clearer context.

This section draws data primarily from the Intrusion Events and Connection Events tables.

The Indications of Compromise Section

The Indications of Compromise (IOC) section of the Context Explorer contains two interactive sections that provide an overall picture of potentially compromised hosts on your monitored network: a proportional view of the most prevalent IOC types triggered, as well as a view of hosts by number of triggered indications.

For more information about IOCs, see Indications of Compromise Data, on page 892.

The Hosts by Indication Graph

The Hosts by Indication graph, in donut form, displays a proportional view of the Indications of Compromise (IOC) triggered by hosts on your monitored network. The inner ring divides by IOC category (such as `CnC Connected` or `Malware Detected`), while the outer ring further divides that data by specific event type (such as `Impact 2 Intrusion Event — attempted-admin` or `Threat Detected in File Transfer`).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts and Host Indications of Compromise tables.

The Indications by Host Graph

The Indications by Host graph, in bar form, displays counts of unique Indications of Compromise (IOC) triggered by the 15 most IOC-active hosts on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts and Host Indications of Compromise tables.

The Network Information Section

The Network Information section of the Context Explorer contains six interactive graphs that display an overall picture of connection traffic on your monitored network: sources, destinations, users, and security zones associated with traffic, a breakdown of operating systems used by hosts on the network, as well as a proportional view of access control actions that have been performed on network traffic.

The Operating Systems Graph

The Operating Systems graph, in donut form, displays a proportional representation of operating systems detected on hosts on your monitored network. The inner ring divides by OS name (such as `Windows` or `Linux`), while the outer ring further divides that data by specific operating system version (such as `Windows Server 2008` or `Linux 11.x`). Some closely related operating systems (such as Windows 2000, Windows XP, and

Windows Server 2003) are grouped together. Very scarce or unrecognized operating systems are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Hosts table.

The Traffic by Source IP Graph

The Traffic by Source IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source IP addresses on your monitored network. For each source IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note If you filter on intrusion event information, the Traffic by Source IP graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Source User Graph

The Traffic by Source User graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active source users on your monitored network. For each source IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note If you filter on intrusion event information, the Traffic by Source User graph is hidden.

This graph draws data primarily from the Connection Events table. It displays authoritative user data.

The Connections by Access Control Action Graph

The Connections by Access Control Action graph, in pie form, displays a proportional view of access control actions (such as `Block` or `Allow`) taken on monitored traffic.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note If you filter on intrusion event information, the Traffic by Source User graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Destination IP Graph

The Traffic by Destination IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most active destination IP addresses on your monitored network. For each destination IP address listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note If you filter on intrusion event information, the Traffic by Destination IP graph is hidden.

This graph draws data primarily from the Connection Events table.

The Traffic by Ingress/Egress Security Zone Graph

The Traffic by Ingress/Egress Security Zone graph, in bar form, displays counts of incoming or outgoing network traffic (in kilobytes per second) and unique connections for each security zone configured on your monitored network. You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

For each security zone listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information

Tip To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

Note If you filter on intrusion event information, the Traffic by Ingress/Egress Security Zone graph is hidden.

This graph draws data primarily from the Connection Events table.

The Application Information Section

The Application Information section of the Context Explorer contains three interactive graphs and one table-format list that display an overall picture of application activity on your monitored network: traffic, intrusion events, and hosts associated with applications, further organized by the estimated risk or business relevance assigned to each application. The Application Details list provides an interactive list of each application and its risk, business relevance, category, and host count.

For all instances of "application" in this section, the Application Information graph set, by default, specifically examines application protocols (such as DNS or SSH). You can also configure the Application Information section to specifically examine client applications (such as PuTTY or Firefox) or web applications (such as Facebook or Pandora).

Focusing the Application Information Section

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Hover your pointer over the **Application Protocol Information** section.

Note
If you previously changed this setting in the same Context Explorer session, the section title may appear as **Client Application Information** or **Web Application Information** instead.

Step 3 Click **Application Protocol**, **Client Application**, or **Web Application**.

The Traffic by Risk/Business Relevance and Application Graph

The Traffic by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of application traffic detected on your monitored network, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring further divides that data by specific application (such as `SSH` or `NetBIOS`). Scarcely detected applications are grouped under **Other**.

Note that this graph reflects all available data regardless of date and time constraints. If you change the explorer time range, the graph does not change.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays traffic by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

Note If you filter on intrusion event information, the Traffic by Risk/Business and Application graph is hidden.

This graph draws data primarily from the Connection Events and Application Statistics tables.

The Intrusion Events by Risk/Business Relevance and Application Graph

The Intrusion Events by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of intrusion events detected on your monitored network and the applications associated with those events, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring

further divides that data by specific application (such as `SSH` or `NetBIOS`). Scarcely detected applications are grouped under **Other**.

Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information, or (where applicable) to view application information.

Tip To constrain the graph so it displays intrusion events by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

This graph draws data primarily from the Intrusion Events and Application Statistics tables.

The Hosts by Risk/Business Relevance and Application Graph

The Hosts by Risk/Business Relevance and Application graph, in donut form, displays a proportional representation of hosts detected on your monitored network and the applications associated with those hosts, arranged by the applications' estimated risk (the default) or estimated business relevance. The inner ring divides by estimated risk/business relevance level (such as `Medium` or `High`), while the outer ring further divides that data by specific application (such as `SSH` or `NetBIOS`). Very scarce applications are grouped under **Other**.

Hover your pointer over any part of the donut graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays hosts by business relevance and application, hover your pointer over the graph, then click **Business Relevance** on the toggle button that appears. Click **Risk** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Risk view.

This graph draws data primarily from the Applications table.

The Application Details List

At the bottom of the Application Information section is the Application Details List, a table that provides estimated risk, estimated business relevance, category, and hosts count information for each application detected on your monitored network. The applications are listed in descending order of associated host count.

The Application Details List table is not sortable, but you can click on any table entry to filter or drill down on that information, or (where applicable) to view application information. This table draws data primarily from the Applications table.

Note that this list reflects all available data regardless of date and time constraints. If you change the explorer time range, the list does not change.

The Security Intelligence Section

The Security Intelligence section of the Context Explorer contains three interactive bar graphs that display an overall picture of traffic on your monitored network that is blocked or monitored by Security Intelligence. The graphs sort such traffic by category, source IP address, and destination IP address, respectively; both the amount of traffic (in kilobytes per second) and the number of applicable connections appear.

The Security Intelligence Traffic by Category Graph

The Security Intelligence Traffic by Category graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top Security Intelligence categories of traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Security Intelligence Traffic by Category graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Security Intelligence Traffic by Source IP Graph

The Security Intelligence Traffic by Source IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top source IP addresses of Security Intelligence-monitored traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Security Intelligence Traffic by Source IP graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Security Intelligence Traffic by Destination IP Graph

The Security Intelligence Traffic by Destination IP graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top destination IP addresses of Security Intelligence-monitored traffic on your monitored network. For each category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Security Intelligence Traffic by Destination IP graph is hidden.

This graph draws data primarily from the Security Intelligence Events table.

The Intrusion Information Section

The Intrusion Information section of the Context Explorer contains six interactive graphs and one table-format list that display an overall picture of intrusion events on your monitored network: impact levels, attack sources, target destinations, users, priority levels, and security zones associated with intrusion events, as well as a detailed list of intrusion event classifications, priorities, and counts.

The Intrusion Events by Impact Graph

The Intrusion Events by Impact graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated impact level (from 0 to 4).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the intrusion detection (IDS Statistics) and Intrusion Events tables.

The Top Attackers Graph

The Top Attackers graph, in bar form, displays counts of intrusion events for the top attacking host IP addresses (causing those events) on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Users Graph

The Top Users graph, in bar form, displays users on your monitored network that are associated with the highest intrusion event counts, by event count.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the intrusion detection (IDS) User Statistics and Intrusion Events tables. It displays authoritative user data.

The Intrusion Events by Priority Graph

The Intrusion Events by Priority graph, in pie form, displays a proportional view of intrusion events on your monitored network, grouped by estimated priority level (such as `High`, `Medium`, or `Low`).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Targets Graph

The Top Targets graph, in bar form, displays counts of intrusion events for the top target host IP addresses (targeted in the connections causing those events) on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

This graph draws data primarily from the Intrusion Events table.

The Top Ingress/Egress Security Zones Graph

The Top Ingress/Egress Security Zones graph, in bar form, displays counts of intrusion events associated with each security zone (ingress or egress, depending on graph settings) configured on your monitored network.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only traffic by egress security zone, hover your pointer over the graph, then click **Egress** on the toggle button that appears. Click **Ingress** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Ingress view.

This graph draws data primarily from the Intrusion Events table.

You can configure this graph to display either ingress (the default) or egress security zone information, according to your needs.

The Intrusion Event Details List

At the bottom of the Intrusion Information section is the Intrusion Event Details List, a table that provides classification, estimated priority, and event count information for each intrusion event detected on your monitored network. The events are listed in descending order of event count.

The Intrusion Event Details List table is not sortable, but you can click on any table entry to filter or drill down on that information. This table draws data primarily from the Intrusion Events table.

The Files Information Section

The Files Information section of the Context Explorer contains six interactive graphs that display an overall picture of file and malware events on your monitored network.

Five of the graphs display data related to malware defense (formerly called AMP for Firepower): the file types, file names, and malware dispositions of the files detected in network traffic, as well as the hosts sending (uploading) and receiving (downloading) those files. The final graph displays all malware threats detected in your organization, whether by 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, in donut form, displays a proportional view of the file types detected in network traffic (outer ring), grouped by file category (inner ring).

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events table.

The Top File Names Graph

The Top File Names graph, in bar form, displays counts of the top unique file names detected in network traffic.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events table.

The Files by Disposition Graph

The Top File Types graph, in pie form, displays a proportional view of the malware dispositions for files detected by the malware defense feature (formerly called AMP for Firepower). Note that only files for which the 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.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events table.

The Top Hosts Sending Files Graph

The Top Hosts Sending Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-sending host IP addresses.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only hosts sending malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events table.

The Top Hosts Receiving Files Graph

The Top Hosts Receiving Files graph, in bar form, displays counts of the number of files detected in network traffic for the top file-receiving host IP addresses.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only hosts receiving malware, hover your pointer over the graph, then click **Malware** on the toggle button that appears. Click **Files** to return to the default files view. Note that navigating away from the Context Explorer also returns the graph to the default files view.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events table.

The Top Malware Detections Graph

The Top Malware Detections graph, in bar form, displays counts of the top malware threats detected in your organization, whether by malware defense or Secure Endpoint.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Note that you must have a Malware Defense license to for this graph to display malware defense data.

This graph draws data primarily from the File Events and Malware Events tables.

The Geolocation Information Section

The Geolocation Information section of the Context Explorer contains three interactive donut graphs that display an overall picture of countries with which hosts on your monitored network are exchanging data: unique connections by initiator or responder country, intrusion events by source or destination country, and file events by sending or receiving country.

The Connections by Initiator/Responder Country Graph

The Connections by Initiator/Responder Country graph, in donut form, displays a proportional view of the countries involved in connections on your network as either the initiator (the default) or the responder. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only countries acting as the responder in connections, hover your pointer over the graph, then click **Responder** on the toggle button that appears. Click **Initiator** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Initiator view.

This graph draws data primarily from the Connection Summary Data table.

The Intrusion Events by Source/Destination Country Graph

The Intrusion Events by Source/Destination Country graph, in donut form, displays a proportional view of the countries involved in intrusion events on your network as either the source of the event (the default) or the destination. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only countries acting as the destinations of intrusion events, hover your pointer over the graph, then click **Destination** on the toggle button that appears. Click **Source** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Source view.

This graph draws data primarily from the Intrusion Events table.

The File Events by Sending/Receiving Country Graph

The File Events by Sending/Receiving Country graph, in donut form, displays a proportional view of the countries detected in file events on your network as either sending (the default) or receiving files. The inner ring groups these countries together by continent.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to filter or drill down on that information.

Tip To constrain the graph so it displays only countries receiving files, hover your pointer over the graph, then click **Receiver** on the toggle button that appears. Click **Sender** to return to the default view. Note that navigating away from the Context Explorer also returns the graph to the default Sender view.

This graph draws data primarily from the File Events table.

The URL Information Section

The URL Information section of the Context Explorer contains three interactive bar graphs that display an overall picture of URLs with which hosts on your monitored network are exchanging data: traffic and unique connections associated with URLs, sorted by individual URL, URL category, and URL reputation. You cannot filter on URL information.

Note If you filter on intrusion event information, the entire URL Information Section is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

The Traffic by URL Graph

The Traffic by URL graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the top 15 most requested URLs on your monitored network. For each URL listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the Connection Events table.

The Traffic by URL Category Graph

The Traffic by URL Category graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL categories (such as `Search Engines` or `Streaming Media`) on your monitored network. For each URL category listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL Category graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the URL Statistics and Connection Events tables.

The Traffic by URL Reputation Graph

The Traffic by URL Reputation graph, in bar form, displays counts of network traffic (in kilobytes per second) and unique connections for the most requested URL reputation groups (such as Trusted or Neutral) on your monitored network. For each URL reputation listed, blue bars represent traffic data and red bars represent connection data.

Hover your pointer over any part of the graph to view more detailed information. Click any part of the graph to drill down on that information.

Note If you filter on intrusion event information, the Traffic by URL Reputation graph is hidden.

Note that you must have a URL Filtering license for this graph to include URL category and reputation data.

This graph draws data primarily from the URL Statistics and Connection Events tables.

Requirements and Prerequisites for the Context Explorer

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Refreshing the Context Explorer

The Context Explorer does not automatically update the information it displays. To incorporate new data, you must manually refresh the explorer.

Note that, although reloading the Context Explorer itself (by refreshing the browser program or navigating away from, then back to, the Context Explorer) refreshes all displayed information, this does not preserve any changes you made to section configuration (such as the Ingress/Egress graphs and the Application Information section) and may cause delays in loading.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Click **Reload** at the upper right.

Reload is dimmed until your refresh is finished.

Setting the Context Explorer Time Range

You can configure the Context Explorer time range to reflect a period as short as the last hour (the default) or as long as the last year. Note that when you change the time range, the Context Explorer does not automatically update to reflect the change. To apply the new time range, you must manually refresh the explorer.

Changes to the time range persist even if you navigate away from the Context Explorer or end your login session.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 From the **Show the last** drop-down list, choose a time range.

Step 3 Optionally, to view data from the new time range, click **Reload**.

Tip
Clicking **Apply Filters** also applies any time range updates.

Minimizing and Maximizing Context Explorer Sections

You can minimize and hide one or more sections of the Context Explorer. This is useful if you want to focus on only certain sections, or if you want a simpler view. You cannot minimize the Traffic and Intrusion Event Counts Time Graph.

Context Explorer sections retain the minimized or maximized states that you configure even if you refresh the page or log out of the appliance.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 To minimize a section, click **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.

Drilling Down on Context Explorer Data

If you want to examine graph or list data in more detail than the Context Explorer allows, you can drill down to the table views of the relevant data. (Note that you cannot drill down on the Traffic and Intrusion Events over Time graph.) For example, drilling down on an IP address in the Traffic by Source IP graph displays the Connections with Application Details view of the Connection Events table, including only data associated with the source IP address you selected.

Depending on the type of data you examine, additional options can appear in the context menu. Data points that are associated with specific IP addresses offer the option to view host or whois information on the IP address you select. Data points associated with specific applications offer the option to view application information on the application you select. Data points associated with a specific user offer the option to view that user's user profile page. Data points associated with an intrusion event message offer the option to view the rule documentation for that event's associated intrusion rule, and data points associated with a specific IP address offer the option to add that address to a Block or Do Not Block list. For more information about these lists, see *Global and Domain Security Intelligence 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.

Filters in the Context Explorer

Beyond the basic, wide-ranging data that the Context Explorer initially displays, you have the option to filter that data for a more granular contextual picture of activity on your network. Filters encompass all types of system data except URL information, support exclusion as well as inclusion, can be applied quickly by clicking on Context Explorer graph data points, and affect the entire explorer. You can apply up to 20 filters at a time.

You can add filters to Context Explorer data in several ways:

- from the Add Filter dialog
- from the context menu, when you select a data point in the explorer
- from the text links that appear in certain detail view pages (Application Detail, Host Profile, Rule Detail, and User Profile). Clicking these links automatically opens and filters the Context Explorer according to the relevant data on the detail view page. For example, clicking the Context Explorer link on a user detail page for the user `jenkins` constrains the explorer to show only data associated with that user

Some filter types are incompatible with others: for example, filters that relate to intrusion events (such as **Device** and **Inline Result**) cannot be applied at the same time as connection event-related filters (such as **Access Control Action**) because the system cannot sort connection event data by intrusion event data. The system automatically prevents incompatible filters from simultaneously applying; when one filter type is more recently activated, filters of the incompatible type are hidden as long as the incompatibility exists.

When multiple filters are active, values for the same data type are treated as OR search criteria: all data that matches at least one of the values appears. Values for different data types are treated as AND search criteria: to appear, data must match at least one value for each filtered data type. For example, data that appears for the filter set of `Application: 2channel`, `Application: Reddit`, and `User: edickinson` must be associated with the user `edickinson` **AND** either the application `2channel` **OR** the application `Reddit`.

In a multidomain deployment, you can filter by multiple descendant domains when viewing the Context Explorer in an ancestor domain. In such cases, use caution when also adding **IP Address** filters. The system builds a separate network map for each leaf domain. Using literal IP addresses to constrain this configuration can have unexpected results.

Note that the data displayed depends on such factors as how you license and deploy your managed devices and whether you configure features that provide the data.

Note Filters function as a simple, agile tool to get the precise data context you need at any given time. They are not intended as permanent configuration settings, and disappear when you navigate away from the Context Explorer or end your session. To preserve filter settings for later use, see Saving Filtered Context Explorer Views, on page 591.

Data Type Field Options

The following table lists the data types available as filters, with examples and brief definitions of each.

Table 64: Filter Data Types

Type	Example Values	Definition
Access Control Action	`Allow`, `Block`	Action taken by your access control policy to allow or block traffic.
Application Category	`web browser`, `email`	General classification of an application's most essential function.
Application Name	`Facebook`, `HTTP`	Name of an application.
Application Risk	`Very High`, `Medium`	Estimated security risk of an application.
Application Tag	`encrypts communications`, `sends mail`	Additional information about an application; applications can have any number of tags, including none.
Application Type	`Client`, `Web Application`	Type of an application: application protocol, client, or web application.
Business Relevance	`Very Low`, `High`	Estimated relevance of an application to business activity (as opposed to recreation).
Continent	`North America`, `Asia`	Continent associated with a routable IP address detected on your monitored network.
Country	`Canada`, `Japan`	Country associated with a routable IP address detected on your monitored network.
Device	`device1.example.com`, `192.168.1.3`	Name or IP address of a device on your monitored network.
Domain	`Asia Division`, `Europe Division`	The domain of the device whose network activity you want to graph. This data type is only present in a multidomain deployment.
Event Classification	`Potential Corporate Policy Violation`, `Attempted Denial of Service`	Capsule description of an intrusion event, determined by the classification of the rule, decoder, or preprocessor that triggered it.
Event Message	`dns response`, `P2P`	Message generated by an event, determined by the rule, decoder, or preprocessor that triggered it.

Type	Example Values	Definition
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.

Creating a Filter from the Add Filter Window

Use this procedure to create filters from scratch with the Add Filter window. (You can also use the context menu to create quick filters.)

The Add Filter window, which you access by clicking **Plus** ( ) under **Filters** at the top left of the Context Explorer, contains only two fields:

- The **Data Type** drop-down list contains many different types of data you can use to constrain the Context Explorer. After you select a data type, you then enter a specific value for that type in the **Filter** field (for example, a value of `Asia` for the type **Continent**). To assist you, the Filter field presents several grayed-out example values for the data type you select. (These are erased when you enter data in the field.)
- In the **Filter** field, you can input special search parameters such as `*` and `!` essentially as you can in event searches. You can create exclusionary filters by prefixing filter parameters with the `!` symbol.

Note Filters that you add are not automatically applied; you must click **Apply Filters** to see the filtering in the Context Explorer.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click **Plus** ( ).

Step 3 From the **Data Type** drop-down list, choose the data type you want to filter on.

Step 4 In the **Filter** field, enter the data type value you want to filter on.

Step 5 Click **OK**.

Step 6 Optionally, repeat the previous steps to add more filters until you have the filter set you need.

Step 7 Click **Apply Filters**.

Related Topics

Data Type Field Options, on page 588
Search Constraints, on page 685

Creating a Quick Filter from the Context Menu

While exploring Context Explorer graph and list data, you can click on data points, then use the context menu to quickly create a filter based on that data, either inclusive or exclusive. If you use the context menu to filter on information of data type Application, User, or Intrusion Event Message, or any individual host, the filter widget includes a widget information that links to the relevant detail page for that data type (such as Application Detail for application data). Note that you cannot filter on URL data.

You can also use the context menu to investigate specific graph or list data in more detail.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 In any explorer section except Traffic and Intrusion Events over Time or sections that contain URL data, click a data point you want to filter on.

Step 3 You have two options:

- To add a filter for this data, click **Add Filter**.
- To add an exclusion filter for this data, click **Add Exclude Filter**. The filter, when applied, displays all data **not** associated with the excluded value. Exclude filters display an exclamation point (!) before the filter value.

Saving Filtered Context Explorer Views

To preserve filter settings in the Context Explorer after you navigate away from the Context Explorer or end your session, create a browser bookmark of the Context Explorer with your preferred filters applied. Because applied filters are incorporated in the Context Explorer page URL, loading a bookmark of that page also loads the corresponding filters.

Procedure

Create a browser bookmark of the Context Explorer with your preferred filters applied.

Viewing Filter Data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Click **Information** on any eligible filter widget.

Deleting a Filter

Procedure

Step 1 Choose **Analysis** > **Context Explorer**.

Step 2 Under **Filters** at the top left, click **Close** (×) to delete the filter widget individually.

Tip
If you want to delete all filters at once, you can click **Clear**.

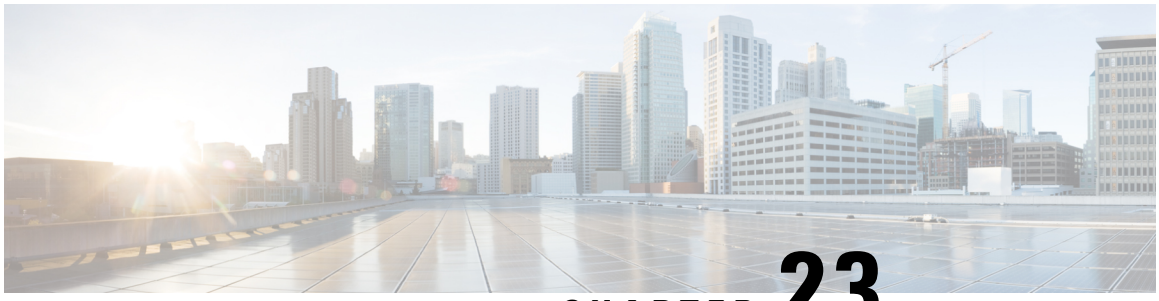

CHAPTER 23

Unified Events

The following topics describe how to use the Unified Events:

About the Unified Events

Unified Events provide you a single-screen view of multiple types (connection, intrusion, file, malware, and some security-related connection events) of firewall events. Events associated with each other are stacked together in the table to provide a unified view and more context about the security event. If you have an intrusion event on the Unified Events table, click the intrusion event to highlight the associated connection event. You can then correlate the connection event with the intrusion event to better understand and troubleshoot the network issues, without toggling between multiple event viewers.

The Unified Events table is highly customizable. You can create and apply custom filters to fine-tune the information displayed on the event viewer. Unified events also has option to save the custom filters that you use often for specific needs, and then quickly load the saved filters. Also, you can make a tailored event viewer table by adding or removing columns, pin columns, or drag and re-order the columns.

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. For example, if you are a firewall administrator, viewing event updates in real time after you make a policy change can help you to ensure that the policy changes are correctly enforced on your network.

Requirements and Prerequisites for the Unified Events

Model Support

Any.

Supported Domains

Any.

User Roles

- Admin
- Security Analyst

Working 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, see Set a Time Range in Unified Events, on page 597.

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, see Filters in Unified Events, on page 598.

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 indicates 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. • Save column sets to quickly reload your customized view later. For more information, see Save a Column Set, on page 600 topic. 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, see Use the Packet Tracer, on page 445.
View events in real time	Click **Go Live**. For more information, see View Live Events in Unified Events, on page 598. 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, see Event Investigation Using Web-Based Resources, on page 620.
Open multiple unified events windows	• You can display different views of the unified events table using multiple browser tabs or windows. • Each new tab or window has the characteristics of the most recently modified tab/window. • To make any open tab/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.
Save searches	Save custom searches as your favorites and quickly load them later. For more information, see Save a Search in Unified Events, on page 599.

To Do This ...	Do This
Bookmark or share query results	Bookmark or copy-paste the URL in the browser window. • The URL retrieves different events later if it used 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

Configure time range in unified events to view firewall events for a specific period. 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.

Important 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-Related Connection 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**.

  Tip
  Click **Now** to quickly set the current time as the **End time**.

- If you want to configure a sliding default time window of the length you specify, click **Sliding Time Range**.

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

Step 4 Click **Apply**.

View Live Events in Unified Events

Configure the unified events to display firewall events in real time without manually refreshing the event viewer. In the **Live View** mode, the event logs appear in real time as the security event occurs in your network which helps you to better troubleshoot the issues.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

By default, the unified events table displays events from the last hour.

Step 2 To see live event updates, click **Go Live**.

New events get populated at the top of the event table. The time range section displays a timer to inform you for how long the unified events table is live.

Note
When using the **Go Live** feature, the following limitation applies for the UDP traffic:

- By default, the **Go Live** feature in management center considers traffic data from the last 30 seconds, which is shorter than the 120 seconds required for UDP connections to be processed into unified events. This may result in incomplete event logging for UDP traffic.
- To improve visibility, configure logging at the beginning of the connection for UDP traffic.

What to do next

To exit the live view mode, click **Live**.

Filters in Unified Events

The unified events table initially displays multiple types of firewall events from the past hour. You can filter the default view of unified events for a more granular contextual picture of activity on your network. Filters support exclusion as well as inclusion filter criteria.

Filters help you to provide quick access to critical information. For example, if you are a firewall administrator and you want to allow or deny specific application access to some users, you can set user search criteria to scan through the firewall logs. The event viewer displays event logs that match the search criteria.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Enter the filter criteria:

- To manually enter the filter criteria, type the exact criteria in the search text field, or select the criteria from the drop-down list. Then, provide the filter criteria value. While typing in the values, you are prompted with suggestions in the drop-down list whenever possible.
- Click the dots in a cell for an event in the table 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 important information about wildcards and search behavior, see Event Searches, on page 685.
- Include operators (such as <, >, !, and so on) 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.

What to do next

To save a custom filter, see Save a Search in Unified Events, on page 599 topic.

Save a Search in Unified Events

Save custom searches as your favorites and quickly load them later. Note that this option is not available for the **Troubleshooting** table.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the **Events** tab.

Step 3 Establish a search criteria as described in the Filters in Unified Events, on page 598 topic.

Step 4 Click the **Search** (🔍) icon.

Step 5 Click **+ Save search**.

Step 6 Specify a name for the search and click the **OK** (✔) icon.

What to do next

To load a saved search, see Load a Saved Search in Unified Events, on page 600 topic.

Load a Saved Search in Unified Events

Before you begin

Establish a saved search as described in the Save a Search in Unified Events, on page 599 topic.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the **Search** (🔍) icon on the search text box and choose the saved search that you want to load.

Step 3 Alternatively, click the **Open saved search details panel** icon and choose the saved search to view the filter conditions for the selected search. Click **Apply search** to load the search.

Save a Column Set

Save custom column sets as your favorites to load them later or quickly toggle between custom tables. Note that this option is not available for the **Troubleshooting** table.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the column picker Icon (▥) and choose the set of columns that you want to save.

Step 3 From the **Saved Column Sets** drop-down, click **+ Create column set from current selection**.

Step 4 Specify a name for the column set and click the **OK** (✔) icon.

What to do next

To load a saved column set, see Load a Saved Column Set, on page 600 topic.

Load a Saved Column Set

Before you begin

Save a favorite column set as described in the Save a Column Set, on page 600 topic.

Procedure

Step 1 Choose **Analysis** > **Unified Events**.

Step 2 Click the column picker icon ( ).

Step 3 From the **Saved Column Sets** drop-down, choose the column set that you want to load.

Unified Events Column Descriptions

Values in some fields depend on the event type. Field correspondences for the default fields are as follows:

Unified Events Field Name	Connection or Security Intelligence Event Field Name	Intrusion Event Field Name	File Event Field Name	Malware Event Field Name
Time	First Packet See note below.	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
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

Click the column picker ( ) icon to see all event fields and their correspondences.

For field descriptions, see the following topics:

- Connection and Security-Related Connection Event Fields, on page 731

- Intrusion Event Fields, on page 766
- File and Malware Event Fields, on page 815

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

Note Even if you have not enabled logging at the beginning of the connection, the system has and uses this value as the time field in the unified events table. To determine whether a connection event was logged at the beginning and end of the connection, expand the event's row to view details. If both ends of the connection were logged, you see a **Last Packet** field.

History for Unified Events

Feature	Minimum Management Center	Minimum 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**(⋮)) (Expand) 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 Verion 7.3.x or 7.4.0.
Save your favorite searches	7.3	Any	Save column sets and searches as your favorites and later launch them quickly.
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: management center

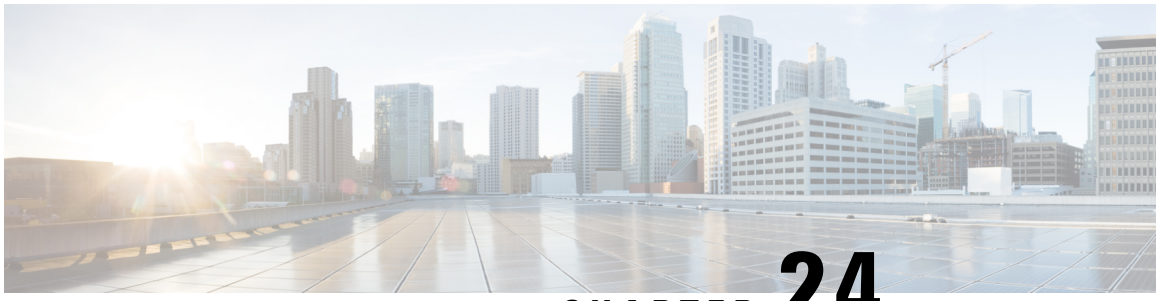

CHAPTER 24

Network Map

The following topics describe how to use the network map:

Requirements and Prerequisites for the Network Map

Model Support

Any.

Supported Domains

Leaf

User Roles

- Admin
- Discovery Admin

The Network Map

The system monitors traffic traveling over your network, decodes the traffic data, and then compares the data to established operating systems and fingerprints. The system then uses this data to build a detailed representation of your network, called a *network map*. In multidomain deployments, the system creates an individual network map for each leaf domain.

The system gathers data from the managed devices identified for monitoring in the network discovery policy. The managed devices detect network assets directly from monitored traffic and indirectly from processed NetFlow records. If multiple devices detect the same network asset, the system combines the information into a composite representation of the asset.

To augment data from passive detection, you can:

- Actively scan hosts using the open-source scanner, Nmap™, and add the scan results to your network map.
- Manually add host data from a third-party application using the host input feature.

The network map displays your network topology in terms of detected hosts and network devices.

You can use the network map to:

- Obtain a quick, overall view of your network.
- Select different views to suit the analysis you want to perform. Each view of the network map has the same format: a hierarchical tree with expandable categories and sub-categories. When you click a category, it expands to show you the sub-categories beneath it.
- Organize and identify subnets via the custom topology feature. For example, if each department in your organization uses a different subnet, you can assign familiar labels to those subnets using the custom topology feature.
- View detailed information by drilling down to any monitored host's *host profile*.
- Delete an asset if you are no longer interested in investigating it.

Note If the system detects activity associated with a host you deleted from a network map, it re-adds the host to the network map. Similarly, deleted applications are re-added to the network map if the system detects a change in the application (for example, if an Apache web server is upgraded to a new version). Vulnerabilities are reactivated on specific hosts if the system detects a change that makes the host vulnerable.

Tip If you want to permanently exclude a host or subnet from the network map, modify the network discovery policy. You may wish to exclude load balancers and NAT devices from monitoring if you find that they are generating excessive or irrelevant events.

The Hosts Network Map

The network map on the Hosts tab displays a host count and a list of host IP addresses and primary MAC addresses. Each address or partial address is a link to the next level. This network map view provides a count of all unique hosts detected by the system, regardless of whether the hosts have one IP address or multiple IP addresses.

Use the hosts network map to view the hosts on your network, organized by subnet in a hierarchical tree, as well as to drill down to the host profiles for specific hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data.

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 845.

The Network Devices Network Map

The network map on the Network Devices tab displays the network devices (bridges, routers, NAT devices, and load balancers) that connect one segment of your network to another. The map contains two sections listing devices identified by an IP address and devices identified by a MAC address.

The map also provides a count of all unique network devices detected by the system, regardless of whether the devices have one IP address or multiple IP addresses.

If you create a custom topology for your network, the labels you assign to your subnets appear in the network devices network map.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their types (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a network device communicates using CDP, it may have one or more IP addresses. If it communicates using STP, it may only have a MAC address.

You cannot delete network devices from the network map, because the system uses their locations to determine network topology.

The host profile for a network device has a Systems section rather than an Operating Systems section, which includes a Hardware column that reflects the hardware platform for any mobile devices detected behind the network device. If a value for a hardware platform is listed under Systems, that system represents a mobile device or devices detected behind the network device. Note that mobile devices may or may not have hardware platform information, but hardware platform information is never detected for systems that are not mobile devices.

The Mobile Devices Network Map

The network map on the Mobile Devices tab displays mobile devices attached to your network. This network map also provides a count of all unique mobile devices detected by the system, regardless of whether the devices have one IP address or multiple IP addresses.

Each address or partial address is a link to the next level. You can also delete a subnet or IP address; if the system rediscovers the device, it re-adds the device to the network map.

You can also drill down to view the host profiles for the mobile devices.

To identify mobile devices, the system:

- analyzes User-Agent strings in HTTP traffic from the mobile device's mobile browser
- monitors the HTTP traffic of specific mobile applications

If you create a custom topology for your network, the labels you assign to your subnets appear in the mobile devices network map.

The Indications of Compromise Network Map

The network map on the Indications of Compromise tab displays the compromised hosts on your network, organized by IOC category. Affected hosts are listed beneath each category. Each address or partial address is a link to the next level.

From the indications of compromise network map, you can view the host profile of each host determined to have been compromised in a specific way. You can also delete (mark as resolved) any IOC category or any specific host, which removes the IOC tag from the relevant hosts. For example, you can delete an IOC category from the network map if you have determined that the issue is addressed and unlikely to recur.

Marking a host or IOC category resolved from the network map does not remove it from your network. A resolved host or IOC category reappears in the network map if your system newly detects information that triggers that IOC.

For more information about how the system determines indications of compromise, see Indications of Compromise Data, on page 892 and subtopics.

The Application Protocols Network Map

The network map on the Application Protocols tab displays the applications running on your network, organized in a hierarchical tree by application name, vendor, version, and finally by the hosts running each application.

The applications that the system detects may change with system software and VDB updates, and if you import any add-on detectors. The release notes or advisory text for each system or VDB update contains information on any new and updated detectors. For a comprehensive up-to-date list of detectors, see the Cisco Support Site (http://www.cisco.com/cisco/web/support/index.html).

From this network map, you can view the host profile of each host that runs a specific application.

You can also delete any application category, any application running on all hosts, or any application running on a specific host. For example, you can delete an application from the network map if you know it is disabled on the host and you want to make sure the system does not use it for impact level qualification.

Deleting an application from the network map does not remove it from your network. A deleted application reappears in the network map if your system detects a change in the application (for example, if an Apache web server is upgraded to a new version) or if you restart your system's discovery function.

Depending on what you delete, the behavior differs:

- Application Category — Deleting removes the application category from the network map. All applications that reside beneath the category are removed from any host profile that contains the applications.

For example, if you delete **http**, all applications identified as **http** are removed from all host profiles and **http** no longer appears in the applications view of the network map.

- Specific Application, Vendor, or Version — Deleting removes the affected application from the network map and from any host profiles that contain it.

  For example, if you expand the **http** category and delete **Apache**, all applications listed as Apache with any version listed beneath Apache are removed from any host profiles that contain them. Similarly, if instead of deleting **Apache**, you delete a specific version (**1.3.17**, for example), only the version you selected will be deleted from affected host profiles.

- Specific IP Address — Deleting the IP address removes it from the application list and removes the application itself from the host profile of the IP address you selected.

  For example, if you expand **http**, **Apache**, **1.3.17 (Win32)**, and then delete **172.16.1.50:80/tcp**, the Apache 1.3.17 (Win32) application is deleted from the host profile of IP address 172.16.1.50.

The Vulnerabilities Network Map

The network map on the Vulnerabilities tab displays vulnerabilities that the system has detected on your network, organized by legacy vulnerability ID (SVID), CVE ID, or Snort ID.

From this network map, you can view the details of specific vulnerabilities, as well as the host profile of any host subject to a specific vulnerability. This information can help you evaluate the threat posed by that vulnerability to specific affected hosts.

If you determine that a specific vulnerability is not applicable to the hosts on your network (for example, you have applied a patch), you can deactivate the vulnerability. Deactivated vulnerabilities still appear on the network map, but the IP addresses of their previously affected hosts appear in gray italics. The host profiles for those hosts show deactivated vulnerabilities as invalid, though you can manually mark them as valid for individual hosts.

If there is an identity conflict for an application or operating system on a host, the system lists the vulnerabilities for both potential identities. When the identity conflict is resolved, the vulnerabilities remain associated with the current identity.

By default, the network map displays the vulnerabilities of a detected application only if the packet contains the application's vendor and version. However, you can configure the system to list the vulnerabilities for applications lacking vendor and version data by enabling the vulnerability mapping setting for the application in the management center configuration.

The numbers next to a vulnerability ID (or range of vulnerability IDs) represent two counts:

Affected Hosts

The first number is a count of non-unique hosts that are affected by a vulnerability or vulnerabilities. If a host is affected by more than one vulnerability, it is counted multiple times. Therefore, it is possible for the count to be higher than the number of hosts on your network. Deactivating a vulnerability decrements this count by the number of hosts that are potentially affected by the vulnerability. If you have not deactivated any vulnerabilities for any of the potentially affected hosts for a vulnerability or range of vulnerabilities, this count is not displayed.

Potentially Affected Hosts

The second number is a count of the total number of non-unique hosts that the system has determined are *potentially* affected by a vulnerability or vulnerabilities.

Deactivating a vulnerability renders it inactive only for the hosts you designate. You can deactivate a vulnerability for all hosts that have been judged vulnerable or for a specified individual vulnerable host. After a vulnerability is deactivated, the applicable hosts' IP addresses appear in gray italics in the network map. In addition, host profiles for those hosts show deactivated vulnerabilities as invalid.

If the system subsequently detects the vulnerability on a host where it has not been deactivated (for example, on a new host in the network map), the system activates the vulnerability for that host. You have to explicitly deactivate the newly discovered vulnerability. Also, if the system detects an operating system or application change for a host, it may reactivate associated deactivated vulnerabilities.

The Host Attributes Network Map

The network map on the Host Attributes tab displays the hosts on your network organized by either user-defined or compliance allow list host attributes. You cannot organize hosts using predefined host attributes in this display.

When you choose the host attribute you want to use to organize your hosts, the management center lists the possible values for that attribute in the network map and groups hosts based on their assigned values. For example, if you choose to organize your hosts by allow list host attributes, the system displays them in categories of Compliant, Non-Compliant, and Not Evaluated.

You can also view the host profile of any host assigned a specific host attribute value.

Related Topics

Host Attributes in the Host Profile, on page 859

Viewing Network Maps

You must be an Admin or Security Analyst user 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, choose a leaf domain from the **Domain** drop-down list.
- Filter Hosts — If you want to filter by IP or MAC addresses, enter an address into the search field. To clear the search, click **Clear** ( ).
- Drill Down — If you want to investigate a category or host profile, drill down through the categories or subnets in the map. If you have defined a custom topology, click **(topology)** from **Hosts** to view it, then click on **(hosts)** if you want to toggle back to the default view.
- Delete — Click **Delete** ( ) next to the appropriate element to:
  - Remove an element from the map on **Hosts**, **Network Devices**, **Mobile Devices**, or **Application Protocols**.
  - Mark an IOC category, compromised host, or group of compromised hosts resolved on **Indications of Compromise**.

- Deactivate a vulnerability for all hosts or a single host on **Vulnerabilities**.

- Specify Vulnerabilities Class — On **Vulnerabilities**, choose the class of vulnerabilities you want to view from the **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 609
Host Profiles, on page 844

Custom Network Topologies

Use the custom topology feature to help you organize and identify subnets in your hosts and network devices network maps.

For example, if each department within your organization uses a different subnet, you can label those subnets using the custom topology feature.

You can also view the hosts network map according to the organization you specified in the custom topology.

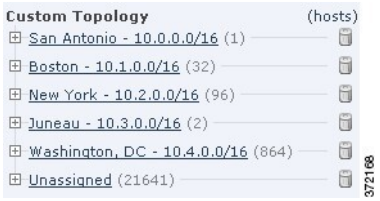

You can specify a custom topology's networks using any or all of the following strategies:

- You can import networks from the network discovery policy to add the networks that you configured the system to monitor.
- You can add networks to your topology manually.

The Custom Topology page lists your custom topologies and their status. If the light bulb icon next to the policy name is lit, the topology is active and affects your network map. If it is dimmed, the topology is inactive.

Related Topics

The Hosts Network Map, on page 604
The Network Devices Network Map, on page 605

Creating Custom Topologies

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology** in the toolbar.

Step 3 Click **Create Topology**.

Step 4 Enter a **Name**.

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

Step 6 Add networks to your topology. You can use any or all of the following strategies:

- Import networks from a network discovery policy as described in Importing Networks from the Network Discovery Policy, on page 610.
- Manually add networks as described in Manually Adding Networks to Your Custom Topology, on page 611.

Step 7 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 611.

Importing Networks from the Network Discovery Policy

Procedure

Step 1 Access the custom topology to which you want to import the network:

- Create a custom topology; see Creating Custom Topologies, on page 609.
- Edit an existing custom topology; see Editing Custom Topologies, on page 611.

Step 2 Click **Import Policy Networks**.

Step 3 Click **Load**. The system displays the topology information for the network discovery policy.

Step 4 Refine your topology:

- Rename a network in the topology by clicking **Edit** ( ) next to the network, typing a name, and clicking **Rename**.
- Remove a network from the topology by clicking **Delete** ( ) and then clicking **OK** to confirm.

Step 5 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 611.

Manually Adding Networks to Your Custom Topology

Procedure

Step 1 Access the custom topology where you want to add the network:

- Create a custom topology; see Creating Custom Topologies, on page 609.
- Edit an existing custom topology; see Editing Custom Topologies, on page 611.

Step 2 Click **Add Network**.

Step 3 If you want to add a custom label for the network in the hosts and network devices network maps, type a **Name**.

Step 4 Enter the **IP Address** and **Netmask** (IPv4) that represent the network you want to add.

Step 5 Click **Add**.

Step 6 Click **Save**.

What to do next

- Activate the topology as described in Activating and Deactivating Custom Topologies, on page 611.

Related Topics

IP Address Conventions, on page 25

Activating and Deactivating Custom Topologies

Note Only one custom topology can be active at any time. If you have created multiple topologies, activating one automatically deactivates the currently active topology.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Choose **Custom Topology**.

Step 3 Click the slider next to a topology to activate or deactivate it.

Editing Custom Topologies

Changes you make to an active topology take effect immediately.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 2 Click **Custom Topology**.

Step 3 Click **Edit** ( ) next to the topology you want to edit.

Step 4 Edit the topology as described in Creating Custom Topologies, on page 609.

Step 5 Click **Save**.

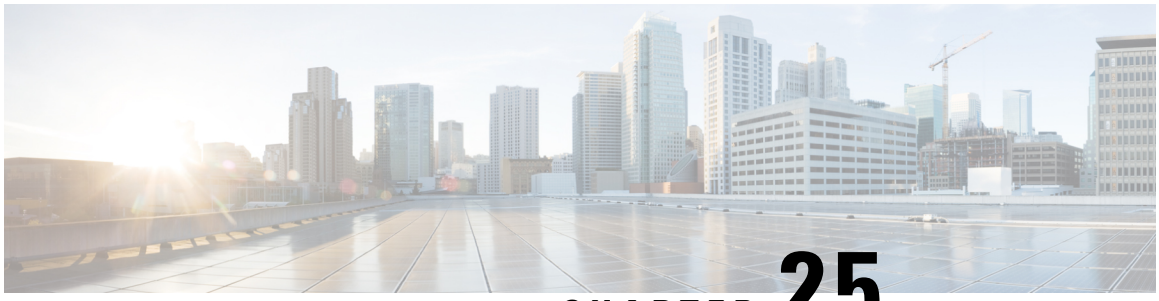

CHAPTER 25

Lookups

The following topics explain how to look up information about entities that may or may not be known to the system:

Introduction to Lookups

If your management center is connected to the Internet, you can use manual lookup features to find the following information:

- Regional Information Registries (RIR) information (whois) for any IP address.
- URL category and reputation as classified by the URL Filtering feature.
- Geolocation information for any IP address: country name, country code, and continent name. (To ensure that you are using up-to-date geolocation information, Cisco strongly recommends that you regularly update the Geolocation Database (GeoDB) on your management center.)

Performing Whois Lookups

Before you begin

- Ensure that the management center has Internet access; see Security, Internet Access, and Communication Ports, on page 1027.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Whois**.

Step 2 Enter an IP address and click **Search**.

Finding URL Category and Reputation

You can manually look up category and reputation of URLs. Use this feature to see how particular URLs are evaluated in order to plan, adjust, or troubleshoot policy processing, or to investigate potentially problematic URLs that come to your attention via sources outside your Cisco solution. The categories and reputations in these results are the same as those that are used by the URL Filtering feature.

Before you begin

- The management center must have Internet access; see Security, Internet Access, and Communication Ports, on page 1027.
- 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 management center and have a valid URL Filtering license assigned to it.
- You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Select **Analysis** > **Advanced** > **URL**.

Step 2 Enter up to 250 URLs and public, routable IP addresses, in any common format (for example, URLs may be with or without "http", "www", or a subdomain, or may be shortened). Separate each entity with a space or a return.

Wildcards such as asterisks (*) are not supported.

Step 3 Click **Search**.

If you enter many URLs and your network is slow, processing may take several minutes.

If you see an error message that the URL is not valid, check your spelling or try a different variation of the URL. For example, add or omit the "www" or "http" or "https" prefix.

A URL may belong to up to six categories but has only one reputation.

Step 4 (Optional) Sort the results by clicking a column heading.

Step 5 (Optional) To save the results as a CSV file, click **Export CSV**.

An additional column for reputation level is included in the CSV file so you can sort by risk. Zero (0) represents an unknown risk, for a URL for which the system has insufficient risk data.

What to do next

If you want to view lists of possible categories and reputations, go to **Policies > Access Control > Access Control**, click a policy or add a new one, click **Add Rule**, then click **URLs**.

Finding Geolocation Information for an IP Address

You can use the geolocation lookup feature to find the country name, ISO 3166-1 three-digit country code, and continent name associated with any IP address.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Geolocation**.

Step 2 To view the geolocation information for one or more IP addresses, enter the address or addresses and click **Search**. You may specify IPv4 addresses, IPv6 addresses, or both. Use a comma, semicolon, return, or any white space character to separate multiple addresses.

Tip
Click **Clear** to clear the text box.

Step 3 Optionally, click the column titles to sort the data. You can sort by any field except IP Address.

Step 4 (Optional) To save the results as a CSV file, click **Export CSV**.

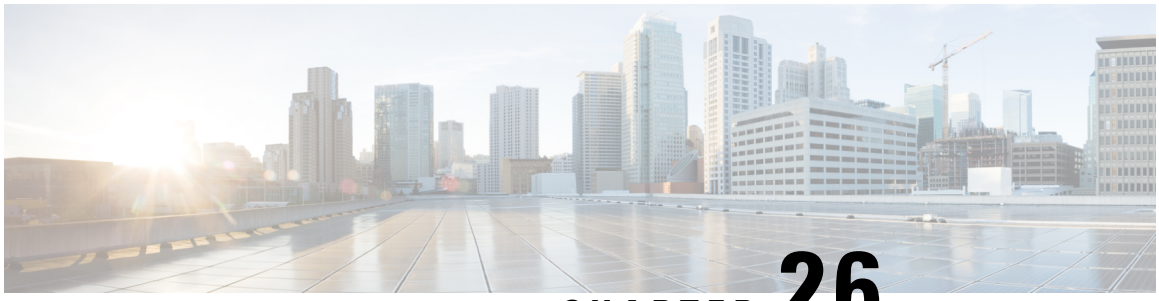

CHAPTER 26

Event Analysis Using External Tools

Cisco Cloud Event Settings

Sending firewall events to the cloud allows you to use external tools to investigate the firewall incidents. The devices send firewall events to the Security Services Exchange (SSE), from where they can be forwarded to various cloud services to unify visibility and enhance your threat investigations.

To allow your devices to send firewall events to Cisco Security Cloud, you must either register the management center with the smart license (**System** (⚙) > **Smart License**) or enable SecureX integration. Cisco Security Cloud integration associates the 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 management center with Cisco Security Cloud, see Enable SecureX Integration, on page 40.

Security Services Exchange Event Consolidation

The Security Services Exchange does not display the complete list of events from the management center. Instead, it correlates and consolidates events, presenting only unique events. This approach reduces redundancy of events and enhances clarity. The current categorization parameters used for this consolidation are detailed as follows:

- For identifying duplication of intrusion events, the following elements are considered: Initiator IP, Initiator IP, SID, and GID.
- For identifying duplication of connection events and security-related connection events, the following elements are considered: Initiator IP, Initiator IP, and Security Intelligence Category.
- For identifying duplication of file and malware events, all elements except Event Second are considered.

Enable Sending Events to the Cisco Security Cloud

Configure your management center to have the managed threat defense devices send events directly to Cisco Security Cloud. The cloud region and event types that you configure in this page can be used for multiple integrations when applicable and enabled.

Before you begin

- Ensure that you register the management center with the Smart License (**System** (⚙) > **Smart License**) or enable Cisco Security Cloud integration to allow your devices to send firewall events to the Cisco cloud.
- In the management center:
  - Go to the **System > Configuration** page and give your management center a unique name to clearly identify it in the **Devices** list in the cloud.
  - Add your threat defense devices to the management center, assign licenses to them, and ensure that the system is working correctly. Ensure that you have created the necessary policies and the generated events are displayed as expected in the management center UI under the **Analysis** menu.
- Ensure that you have your Cisco 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, see Regional Clouds.
- If you are currently sending events to the cloud using syslog, disable it to avoid duplication.

Procedure

Step 1 Determine the regional cloud you want to use for sending firewall events. For more information for choosing a regional cloud, see Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide.

Note
If SecureX integration is enabled and the 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 management center, click **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 the following table.

Integration	Supported Event Options	Notes
Cisco Security Analytics and Logging (SaaS)	All	High-priority connection events include: • Security-related 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, see Cisco XDR Licenses.

Note

- When you enable **Intrusion Events**, the 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 threat defense devices, the 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, and enables security teams to detect, prioritize, and respond to the most sophisticated threats. Integrate threat defense with Cisco XDR to connect Cisco's integrated security portfolio and your firewall deployment for a consistent experience that unifies visibility, enables automation, and strengthens your security across network.

For more information about Cisco XDR, see 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, see 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 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 threat defense with Cisco XDR, see the 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, see the Frequently Asked Questions.

Event Investigation Using Web-Based Resources

Use the contextual cross-launch feature to quickly find more information about potential threats in web-based resources outside of the Secure Firewall Management Center. For example, you might:

- Look up a suspicious source IP address in a Cisco or third-party cloud-hosted service that publishes information about known and suspected threats, or
- 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.

For example, 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 web site opens to a page with information about this specific IP address.

You can choose from a set of pre-defined links to commonly used Cisco and third-party threat intelligence services, and add custom links to other web-based services, and to SIEMs or other products that have a web interface. Note that some resources may require an account or a product purchase.

About Managing 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 623.

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, see Add Contextual Cross-Launch Resources, on page 621.

Requirements for Custom Contextual Cross-Launch Resources

When adding custom contextual cross-launch resources:

- 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

You can add contextual cross-launch resources such as threat intelligence services and Security Information and Event Management (SIEM) tools.

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).
- See Requirements for Custom Contextual Cross-Launch Resources, on page 621.
- 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

Before you begin

If the resource you will access requires credentials, make sure you have those credentials.

Procedure

Step 1 Navigate to one of the following pages in the Secure Firewall Management Center that shows events:

- A dashboard (**Overview > Dashboards**), 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 threat defense to related data in your Secure Network Analytics appliance. For more information about the Secure Network Analytics product, see Cisco Security Analytics and Logging product page.

For general information about contextual cross-launching, see Investigate Events Using Contextual Cross-Launch, on page 622.

Use this procedure to configure a set of cross-launch links to your Secure Network Analytics appliance.

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 621, but those links remain independent of the auto-created resources and you must manage them manually.

Before you begin

- You must have a deployed and running Secure Network Analytics appliance.
- If you are currently using syslog to send events to Secure Network Analytics from device versions that support sending events directly, disable syslog for those devices (or assign those devices an access control policy that does not include syslog configurations) to avoid duplicate events on the remote volume.
- You must have the following:
  - 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 threat defense data to your Secure Network Analytics appliance using Security Analytics and Logging (On Premises), see Remote Data Storage on a Secure Network Analytics Appliance, on page 519.

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, see the 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.

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 management center event viewer or dashboard, right-click a relevant event's table cell and choose the appropriate option.

It may take some time to process the queries, depending on the amount of data to process, speed of and demand on the Secure Network Analytics Manager, and so on.

About Sending Syslog Messages for Security Events

You can send 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.

These events are also sometimes referred to as Snort® events.

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.

About Configuring the System to Send Security Event Data to Syslog

In order to configure the system to send security event syslogs, you will need to know the following:

- Best Practices for Configuring Security Event Syslog Messaging, on page 625
- Configuration Locations for Security Event Syslogs, on page 629

- *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 changes to take effect.

Best Practices for Configuring Security Event Syslog Messaging

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 threat defense platform settings: (**Devices > Platform Settings > Threat Defense Settings > Syslog**.) **a.** Click **Devices** > **Platform Settings**. **b.** Edit the threat defense settings policy. **c.** In the left navigation pane, click **Syslog**. See also *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 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, see Send Security Event Syslog Messages from Threat Defense Devices, on page 625.
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, see Send Security Event Syslog Messages from Classic Devices, on page 628.

Send Security Event Syslog Messages from Threat Defense Devices

This procedure documents the best practice configuration for sending syslog messages for security events (connection, Security Intelligence, intrusion, file, and malware events) from threat defense devices.

Note Many 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 important information about connection logging, see the chapter on Connection Logging, on page 713.

Procedure

Step 1 Configure syslog settings for your threat defense device:

a) Click **Devices > Platform Settings**.
b) **Edit** the platform settings policy associated with your threat defense device.
c) In the left navigation pane, click **Syslog**.
d) Click **Syslog Servers** and click **Add** (+) to enter server, protocol, interface, and related information.

If you have questions about options on this page, see 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**.
b) Edit the applicable access control policy.
c) Click **More** > **Logging**.
d) 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 the **Access Control** > **Add Rule**.
b) Select a rule to edit.
c) Click the **Logging** tab in the rule.
d) 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.)

e) If you want to log file events, select **Log Files**.
f) Enable **Syslog Server**.
g) Verify that the rule is "**Using default syslog configuration in Access Control Logging**."
h) Click **Confirm**.
i) 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, click **Advanced Settings** > **Syslog Alerting** > **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, see Syslog Alert Facilities, on page 555.
Severity	This setting is applicable only if you specify a Logging Host on this page. For descriptions, see Syslog Severity Levels, on page 556.

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.

  See the applicable table rows in Configuration Locations for Syslogs for Connection and Security Intelligence Events (All Devices), on page 629.

  These settings will require syslog alert responses, which are configured as described in Creating a Syslog Alert Response, on page 554. They do not use the platform settings you configured in this procedure.

- To configure security event syslog logging for Classic devices, see Send Security Event Syslog Messages from Classic Devices, on page 628.

- If you are done making changes, deploy your changes to managed devices.

Send Security Event Syslog Messages from Classic Devices

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, see the chapter on Connection Logging, on page 713.

Procedure

Step 1 Configure an alert response for your Classic devices:

See Creating a Syslog Alert Response, on page 554.

Step 2 Configure syslog settings in the access control policy:

a) Click **Policies** > **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 you will send file and malware events:

a) Select **Send Syslog messages for File and Malware events**.
b) Click **Save**.

Step 4 If you will send intrusion events:

a) Navigate to the intrusion policy associated with your access control policy.
b) In your intrusion policy, click **Advanced Settings** > **Syslog Alerting** > **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. See Syslog Alert Facilities, on page 555.
Severity	This setting is applicable only if you specify a Logging Host on this page. See Syslog Severity Levels, on page 556.

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 access control rules. See the applicable table rows in Configuration Locations for Syslogs for Connection and Security Intelligence Events (All Devices), on page 629. These settings will require syslog alert responses, which are configured as described in Creating a Syslog Alert Response, on page 554. They do not use the settings you configured above.
- To configure security event syslog logging for threat defense devices, see Send Security Event Syslog Messages from Threat Defense Devices, on page 625.

Configuration Locations 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 713.

Configuration Location	Description and More Information
Devices > Platform Settings, Threat Defense Settings policy, **Syslog**	This option applies only to 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, <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 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 Creating a Syslog Alert Response, on page 554. For more information about the settings on the Logging tab, see Cisco Secure Firewall Management Center Device Configuration Guide.
Policies > 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 727.
Policies > 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.
Policies > 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 > 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 727.

Configuration Location	Description and More Information
Policies > SSL, <each policy>, <each rule>, **Logging**	Logging settings for SSL rules. See Cisco Secure Firewall Management Center Device Configuration Guide.
Policies > 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 727.
Policies > 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 > 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 threat defense cluster configurations:	The Cisco Secure Firewall Management Center Device Configuration Guide has multiple references to syslog; search the chapter for "syslog."

Configuration Locations for Syslogs for Intrusion Events (Threat Defense Devices)

You can specify syslog settings for intrusion policies in various places and, optionally, inherit settings from the access control policy or the Threat Defense Platform Settings or both.

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. See Cisco Secure Firewall Management Center Device Configuration Guide.
Policies > 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.

Configuration Location	Description and More Information
Policies > 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, see Configuring Syslog Alerting for Intrusion Events, on page 564. 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** > **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 Threat Defense)

- (Default) Access control policy Cisco Secure Firewall Management Center Device Configuration Guide, IF you specify a syslog alert (See Creating a Syslog Alert Response, on page 554.)
- Or see Configuring Syslog Alerting for Intrusion Events, on page 564.

By default, the intrusion policy uses the settings in the Logging tab of the access control policy. If settings applicable to devices other than threat defense are not configured there, syslogs will not be sent for devices other than threat defense and no warning appears.

Configuration Locations for Syslogs for File and Malware Events

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 Threat Defense Platform Settings, you must also create an alert response. See Creating a Syslog Alert Response, on page 554.
In Threat Defense Platform Settings: **Devices > Platform Settings**, Threat Defense Settings policy, **Syslog**	These settings apply only to threat defense devices running supported versions, and only if you configure the Logging tab in the access control policy to use 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 Threat Defense Platform Settings, you must also create an alert response. See Creating a Syslog Alert Response, on page 554.

Anatomy of Security Event Syslog Messages

Example security event message from Threat Defense (Intrusion Event)

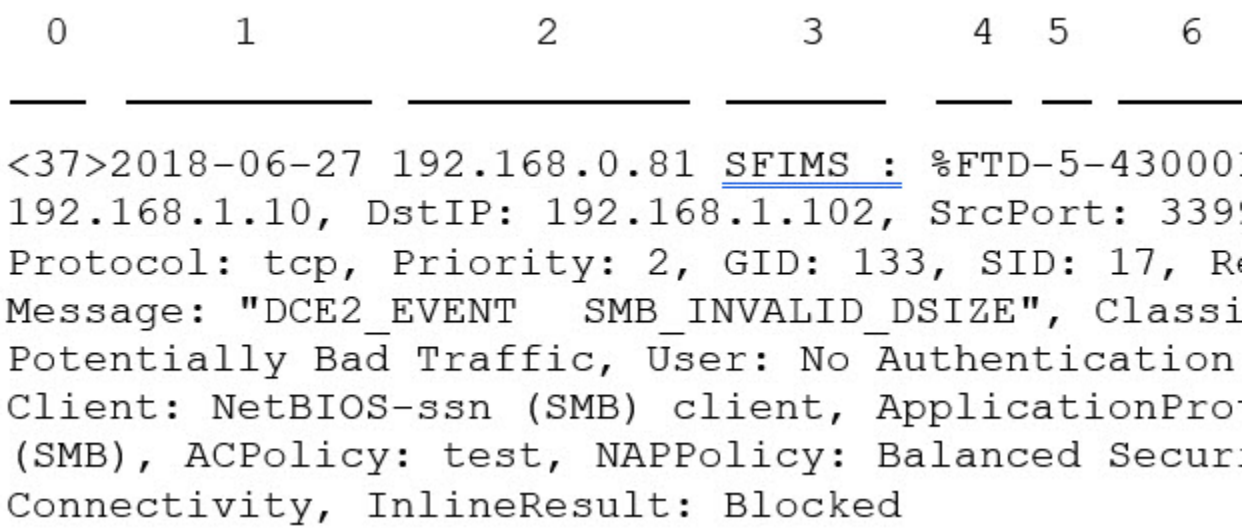

Table 65: Components of Security Event Syslog Messages

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 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, see Cisco Secure Firewall Management Center Device Configuration Guide. For information on PRI, see RFC5424.

Item Number in Sample Message	Header Element	Description
1	Timestamp	Date and time the syslog message was sent from the device. • (Syslogs sent from 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, see 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 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, see Cisco Secure Firewall Management Center Device Configuration Guide.
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. (See Creating a Syslog Alert Response, on page 554.) Otherwise, this element is not present in the header.
4	%FTD	Type of device that sent the message. %FTD is Secure Firewall Threat Defense

Item Number in Sample Message	Header Element	Description
5	Severity	The severity specified in the syslog settings for the policy that triggered the message. For severity descriptions, see *Severity Levels* in the Cisco Secure Firewall Management Center Device Configuration Guide or Syslog Severity Levels, on page 556.
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	See Facility in Security Event Syslog Messages, on page 635.
--	Remainder of message	Fields and values separated by colons. Fields with empty or unknown values are omitted from messages. For field descriptions, see: • Connection and Security-Related Connection Event Fields, on page 731. • Intrusion Event Fields, on page 766 • File and Malware Event Fields, on page 815 **Note** Field description lists include both syslog fields and fields visible in the event viewer (menu options under the Analysis menu in the 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 in Security Event Syslog Messages

Facility values are not generally relevant in syslog messages for security events. However, if you require Facility, use the following table:

Device	To Include Facility in Connection Events	To Include Facility in Intrusion Events	Location in Syslog Message
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. See Configuring Syslog Alerting for Intrusion Events, on page 564.	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 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, see Facilities and Severities for Intrusion Syslog Alerts, on page 565 and Creating a Syslog Alert Response, on page 554.

Secure Firewall Syslog Message Types

Secure Firewall can send multiple syslog data types, as described in the following table:

Syslog Data Type	See
Audit logs from management center	Stream Audit Logs to Syslog, on page 48 and the Audit and Syslog, on page 405 chapter
Device health and network-related logs from threat defense devices	Cisco Secure Firewall Management Center Device Configuration Guide
Connection, security intelligence, and intrusion event logs from threat defense devices	About Configuring the System to Send Security Event Data to Syslog, on page 624.
Connection, security intelligence, and intrusion event logs from Classic devices	About Configuring the System to Send Security Event Data to Syslog, on page 624
Logs for file and malware events	About Configuring the System to Send Security Event Data to Syslog, on page 624
IPS Settings	Send Syslog messages for IPS events. Configuration Locations for Syslogs for Intrusion Events (Threat Defense Devices), on page 631

Limitations of 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

The Event Streamer (eStreamer) allows you to stream several kinds of event data from a Secure Firewall Management Center to a custom-developed client application. For more information, see *Secure Firewall Management Center Event Streamer Integration Guide*.

Before the appliance you want to use as an eStreamer server can begin streaming eStreamer events to an external client, you must configure the eStreamer server to send events to clients, provide information about the client, and generate a set of authentication credentials to use when establishing communication. You can perform all of these tasks from the appliance's user interface. Once your settings are saved, the events you selected will be forwarded to eStreamer clients when requested.

You can control which types of events the eStreamer server is able to transmit to clients that request them.

Table 66: 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 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
Sends directly from devices	Sends from 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 637.
Some event data can be sent only from management center. See Data Sent Only via eStreamer, Not via Syslog, on page 638.	Includes data that cannot be sent via syslog directly from devices. See Data Sent Only via eStreamer, Not via Syslog, on page 638.

Data Sent Only via eStreamer, Not via Syslog

The following data 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 a description, see eStreamer Server Streaming, on page 637.
- 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
- The following fields:
  - Impact and ImpactFlag fields

    For a description, see eStreamer Server Streaming, on page 637.

- the IOC_Count field

- Most raw IDs and UUIDs.

  Exceptions:

  - Syslogs for connection events do include the following: 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

Choosing eStreamer Event Types

The **eStreamer Event Configuration** check boxes control which events the eStreamer server can transmit. Your client must still specifically request the types of events you want it to receive in the request message it sends to the eStreamer server. For more information, see the *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 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 637.

Step 4 Click **Save**.

Configuring eStreamer Client Communications

Before eStreamer can send eStreamer events to a client, you must add the client to the eStreamer server's peers database from the eStreamer page. You must also copy the authentication certificate generated by the eStreamer server to the client. After completing these steps you do not need to restart the eStreamer service to enable the client to connect to the eStreamer server.

In a multidomain deployment, you can create an eStreamer client in any domain. The authentication certificate allows the client to request events only from the client certificate's domain and any descendant domains. The eStreamer configuration page shows only clients associated with the current domain, so if you want to download or revoke a certificate, switch to the domain where the client was created.

You must be an Admin or Discovery Admin user to perform this task, for 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.

Note
If you have not configured DNS resolution, use an IP address.

Step 5 If you want to encrypt the certificate file, enter a password in the **Password** field.

Step 6 Click **Save**.
The eStreamer server now allows the host to access port 8302 on the eStreamer server and creates an authentication certificate to use during client-server authentication.

Step 7 Click **Download** ( ) next to the client hostname to download the certificate file.

Step 8 Save the certificate file to the appropriate directory used by your client for SSL authentication.

Step 9 To revoke access for a client, click **Delete** ( ) next to the host you want to remove.

Note that you do not need to restart the eStreamer service; access is revoked immediately.

Event Analysis in Splunk

You can use the Cisco Secure Firewall (f.k.a. Firepower) app for Splunk (formerly known as the Cisco Firepower App for Splunk) as an external tool to display and work with Secure Firewall event data, to hunt and investigate threats on your network. To use the Splunk tool, eStreamer is required. This is an advanced functionality. See eStreamer Server Streaming, on page 637. For more information, see User Guide for Cisco Secure Firewall (f.k.a. Firepower) App for Splunk.

Event Analysis in IBM QRadar

You can use the Cisco Firepower app for IBM QRadar as an alternate way to display event data and help you analyze, hunt for, and investigate threats to your network.

eStreamer is required. This is an advanced functionality. See eStreamer Server Streaming, on page 637.

For more information, see 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

Feature	Minimum Management Center	Minimum 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 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 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, see Frequently Asked Questions.
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 management center, see 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 page

Feature	Minimum Management Center	Minimum Threat Defense	Details
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. See Event Analysis in IBM QRadar, on page 641.

Feature	Minimum Management Center	Minimum 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.
Integration with Splunk	Supports all 6.x versions	Any	Splunk users can use a new, separate Splunk app, Cisco Secure Firewall (f.k.a. Firepower) app for Splunk, to analyze events. Available functionality is affected by your Firepower version. See Event Analysis in Splunk, on page 640.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Integration with Cisco Security Packet Analyzer	6.3	Any	Feature introduced: 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: threat defense devices running version 6.3. • Some new functionality: Non-threat defense devices running version 6.3. • Less new functionality: All devices running versions older than 6.3. For more information, see the topics under About Sending Syslog Messages for Security Events, on page 624 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

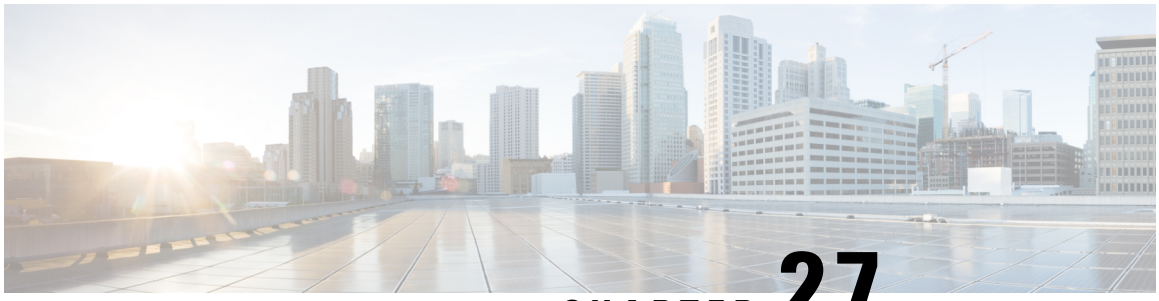

CHAPTER 27

Workflows

The following topics describe how to use workflows:

Overview: Workflows

A workflow is a tailored series of data pages on the management center web interface that analysts can use to evaluate events generated by the system.

The following types of workflows are available on the 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 management center. You can edit, delete, and copy these workflows.

Custom Workflows

Workflows that you create and customize for your specific needs, or that the system generates automatically when you create custom tables. You can edit, delete, and copy these workflows.

The data displayed in a workflow often depends on such factors as how you license and deploy your managed devices, and whether you configure features that provide the data.

Predefined Workflows

The predefined workflows described in the following sections are delivered with the system. You cannot edit or delete a predefined workflow, but you can copy a predefined workflow and use it as the basis for a custom workflow.

Predefined Intrusion Event Workflows

The following table describes the predefined intrusion event workflows included with the system.

Table 67: 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.

Workflow Name	Description
Source Port	This workflow indicates which servers are generating the most alerts. You can use this information to identify areas that require tuning, and to decide which servers require attention.
Source and Destination	This workflow identifies host IP addresses sharing high levels of alerts. Pairs at the top of the list could be false positives, and may identify areas that require tuning. You can check pairs at the bottom of the list for targeted attacks, for users accessing resources they should not be accessing, or for hosts that do not belong on the network.

Predefined Malware Workflows

The following table describes the predefined malware workflows included on the management center. All predefined malware workflows use the table view of malware events.

Table 68: 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

The following table describes the predefined file event workflows included on the management center. All the predefined file event workflows use the table view of file events.

Table 69: Predefined File Workflows

Workflow Name	Description
File Summary	This workflow provides a quick breakdown of the different file event categories and types, along with any associated malware dispositions.
Hosts Receiving Files	This workflow provides a list of host IP addresses that have received files, grouped by the associated malware dispositions for those files.
Hosts Sending Files	This workflow provides a list of host IP addresses that have sent files, grouped by the associated malware dispositions for those files.

Predefined Captured File Workflows

The following table describes the predefined captured file workflows included on the management center. All predefined captured file workflows use the table view of captured files.

Table 70: Predefined Captured File Workflows

Workflow Name	Description
Captured File Summary	This workflow provides a breakdown of captured files based on type, category, and threat score.
Dynamic Analysis Status	This workflow provides a count of captured files based on whether they have been submitted for dynamic analysis.

Predefined Connection Data Workflows

The following table describes the predefined connection data workflows included on the management center. All the predefined connection data workflows use the table view of connection data.

Table 71: 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

The following table describes the predefined Security Intelligence workflows included on the management center. All the predefined Security Intelligence workflows use the table view of Security Intelligence events.

Table 72: 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

The following table describes the predefined workflows that you can use with host data.

Table 73: 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

The following table describes the predefined workflows that you can use with IOC (Indications of Compromise) data.

Table 74: Predefined Indications of Compromise Workflows

Workflow Name	Description
Host Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis > Hosts** menu.
Indications of Compromise by Host	You can use this workflow to gauge which hosts on your network are most likely to be compromised (based on IOC data). Access this workflow via the **Analysis > Hosts** menu.
User Indications of Compromise	This workflow begins with a summary view of IOC data grouped by count and category, and provides a detail view that further subdivides the summary data by event type. Access this workflow via the **Analysis > Users** menu.
Indications of Compromise by User	Use this workflow to gauge which users on your network are most likely to be involved in potential compromises (based on IOC data.) Access this workflow via the **Analysis > Users** menu.

Predefined Applications Workflows

The following table describes the predefined workflows that you can use with application data.

Table 75: 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

The following table describes the predefined workflows that you can use with application detail and client data.

Table 76: Predefined Application Details Workflows

Workflow Name	Description
Application Details	You can use this workflow to analyze the client applications on your network in more detail. The workflow then provides a table view of client applications, followed by the host view.
Clients	This workflow contains a table view of client applications, followed by the host view.

Predefined Servers Workflows

The following table describes the predefined workflows that you can use with server data.

Table 77: 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.

Workflow Name	Description
Server Details	You can use this workflow to analyze the vendors and versions of detected server application protocols in detail.
Servers	This workflow contains a table view of applications followed by the host view.

Predefined Host Attributes Workflows

The following table describes the predefined workflow that you can use with host attribute data.

Table 78: Predefined Host Attributes Workflows

Workflow Name	Description
Attributes	You can use this workflow to monitor IP addresses of hosts on your network and the hosts' status.

The Predefined Discovery Events Workflow

The following table describes the predefined workflow that you can use to view discovery and identity data.

Table 79: Predefined Discovery Event Workflows

Workflow Name	Description
Discovery Events	This workflow provides a detailed list, in table view form, of discovery events, followed by the host view.

Predefined User Workflows

The following table describes the predefined workflow that you can use to view user discovery and user identity data.

Table 80: 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

The following table describes the predefined vulnerabilities workflow included on the management center.

Table 81: Predefined Vulnerabilities Workflows

Workflow Name	Description
Vulnerabilities	You can use this workflow to review vulnerabilities in the database, including a table view of only those active vulnerabilities that apply to the detected hosts on your network. The workflow provides a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.

Predefined Third-Party Vulnerabilities Workflows

The following table describes the predefined third-party vulnerabilities workflows included on the management center.

Table 82: 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 83: 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 84: Additional Predefined Workflows

Workflow Name	Description
Audit Log	This workflow contains a table view of the audit log that lists audit events.
Health Events	This workflow displays events triggered by the health monitoring policy.
Rule Update Import Log	This workflow contains a table view listing information about both successful and failed rule update imports.
Scan Results	This workflow contains a table view listing each completed scan.

Custom Table Workflows

You can use the custom tables feature to create tables that use the data from two or more types of events. This is useful because you can, for example, create tables and workflows that correlate intrusion event data with discovery data to allow simple searches for events that affect critical systems.

When you create a custom table, the system automatically creates a workflow that you can use to view the events associated with the table. The features in the workflow differ depending on which type of table you use. For example, custom table workflows based on the intrusion event table always end with the packet view. However, custom table workflows based on discovery events end with the host view.

Unlike workflows based on the predefined event tables, workflows based on custom tables do not have links to other types of workflows.

Using Workflows

Procedure

Step 1 Choose the appropriate menu path and option as described in Workflow Selection, on page 658.

Step 2 Navigate within the current workflow:

- To view all of the columns available in your chosen event data type, use table view pages; see Using Table View Pages, on page 664.
- To view a subset of the columns available in your chosen event data type, use drill-down pages; see Using Drill-Down Pages, on page 664.
- To display the corresponding row in the next page of the workflow, click **Down-Arrow** ( ).

- To move among the pages of a multipage workflow, use the tools at the bottom of each page; see Workflow Page Traversal Tools, on page 661.
- To view the same constraints applied within a workflow for a different type of event, click **Jump to** and choose the event view from the drop-down list.

Step 3 Modify the display of the current workflow:

- Check the check boxes by one or more rows on a page to indicate which row(s) you want to affect, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all selected rows.
- Check the check box at the top of the row to select all the rows on the page, then click one of the buttons at the bottom of the page (for example, **View**) to perform that action for all rows on the page.
- Constrain the columns in the display by clicking **Close** () in the column heading that you want to hide. In the pop-up window that appears, click **Apply**

  Tip
  To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the expand arrow to expand the search constraints, then click the column name under Disabled Columns.

- Constrain the data view by selected values for selected fields. For information, see Event View Constraints, on page 678 and Compound Event View Constraints, on page 679.
- Change the time constraints on the event view. The date range located in the upper right corner of the page sets a time range for events to include in the workflow; for information, see Event Time Constraints, on page 672.

  Note
  Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- To sort data by columns, click the name of a column. To reverse the sort order, click the column name again. The direction indicates which column the data is sorted by, and whether the sort is **Ascending** or **Descending**.
- Click a workflow page link to display that page using any active constraints. Workflow page links appear in the upper left corner of predefined workflow table views and drill-down pages, above events and below the workflow name.

Step 4 View additional data within the current workflow:

- To view the file's trajectory map in a new window, click network file trajectory in file name and SHA-256 hash value columns. The icon is different depending on the file status; see File Trajectory Icons, on page 662.
- To display a pop-up window of the host profile associated with an IP address, click host profile in any IP address column. The icon is different depending on the file status; see Host Profile Icons, on page 662.
- To view the Dynamic Analysis Summary report for the highest threat score associated with a file, click threat score in any threat score column. The icon is different depending on the file's highest threat score; see Threat Score Icons, on page 662.

- 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 666.

Step 5 Navigate to a different workflow:

To view the same event type using a different workflow, click **(switch workflow)** next to the workflow title, then choose the workflow you want to use. Note that you **cannot** use a different workflow for scan results.

Workflow Access by User Role

Access to a workflow is determined by the user's role. See the table below for more information.

User Role	Accessible Workflows
Administrator	Can access any workflow, and are the only users who can access the audit log, scan results, and the rule update import log.
Maintenance User	Can access health events.
Security Analyst and Security Analyst (Read Only)	Can access intrusion, malware, file, connection, discovery, vulnerability, correlation, and health workflows.

Workflow Selection

The system provides predefined workflows for the types of data listed in the following table.

Table 85: Features Using Workflows

Feature	Menu Path	Option
Connection events	Analysis > Connections	Events
Security Intelligence events	Analysis > Connections	Security Intelligence Events
Correlation events	Analysis > Correlation	Correlation Events Allow List Events Allow List Violations Status

Feature	Menu Path	Option
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
Rule Update Import Log	System > Updates	Rule Updates

When you view any of the kinds of data described in the above table, events appear on the first page of the default workflow for that data. You can specify a different default workflow by configuring your event view settings. Note that workflow access depends on your user role.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Related Topics

Configuring Event View Settings, on page 206

Workflow Pages

Although the data in each type of workflow is different, all workflows share a common set of features. Workflows can include several types of pages. The actions you can perform on a workflow page depend on the type of page.

Drill-down and table view pages in workflows allow you to quickly narrow your view of the data so you can zero in on events that are significant to your analysis. Table view pages and drill-down pages both support many features you can use to constrain the set of events you want to view or to navigate the workflow. When viewing data on drill-down pages or in the table view in a workflow, you can sort the data in ascending or descending order based on any available column. If the database contains more events than can be displayed on a single workflow page, you can click the links at the bottom of the page to display more events. When you click one of these links, the time window automatically pauses so that you do not see the same events twice; you can unpause the time window when you are ready.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Table Views

Table views include a column for each of the fields in the database on which your workflow is based if the page is enabled by default.

For best performance, display only the columns you need. The more columns are displayed, the more resources are required to display the data.

Note that when you disable a column on a table view, the 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

Generally, drill-down pages are intermediate pages that you use to narrow your investigation to a few events before moving to a table view page. Drill-down pages contain a subset of columns that are available in the database.

For example, a drill-down page for discovery events might include only the IP Address, MAC Address, and Time columns. A drill-down page for intrusion events, on the other hand, might include the Priority, Impact Flag, Inline Result, and Message columns.

Drill-down pages allow you to narrow the scope of events you are viewing and to move forward in the workflow. If you click on a value in a drill-down page, for example, you constrain by that value and move to the next page in the workflow, focusing more closely on events that match your selected values. Clicking a value in a drill-down page does not disable the column where the value is, even if the page you advance to is a table view. Note that drill-down pages for predefined workflows always have a Count column. When you create a custom workflow, you add a drill-down page to it by clicking **Add Page**.

Graphs

Workflows based on connection data can include graph pages, also called *connection graphs*.

For example, a connection graph might display a line graph that shows the number of connections detected by the system over time. Generally, connection graphs are, like drill-down pages, intermediate pages that you use to narrow your investigation.

Final Pages

The final page of a workflow depends on the type of event on which the workflow is based:

- The host view is the final page for workflows based on applications, application details, discovery events, hosts, indications of compromise (IOC), servers, 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.
- The user detail view is the final page for workflows based on users, user activity, and user indications of compromise.
- The vulnerability detail view is the final page for workflows based on Cisco vulnerabilities.
- The packet view is the final page for workflows based on intrusion events.

Workflows based on other kinds of events (for example, audit log events or malware events) do not have final pages.

On the final page of a workflow, you can expand detail sections to view specific information about each object in the set you focused on over the course of the workflow. Although the web interface does not list the constraints on the final page of a workflow, previously set constraints are retained and applied to the set of data.

Workflow Page Navigation Tools

Workflow pages provide visual cues to facilitate navigating among them and choosing what information to display during event analysis.

Workflow Page Traversal Tools

If a workflow contains multiple pages of data, the bottom of each page displays the number of pages in the workflow, as well as the tools listed in the table below which you may use to navigate among the pages:

Table 86: 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 87: 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 88: Host Profile Icons

Host Profile Icon	Host Status
	Host is not tagged as potentially compromised.
	Host is tagged as potentially compromised by triggered indications of compromise (IOC) rules.
	Added to Block List (Appears only if you are performing traffic filtering based on Security Intelligence data.)
	Added to Block List, set to monitor (Appears only if you are performing traffic filtering based on Security Intelligence data.)

Threat Score Icons

When a workflow page provides the opportunity to view a Dynamic Analysis Summary report for the highest threat score associate with a file, a threat score icon appears. The icon differs depending on the file's highest threat score.

Table 89: Threat Score Icons

Threat Score Icon	Threat Score Level
Low	Low

Threat Score Icon	Threat Score Level
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 90: User Icons

User Icon	User Status
User	User is not associated with any indications of compromise.
Red User	User is associated with one or more indications of compromise.

The Workflow Toolbar

Each page in a workflow includes a toolbar that offers quick access to related features. The following table describes each of the links on the toolbar.

Table 91: Workflow Toolbar Links

Feature	Description
Bookmark This Page	Bookmarks the current page so you can return to it later. Bookmarking captures the constraints in effect on the page you are viewing so you can return to the same data (assuming the data still exists) at a later time.
Report Designer	Opens the report designer with the currently constrained workflow as the selection criteria.
Dashboard	Opens a dashboard relevant to your current workflow. For example, Connection Events workflows link to the Connection Summary dashboard.
View Bookmarks	Displays a list of saved bookmarks from which you can select.
Search	Displays a Search page where you can perform advanced searches on data in the workflow. You can also click the down arrow icon to select and use a saved search.

Related Topics

Creating a Report Template from an Event View, on page 529
About Dashboards, on page 339
Event Searches, on page 685
Bookmarks, on page 682
Creating Bookmarks, on page 682
Viewing Bookmarks, on page 682

Using Drill-Down Pages

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Features Using Workflows.

Step 2 In any workflow, you have the following options:

- To drill down to the next workflow page constraining on a specific value, click a value within a row. Note that this works only on drill-down pages. Clicking a value within a row in a table view only constrains the table view and does not drill down to the next page.
- To drill down to the next workflow page constraining on some events, check the check boxes next to the events you want to view on the next workflow page, then click **View**.
- To drill down to the next workflow page keeping the current constraints, click **View All**.

Tip
Table views always include "Table View" in the page name.

Using Table View Pages

Table view pages provide some features not available on drill-down, host view, packet view, or vulnerability detail pages. Use these features as described below:

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 658.

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.

See Work in Secure Firewall Management Center with Connection Events Stored on a Secure Network Analytics Appliance, on page 665.

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 Secure 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 management center's event viewer and context explorer, and include them when generating reports. You can also cross-launch from an event in 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.

Important 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 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 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:

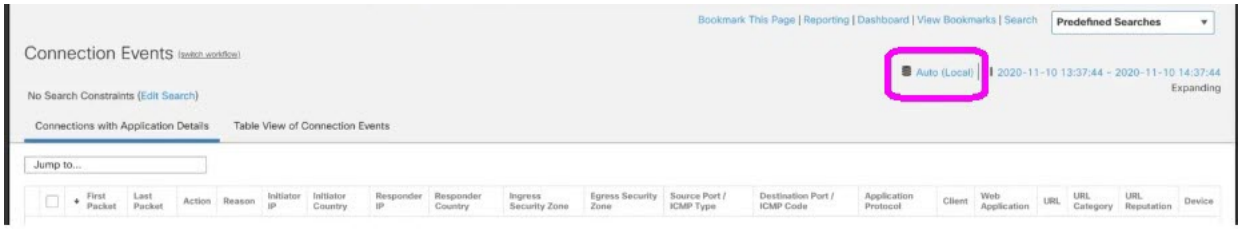

Caution
If you select **Local**, the system displays only the data available on the management center, even if local data is not available for the entire time range selected. You will not be notified that this situation is occurring.

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

You can view and filter traffic based on country and continent by leveraging a geolocation database (GeoDB) that maps IP addresses to countries/continents. Note that for mobile devices and other hosts detected moving from country to country, the system may report a continent instead of a specific country. We issue periodic updates to the GeoDB. You must regularly update the GeoDB to have accurate geolocation information; see Update the Geolocation Database (GeoDB), on page 228.

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.

Related Topics

Network Conditions
Geolocation
Introduction to Correlation Policies and Rules, on page 953
Traffic Profile Conditions, on page 993
Update the Geolocation Database (GeoDB), on page 228

Connection Event Graphs

In addition to workflows that use tabular drill-down pages and a final table view of events, the system can present certain connection data graphically, using data aggregated over five-minute intervals. Note that you can graph only the information used to aggregate data: source and destination IP addresses (and those hosts' associated users), destination port, transport protocol, and application protocol.

Tip You cannot graph Security Intelligence events separately from their associated connection events. For a graphical overview of Security Intelligence filtering activity, use dashboards and the Context Explorer.

There are three different types of connection graphs:

- *Pie charts* display data from one dataset grouped into discrete categories.
- *Bar graphs* display data from one or more datasets grouped into discrete categories.
- *Line graphs* plot data from one or more datasets over time, using either a standard or a velocity (rate of change) view.

Note The system displays traffic profiles as line graphs, which you can manipulate in the same way as you would any other connection graph, with some restrictions. To view traffic profiles, you must have Administrator access.

Like workflow tables, you can drill down and constrain workflow graphs to focus your analysis.

Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and responders. Pie charts can only display one dataset.

You can display different data and datasets on a connection graph by changing either the x-axis, the y-axis, or both. On a pie chart, changing the x-axis changes the independent variable and changing the y-axis changes the dependent variable.

Related Topics

Connection Summaries (Aggregated Data for Graphs), on page 730

Using Connection Event Graphs

On the management center, you can view connection event graphs and manipulate them depending on the information you are looking for.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of connection events. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Connections** > **Events**.

Note
If a connection event table appears instead of a graph, or to view a different graph, click **(switch workflow)** by the workflow title and choose a predefined workflow that includes graphs, or a custom workflow. Note that all predefined connection event workflows—including connection graphs—terminate in a table view of connections.

Step 2 You have the following options:

- Time Range — To adjust the time range, which is useful if the graph is blank, see Changing the Time Window, on page 675.
- Field Names — To learn more about the data you can graph, see Connection and Security-Related Connection Event Fields, on page 731.
- Host Profile — To view the host profile for an IP address, on a graph displaying connection data by initiator or responder, click either a bar on a bar graph or a wedge on a pie chart and choose **View Host Profile**.

- User Profile — To view user profile information, on a graph displaying connection data by initiator user, click either a bar on a bar graph or a wedge on a pie chart and choose **View User Profile**.
- Other Information — To learn more information about the graphed data, position your cursor over a point on a line graph, a bar in a bar graph, or a wedge in a pie chart.
- Constrain — To constrain a connection graph by any x-axis (independent variable) criterion without advancing the workflow to the next page, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, and choose a **View by...** option.
- Data Selection — To change the data displayed on the graph, click **X-Axis** or **Y-Axis** and choose the new data to graph. Note that changing the x-axis to or from **Time** also changes the graph type; changing the y-axis affects the displayed datasets.
- Datasets — To change the graph's dataset, click **Datasets** and choose a new dataset.
- Detach — To detach a connection graph so you can perform further analysis without affecting the default time range, click **Detach**.

  Tip
  Click **New Window** in a detached graph to create a copy. You can then perform different analyses on each of the detached graphs. Note that traffic profiles are detached graphs.

- Drill Down — To drill down to the next page in the workflow, click a point on a line graph, a bar on a bar graph, or a wedge on a pie chart, then choose **Drill-down**. Clicking a point on a line graph changes the time range on the next page to a 10-minute span, centered on the point you clicked. Clicking a bar on a bar graph or a wedge on a pie chart constrains the next page based on the criterion represented by the bar or wedge.
- Export — To export the connection data for a graph as a CSV (comma-separated values) file, **Export Data**. Then, click **Download CSV File** and save the file.
- Graph Type: Line — To switch between a standard and velocity (rate of change) line graph, click **Velocity**, then choose **Standard** or **Velocity**.
- Graph Type: Bar and Pie — To switch between a bar graph and pie chart, click **Switch to Bar** or **Switch to Pie**. Because you cannot display multiple datasets on a pie char, if you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When choosing which dataset to display, the management center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.
- Navigate Between Pages — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views — To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.
- Recenter — To recenter a line graph around a point in time without changing the length of the time range, click that point, then choose **Recenter**.
- Zoom -— To recenter a line graph around a point in time while zooming in or out, click that point, choose **Zoom**, then choose a new time span.

Note

Unless you are working with a detached graph, constraining, recentering, and zooming changes the default time range for the management center.

Example

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

Using Workflows, on page 656
Configuring Event View Settings, on page 206

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 92: 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

X-Axis Option	Graph Type	Graphs This Data
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 93: Y-Axis Options

Y-Axis Option	Graphs This Data Using The X-Axis Criterion
Bytes	bytes transmitted
Connections	number of connections
KBytes	kilobytes transmitted
KBytes Per Second	kilobytes per second
Packets	number of packets transmitted
Unique Hosts	number of unique hosts detected
Unique Application Protocols	number of unique application protocols
Unique Users	number of unique users

Connection Graphs with Multiple Datasets

Both bar graphs and line graphs can display multiple datasets; that is, they can display several values on the y-axis for each x-axis data point. For example, you could display the total number of unique initiators and responders.

Note You **cannot** display multiple datasets on a pie chart. If you switch to a pie chart from a bar graph that has multiple datasets, the pie chart shows only one dataset, which is selected automatically. When selecting which dataset to display, the management center favors total statistics over initiator and responder statistics, and favors initiator statistics over responder statistics.

On line graphs, multiple datasets appear as multiple lines, each with a different color. For example, the following graphic displays the total number of unique initiators and the total number of unique responders detected on a monitored network over a one hour interval.

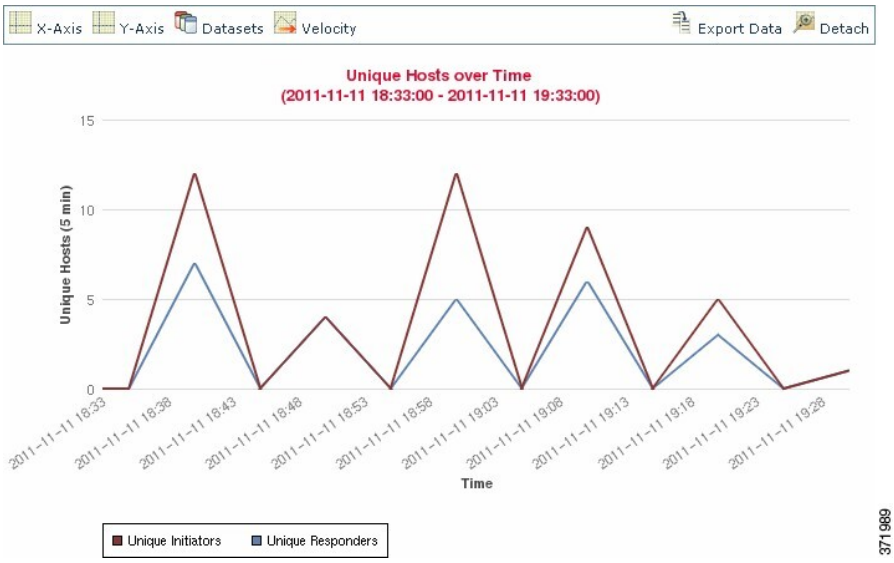

On bar graphs, multiple datasets appear as a set of colored bars for each x-axis data point. For example, the following bar graph displays the total packets transmitted on a monitored network, packets transmitted by initiators, and packets transmitted by responders.

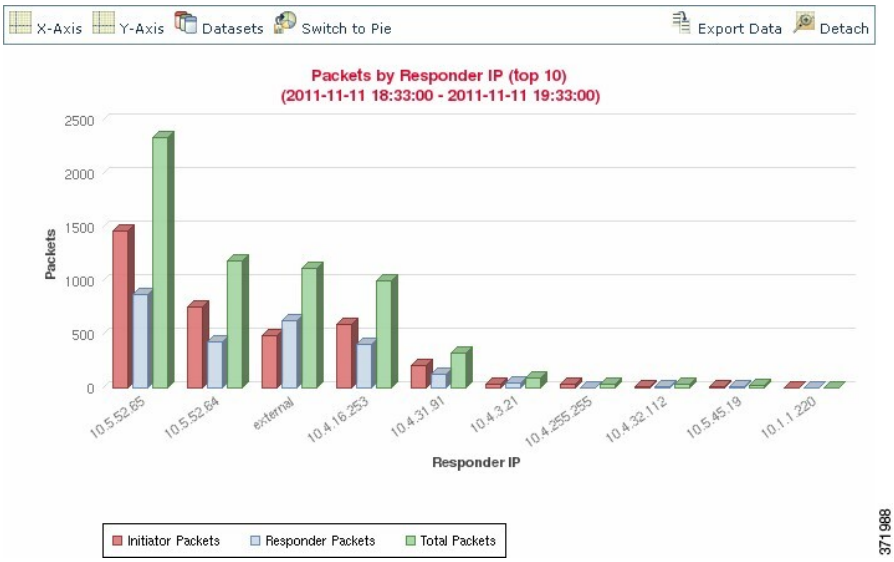

Connection Graph Dataset Options

The following table describes the datasets you can display on the x-axis of a connection graph.

Table 94: 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.

If the y-axis displays...	You can select as datasets...
KBytes	combinations of: • the total kilobytes transmitted on the monitored network (**Total KBytes**) • the number of kilobytes transmitted from host IP addresses on the monitored network (**Initiator KBytes**) • the number of kilobytes received by host IP addresses on the monitored network (**Responder KBytes**)
KBytes Per Second	the default only, which is the total kilobytes per second transmitted on the monitored network (**Total KBytes Per Second**)
Packets	combinations of: • the total packets transmitted on the monitored network (**Total Packets**) • the number of packets transmitted from host IP addresses on the monitored network (**Initiator Packets**) • the number of packets received by host IP addresses on the monitored network (**Responder Packets**)
Unique Hosts	combinations of: • the number of unique session initiators on the monitored network (**Unique Initiators**) • the number of unique session responders on the monitored network (**Unique Responders**)
Unique Application Protocols	the default only, which is the number of unique application protocols on the monitored network (**Unique Application Protocols**)
Unique Users	the default only, which is the number of unique users logged into session initiators on the monitored network (**Unique Initiator Users**)

Event Time Constraints

Each event has a time stamp that indicates when the event occurred. You can constrain the information that appears in some workflows by setting the time window, sometimes called the time range.

Workflows based on events that can be constrained by time include a time range line at the top of the page.

By default, workflows use an expanding time window set to the past hour. For example, if you log in at 11:30 AM, you will see events that occurred between 10:30 AM and 11:30 AM. As time moves forward, the time window expands. At 12:30 PM, you will see events that occurred between 10:30 AM and 12:30 PM.

You can change this behavior by setting your own default time window in the event view settings. This governs three properties:

- time window type (static, expanding, or sliding)

- time window length
- the number of time windows (either multiple time windows or a single global time window)

Regardless of the default time window setting, you can manually change the time window during your event analysis by clicking the time range at the top of the page, which displays the Date/Time pop-up window. Depending on the number of time windows you configured and the type of appliance you are using, you can also use the Date/Time window to change the default time window for the type of event you are viewing.

Finally, you can pause the time window while looking at a sliding or expanding workflow. See Pause the Time Window to Temporarily Freeze the Data Set, on page 675.

Related Topics

Configuring Event View Settings, on page 206
Using Connection and Security-Related Connection Event Tables, on page 755

Per-Session Time Window Customization for Events

Regardless of the default time window, you can manually change the time window during your event analysis.

Note Manual time window settings are valid for only the current session. When you log out and then log back in, time windows are reset to the default.

Depending on the number of time windows you configured, changing the time window for one workflow may affect other workflows on the appliance. For example, if you have a single, global time window, changing the time window for one workflow changes it for all other workflows on the appliance. On the other hand, if you are using multiple time windows, changing the audit log or health event workflow time windows has no effect on any other time window, while changing the time window for other kinds of events affects all events that can be constrained by time (with the exception of audit events and health events).

Note that because not all workflows can be constrained by time, time window settings have no effect on workflows based on hosts, host attributes, applications, application details, vulnerabilities, users, or allow list violations.

Use the Time Window tab on the Date/Time window to manually configure a time window. Depending on the number of time windows you configured in your default time window settings, the tab's title is one of the following:

- **Events Time Window**, if you configured multiple time windows and are setting the time window for a workflow other than the audit log or health events workflow
- **Health Monitoring Time Window**, if you configured multiple time windows and are setting the time window for the health events workflow
- **Audit Log Time Window**, if you configured multiple time windows and are setting the time window for the audit log
- **Global Time Window**, if you configured a single time window

The first decision you must make when configuring a time window is the type of time window you want to use:

- A *static* time window displays all the events generated from a specific start time to a specific end time.

- An *expanding* time window displays all the events generated from a specific start time to the present; as time moves forward, the time window expands and new events are added to the event view.
- A *sliding* time window displays all the events generated from a specific start time (for example, one week ago) to the present; when you refresh the page, the time window "slides" so that you see only the events in the time range you configured (in this example, for the last week). To temporarily prevent the data set from updating while you are examining it, see Pause the Time Window to Temporarily Freeze the Data Set, on page 675.

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

The following table explains the various settings you can configure on the Time Window tab.

Table 95: Time Window Settings

Setting	Time Window Type	Description
time window type drop-down list	n/a	Select the type of time window you want to use: static, expanding, or sliding. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
Start Time calendar	static and expanding	Specify a start date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Instead of using the calendar, you can use the Presets options, described below.
End Time calendar	static	Specify an end date and time for your time window. The maximum time range for all time windows is from midnight on January 1, 1970 (UTC) to 3:14:07 AM on January 19, 2038 (UTC). Note that if you are using an expanding time window, the End Time calendar is grayed out and specifies that the end time is "Now." Instead of using the calendar, you can use the Presets options, described below.
Show the Last field and drop-down list	sliding	Configure the length of the sliding time window.
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.

Setting	Time Window Type	Description
Presets: Current	static and expanding	Click one of the time ranges in the list to change the time window, based on the local time and date of the appliance. Clicking a preset changes the calendars to reflect the preset you choose. Note that: • the current day begins at midnight • the current week begins at midnight Sunday • the current month begins at midnight on the first of the month
Presets: Synchronize with	all (not available if you are using a global time window)	Click one of: • **Events Time Window** to synchronize the current time window with the events time window • **Health Monitoring Time Window** to synchronize the current time window with the health monitoring time window • **Audit Log Time Window** to synchronize the current time window with the audit log time window

Changing the Time Window

Procedure

Step 1 On a workflow constrained by time, click **Time Range** ( ) to go to the Date/Time window.

Step 2 On **Events Time Window**, set the time window as described in Time Window Settings, on page 674.

Tip
Click **Reset** to change the time window back to the default settings.

Step 3 Click **Apply**.

Pause the Time Window to Temporarily Freeze the Data Set

If you are using a sliding or expanding time window, you can pause the time window to examine a snapshot of the data provided by the workflow. This is useful because when an unpaused workflow updates, it may remove events that you want to examine or add events that you are not interested in.

The time window automatically pauses when you click a link at the bottom of the page to display another page of events; you can unpause the time window when you are ready.

When you are finished with your analysis, you can unpause the time window. Unpausing the time window updates it according to your preferences, and also updates the event view to reflect the unpaused time window.

Pausing an event time window has no effect on dashboards, nor does pausing a dashboard have any effect on pausing an event time window.

Procedure

On a workflow constrained by time, choose the desired time range control:

- To pause the time window, click time range control **Pause** ( ).
- To unpause the time window, click time range control **Play** ( ).

The Default Time Window for Events

During your event analysis, you can use the Preferences tab on the Date/Time window to change the default time window for the type of event you are viewing without having to use the event view settings.

Keep in mind that changing the default time window in this way changes the default time window for only the type of event you are viewing. For example, if you configured multiple time windows, changing the default time window on the Preferences tab changes the settings for either the events, health monitoring, or audit log window, in other words, whichever time window is indicated by the first tab. If you configured a single time window, changing the default time window on the Preferences tab changes the default time window for all types of events.

Related Topics

Default Time Windows, on page 208

Default Time Window Options for Event Types

The following table explains the various settings you can configure on the Preferences tab.

Table 96: Time Window Preferences

Preference	Description
Refresh Interval	Sets the refresh interval for event views, in minutes. Entering zero disables the refresh option.
Number of Time Windows	Specify how many time windows you want to use: • Select **Multiple** to configure separate default time windows for the audit log, for health events, and for workflows based on events that can be constrained by time. • Select **Single** to use a global time window that applies to all events.
Default Time Window: Show the Last - Sliding	This setting allows you to configure a sliding default time window of the length you specify. The appliance displays all the events generated from a specific start time (for example, 1 hour ago) to the present. As you change event views, the time window “slides” so that you always see events from the last hour.

Preference	Description
Default Time Window: Show the Last - Static/Expanding	This setting allows you to configure either a static or expanding default time window of the length you specify. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from a specific start time (for example, 1 hour ago), to the present. As you change event views, the time window expands to the present time.
Default Time Window: Current Day - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current day. The current day begins at midnight, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 24 hours before you log out, this time window can be more than 24 hours.
Default Time Window: Current Week - Static/Expanding	This setting allows you to configure either a static or expanding default time window for the current week. The current week begins at midnight on the previous Sunday, based on the time zone setting for your current session. For **static** time windows (enable the **Use End Time** check box), the appliance displays all the events generated from midnight to the time when you first viewed the events. As you change event views, the time window stays fixed so that you see only the events that occurred during the static time window. For **expanding** time windows (disable the **Use End Time** check box), the appliance displays all the events generated from midnight Sunday to the present. As you change event views, the time window expands to the present time. Note that if your analysis continues for over 1 week before you log out, this time window can be more than 1 week.

Changing the Default Time Window for Your Event Type

Procedure

Step 1 On a workflow constrained by time, click **Time Range** ( ) to go to the Date/Time window.

Step 2 Click **Preferences** and change your preferences, as described in Default Time Window Options for Event Types, on page 676.

Step 3 Click **Save Preferences**.

Step 4 You have two options:

- To apply your new default time window settings to the event view you are using, click **Apply** to close the Date/Time window and refresh the event view.
- To continue with your analysis without applying the default time window settings, close the Date/Time window without clicking **Apply**.

Event View Constraints

The information that you see on a workflow page is determined by the constraints that you impose. For example, when you initially open an event workflow, the information is constrained to events that were generated in the previous hour.

To advance to the next page in the workflow and constrain the data you are viewing by specific values, select the rows with those values on the page and click **View**. To advance to the next page in the workflow retaining the current constraints and carrying forward all events, select **View All.**

Note If you select a row with multiple non-count values and click **View**, you create a compound constraint.

There is a third method for constraining data in a workflow. To constrain the page to the rows with values that you selected and also add the selected value to the list of constraints at the top of the page, click a value within a row on the page. For example, if you are viewing a list of logged connections and want to constrain the list to only those you allowed using access control, click **Allow** in the **Action** column. As another example, if you are viewing intrusion events and want to constrain the list to only events where the destination port is 80, click **80 (http)/tcp** in the **Destination Port/ICMP Code** column.

Tip The procedure for constraining connection events based on Monitor rule criteria is slightly different and you may need to take some extra steps. Additionally, you cannot constrain connection events by associated file or intrusion information.

You can also use searches to constrain the information in a workflow. Use this feature when you want to constrain against multiple values in a single column. For example, if you want to view the events related to two IP addresses, click **Edit Search**, then modify the appropriate IP address field on the Search page to include both addresses, and then click **Search**.

The search criteria you enter on the search page are listed as the constraints at the top of the page, with the resulting events constrained accordingly. On the 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 management center displays connection data as if you had not constrained it at all. Invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough.

Constraining Events

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 658.

Step 2 In any workflow, you have the following options:

- To constrain the view to events that match a single value, click the desired value within a row on the page.
- To constrain the view to events that match multiple values, check the check boxes for events with those values, and click **View**.

  Note
  A compound constraint is added if the row contains multiple non-count values.

- To remove a constraint, click the Search Constraints **Expand Arrow** ( ) and click the name of the constraint in the expanded Search Constraints list.
- To edit constraints using the Search page, click **Edit Search.**
- To save constraints as a saved search, click **Save Search** and give the query a name.

  Note
  You cannot save queries containing compound constraints.

- To use the same constraints with another event view, click **Jump to** and choose the event view.

  Note
  You do not retain compound constraints when you switch to another workflow.

- To toggle the display of constraints click the Search Constraints **Expand Arrow** ( ) or the Search Constraints **Collapse Arrow** ( ). This is useful when the list of constraints is large and takes up most of the screen.

Compound Event View Constraints

Compound constraints are based on all non-count values for a specific event. When you select a row with multiple non-count values, you set a compound constraint that retrieves only events matching all the non-count values in that row on that page. For example, if you select a row that has a source IP address of `10.10.31.17` and a destination IP address of `10.10.31.15` and a row that has a source IP address of `172.10.10.17` and a destination IP address of `172.10.10.15`, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15

  OR

- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15

When you combine compound constraints with simple constraints, the simple constraints are distributed across each set of compound constraints. If, for example, you added a simple constraint for a protocol value of `tcp` to the compound constraints listed above, you retrieve all of the following:

- Events that have a source IP address of 10.10.31.17 AND a destination IP address of 10.10.31.15 AND a protocol of tcp

  OR

- Events that have a source IP address of 172.10.31.17 AND a destination IP address of 172.10.31.15 AND a protocol of tcp

You cannot perform a search or save a search on a compound constraint. You also cannot retain compound constraints when you use the event view links or click **(switch workflow)** to switch to another workflow. If you bookmark an event view with compound constraints applied, the constraints are not saved with the bookmark.

Using Compound Event View Constraints

Procedure

Step 1 Access a workflow by choosing the appropriate menu path and option as described in Workflow Selection, on page 658.

Step 2 To manage compound constraints, you have the following options:

- To create a compound constraint, choose one or more rows with multiple non-count values and click **View**.
- To clear compound constraints, click the Search Constraints **Expand Arrow** ( ) and click **Compound Constraints**.

Inter-Workflow Navigation

You can navigate to other workflows using the links in the **Jump to...** drop-down list on a workflow page. Select the drop-down list to view and select additional workflows.

When you select a new workflow, properties shared by the rows you select and the constraints you set are used in the new workflow, if they are applicable. If configured constraints or event properties do not map to fields in the new workflow, they are dropped. In addition, compound constraints are not retained when you switch from one workflow to another. In addition, constraints from the captured files workflow only transfer to file and malware event workflows.

Note When you view event counts over a time range, the total number of events may not reflect the number of events for which more detailed data is available. This occurs because the system sometimes prunes older event details to manage disk space usage. To minimize the occurrence of event detail pruning, you can fine-tune event logging to log only those events most important to your deployment.

Note that unless you have either paused the time window or have configured a static time window, the time window changes when you change workflows.

This feature enhances your ability to investigate suspicious activity. For example, if you are viewing connection data and notice that an internal host is transmitting an abnormally large amount of data to an external site, you can select the responder IP address and the port as constraints and then jump to the **Applications** workflow. The applications workflow will use the responder IP address and port as IP Address and Port constraints and display additional information about the application, such as what kind of application it is. You can also click **Hosts** at the top of the page to view the host profile for the remote host.

After finding more information about the application, you can select **Correlation Events** to return to the connection data workflow, remove the Responder IP from the constraints, add the Initiator IP to constraints, and select **Application Details** to see what client the user on the initiating host used when transferring data to the remote host. Note that the Port constraint is not transferred to the Application Details page. While keeping the local host as a constraint, you can also use other navigation buttons to find additional information:

- To discover if any policies have been violated by the local host, keep the IP address as a constraint and select **Correlation Events** from the **Jump to** drop-down list.
- To find out if an intrusion rule triggered against the host, indicating a compromise, select **Intrusion Events** from the **Jump to** drop-down list.
- To view the host profile for the local host and determine if the host is susceptible to any vulnerabilities that may have been exploited, select **Hosts** from the **Jump to** drop-down list.

Working 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.

Use the unified event viewer to:

- Look for relationships between events of different types
- See the effects of policy changes in real time

Procedure

Step 1 Select **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

Create a bookmark if you want to return quickly to a specific location and time in an event analysis. Bookmarks retain information about:

- the workflow you are using
- the part of the workflow you are viewing
- the page number within the workflow
- any search constraints
- any disabled columns
- the time range you are using

The bookmarks you create are available to all user accounts with bookmark access. This means that if you uncover a set of events that require more in-depth analysis, you can easily create a bookmark and turn over the investigation to another user with the appropriate privileges.

Note If the events that appear in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

Creating Bookmarks

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

Step 1 During an event analysis, with the events of interest displayed, click **Bookmark This Page**.
Step 2 In the **Bookmark Name** field, enter a name.
Step 3 Click **Save Bookmark**.

Viewing Bookmarks

In a multidomain deployment, you can only view bookmarks created in the current domain.

Procedure

From any event view, you have two options:

- Hover your pointer over **View Bookmarks**, and click on the desired bookmark in the drop-down menu.

- Click on click **View Bookmarks** and on the View Bookmarks page, click on the desired bookmark name or **View** ( ) next to it.

  Note
  If the events that originally appeared in a bookmark are deleted (either directly by a user or by automatic database cleanup), the bookmark no longer displays the original set of events.

History for Workflows

Table 97:

Feature	Minimum Management Center	Minimum 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 intelligence), intrusion, file, and malware. New/modified screens: **Analysis** > **Unified Events**
Work with events stored remotely.	7.0	Any	You can use the FMC 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: event viewer.

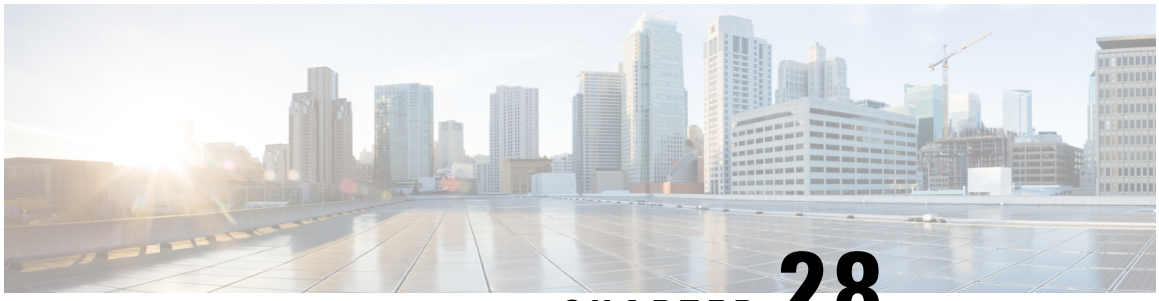

CHAPTER 28

Event Search

The following topics describe how to search for events within a workflow:

Event Searches

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

Searching Custom Tables, on page 708

Search Constraints

Each database table has its own search page where you can enter search constraint values to apply to fields defined for the table. Depending on the type of field, special syntax may be used to specify criteria such as wildcard characters or a range of numeric values.

Search results appear on workflow pages displaying each table field in columnar layout. Some database tables can additionally be searched using fields that are not displayed as columns on workflow pages. To determine whether such a constraint applies to your search results when viewing the results on a workflow page, click **Expand Arrow** ( ) to view the active search constraints.

General Search Constraints

When searching for events, observe the following general guidelines:

- Many fields require wildcards for partial-match searches. All fields accept wildcards for these searches.

  See Wildcards and Symbols in Searches, on page 686.

- All fields accept negation (`!`).
- All fields accept comma-separated lists of search values. Records that contain any of the listed values in the specified field match that search criteria.
- All fields accept comma-separated lists enclosed in quotation marks as search values.
  - For fields that may contain only a single value, records with the specified field containing the exact string specified within the quotation marks match the search criteria. For instance, a search for `A, B, "C, D, E"` will match records where the specified field contains `"A"` or `"B"` or `"C, D, E"`. This permits matching on fields that include the comma in possible values.
  - For fields that may contain multiple values at the same time, records with the specified fields containing all of the values in the quote-enclosed comma-separated list match that search criteria.
  - For fields that may contain multiple values at the same time, search criteria may include single values as well as quote-enclosed comma-separated lists. For instance, a search for `A, B, "C, D, E"` on a field that may contain one of more of these letters matches records where the specified field contains `A` or `B`, or all of `C`, `D`, and `E`.
- Specify `n/a` in any field to identify events where information is not available for that field; use `!n/a` to identify the events where that field is populated.
- You can precede many numeric fields with greater than (`>`), greater than or equal to (`>=`), less than (**<**), less than or equal to (`<=`), equal to (`=`), or not equal to (`<>`) operators.

Tip When searching a field with long complicated values (such as SHA-256 hash values), you can copy the search criteria value from source material and paste it into the appropriate field on the search page.

Wildcards and Symbols in Searches

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`.

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.

You can click **Object** (+) that appears next to a search field where you can use an object as a search criterion.

Related Topics

The Object Manager

Time Constraints in Searches

The formats accepted by search criteria fields that take a time value are shown in the following table.

Table 98: Time Specification in Search Fields

Time Formats	Example
`today [at HH:MMam|pm]`	`today` `today at 12:45pm`
`YYYY-DDMM- HH:MM:SS`	`2006-03-22 14:22:59`

You can precede a time value with one of the following operators:

Table 99: 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 100: Acceptable IP Address Syntax

To specify...	Type...	For example...
a single IP address	the IP address.	`192.168.1.1` `2001:db8::abcd`
multiple IP addresses using a list	a comma-separated list of IP addresses. Do **not** add a space before or after the commas.	`192.168.1.1,192.168.1.2` `2001:db8::b3ff,2001:db8::0202`
a range of IP addresses that can be specified with a CIDR block or prefix length	the IP address block in IPv4 CIDR or IPv6 prefix length notation.	`192.168.1.0/24` This specifies any IP in the 192.168.1.0 network with a subnet mask of 255.255.255.0, that is, 192.168.1.0 through 192.168.1.255.
a range of IP addresses that cannot be specified with a CIDR block or prefix	the IP address range using a hyphen. Do **not** add a space before or after the hyphen.	`192.168.1.1-192.168.1.5` `2001:db8::0202-2001:db8::8329`

negation of any of the other ways to specify IP addresses or ranges of IP addresses	an exclamation point in front of the IP address, block, or range.	`192.168.0.0/32,!192.168.1.10`
		`!2001:db8::/32`
		`!192.168.1.10,!2001:db8::/32`

hosts that are blocked or monitored (but would have been blocked) See Host Profile Icons, on page 662.	In connection and Security Intelligence events, in Initiator IP and Responder IP fields: • block • monitor	--

Related Topics

IP Address Conventions, on page 25

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

If you group devices—whether just on the 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 system accepts specific syntax for port numbers in searches. You can enter:

- a single port number
- a comma-separated list of port numbers
- two port numbers separated by a dash to represent a range of port numbers
- a port number followed by a protocol abbreviation, separated by a forward slash (only when searching for intrusion events)
- a port number or range of port numbers preceded by an exclamation mark to indicate a negation of the specified ports

Note Do **not** use spaces when specifying port numbers or ranges.

Table 101: Port Syntax Examples

Example	Description
`21`	Returns all events on port 21, including TCP and UDP events.
`!23`	Returns all events except those on port 23.
`25/tcp`	Returns all TCP-related intrusion events on port 25.
`21/tcp,25/tcp`	Returns all TCP-related intrusion events on ports 21 and 25.
`21-25`	Returns all events on ports 21 through 25.

Event Fields in Searches

When searching for events, you can use the following fields as search criteria:

- Audit Log Workflow Fields, on page 408
- Application Data Fields, on page 900
- Application Detail Data Fields, on page 902
- Captured File Fields, on page 832

- Allow List Event Fields, on page 930
- Connection and Security-Related Connection Event Fields, on page 731
- Correlation Event Fields, on page 926
- Discovery Event Fields, on page 883
- The Health Events Table, on page 395
- Host Attribute Data Fields, on page 890
- Host Data Fields, on page 885
- File and Malware Event Fields, on page 815
- Intrusion Event Fields, on page 766
- Intrusion Rule Update Log Details, on page 234
- Remediation Status Table Fields, on page 934
- See *Nmap Scan Results Fields* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Server Data Fields, on page 897
- Third-Party Vulnerability Data Fields, on page 908
- User-Related Fields, on page 910
- Vulnerability Data Fields, on page 904
- Allow List Violation Fields, on page 932

Performing a Search

You must have Admin or Security Analyst privileges to perform a search.

Procedure

Step 1 Select **Analysis** > **Search**.

Tip
You may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, select the type of event or data to search.

Step 3 Enter your search criteria in the appropriate fields. See the following sections for detailed information on the search criteria you can use:

- Search Constraints, on page 685
- Audit Log Workflow Fields, on page 408
- Application Data Fields, on page 900

- Application Detail Data Fields, on page 902
- Captured File Fields, on page 832
- Allow List Event Fields, on page 930
- Connection and Security-Related Connection Event Fields, on page 731
- Correlation Event Fields, on page 926
- Discovery Event Fields, on page 883
- The Health Events Table, on page 395
- Host Attribute Data Fields, on page 890
- Host Data Fields, on page 885
- File and Malware Event Fields, on page 815
- Intrusion Event Fields, on page 766
- Intrusion Rule Update Log Details, on page 234
- Remediation Status Table Fields, on page 934
- See *Nmap Scan Results Fields* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Server Data Fields, on page 897
- Third-Party Vulnerability Data Fields, on page 908
- User Data Fields
- User Activity Data Fields
- Vulnerability Data Fields, on page 904
- Allow List Violation Fields, on page 932

Step 4 If you want to use the search again in the future, save the search as described in Saving a Search, on page 691.

Step 5 Click **Search** to start the search. Your search results appear in the default workflow for the table you are searching, constrained by time (if applicable).

What to do next

- To analyze the search results using workflows, see Using Workflows, on page 656.

Related Topics

Configuring Event View Settings, on page 206

Saving a Search

You must have Admin or Security Analyst privileges to save a search.

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

Before you begin

- Establish search criteria as described in Performing a Search, on page 690, or load a saved search as described in Loading a Saved Search, on page 692.

Procedure

Step 1 From the Search page, if you want to save the search as private so only you can access it, check the **Private** checkbox.

Tip
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

Step 2 You have two options:

- If you want to save a new version of a loaded search, click **Save As New**.
- If you want to save a new search, or overwrite a custom search using the same name, click **Save**. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Loading a Saved Search

You must have Admin or Security Analyst privileges to load a saved search.

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

Procedure

Step 1 Choose **Analysis** > **Search**.

Tip
You may also click **Search** from any page on a workflow.

Step 2 From the table drop-down list, choose the type of event or data to search.

Step 3 Choose the search you want to load from the **Custom Searches** list or the **Predefined Searches** list.

Step 4 If you want to use different search criteria, change the search constraints.

Step 5 If you want to use a changed search again in the future, save the search as described in Saving a Search, on page 691.

Step 6 Click **Search**.

Query Overrides Via the Shell

System administrators can use a Linux shell-based query management tool to locate and stop long-running queries.

The query management tool allows you to locate queries running longer than a specified number of minutes and stop those queries. The tool logs an event to the audit log and to syslog when you stop a query.

Note that the `admin` internal user can access the management center CLI. If you use an external authentication object which grants CLI access, users matching the shell access filter can also log into the CLI.

Note Leaving the search page in the web interface does not stop a query. Queries that take a long time to return results impact overall system performance while the query is running.

Shell-Based Query Management Syntax

Use the following syntax to manage long-running queries:

```
query_manager [-v] [-l [minutes]] [-k query_id [...]] [--kill-all minutes]
```

Table 102: `query_manager` Options

Option	Description
`-h, --help`	Prints a brief help message.
`-l, --list [minutes]`	Lists all queries taking longer than passed-in minutes. By default it will show all queries taking longer than 1 minute.
`-k, --kill query_id [...]`	Kills the query with the passed-in id. The option can take multiple ids.
`--kill-all minutes`	Kills all queries taking longer than passed-in minutes.
`-v, --verbose`	Verbose output including full SQL queries.

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

Stopping Long-Running Queries

You must be the **admin** user or externally authenticated user with CLI access

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 693.

History for Searching for Events

Feature	Minimum Management Center	Minimum 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 Intelligence 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: management center

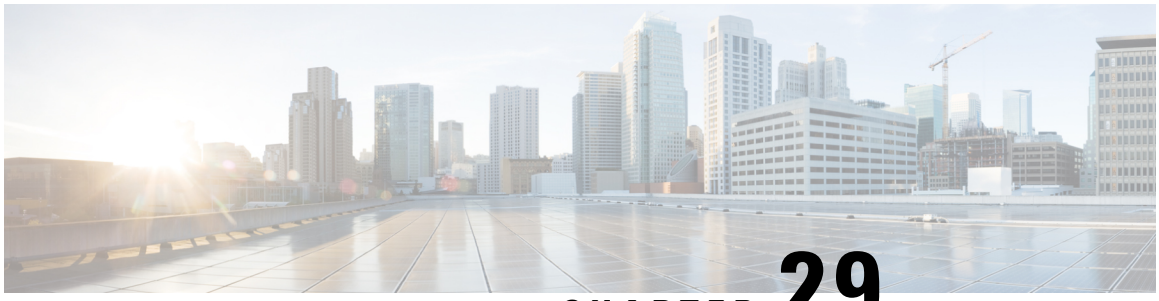

CHAPTER 29

Custom Workflows

The following topics describe how to use custom workflows:

Introduction to Custom Workflows

If the predefined and Cisco-provided custom workflows do not meet your needs, you can create and manage custom workflows.

Custom workflows are workflows that you create to meet the unique needs of your organization. When you create a custom workflow, you choose the kind of event (or database table) on which the workflow is based. On the management center, you can base a custom workflow on a custom table. You can also choose the pages a custom workflow contains; custom workflows can contain drill-down, table view, and host or packet view pages.

If your event evaluation process changes, you can edit custom workflows to meet your new needs. Note that you cannot edit any of the predefined workflows.

Tip You can set a custom workflow as the default workflow for any event type.

Saved Custom Workflows

In addition to predefined workflows, which cannot be modified, the 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 103: 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.

Custom Workflow Creation

If the predefined and Cisco-provided custom workflows do not meet your needs, you can create custom workflows.

Tip Instead of creating a new custom workflow, you can export a custom workflow from another appliance and then import it onto your appliance. You can then edit the imported workflow to suit your needs.

When you create a custom workflow, you:

- Select a table to be the source of the workflow
- Provide a workflow name
- Add drill-down pages and table view pages to the workflow

For each drill-down page in the workflow, you can:

- Provide a name that appears at the top of the page in the web interface
- Include up to five columns per page
- Specify a default sort order, ascending or descending

You can add table view pages in any position in the sequence of workflow pages. They do not have any editable properties, such as a page name, sort order, or user-definable column positions.

Note You must add at least one drill-down page or a table view of events to a custom workflow.

Note If you selected **Vulnerabilities** as the table type, then add **IP Address** as a table column, the IP Address column does not appear when you are viewing vulnerabilities using your custom workflow, unless you use the search feature to constrain the workflow to view a specific IP address or block of addresses.

The final page of a custom workflow depends on the table on which you base the workflow, as described in the following table. These final pages are added by default when you create the workflow.

Table 104: Custom Workflow Final Pages

Event/Asset Type	Final Page
Discovery events	Hosts
Vulnerabilities	Vulnerability detail
Third-party vulnerabilities	Hosts
Users	Users
Indications of compromise	Hosts or users
Intrusion events	Packets

The system does not add a final page to custom workflows based on other kinds of events (for example, audit log or malware events).

Custom workflows based on connection data are like other custom workflows, except you can include drill-down pages containing connection summary data, and connection data graph pages as well as drill-down pages containing data for individual connections and table view pages.

Creating Custom Workflows Based on Non-Connection Data

You must have Admin or Security Analyst privileges to create a custom workflow based on non-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 Optionally, enter a **Description**.

Step 5 Choose the table you want to include from the **Table** drop-down list.

Step 6 If you want to add one or more drill-down pages to the workflow, click **Add Page**.

Step 7 Enter a name for the page in the **Page Name** field.

Step 8 Under Column 1, choose a sort priority and a table column. This column will appear in the leftmost column of the page.

Example:

For example, to create a page showing the destination ports that are targeted, and to sort the page by count, choose **2** from the **Sort Priority** drop-down list and **Destination Port/ICMP Code** from the **Field** drop-down list.

Step 9 Continue choosing fields to include and setting their sort priority until you have specified all the fields you want to appear on the page.

Step 10 If you want to add a table view page to the workflow, click **Add Table View**.

Step 11 Click **Save**.

Creating Custom Connection Data Workflows

Custom workflows based on connection data are like other custom workflows, except you can include connection data graph pages as well as drill-down pages and table view pages. You can include as many of each type of page in the workflow as you want, in any order. Each connection data graph page contains a single graph, which can be a line graph, bar graph, or pie chart. On line and bar graphs, you may include more than one dataset.

You must have Admin privileges to create a custom workflow based on connection data.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 2 Click **Create Custom Workflow**.

Step 3 Enter a name for the workflow in the **Name** field.

Step 4 Optionally, enter a **Description**.

Step 5 From the **Table** drop-down list, choose **Connection Events**.

Step 6 If you want to add one or more drill-down pages to the workflow, you have two options:

- Click **Add Page** to add a drill-down page that contains data on individual connections,
- Click **Add Summary Page** to add a drill-down page that contains connection summary data.

Step 7 Enter a name for the page in the **Page Name** field.

Step 8 Under **Column 1**, choose a sort priority and a table column. This column will appear in the leftmost column of the page.

Step 9 Continue choosing fields to include and setting their sort priority until you have specified all the fields you want to appear on the page.

Example:

For example, to create a page showing the amount of traffic transmitted over your monitored network and to sort the page by the responders that transmitted the most traffic, choose **1** from the **Sort Priority** drop-down list and **Responder Bytes** from the **Field** drop-down list.

Step 10 If you want to add one or more graph pages to the workflow, click **Add Graph**.

Step 11 Enter a name for the page in the **Graph Name** field.

Step 12 Choose the type of graph you want to include on the page:

- line graph (**Line chart** ())
- bar graph(**Bar chart** ())
- pie chart (**Pie chart** ())

Step 13 Specify what kind of data you want to graph by choosing the x- and y-axes of the graph.

On a pie chart, the x-axis represents the independent variable and the y-axis represents the dependent variable.

Step 14 Choose the datasets you want to include on the graph.

Note that pie charts can include only one data set.

Step 15 If you want to add a table view of connection data, click **Add Table View**.

Table views are not configurable.

Step 16 Click **Save**.

Custom Workflow Use and Management

The method you use to view a workflow depends on whether the workflow is based on one of the predefined event tables or on a custom table.

If your custom workflow is based on a predefined event table, access it in the same way that you would access a workflow that ships with the appliance. For example, to access a custom workflow based on the Hosts table, choose **Analysis** > **Hosts heading** > **Hosts**. If, on the other hand, your custom workflow is based on a custom table, you must access it from the Custom Tables page.

If your event evaluation process changes, you can edit custom workflows to meet your new needs. Note that you cannot edit any of the predefined workflows.

Tip You can set a custom workflow as the default workflow for any event type.

Viewing Custom Workflows Based on Predefined Tables

You must have Admin, Maintenance, or Security Analyst privileges to view a custom workflow.

Procedure

Step 1 Choose the appropriate menu path and option for the table on which you based your custom workflow, as described in the Workflow Selection, on page 658.

Step 2 To use a different workflow, including a custom workflow, click **(switch workflow)** next to the current workflow title.

Step 3 If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see Event Time Constraints, on page 672.

Viewing Custom Workflows Based on Custom Tables

You must have Admin or Security Analyst privileges to view a custom workflow that is based on custom tables.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table you want to view, or click the name of the custom table.

Step 3 To use a different workflow, including a custom workflow, click **(switch workflow)** beside the current workflow title.

Step 4 If no events appear and the workflow can be constrained by time, you may need to adjust the time range; see Event Time Constraints, on page 672.

Editing Custom Workflows

You must have Admin or Security Analyst privileges to edit a custom workflow.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Workflows**.

Step 2 Click **Edit** ( ) next to the name of the workflow that you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Make any changes that you want to the workflow.

Step 4 Click **Save**.

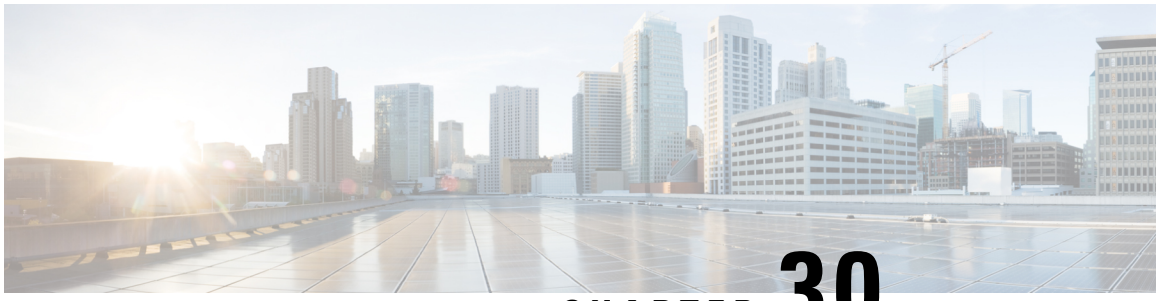

CHAPTER 30

Custom Tables

The following topics describe how to use custom tables:

Introduction to Custom Tables

As the system collects information about your network, the management center stores it in a series of database tables. When you use a workflow to view the resulting information, the management center pulls the data from one of these tables. For example, the columns on each page of the Network Applications by Count workflow are taken from the fields in the Applications table.

If you determine that your analysis of the activity on your network would be enhanced by combining fields from different tables, you can create a custom table.

Note that you can create custom workflows for either predefined or custom tables.

Predefined Custom Tables

Custom tables contain fields from two or more predefined tables. The system is delivered with a number of system-defined custom tables, but you can create additional custom tables that contain only information that matches your specific needs.

For example, the system is delivered with system-defined custom tables that correlate intrusion event data with host data, so you can search for events that impact critical systems and view the results of that search in one workflow.

In a multidomain deployment, the predefined custom tables belong to the Global domain and cannot be modified in lower domains.

The following table describes the custom tables provided with the system.

Table 105: System-Defined Custom Tables

Table	Description
Hosts with Servers	Includes fields from the Hosts and Servers tables, providing you with information about the detected applications running on your network, as well as basic operating system information about the hosts running those applications.

Possible Table Combinations

When you create a custom table, you can combine fields from predefined tables that have related data. The following table lists the predefined tables you can combine to create a new custom table. Keep in mind that you can create a custom table that combines fields from more than two predefined custom tables.

Table 106: 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
Host Indications 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

Tip Instead of creating a new custom table, you can export a custom table from another management center, then import it onto your management center.

To create a custom table, decide which predefined tables contain the fields you want to include in your custom table. You can then choose which fields you want to include and, if necessary, configure field mappings for any common fields.

Tip Data involving the Hosts table allows you to view data associated with all IP addresses from one host, rather than one specific IP address.

For example, consider a custom table that combines fields from the Correlation Events table and the Hosts table. You can use this custom table to get detailed information about the hosts involved in violations of any of your correlation policies. Note that you must decide whether to display data from the Hosts table that matches the source IP address or the destination IP address in the Correlation Events table.

If you view the table view of events for this custom table, it displays correlation events, one per row. You can configure the custom table to include the following information:

- the date and time the event was generated
- the name of the correlation policy that was violated
- the name of the rule that triggered the violation
- the IP address associated with the source, or initiating, host involved in the correlation event
- the source host's NetBIOS name
- the operating system and version the source host is running
- the source host criticality

Tip You could create a similar custom table that displays the same information for destination, or responding, hosts.

Creating a Custom Table

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **Create Custom Table**.

Step 3 In the **Name** field, enter a name for the custom table.

Example:

For example, you might enter `Correlation Events with Host Information (Src IP).`

Step 4 From the **Tables** drop-down list, choose **Correlation Events**.

Step 5 Under **Fields**, choose **Time** and click **Add** to add the date and time when a correlation event was generated.

Step 6 Repeat step 5 to add the **Policy** and **Rule** fields.

Tip
You can use Ctrl or Shift while clicking to choose multiple fields. You can also click and drag to choose multiple adjacent values. However, if you want to specify the order the fields appear in the table view of events associated with the table, add the fields one at a time.

Step 7 From the **Tables** drop-down list, choose **Hosts**.

Step 8 Add the **IP Address**, **NetBIOS Name**, **OS Name**, **OS Version**, and **Host Criticality** fields to the custom table.

Step 9 Under **Common Fields**, next to **Correlation Events**, choose **Source IP**.

Your custom table is configured to display the host information you chose in step 8 for the source, or initiating, hosts involved in correlation events.

Tip
You can create a custom table that displays detailed host information for the destination, or responding, hosts involved in a correlation event by following this procedure but choosing **Destination IP** instead of **Source IP**.

Step 10 Click **Save**.

Modifying a Custom Table

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **Edit** ( ) next to the table you want to edit.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Optionally, remove fields from the table by clicking **Delete** ( ) next to the fields you want to remove.

Note
If you delete fields currently in use in reports, the system will prompt you to confirm that you want to remove the sections using those fields from those reports.

Step 4 Make other changes as needed.

Step 5 Click **Save**.

Deleting a Custom Table

In a multidomain deployment, the system displays custom tables created in the current domain, which you can delete. It also displays custom tables created in ancestor domains, which you cannot delete. To delete custom tables in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **Delete** ( ) next to the custom table you want to delete.

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

Viewing a Workflow Based on a Custom Table

When you create a custom table, the system automatically creates a default workflow for it. The first page of this workflow displays a table view of events. If you include intrusion events in your custom table, the second page of the workflow is the packet view. Otherwise, the second page of the workflow is a hosts page. You can also create your own custom workflows based on your custom table.

Tip If you create a custom workflow based on a custom table, you can specify it as the default workflow for that table.

You can use the same techniques to view events in your custom table that you use for event views based on predefined tables.

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table related to the workflow you want to see.

Searching Custom Tables

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

Procedure

Step 1 Choose **Analysis** > **Advanced** > **Custom Tables**.

Step 2 Click **View** ( ) next to the custom table you want to search.

Tip
To use a different workflow, including a custom workflow, click **(switch workflow)** next the workflow title.

Step 3 Click **Search**.

Tip
To search the database for a different kind of event or data, choose it from the table drop-down list.

Step 4 Enter your search criteria in the appropriate fields.

If you enter criteria for multiple fields, the search returns only the records that match search criteria specified for all fields.

Tip
Click **Object** ( ) next to a search field to use an object as a search criterion.

Step 5 Optionally, if you plan to save the search, you can check the **Private** check box to save the search as private so only you can access it. Otherwise, leave the check box clear to save the search for all users.

Tip
If you want to use the search as a data restriction for a custom user role, you **must** save it as a private search.

Step 6 Optionally, you can save the search to be used again in the future. You have the following options:

- Click **Save** to save the search criteria. The search is visible only to your account if you checked the **Private** check box.
- Click **Save As New** to save a new search or assign a name to a search you created by altering a previously-saved search. The search is saved and visible only to your account if you checked the **Private** check box.

Step 7 Click **Search** to start the search.

Your search results appear in the default workflow for the custom table, constrained by the current time range (if applicable).

History for Custom Tables

Feature	Minimum Management Center	Minimum 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. Affected Platforms: management center

PART VIII

Events and Assets

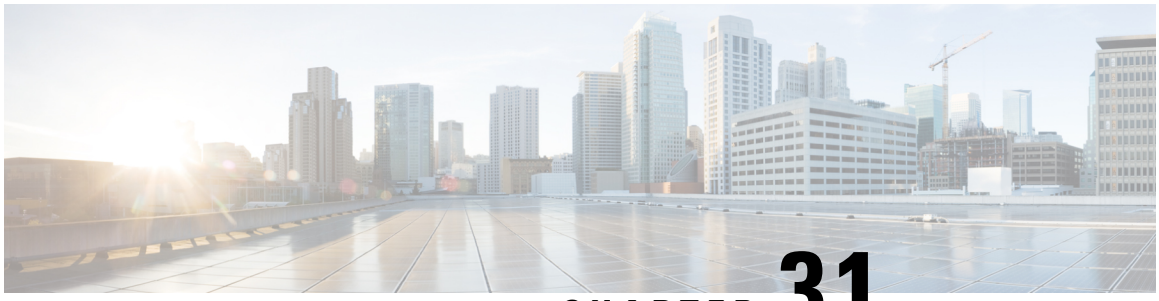

CHAPTER 31

Connection Logging

The following topics describe how to configure the system to log connections made by hosts on your monitored network:

About Connection Logging

The system can generate logs of the connections its managed devices detect. These logs are called *connection events*. Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data. Special connection events, called *security-related connection events*, represent connections that were blocked by the reputation-based Security Intelligence feature.

Connection events contain data about the detected sessions. The information available for any individual connection event depends on several factors, but in general includes:

- Basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- Additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- Metadata about why the connection was logged: which configuration handled the traffic, whether the connection was allowed or blocked, details about encrypted and decrypted connections, and so on

Log connections according to the security and compliance needs of your organization. When setting up connection logging, keep in mind that the system can log a connection for multiple reasons, and that disabling logging in one place does not mean that matching connections will not be logged.

The information in a connection event depends on several factors, including traffic characteristics, the configuration that ultimately handled the connection, and so on.

Note You can supplement the connection logs gathered by your managed devices with connection data generated from exported NetFlow records. This is especially useful if you have NetFlow-enabled routers or other devices deployed on networks that your managed devices cannot monitor.

Connections That Are Always Logged

Unless you disable connection event storage, the system automatically saves the following end-of-connection events to the management center database, regardless of any other logging configurations.

Connections Associated with Intrusions

The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action.

When an intrusion policy associated with the access control default action generates an intrusion event, the system does *not* automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging. This is useful for intrusion prevention-only deployments where you do not want to log any connection data.

However, if you enable beginning-of-connection logging for the default action, the system *does* log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

Connections Associated with File and Malware Events

The system automatically logs connections associated with file and malware events.

Note File events generated by inspecting NetBIOS-SSN (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

Connections Associated with Intelligent Application Bypass

The system automatically logs bypassed and would-have-bypassed connections associated with IAB.

Monitored Connections

The system always logs the ends of connections for monitored traffic, even if the traffic matches no other rules and you do not enable default action logging. For more information, see Logging for Monitored Connections, on page 716.

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.

Decryption Policy: Rules and Default Action

You can log connections that match a decryption rule or decryption policy default action.

For blocked connections, the system immediately ends the session and generates an event. For monitored connections and connections that you pass to access control rules, the system generates an event when the session ends.

Access Control Policy: Security Intelligence Decisions

You can log a connection whenever it is blocked by the reputation-based Security Intelligence feature.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blocked by Security Intelligence, but still log the match. Security Intelligence monitoring also allows you to create traffic profiles using Security Intelligence information.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. So that you can identify the matching IP address in the connection, host icons beside blocked and monitored IP addresses look slightly different in the tables on the pages under the **Analysis** > **Connections** menus.

Access Control Policy: Rules and Default Action

You can log connections that match an access control rule or access control policy default action.

Related Topics

How Rules and Policy Actions Affect Logging, on page 715

How Rules and Policy Actions Affect Logging

Connection events contain metadata about why the connection was logged, including which configurations handled the traffic. Where you can configure connection logging, rule actions, and policy default actions determine not only how the system inspects and handles matching traffic, but also when and how you can log details about matching traffic.

Related Topics

Connection and Security-Related Connection Event Fields, on page 731

Logging for Fastpathed Connections

You can log fastpathed connections and non-encrypted tunnels, which includes traffic matching the following rules and actions in the prefilter policy:

- Tunnel rules—**Fastpath** action (logs the outer session)
- Prefilter rules—**Fastpath** action

Fastpathed traffic bypasses the rest of access control and QoS, so connection events for fastpathed connections contain limited information.

Logging for Monitored Connections

The system always logs the ends of connections for traffic matching the following configurations, even if the traffic matches no other rules and you do not enable default action logging:

- Security Intelligence—Block lists set to monitor (also generates a Security Intelligence event)
- SSL rules—**Monitor** action
- Access control rules—**Monitor** action

The system does not generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching SSL Monitor rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Logging for Trusted Connections

You can log the beginnings and ends of trusted connections, which includes traffic matching the following rules and actions:

- Access control rules—**Trust** action
- Access control default action—**Trust All Traffic**

Note Although you *can* log trusted connections, we recommend you do not do so because trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

TCP connections detected by a Trust rule on the first packet generate only an end-of-connection event. The system generates the event one hour after the final session packet.

Logging for Blocked Connections

You can log blocked connections, which includes traffic matching the following rules and actions:

- Tunnel rules—**Block**

- Prefilter rules—**Block**
- Prefilter default action—**Block all tunnel traffic**
- Security Intelligence—Block lists not set to Monitor (also generates a Security Intelligence event)
- Decryption rules—**Block** and **Block with reset**
- SSL default action—**Block** and **Block with reset**
- Access control rules—**Block**, **Block with reset**, and **Interactive Block**
- Access control default action—**Block All Traffic**

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for an Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Beginning vs End-of-Connection Logging for Blocked Connections

When you log a blocked connection, how the system logs it depends on why the connection was blocked; this is important to keep in mind when configuring correlation rules based on connection logs:

- For decryption rules and decryption policy default actions that block encrypted traffic, the system logs **end**-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session.
- For other blocking actions, the system logs **beginning**-of-connection events. Matching traffic is denied without further inspection.

Logging Bypassed Interactive Blocks

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, allow you to configure end-of-connection logging. This is because if the user clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and log.

Therefore, for packets that match an Interactive Block or Interactive Block with Reset rule, the system can generate the following connection events:

- A beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of `Interactive Block` or `Interactive Block with Reset`
- Multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of `Allow` and a reason of `User Bypass`

  The following figure shows an example of an interactive block followed by allow.

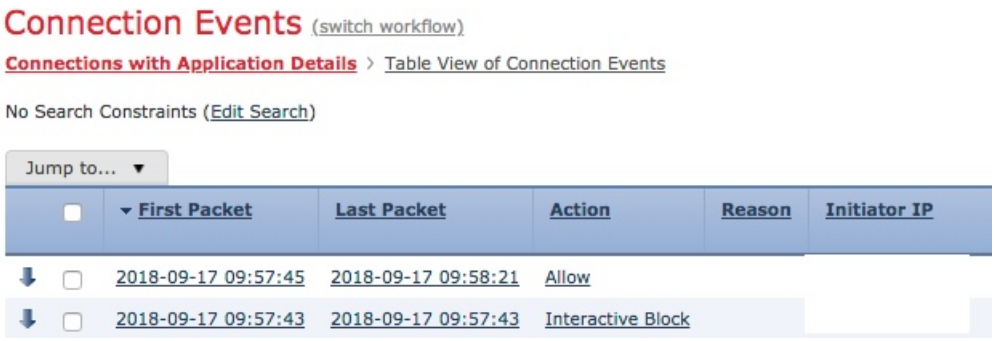

Logging for Allowed Connections

You can log allowed connections, which includes traffic matching the following rules and actions:

- SSL rules—**Decrypt** action
- SSL rules—**Do not decrypt** action
- SSL default action—**Do not decrypt**
- Access control rules—**Allow** action
- Access control default action—**Network Discovery Only** and any intrusion prevention option

Enabling logging for these configurations ensures the connection is logged, while also permitting (or specifying) the next phase of inspection and traffic handling. SSL logging is always end-of-connection; access control configurations also allow beginning-of-connection logging.

Although the **Analyze** action in tunnel and prefilter rules also allows connections to continue with access control, logging is disabled for rules with this action. Matching connections may still be logged by other configurations. Allowed tunnels might have their encapsulated sessions evaluated and logged individually.

When you allow traffic with an access control rule or default action, you can use an associated intrusion policy to further inspect traffic and block intrusions. For access control rules, you can also use a file policy to detect and block prohibited files, 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, see Connections That Are Always Logged, on page 714.

Connections with encrypted payloads are not subject to deep inspection, so connection events for encrypted connections contain limited information.

File and Malware Event Logging for Allowed Connections

When a file policy detects or blocks a file, it logs one of the following events to the management center database:

- *File events*, which represent detected or blocked files, including malware files.
- *Malware events*, which represent detected or blocked malware files only.
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes.

You can disable this logging on a per-access-control-rule basis. You can also disable file and malware event storage entirely.

Note We recommend you leave file and malware event logging enabled.

Beginning vs End-of-Connection Logging

You can log a connection at its beginning or its end, with the following exceptions for blocked traffic:

- Blocked traffic—Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events for blocked traffic. There is no unique end of connection to log.
- Blocked encrypted traffic—When you enable connection logging in a decryption policy, the system logs end-of-connection rather than beginning-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session, and thus cannot immediately block encrypted sessions.

To optimize performance, log either the beginning or the end of any connection, but not both. Monitoring a connection for any reason forces end-of-connection logging. For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session.

The following table details the differences between beginning and end-of-connection events, including the advantages to logging each.

Table 107: Comparing Beginning and End-of-Connection Events

	Beginning-of-Connection Events	**End-of-Connection Events**
Can be generated...	When the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification).	When the system: • Detects the close of a connection. • Does not detect the end of a connection 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 decryption policy.	Most connections.

	Beginning-of-Connection Events	End-of-Connection Events
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 a decryption 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

If you store connection and Security Intelligence event logs on the 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 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 the following:

- For external logging on any device:

  A connection you configure called an *alert response*.

- For external logging on threat defense devices:

  See *About Configuring Syslog* and *Configure SNMP Traps* in the Cisco Secure Firewall Management Center Device Configuration Guide.

- For additional options related to external logging:

  See Event Analysis Using External Tools, on page 617.

Related Topics

Secure Firewall Management Center Alert Responses, on page 551

Limitations of Connection Logging

You cannot log:

- 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 Overview, on page 441).

If a connection event does not contain the information you think it should, see Requirements for Populating Connection Event Fields, on page 749 and Information Available in Connection Event Fields, on page 751.

When Events Appear in the Event Viewer

The following points are applicable to all types of events:

- If you are looking at a page under the Analysis menu, you must refresh the page to display new events.
- Events generally are available for viewing within a few seconds of the time the traffic was detected. However, there can be an arbitrary delay under situations such as: Exceptionally heavy traffic conditions; the 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

Use the following best practices to ensure that you log *only* the connections you want to log.

So that you log only critical connections, enable connection logging on a per-access-control-rule basis.

Connections that are always logged

The system automatically logs the following:

- Some connections associated with detected files, malware, intrusions, and Intelligent Application Bypass (IAB).

  For more information, see Connections That Are Always Logged, on page 714.

- Monitored connections.

  For more information, see Logging for Monitored Connections, on page 716.

Connections to never log

Do not enable logging for the following:

- Access control rules with a Trust action.

  Trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

- Do not enable logging for Block rules in passive deployments. To log connections that the system *would have blocked* if your devices were deployed inline, use a Monitor rule instead of a Block rule.

  Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

- Traffic you're not interested in. Examples follow:

  - Specific allowed traffic, such as DNS requests to a trusted DNS host.
  - Infrastructure traffic that is not related to your service offering.

  (As previously mentioned, you can still monitor this traffic for threats.)

As discussed in Connections That Are Always Logged, on page 714, even if you disable logging for the preceding, intrusion events, malware, and IAB are still logged.

Avoid logging what's being logged elsewhere

If another device or service is logging connection data for a network segment, disable logging for that segment's data in the management center. Examples follow:

- If a router logs connection events on the same network segment as the management center, avoid logging the same connections on the management center unless you need those connection events for something else, such as correlation policies or traffic profiles.

  For more information about correlation policies, see Introduction to Correlation Policies and Rules, on page 953. For more information about traffic profiles, see Introduction to Traffic Profiles, on page 991.

- 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, see Beginning vs End-of-Connection Logging, on page 719.

Logging for blocked traffic

Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events.

For more information, see Logging for Blocked Connections, on page 716.

Log events to an external location

If your company's security policy permits it, you can save disk space on your management center by streaming logs to an external source using any of the following:

- eStreamer, which enables you to stream logs from a management center to a custom-developed client application. For more information, see the *Secure Firewall Management Center Event Streamer Integration Guide* .
- Syslog or SNMP trap, which are referred to as *alert responses*. For more information, see Secure Firewall Management Center Alert Responses, on page 551.

Specify the maximum number of event records

Consider the minimum and maximum number of records that can be stored in the database. For example, a virtual management center by default stores 10 million events but the maximum number of events is 50 million. Go to **System** > **Configuration** > **Database** to adjust the size to meet your needs.

For a list of all management center models and their event database sizes, see Database Event Limits, on page 59.

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 management center and click **User Preferences** > **Event View Settings**. The maximum you can set is 1000 events per page.

Set up connection event reports

To make sure you do not miss connection events, you can set up automated reports in .csv format and optionally schedule them to occur at a regular interval. For more information, see the following:

- Use the report designer (**Analysis** > **Connection** > **Events** > **Report Designer**): About Designing Reports, on page 526.
- Schedule tasks (**System** > **Tools** > **Scheduling**): About Task Scheduling, on page 487.

Requirements and Prerequisites for Connection Logging

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configure Connection Logging

The following sections describe how to set up connection logging to match various rules and conditions.

Logging Connections with Tunnel and Prefilter Rules

The prefilter policy applies to 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.

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.

Logging Decryptable Connections with TLS/SSLDecryption Rules

Procedure

Step 1 In the decryption 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 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 decryption policy.

What to do next

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

Logging Connections with Security Intelligence

The Security Intelligence policy requires the Threat Smart License or Protection Classic License.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

Step 2 Click 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-Related connection events.

Send events to the event viewer if you want to perform 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; see How Rules and Policy Actions Affect Logging, on page 715.

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 **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.

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 management center (or other device manager) if you want to perform 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 **Confirm**.

Step 7 Click **Apply** 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 decryption policies, which match and log—but do not handle or inspect—traffic).

Logging settings for the decryption 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 **Default Logging and Inspection** ⚙ 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 management center-based analysis on these connection events.

Step 4 Click **Apply**.

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 **More** > **Advanced Settings**, 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.

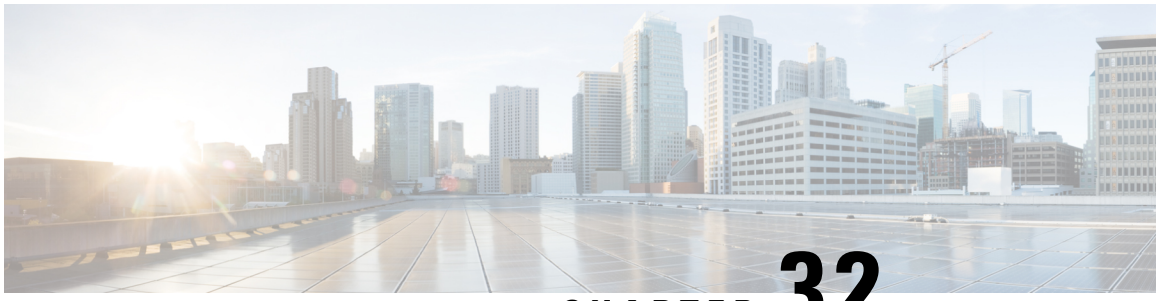

CHAPTER 32

Connection and Security-Related Connection Events

The following topics describe how to use connection and security events tables.

About Connection Events

The system can generate logs of the connections its managed devices detect. These logs are called *connection events*. Connection events include *Security-Related connection events* (connections blocked by the reputation-based Security Intelligence feature.)

Connection events generally include transactions detected by:

- Access control policies
- Decryption policies
- Prefilter policies (captured by prefilter or tunnel rules)
- DNS Block lists
- URL Block lists
- Network (IP address) Block lists

Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data.

For detailed information, see Connection Logging, on page 713.

Related Topics

About Security Intelligence

Connection vs. Security-Related Connection Events

A *Security-Related connection events* is a connection event that is generated whenever a session is blocked or monitored by the reputation-based Security Intelligence feature.

However, for every Security-Related connection event, there is an identical connection event. You can view and analyze Security-Related connection events independently. The system also stores and prunes Security-Related connection events separately.

Note that the system enforces Security Intelligence before more resource-intensive evaluations. When a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.

Note In this guide, information about connection events also pertains to Security-Related connection events, unless otherwise noted.

NetFlow Connections

To supplement the connection data gathered by your managed devices, you can use records broadcast by NetFlow exporters to generate connection events. This is especially useful if the NetFlow exporters are monitoring different networks than those monitored by your managed devices.

The system logs NetFlow records as unidirectional end-of-connection events in the Secure Firewall Management Center database. The available information for these connections differs somewhat from connections detected by your access control policy; see Differences between NetFlow and Managed Device Data.

Related Topics

NetFlow Data

Connection Summaries (Aggregated Data for Graphs)

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.

Note that there are no connection summaries specifically for Security-Related connection events, although corresponding end-of-connection events can be aggregated into connection summary data.

To be aggregated, multiple connections must:

- represent the end of connections
- have the same source and destination IP addresses, and use the same port on the responder (destination) host
- use the same protocol (TCP or UDP)
- use the same application protocol
- either be detected by the same managed device or by the same NetFlow exporter

Each connection summary includes total traffic statistics, as well as the number of connections in the summary. Because NetFlow exporters generate unidirectional connections, a summary's connection count is incremented by two for every connection based on NetFlow data.

Note that connection summaries do not contain all of the information associated with the summaries' aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Long-Running Connections

If a monitored session spans two or more five-minute intervals over which connection data is aggregated, the connection is considered a *long-running connection*. When calculating the number of connections in a connection summary, the system increments the count only for the five-minute interval in which a long-running connection was initiated.

Also, when calculating the number of packets and bytes transmitted by the initiator and responder in a long-running connection, the system does not report the number of packets and bytes that were actually transmitted during each five-minute interval. Instead, the system assumes a constant rate of transmission and calculates estimated figures based on the total number of packets and bytes transmitted, the length of the connection, and what portion of the connection occurred during each five-minute interval.

Combined Connection Summaries from External Responders

To reduce the space required to store connection data and speed up the rendering of connection graphs, the system combines connection summaries when:

- one of the hosts involved in the connection is not on your monitored network
- other than the IP address of the external host, the connections in the summaries meet the summary aggregation criteria

When viewing connection summaries in the Analysis > Connections submenu pages, and when working with connection graphs, the system displays `external` instead of an IP address for the non-monitored hosts.

As a consequence of this aggregation, if you attempt to drill down to the table view of connection data (that is, access data on individual connections) from a connection summary or graph that involves an external responder, the table view contains no information.

Connection and Security-Related Connection Event Fields

Note You cannot use the connection/Security-Related connection 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 Secure 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 Secure 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 Secure 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 can be **AppID** or **Encrypted Visibility**.

Destination Port/ICMP Code (Syslog: Separate fields - DstPort, ICMPCode)

In the Secure 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.

Detection Type

This field shows the source of detection of a client.

Device

In the Secure 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, see DNS Filtering: Identify URL Reputation and Category During DNS Lookup.

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 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)

The 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 Secure 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 Secure 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

In the Secure 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 Secure 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

Intrusion Events (Syslog: IPSCount)

The number of intrusion events, if any, associated with the connection.

In the Secure 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 Secure 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, see Connection Event Reasons, on page 747.

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 Secure 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, see Security Intelligence Categories.

Source Device

In the Secure 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 Secure 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
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 Secure 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 Secure 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 Secure 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, see 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, see About 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 Secure 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 Secure 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 Secure 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.

About Connection and Security-Related Connection Event Fields

In the Secure Firewall Management Center web interface, you can view and search connection and Security-Related connection events using tabular and graphical workflows under the **Analysis > Connections** submenus.

Note For each Security-Related connection event, there is an identical, separately stored connection event. All Security-Related connection event have a populated **Security Intelligence Category** field.

The information available for any individual event can vary depending on how, why, and when the system logged the connection.

Search Constraints

Fields marked with an asterisk (*****) on search pages constrain connection graphs and connection summaries. Because connection graphs are based on connection summaries, the same criteria that constrain connection summaries also constrain connection graphs. If you search connection summaries using invalid search constraints and view your results using a connection summary page in a custom workflow, the invalid constraints are labeled as not applicable (N/A) and are marked with a strikethrough.

Syslog Fields

Most fields appear both in the 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.

A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields

Table 108: Comparison of Terms

Fields	Event Type	Description
Initiator/Responder	Connection	Initiator/responder of the connection. The initiator of a connection is not necessarily the same as the source of an intrusion or the sender of a malware file.
Source/Destination	Intrusion	Source/destination of the attack. The source of an intrusion event can be the initiator or the responder of the connection.
Sender/Receiver (Sending..., Receiving...)	File, Malware	Sender/receiver of a file or malware. The sender of a file is not necessarily the initiator of the connection, as a file may be uploaded or downloaded.

Connection Event Reasons

The Reason field in a connection event displays the reason or reasons the connection was logged, in the following situations:

Reason	Description
Content Restriction	The system modified the packet to enforce content restrictions related to 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 threat defense CLI using the **system support elephant-flow-detection** command. For more information, see 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, see the *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

The information available for a connection event, Security-Related connection events, or connection summary depends on several factors.

Appliance Model and License

Many features require that you enable specific licensed capabilities on target devices, and many features are only available on some models.

Traffic Characteristics

The system only reports information present (and detectable) in network traffic. For example, there could be no user associated with an initiator host, or no referenced host detected in a connection where the protocol is not DNS, HTTP, or HTTPS.

Origin/Detection Method: Traffic-Based Detection vs NetFlow

With the exception of NetFlow-only fields, the information available in NetFlow records is more limited than the information generated by traffic-based detection; see Differences between NetFlow and Managed Device Data.

Evaluation Stage

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance.

For example, the system enforces Security Intelligence before more resource-intensive evaluations. When a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.

Logging Method: Beginning or End of Connection

When the system detects a connection, whether you can log it at its beginning or its end (or both) depends on how you configure the system to detect and handle it.

Beginning-of-connection events do not have information that must be determined by examining traffic over the duration of the session (for example, the total amount of data transmitted or the timestamp of the last packet in the connection). Beginning-of-connection events are also not guaranteed to have information about application or URL traffic in the session, and do not contain any details about the session's encryption. Beginning-of-connection logging is usually the only option for blocked connections.

Connection Event Type: Individual vs Summary

Connection summaries do not contain all of the information associated with their aggregated connections. For example, because client information is not used to aggregate connections into connection summaries, summaries do not contain client information.

Keep in mind that connection graphs are based on connection summary data, which use only end-of-connection logs. If your system is configured to log only beginning-of-connection data, connection graphs and connection summary event views contain no data.

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 a decryption 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 in this topic indicates when the system can populate connection and Security Intelligence fields. The columns in the table represent the following event types:

- Origin: Direct—Events that represent connections detected and handled by a System managed device.
- Origin: NetFlow—Events that represent connections exported by a NetFlow exporter.
- Logging: Start—Events that represent connections logged at their beginning.
- Logging: End—Events that represent connections logged at their end.

A "yes" in the table does not mean that the system must populate a connection event field, rather, that it can. The system only reports information present (and detectable) in network traffic. For example, TLS/SSL-related fields are populated only for records of encrypted connections handled by a decryption policy.

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Access Control Policy	yes	no	yes	yes
Access Control Rule	yes	no	yes	yes
Action	yes	no	yes	yes
Application Protocol	yes	yes	if available	yes
Application Protocol Category & Tag	yes	no	if available	yes
Application Risk	yes	no	if available	yes
Business Relevance	yes	no	if available	yes
Client	yes	no	if available	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Client Category & Tag	yes	no	if available	yes
Client Version	yes	no	if available	yes
Connections	yes	yes	no	yes
Count	yes	yes	yes	yes
Destination Port/ICMP Type	yes	yes	yes	yes
Destination SGT	yes	no	yes	yes
Device	yes	yes	yes	yes
Domain	yes	yes	yes	yes
DNS Query	yes	no	yes	yes
DNS Record Type	yes	no	yes	yes
DNS Response	yes	no	yes	yes
DNS Sinkhole Name	yes	no	yes	yes
DNS TTL	yes	no	yes	yes
Egress Interface	yes	no	yes	yes
Egress Security Zone	yes	no	yes	yes
Endpoint Location	yes	no	yes	yes
Endpoint Profile	yes	no	yes	yes
Files	yes	no	no	yes
First Packet	yes	yes	yes	yes
HTTP Referrer	yes	no	no	yes
HTTP Response Code	yes	no	yes	yes
Ingress Interface	yes	no	yes	yes
Ingress Security Zone	yes	no	yes	yes
Initiator Bytes	yes	yes	not useful	yes
Initiator Country	yes	no	yes	yes
Initiator IP	yes	yes	yes	yes
Initiator Packets	yes	yes	not useful	yes
Initiator User	yes	yes	yes	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Intrusion Events	yes	no	no	yes
Intrusion Policy	yes	no	yes	yes
IOC (Indication of Compromise)	yes	no	yes	yes
Last Packet	yes	yes	no	yes
NetBIOS Domain	yes	no	yes	yes
NetFlow Source/Destination Autonomous System	no	yes	no	yes
NetFlow Source/Destination Prefix	no	yes	no	yes
NetFlow Source/Destination TOS	no	yes	no	yes
NetFlow SNMP Input/Output	no	yes	no	yes
Network Analysis Policy	yes	no	yes	yes
Original Client Country	yes	no	yes	yes
Original Client IP	yes	no	yes	yes
Prefilter Policy	yes	no	yes	yes
QoS-Applied Interface	yes	no	no	yes
QoS-Dropped Initiator Bytes	yes	no	no	yes
QoS-Dropped Initiator Packets	yes	no	no	yes
QoS-Dropped Responder Bytes	yes	no	no	yes
QoS-Dropped Responder Packets	yes	no	no	yes
QoS Policy	yes	no	no	yes
QoS Rule	yes	no	no	yes
Reason	yes	no	yes	yes
Referenced Host	yes	no	no	yes
Responder Bytes	yes	yes	not useful	yes
Responder Country	yes	no	yes	yes
Responder IP	yes	yes	yes	yes
Responder Packets	yes	yes	not useful	yes
Security Context (ASA only)	yes	no	yes	yes

Connection Event Field	Origin: Direct	Origin: NetFlow	Logging: Start	Logging: End
Security Intelligence Category	yes	no	yes	yes
Source Device	yes	yes	yes	yes
Source Port/ICMP Type	yes	yes	yes	yes
Source SGT	yes	no	yes	yes
SSL Certificate Status	yes	no	no	yes
SSL Cipher Suite	yes	no	no	yes
SSL Flow Error	yes	no	no	yes
SSL Flow Flags	yes	no	no	yes
SSL Flow Messages	yes	no	no	yes
Decryption Policy	yes	no	no	yes
Decryption Rule	yes	no	no	yes
SSL Session ID	yes	no	no	yes
SSL Status	yes	no	no	yes
SSL Version	yes	no	no	yes
TCP Flags	no	yes	no	yes
Time	yes	yes	no	yes
Tunnel/Prefilter Rule	yes	no	yes	yes
URL	yes	no	if available	yes
URL Category	yes	no	if available	yes
URL Reputation	yes	no	if available	yes
User Agent	yes	no	no	yes
VLAN ID	yes	no	yes	yes
Web Application	yes	no	if available	yes
Web Application Category & Tag	yes	no	if available	yes

Using Connection and Security-Related Connection Event Tables

You can use the Secure Firewall Management Center to view a table of connection or Security-Related connection events. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access connection graphs differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

When you are using a connection or Security Intelligence workflow table, you can perform many common actions.

Note that when you constrain connection events on a drill-down page, the packets and bytes from identical events are summed. However, if you are using a custom workflow and did not add a **Count** column to a drill-down page, the events are listed individually and packets and bytes are not summed.

Note that **Connection Events** table view displays **1 of Many** instead of how many pages of events are available if your system generates more than 25 connection events.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose either of the following:

- **Analysis** > **Connections** > **Events** (for connection events)
- **Analysis** > **Connections** > **Security-Related Events**

Note
If a connection graph appears instead of a table, click **(switch workflow)** by the workflow title, and choose the predefined **Connection Events** workflow, or a custom workflow. Note that all predefined connection event workflows—including connection graphs—terminate in a table view of connections.

Step 2 You have the following choices:

- Time Range — To adjust the time range, which is useful if no events appear, see Changing the Time Window, on page 675.
- 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, see Work in Secure Firewall Management Center with Connection Events Stored on a Secure Network Analytics Appliance, on page 665.
- Field Names — To learn more about the contents of the columns in the table, see Connection and Security-Related Connection Event Fields, on page 731.

Tip

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 information, see Event Investigation Using Web-Based Resources, on page 620
- 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, see Event Investigation Using Web-Based Resources, on page 620.
- Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, **Compromised Host** that appears next to the IP address.
- User Profile — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- Files and Malware —To view the files, including malware, detected or blocked in a connection, click **View Files** and proceed as described in Viewing Files and Malware Detected in a Connection, on page 757.
- Intrusion Events — To view the intrusion events associated with a connection, as well as their priority and impact, click **Intrusion Events** in the **Intrusion Events** column and proceed as described in Viewing Intrusion Events Associated with a Connection, on page 758.

  Tip
  To quickly view intrusion, file, or malware events associated with one or more connections, check the connections using the check boxes in the table, then choose the appropriate option from the **Jump to** drop-down list. Note that because they are blocked before access control rule evaluation, there can be no files or intrusions associated with connections blocked by Security Intelligence. You can only see this information for a Security Intelligence event if you configured Security Intelligence to monitor, rather than block, connections.

- Certificate — To view details about an available certificate used to encrypt a connection, click **Enabled Lock** in the **SSL Status** column.
- Constrain — To constrain the columns that appear, click **Close** ( ) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip
  To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Delete Events — (Security-Related connection 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 — See Using Drill-Down Pages, on page 664.

  Tip

To drill down using one of several Monitor rules that matched a logged connection, click an *N* **Monitor Rules** value. In the pop-up window that appears, click the Monitor rule you want to use to constrain connection events.

- Navigate This Page — See Workflow Page Traversal Tools, on page 661.
- Navigate Between Pages — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate Between Event Views — To navigate to other event views to view associated events, click **Jump to** and choose the event view from the drop-down list.
- Sort — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.

Related Topics

Overview: Workflows, on page 647
Configuring Event View Settings, on page 206

Viewing Files and Malware Detected in a Connection

If you associate a file policy with one or more access control rules, the system can detect files (including malware) in matching traffic. Use the Analysis > Connections menu options to see the file events, if any, associated with the connections logged by those rules. Instead of a list of files, the 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.

Procedure

Step 1 Go to **Analysis > Connections** and choose the relevant option.

Step 2 While using a connection event table, click **View Files**.
A pop-up window appears with a list of the files detected in the connection as well as their types, and if applicable, their malware dispositions.

Step 3 You have the following choices:

- View — To view a table view of file events, click a **File's View**.
- View — To view details in a table view of malware events, click a **Malware File's View**.
- Track — To track the file's transmission through your network, click a **File's Trajectory**.
- View — To view details on all of the connection's detected file or malware events detected by malware defense ("network-based malware events"), click **View File Events** or **View Malware Events**.

Related Topics

Overview: Workflows, on page 647
Configuring Event View Settings, on page 206

Viewing Intrusion Events Associated with a Connection

If you associate an intrusion policy with an access control rule or default action, the system can detect exploits in matching traffic. Use the Analysis > Connections menu options to see the intrusion events, if any, associated with logged connections, as well as their priority and impact.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Go to **Analysis > Connections** and choose the relevant option.

Step 2 While using a connection event table, click **Intrusion Events** in the **Intrusion Events** column.

Step 3 In the pop-up window that appears, you have the following options:

- Click a **Listed Event's View** to view details in the packet view.
- Click **View Intrusion Events** to view details on all of the connection's associated intrusion events.

Related Topics

Overview: Workflows, on page 647
Configuring Event View Settings, on page 206

Encrypted Connection Certificate Details

You can use options under the Analysis > Connections menu to display the public key certificate (if available) used to encrypt a connection handled by the system. The certificate contains the following information.

Table 109: 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

Overview: Workflows, on page 647
Configuring Event View Settings, on page 206

Viewing the Connection Summary Page

The Connection Summary page is visible only to users who have custom roles that are restricted by searches on connection events and who have been granted explicit menu-based access to the Connection Summary page. This page provides graphs of the activity on your monitored network organized by different criteria. For example, the Connections over Time graph displays the total number of connections on your monitored network over the interval that you choose.

You can perform almost all the same actions on connection summary graphs that you can perform on connection graphs. However, because the graphs on the Connection Summary page are based on aggregated data, you cannot examine the individual connection events on which the graphs are based. In other words, you cannot drill down to a connection data table view from a connection summary graph.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Connection Summary**.

Step 2 From the **Select Device** list, choose the device whose summary you want to view, or choose **All** to view a summary of all devices.

Step 3 To manipulate and analyze the connection graphs, proceed as described in Using Connection Event Graphs, on page 667.

Tip

To detach a connection graph so you can perform further analysis without affecting the default time range, click **View**.

Related Topics

Enable User Role Escalation, on page 200

History for Connection and Security Intelligence Events

Feature	Minimum Management Center	Minimum Threat Defense	Details
New Connection Event Reason - Elephant Flow.	7.1	Any	See Connection Event Reasons, on page 747.
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	See History for Workflows, on page 683.
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. • See also *DNS Filtering and Events* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Feature	Minimum Management Center	Minimum Threat Defense	Details
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: 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, see Data Purge and Storage, on page 515. New/Modified screens: **Analysis > Connections > Events** Platform: 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)

New and changed Security Group Tag fields	6.5	Any	Changes to fields in the 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: 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.

Feature	Minimum Management Center	Minimum Threat Defense	Details
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.

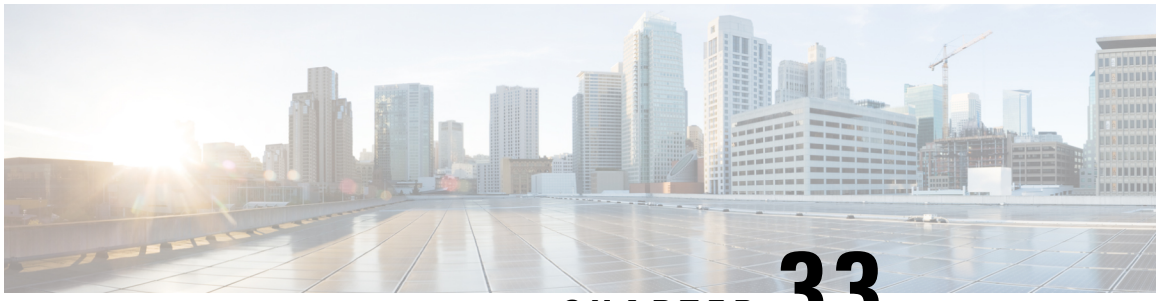

CHAPTER 33

Intrusion Events

The following topics describe how to work with intrusion events.

About Intrusion Events

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

You can use the following tools to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

- An event summary page that gives you an overview of the current activity on your managed devices

- Text-based and graphical reports that you can generate for any time period you choose; you can also design your own reports and configure them to run at scheduled intervals
- An incident-handling tool that you can use to gather event data related to an attack; you can also add notes to help you track your investigation and response
- Automated alerting that you can configure for SNMP, email, and syslog
- Automated correlation policies that you can use to respond to and remediate specific intrusion events
- Predefined and custom workflows that you can use to drill down through the data to identify the events that you want to investigate further
- External tools for managing and analyzing data. You can send data to those tools using syslog or eStreamer. For more information, see Event Analysis Using External Tools, on page 617

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, see https://www.snort.org/rule_docs/.

License Requirements for Intrusion Events

Threat Defense License

IPS

Classic License

Protection

Requirements and Prerequisites for Intrusion Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Viewing Intrusion Events

You view an intrusion event to determine whether there is a threat to your network security.

The initial intrusion events view differs depending on the workflow you use to access the page. You can use one of the predefined workflows, which includes one or more drill-down pages, a table view of intrusion events, and a terminating packet view, or you can create your own workflow. You can also view workflows based on custom tables, which may include intrusion events.

An event view may be slow to display if it contains a large number of IP addresses and you have enabled the **Resolve IP Addresses** event view setting.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 You have the following choices:

- Adjust time range — Adjust the time range for the event view as described in Changing the Time Window, on page 675.
- Change workflows — If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- Constrain — To narrow your view to the intrusion events that are important to your analysis, see Using Intrusion Event Workflows, on page 784.
- Delete event — To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Mark reviewed — To mark intrusion events reviewed, see Marking Intrusion Events Reviewed, on page 779.
- View connection data —To view connection data associated with intrusion events, see Viewing Connection Data Associated with Intrusion Events, on page 778.
- View contents — To view the contents of the columns in the table as described in Intrusion Event Fields, on page 766.

Related Topics

Using the Intrusion Event Packet View, on page 787

About Intrusion Event Fields

When the system identifies a possible intrusion, it generates an *intrusion event*, which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded.

You can view intrusion event data in the 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

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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

Destination Port / ICMP Code (Syslog: DstPort, ICMPCode)

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, this field displays the ICMP code.

Destination User

The username associated with the Responder IP of the connection event. This host may or may not be the host receiving the exploit. This value is typically known only for users on your network.

.

See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

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 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

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 110: 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.

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 111: 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

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** (Decryption 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.

Related Topics

Event Searches, on page 685

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 the impact levels.

Table 112: Impact Levels

Impact Level	Vulnerability	Color	Description
Unknown (0)	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.
Vulnerable (1)	Vulnerable	red	Either: • the source or the destination host is in the network map, and a vulnerability is mapped to the host • the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software
Potentially Vulnerable (2)	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of the following is true: • for port-oriented traffic, the port is running a server application protocol • for non-port-oriented traffic, the host uses the protocol
Currently Not Vulnerable (3)	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of the following is true: • for port-oriented traffic (for example, TCP or UDP), the port is not open • for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
Unknown Target (4)	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.

Viewing Connection Data Associated with Intrusion Events

The system can log the connections where intrusion events are detected. Although this logging is automatic for intrusion policies associated with access control rules, you must manually enable connection logging to see associated connection data for the default action.

Viewing associated data is most useful when navigating between table views of events.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 Choose the intrusion events using the check boxes in the table, then choose **Connections** from the **Jump to** drop-down list.

Tip
You can view the intrusion events associated with particular connections in a similar way. For more information, see Inter-Workflow Navigation, on page 680.

Related Topics

Logging for Allowed Connections, on page 718
Using Intrusion Event Workflows, on page 784
Using Connection and Security-Related Connection Event Tables, on page 755

Marking Intrusion Events Reviewed

If you are confident that an intrusion event is not malicious, you can mark the event reviewed.

If you have examined an intrusion event and are confident that the event does not represent a threat to your network security (for example, because you know that none of the hosts on your network are vulnerable to the detected exploit), you can mark the event reviewed. Reviewed events are stored in the event database and are included in the event summary statistics, but no longer appear in the default intrusion event pages. Your name appears as the reviewer.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

If you perform a backup and then delete reviewed intrusion events, restoring your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under **Intrusion Events**, not under **Reviewed Events**.

Procedure

On a page that displays intrusion events, you have two options:

- To mark one or more intrusion events from the list of events, check the check boxes next to the events and click **Review**.
- To mark all intrusion events from the list of events, click **Review All**.

Related Topics

Using Intrusion Event Workflows, on page 784

Viewing Previously Reviewed Intrusion Events

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Reviewed Events**.

Step 2 You have the following choices:

- Adjust the time range as described in Changing the Time Window, on page 675.
- If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- To learn more about the events that appear, see Intrusion Event Fields, on page 766.

Related Topics

Using Intrusion Event Workflows, on page 784

Marking Reviewed Intrusion Events Unreviewed

You can return a reviewed event to the default intrusion events view by marking the event unreviewed.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

On a page that displays reviewed events, you have two choices:

- To remove individual intrusion events from the list of reviewed events, check the check boxes next to specific events and click **Unreview**.
- To remove all intrusion events from the list of reviewed events, click **Unreview All**.

Preprocessor Events

Preprocessors provide two functions: performing the specified action on the packet (for example, decoding and normalizing HTTP traffic) and reporting the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled. For example, you can enable the `Double Encoding` HTTP Inspect option and the associated preprocessor rule

with the HTTP Inspect Generator (GID) 119 and the Snort ID (SID) 2 to generate an event when the preprocessor encounters IIS double-encoded traffic.

Generating events to report the execution of preprocessors helps you detect anomalous protocol exploits. For example, attackers can craft overlapping IP fragments to cause a DoS attack on a host. The IP defragmentation preprocessor can detect this type of attack and generate an intrusion event for it.

Preprocessor events differ from rule events in that the packet display does not include a detailed rule description for the event. Instead, the packet display shows the event message, the GID, SID, the packet header data, and the packet payload. This allows you to analyze the packet's header information, determine if its header options are being used and if they can exploit your system, and inspect the packet payload. After the preprocessors analyze each packet, the rules engine executes appropriate rules against it (if the preprocessor was able to defragment it and establish it as part of a valid session) to further analyze potential content-level threats and report on them.

Preprocessor Generator IDs

Each preprocessor has its own Generator ID number, or GID, that indicates which preprocessor was triggered by the packet. Some of the preprocessors also have related SIDs, which are ID numbers that classify potential attacks. This helps you analyze events more effectively by categorizing the type of event much the way a rule's Snort ID (SID) can offer context for packets triggering rules. You can list preprocessor rules by preprocessor in the Preprocessors filter group on the intrusion policy Rules page; you can also list preprocessor rules in the preprocessor and packet decoder sub-groupings in the Category filter group.

Note Events generated by standard text rules have a generator ID of 1 (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.

The following table describes the types of events that generate each GID.

Table 113: 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.

ID	Component	Description
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.
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.

ID	Component	Description
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

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 system provides a set of predefined workflows, populated with event data, that you can use to view and analyze intrusion events. Each of these workflows steps you through a series of pages to help you pinpoint the intrusion events that you want to evaluate.

The predefined intrusion event workflows contain three different types of pages, or event views:

- one or more drill-down pages
- the table view of intrusion events
- a packet view

Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information.

When you "drill down" to find more information for one or more destination ports, you automatically select those events and the next page in the workflow appears. In this way, drill-down tables help you reduce the number of events you are analyzing at one time.

The initial *table view* of intrusion events lists each intrusion event in its own row. The columns in the table list information such as the time, the source IP address and port, the destination IP address and port, the event priority, the event message, and more.

When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze.

For example, if you click **Close** ( ) in any column and clear **Time** from the drop-down list, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated from one of the source IP addresses (presumably, a potential attacker), click the IP address in the **Source IP Address** column.

If you select one or more rows in a table view and then click **View**, the packet view appears. A *packet view* provides information about the packet that triggered the rule or the preprocessor that generated the event. Each section of the packet view contains information about a specific layer in the packet. You can expand collapsed sections to see more information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

If the predefined workflows do not meet your specific needs, you can create custom workflows that display only the information you are interested in. Custom intrusion event workflows can include drill-down pages, a table view of events, or both; the system automatically includes a packet view as the last page. You can easily switch between the predefined workflows and your own custom workflows depending on how you want to investigate events.

Using Intrusion Event Workflows

The drill-down views and table view of events share some common features that you can use to narrow a list of events and then concentrate your analysis on a group of related events.

To avoid displaying the same intrusion events on different workflow pages, the time range pauses when you click a link at the bottom of the page to display another page of events, and resumes when you click to take any other action on the subsequent page.

Tip At any point in the process, you can save the constraints as a set of search criteria. For example, if you find that over the course of a few days your network is being probed by an attacker from a single IP address, you can save your constraints during your investigation and then use them again later. You cannot, however, save compound constraints as a set of search criteria.

Procedure

Step 1 Access an intrusion event workflow using **Analysis** > **Intrusions** > **Events**.

Step 2 Optionally, constrain the number of intrusion events that appear on the event views as described in Intrusion Event Drill-Down Page Constraints, on page 785 or Intrusion Event Table View Constraints, on page 786.

Step 3 You have the following choices:

- To learn more about the columns that appear, see Intrusion Event Fields, on page 766.
- 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, see Event Investigation Using Web-Based Resources, on page 620
- 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, see Event Investigation Using Web-Based Resources, on page 620.
- To modify the time and date range for displayed events, see Changing the Time Window, on page 675.

Tip
If no intrusion events appear on the event views, adjusting the specified time range might return results. If you specified an older time range, events in that time range might have been deleted. Adjusting the rule thresholding configuration might generate events.

Note
Events generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- To sort events on the current workflow page or navigate within the current workflow page, see Using Workflows, on page 656.
- 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 delete events from the event database, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All**.
- To mark events reviewed to remove them from intrusion event pages, but not the event database, see Marking Intrusion Events Reviewed, on page 779.
- To download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event, check the check boxes next to events triggered by the packets you want to download, then click **Download Packets**, or click **Download All Packets**. Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.
- To navigate to other event views to view associated events, see Inter-Workflow Navigation, on page 680.
- To temporarily use a different workflow, click **(switch workflow)**.
- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**.
- To view the Intrusion Events section of the Summary Dashboard, click **Dashboards**.
- To navigate to the bookmark management page, click **View Bookmarks**.
- To generate a report based on the data in the current view, see Creating a Report Template from an Event View, on page 529.

Related Topics

Event Searches, on page 685
Bookmarks, on page 682

Intrusion Event Drill-Down Page Constraints

The following table describes how to use the drill-down pages.

Table 114: Constraining Events on Drill-Down Pages

To...	You can...
drill down to the next workflow page constraining on a specific value	click the value. For example, on the Destination Port workflow, to constrain the events to those with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column. The next page of the workflow, Events, appears and contains only port 80/tcp events.
drill down to the next workflow page constraining on selected events	select the check boxes next to the events you want to view on the next workflow page, then click **View**. For example, on the Destination Port workflow, to constrain the events to those with destination ports 20/tcp and 21/tcp, select the check boxes next to the rows for those ports and click **View**. The next page of the workflow, Events, appears and contains only port 20/tcp and 21/tcp events. Note that if you constrain on multiple rows and the table has more than one column (not including a Count column), you build what is called a compound constraint. Compound constraints ensure that you do not include more events in your constraint than you mean to. For example, if you use the Event and Destination workflow, each row that you select on the first drill-down page creates a compound constraint. If you pick event 1:100 with a destination IP address of 10.10.10.100 and you also pick event 1:200 with a destination IP address of 192.168.10.100, the compound constraint ensures that you do not also select events with 1:100 as the event type and 192.168.10.100 as the destination IP address or events with 1:200 as the event type and 10.10.10.100 as the destination IP address.
drill down to the next workflow page keeping the current constraints	click **View All**.

Intrusion Event Table View Constraints

The following table describes how to use the table view.

Table 115: 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**.

To...	You can...
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 786.

Step 2 Optionally, if you chose more than one event, you can page through the packets in the packet view by using the page numbers at the bottom of the page.

Step 3 You also have the following options:

- Adjust — To modify the date and time range in the packet views, see Changing the Time Window, on page 675.

- Configure — To configure the intrusion rule that triggered the event, click the arrow next to Actions and continue as described in Configuring Intrusion Rules within the Packet View, on page 791.
- Delete — To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Download — To download a local copy of the packet (a packet capture file in libpcap format) that triggered the event, click **Download Packet** to save a copy of the captured packet for the event you are viewing or click **Download All Packets** to save copies of the captured packets for all the events whose packets you previously selected. The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers.

  Note
  You cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan view provides all usable packet information. You must have at least 15% available disk space in order to download.

- Mark reviewed — To mark an event reviewed to remove it from event views, but not the event database, click **Review** to mark the event whose packet you are viewing or click **Review All** to mark all the events whose packets you previously selected. For more information, see Marking Intrusion Events Reviewed, on page 779.
- View additional information — To expand or collapse a page section, click the arrow next to the section. For details, see Event Information Fields, on page 788, Frame Information Fields, on page 794, and Data Link Layer Information Fields, on page 795.
- View network layer information — See Viewing Network Layer Information, on page 796.
- View packet byte information — See Viewing Packet Byte Information, on page 801.
- View transport layer information — See Viewing Transport Layer Information, on page 798

Related Topics

Portscan Detection

Event Information Fields

On the packet view, you can view information about the packet in the Event Information section.

Event

The event message. For rule-based events, this corresponds to the rule message. For other events, this is determined by the decoder or preprocessor.

The ID for the event is appended to the message in the format (GID:SID:Rev). GID is the generator ID of the rules engine, the decoder, or the preprocessor that generated the event. SID is the identifier for the rule, decoder message, or preprocessor message. Rev is the revision number of the rule.

Timestamp

The time that the packet was captured, in UTC time zone.

Classification

The event classification. For rule-based events, this corresponds to the rule classification. For other events, this is determined by the decoder or preprocessor.

Priority

The event priority. For rule-based events, this corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other events, this is determined by the decoder or preprocessor.

Ingress Security Zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Egress Security Zone

The egress security zone of the packet that triggered the event. This field is not populated in a passive deployments

Domain

The domain where the managed device belongs. This field is only present if you have ever configured the 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 the following 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

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Configuring Intrusion Rules within the Packet View

Within the packet view of an intrusion event, you can take several actions on the rule that triggered the event. Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 You have the following choices:

- Comment — For standard text rule events, click **Rule Comment** to add a text comment to the rule that generated the event. This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies. You can also add and view rule comments in the intrusion rules editor.
- Disable — 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 Setting Suppression Options within the Packet View, on page 793.

  You can use this option to suppress the rule that triggered this event in all policies that you can edit locally. Alternately, you can suppress the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

- Set threshold options — Expand **Set Thresholding Options** and continue as described in Setting Threshold Options within the Packet View, on page 793.

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.

Setting Threshold Options within the Packet View

You can control the number of events that are generated per rule over time by setting the threshold options in the packet view of an intrusion event. You can set threshold options in all policies that you can edit locally or, when it can be edited locally, only in the in the current policy (that is, the policy that caused the event to be generated).

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Thresholding Options** and choose one of the two possible options:

- **in the current 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**.

Setting Suppression Options within the Packet View

You can use the suppression options to suppress intrusion events altogether, or based on the source or destination IP address. You can set suppression options in all policies that you can edit locally. Alternately, you can set

suppression options only in the current policy (that is, the policy that generated the event) when the current policy can be edited locally.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Suppression Options** and click one of the two possible options:

- **in the current 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**.

Related Topics

IP Address Conventions, on page 25

Frame Information Fields

On the packet view, click the arrow next to **Frame** to view information about the captured frame. The packet view may display a single frame or multiple frames. Each frame provides information about an individual network packet. You would see multiple frames, for example, in the case of tagged packets or packets in reassembled TCP streams.

Frame n

The captured frame, where *n* is 1 for single-frame packets and the incremental frame number for multi-frame packets. The number of captured bytes in the frame is appended to the frame number.

Arrival Time

The date and time the frame was captured.

Time delta from previous captured frame

For multi-frame packets, the elapsed time since the previous frame was captured.

Time delta from previous displayed frame

For multi-frame packets, the elapsed time since the previous frame was displayed.

Time since reference or first frame

For multi-frame packets, the elapsed time since the first frame was captured.

Frame Number

The incremental frame number.

Frame Length

The length of the frame in bytes.

Capture Length

The length of the captured frame in bytes.

Frame is marked

Whether the frame is marked (true or false).

Protocols in frame

The protocols included in the frame.

Related Topics

The tag Keyword
TCP Stream Reassembly

Data Link Layer Information Fields

On the packet view, click the arrow next to the data link layer protocol (for example, **Ethernet II**) to view the data link layer information about the packet, which contains the 48-bit media access control (MAC) addresses for the source and destination hosts. It may also display other information about the packet, depending on the hardware protocol.

Note Note that this example discusses Ethernet link layer information; other protocols may also appear.

The packet view reflects the protocol used at the data link layer. The following listing describes the information you might see for an Ethernet II or IEEE 802.3 Ethernet packet in the packet view.

Destination

The MAC address for the destination host.

Note Ethernet can also use multicast and broadcast addresses as the destination address.

Source

The MAC address for the source host.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IPv6 or ARP datagrams. Note that this item only appears for Ethernet II packets.

Length

For IEEE 802.3 Ethernet packets, the total length of the packet, in bytes, not including the checksum. Note that this item only appears for IEEE 802.3 Ethernet packets.

Viewing Network Layer Information

Procedure

On the packet view, click the arrow next to the network layer protocol (for example, **Internet Protocol**) to view more detailed information about network layer information related to the packet.

Note
Note that this example discusses IP packets; other protocols may also appear.

IPv4 Network Layer Information Fields

The following listing describes protocol-specific information that might appear in an IPv4 packet.

Version

The Internet Protocol version number.

Header Length

The number of bytes in the header, including any IP options. An IP header with no options is 20 bytes long.

Differentiated Services Field

The values for differentiated services that indicate how the sending host supports Explicit Congestion Notification (ECN):

- `0x0` — does not support ECN-Capable Transport (ECT)
- `0x1` and `0x2` — supports ECT
- `0x3` — Congestion Experienced (CE)

Total Length

The length of the IP packet, in bytes, minus the IP header.

Identification

The value that uniquely identifies an IP datagram sent by the source host. This value is used to trace fragments of the same datagram.

Flags

The values that control IP fragmentation, where:

values for the Last Fragment flag indicate whether there are more fragments associated with the datagram:

- `0` — there are no more fragments associated with the datagram
- `1` — there are more fragments associated with the datagram

values for the Don't Fragment flag control whether the datagram can be fragmented:

- `0` — the datagram can be fragmented
- `1` — the datagram must **not** be fragmented

Fragment Offset

The value for the fragment offset from the beginning of the datagram.

Time to Live (ttl)

The remaining number of hops that the datagram can make between routers before the datagram expires.

Protocol

The transport protocol that is encapsulated in the IP datagram; for example, ICMP, IGMP, TCP, or UDP.

Header Checksum

The indicator for whether the IP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an intrusion evasion attempt.

Source/Destination

The IP address or domain name for the source (or destination) host.

Note that to display the domain name, you must enable IP address resolution.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or choose an option to add the address to a global Block list or Do-Not-Block list.

IPv6 Network Layer Information Fields

The following listing describes protocol-specific information that might appear in an IPv6 packet.

Traffic Class

An experimental 8-bit field in the IPv6 header for identifying IPv6 packet classes or priorities similar to the differentiated services functionality provided for IPv4. When unused, this field is set to zero.

Flow Label

A optional 20-bit IPv6 hexadecimal value 1 to FFFFF that identifies a special flow such as non-default quality of service or real-time service. When unused, this field is set to zero.

Payload Length

A 16-bit field identifying the number of octets in the IPv6 payload, which is comprised of all of the packet following the IPv6 header, including any extension headers.

Next Header

An 8-bit field identifying the type of header immediately following the IPv6 header, using the same values as the IPv4 Protocol field.

Hop Limit

An 8-bit decimal integer that each node that forwards the packet decrements by one. The packet is discarded if the decremented value reaches zero.

Source

The 128-bit IPv6 address for the source host.

Destination

The 128-bit IPv6 address for the destination host.

Viewing Transport Layer Information

Procedure

Step 1 On the packet view, click the arrow next to the transport layer protocol (for example, **TCP**, **UDP**, or **ICMP**).

Step 2 Optionally, click **Data** when present to view the first twenty-four bytes of the payload for the protocol immediately above it in the Packet Information section of the packet view.

Step 3 View the contents of the transport layer for TCP, UDP, and ICMP protocols as described in TCP Packet View Fields, on page 798, UDP Packet View Fields, on page 799, or ICMP Packet View Fields, on page 800.

Note
Note that these examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP Packet View Fields

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Sequence number

The value for the first byte in the current TCP segment, keyed to initial sequence number in the TCP stream.

Next sequence number

In a response packet, the sequence number of the next packet to send.

Acknowledgement number

The TCP acknowledgement, which is keyed to the sequence number of the previously accepted data.

Header Length

The number of bytes in the header.

Flags

The six bits that indicate the TCP segment's transmission state:

- `U` — the urgent pointer is valid
- `A` — the acknowledgement number is valid
- `P` — the receiver should push data
- `R` — reset the connection
- `S` — synchronize sequence numbers to start a new connection
- `F` — the sender has finished sending data

Window size

The amount of unacknowledged data, in bytes, that the receiving host will accept.

Checksum

The indicator for whether the TCP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an in evasion attempt.

Urgent Pointer

The position, if present, in the TCP segment where the urgent data ends. Used in conjunction with the `U` flag.

Options

The values, if present, for TCP options.

UDP Packet View Fields

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Length

The combined length of the UDP header and data.

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

ICMP Packet View Fields

This section describes the protocol-specific information for an ICMP packet.

Type

The type of ICMP message:

- 0 — echo reply
- 3 — destination unreachable
- 4 — source quench
- 5 — redirect
- 8 — echo request
- 9 — router advertisement
- 10 — router solicitation
- 11 — time exceeded
- 12 — parameter problem
- 13 — timestamp request
- 14 — timestamp reply
- 15 — information request (obsolete)
- 16 — information reply (obsolete)
- 17 — address mask request
- 18 — address mask reply

Code

The accompanying code for the ICMP message type. ICMP message types 3, 5, 11, and 12 have corresponding codes as described in RFC 792.

Checksum

The indicator for whether the ICMP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

Viewing Packet Byte Information

Procedure

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes that comprise the packet. If the system decrypted traffic, you can view the decrypted packet bytes.

Internally Sourced Intrusion Events

Intrusion events coming from internal sources indicate a compromised host on your network. If the source IP address is on your network, this is a sign that you should investigate this host.

Viewing Intrusion Event Statistics

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated for your network.

Each of the IP addresses, ports, protocols, event messages, and so on shown on the page is a link. Click any link to view the associated event information. For example, if one of the top 10 destination ports is `80 (http)/tcp`, clicking that link displays the first page in the default intrusion events workflow, and lists the events targeting that port. Note that only the events (and the managed devices that generate events) in the current time range appear. Also, intrusion events that you have marked reviewed continue to appear in the statistics. For example, if the current time range is the past hour but the first event was generated five hours ago, when you click the **First Event** link, the resulting event pages will not show the event until you change the time range.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Statistics**.

Step 2 From the two selection boxes at the top of the page, choose the zones and devices whose statistics you want to view, or choose **All Security Zones** and **All Devices** to view statistics for all the devices that are collecting intrusion events.

Step 3 Click **Get Statistics**.

Tip

To view data from a custom time range, click the link in the upper right page area and follow the directions in Changing the Time Window, on page 675.

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 the following:

Time

The current time on the appliance.

Uptime

The number of days, hours, and minutes since the appliance itself was restarted. On the 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 of the information in the intrusion event database.

This information includes details on:

- the top 10 event types
- the top 10 source IP addressees
- the top 10 destination IP addresses
- the top 10 destination ports
- the protocols, ingress and egress security zones, and devices with the greatest number of events

Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view event statistics in an ancestor domain, the system may display multiple instances of that repeated IP address. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

Viewing Intrusion Event Performance Graphs

The intrusion event performance page allows you to generate graphs that depict performance statistics for intrusion events over a specific period of time for a Secure Firewall Management Center or a managed device. Graphs can be generated to reflect number of intrusion events per second, number of megabits per second, average number of bytes per packet, the percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

Note New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs. Each graph displays *average* values in the intervals shown (day, hour, or five minutes) for the selected time period (last month, week, day, or hour). Decimal values are displayed when the average is less than one.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Performance**.

Step 2 From the **Select Device** list, choose the devices whose data you want to view.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Intrusion Event Performance Statistics Graph Types, on page 804.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph**.

Step 6 To save the graph, right-click it and follow the instructions for your browser to save the image.

Intrusion Event Performance Statistics Graph Types

The following table lists the available graph types. Note that graph types display differently if they are populated with data affected by the network analysis policy **Inline Mode** setting. If **Inline Mode** is disabled, the graph types marked with an asterisk (`*`) in the web interface (a `yes` in the column below) populate with data about the traffic the system would have modified or dropped if **Inline Mode** was enabled..

Table 116: 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

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
IPv4 Options Normalizations	enable **Normalize IPv4**	the number of IPv4 packets for which the option octet was set to 1 (No Operation).	yes
IPv4 Reserved Flag Normalizations	enable **Normalize IPv4** and **Normalize Reserved Bit**	the number of IPv4 packets for which the single-bit Reserved subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Resize Normalizations	enable **Normalize IPv4**	the number of IPv4 packets with excessive-length payload that have been truncated to the datagram length specified in the IP header.	yes
IPv4 TOS Normalizations	enable **Normalize IPv4** and **Normalize TOS Bit**	the number of IPv4 packets for which the one-byte Differentiated Services (DS) field (formerly known as the Type of Service (TOS) field) was cleared.	yes
IPv4 TTL Normalizations	enable **Normalize IPv4**, **Maximum TTL**, and **Reset TTL**	the number of IPv4 Time to Live normalizations.	yes
IPv6 Options Normalizations	enable **Normalize IPv6**	the number of IPv6 packets for which the Option Type field in the Hop-by-Hop Options or Destination Options extension header was set to `00` (Skip and continue processing).	yes
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

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
SYN Packets With Data Stripped Normalizations	enable **Remove Data on SYN**	the number of packets for which data was removed from SYN packets when the TCP operating system was not Mac OS.	yes
TCP Header Padding Normalizations	enable **Normalize/Clear Option Padding Bytes**	the number of TCP packets in which option padding bytes were set to 0.	yes
TCP No Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets from which the Time Stamp option was stripped.	yes
TCP NS Flag Normalizations	enable **Explicit Congestion Notification** and select **Packet**	the number of ECN Nonce Sum (NS) option normalizations.	yes
TCP Options Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options (excluding MSS, Window Scale, Time Stamp, and explicitly allowed options) for which the option field is set to No Operation (TCP Option 1).	yes
TCP Packets Blocked By Normalizations	enable **Normalize TCP Payload** (segment reassembly must fail)	the number of packets dropped because the TCP segments could not be properly reassembled.	yes
TCP Reserved Flags Normalizations	enable **Normalize/Clear Reserved Bits**	the number of TCP packets where the Reserved bits have been cleared.	yes
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

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
Total TCP Filtered Packets	configure TCP Stream Preprocessing	the number of packets skipped by the stream because of TCP port filtering.	no
Total UDP Filtered Packets	configure UDP Stream Preprocessing	the number of packets skipped by the stream because of UDP port filtering.	no
Urgent Flag Cleared Normalizations	enable **Clear URG if Urgent Pointer is Not Set**	the number of packets for which the TCP header URG control bit was cleared because the urgent pointer was not set.	yes
Urgent Pointer and Urgent Flag Cleared Normalizations	enable **Clear Urgent Pointer/URG on Empty Payload**	the number of packets for which the TCP header Urgent Pointer field and the URG control bit have been cleared because there was no payload.	yes
Urgent Pointer Cleared Normalizations	enable **Clear Urgent Pointer if URG=0**	the number of packets for which the 16-bit TCP header Urgent Pointer field was cleared because the urgent (URG) control bit was not set.	yes

Related Topics

The Inline Normalization Preprocessor
Preprocessor Traffic Modification in Inline Deployments
Drop Behavior in an Inline Deployment

Viewing 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

Feature	Minimum Management Center	Minimum 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	The following 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.

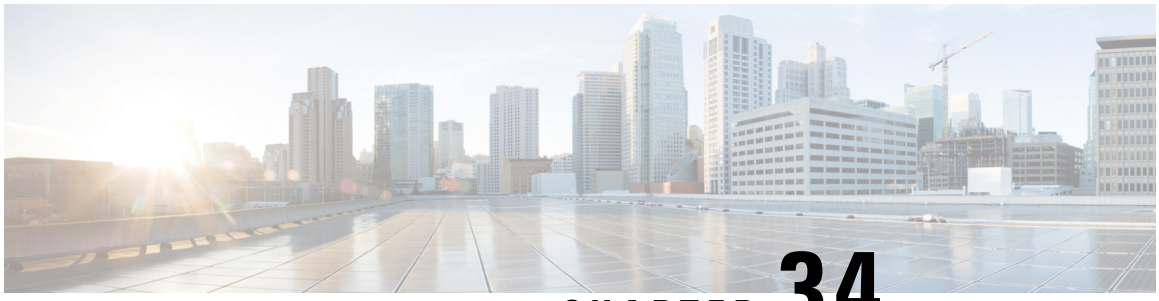

CHAPTER 34

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.

About File/Malware Events and Network File Trajectory

File policies automatically generate file and malware events for matched traffic, and log captured file information. When a file policy generates a file or malware event, or captures a file, the system also automatically logs the end of the associated connection to the Secure Firewall Management Center database. You can analyze this data to address any negative impacts and block future attacks.

Based on the file analysis results, you can review captured files and generated malware and file events using tables on pages available under the Analysis > Files menu. When available, you can examine a file's composition, disposition, threat score, and dynamic analysis summary report for further insight into the malware analysis.

To further target your analysis, you can use a malware file's *network file trajectory* (a map of how the file traversed your network, passing among hosts, as well as various file properties) to track the spread of an individual threat across hosts over time, allowing you to concentrate outbreak control and prevention efforts where most useful.

If you configure local malware analysis or dynamic analysis in a file rule, the system preclassifies files matching the rule and generates a file composition report.

If your organization has deployed *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, see *Network Malware Protection and File Policies* in the Cisco Secure Firewall Management Center Device Configuration Guide.

File and Malware Events

The Secure Firewall Management Center can log various types of file and malware events. The information available for any individual event can vary depending on how and why it was generated:

- *File events* represent files, including malware, detected by the system (malware defense.) File events do not contain Secure Endpoint-related fields.
- *Malware events* 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* 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 Events

The system logs the file events generated when a managed device detects or blocks a file in network traffic, according to the rules in currently deployed file policies.

When the system generates a file event, the system also logs the end of the associated connection to the 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 117: Malware Event Generation Scenarios

When the system detects a file and...	Disposition
successfully queries the AMP cloud (performs a malware cloud lookup) for the file's disposition	Malware, Clean, or Unknown
queries the AMP cloud but cannot establish a connection or the cloud is otherwise unavailable	Unavailable You may see a small percentage of events with this disposition; this is expected behavior.
the threat score associated with a file exceeds the malware threshold threat score defined in the file policy that detected the file, or local malware analysis identifies malware	Malware
it is on the custom detection list (manually marked as malware)	Custom Detection
it is on the on the clean list (manually marked as clean),	Clean

File Disposition and File Action in Malware Events

Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. If you select *Block Malware* or *Malware Cloud Lookup* as file rule action, the system queries the AMP cloud to determine if files traversing your network contain malware, then block files that represent threats. Cloud lookup allows you to obtain and log the file's disposition based on its SHA-256 hash value.

The following table describes the file action that associates with the file disposition returned by the AMP cloud:

Table 118: File Disposition and File Action in Malware Events

File Rule Action Selected	File Disposition	File Action in the Malware Event
• Block Malware • Malware Cloud Lookup	Malware	Block
	• Clean • Unknown • Unavailable • NA	Cloud Lookup **Note** Under the file policy editor Advanced Settings, you can set a threshold threat score for **If AMP Cloud disposition is Unknown, override disposition based upon threat score** option. If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their dynamic analysis score is equal to or worse than the threshold.

Retrospective Malware Events

For malware detected in network traffic, dispositions can change. For example, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the last week, the AMP cloud notifies the system. Then, two things happen:

- The 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

If your organization uses Secure Endpoint, individual users install lightweight connectors on *endpoints*: computers and mobile devices. Connectors can inspect files upon upload, download, execution, open, copy, move, and so on. These connectors communicate with the AMP cloud to determine if inspected files contain malware.

When a file is positively identified as malware, the AMP cloud sends the threat identification to the 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

If your organization has deployed Cisco Secure Endpoint:

- You can configure the system to display malware events detected by Secure Endpoint on 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 the above functionality, see *Integrate DeviceUUID (Syslog Only)Secure Firewall and Secure Endpoint* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Event Data from Secure Endpoint

If your organization has deployed Secure Endpoint for malware protection, you can configure the system to let you work in management center with file and malware data from Secure Endpoint.

However, you should be aware of the differences between file and malware data from Secure Endpoint and file and malware data from the system's malware defense feature.

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.

The following table summarizes the differences between the event data generated by Firepower when using a Malware Defense license, and event data generated by Secure Endpoint.

Table 119: 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	management center-based	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

Use this procedure to view file and malware events in a table and to manipulate the event view depending on the information relevant to your analysis. The page you see when you access events differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose one of the following:

- **Analysis** > **Files** > **File Events**
- **Analysis** > **Files** > **Malware Events**

Tip
In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Tip
To quickly view the connections where specific files were detected, choose the files using the check boxes in the table, then choose **Connections Events** from the **Jump to** drop-down list.

Tip

Right-click an item in the table to see options. (Not every column offers options.)

Related Topics

File and Malware Event Fields, on page 815
Predefined File Workflows, on page 649
Predefined Malware Workflows, on page 649
Configuring Event View Settings, on page 206

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 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** that lists the archive file, right-click on the archive file's table row to open the context menu, then click **View Archive Contents**.

Archive SHA256 (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 that lists the archive file, right-click on 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.

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 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 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

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 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

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 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 management center web interface, the IP address of the host sending the file. See also A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

For syslog equivalents, see **DstIP** and **SrcIP**.

Sending Port

In the 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 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 A Note About Initiator/Responder, Source/Destination, and Sender/Receiver Fields, on page 747.

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** (Decryption 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 following 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 120: 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

Malware Event Subtype/Search Value	malware defense	Secure Endpoint	File Trajectory
Cloud Recall Quarantine Started	no	yes	no
Cloud Recall Restore from Quarantine	no	yes	no
Cloud Recall Restore from Quarantine Failed	no	yes	no
Cloud Recall Restore from Quarantine Started	no	yes	no
Quarantine Failure	no	yes	no
Quarantined Item Restored	no	yes	no
Quarantine Restore Failed	no	yes	no
Quarantine Restore Started	no	yes	no
Scan Completed, No Detections	no	yes	no
Scan Completed With Detections	no	yes	no
Scan Failed	no	yes	no
Scan Started	no	yes	no

Information Available in File and Malware Event Fields

The following table lists whether the system displays information for each file and malware event field.

If your organization has deployed 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.

Table 121: 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

Field	File Event	Malware Events Detected by the System	Retrospective Events Detected by the System	Malware Events Detected by Secure Endpoint
Application Protocol	yes	yes	no	no
Application Protocol Category or Tag	yes	yes	yes	no
Application Risk	yes	yes	yes	no
Archive Depth	yes	yes	no	yes
Archive Name	yes	yes	no	yes
Archive SHA256	yes	yes	no	yes
Business Relevance	yes	yes	yes	no
Category / File Type Category	yes	yes	no	yes
Client	yes	yes	yes	no
Client Category or Tag	yes	yes	yes	no
Count	yes	yes	yes	yes
Detection Name	no	yes	no	no
Detector	no	no	no	yes
Device	yes	yes	yes	yes
Disposition / File Disposition	yes	yes	yes	no
Domain	yes	yes	yes	yes
Event Subtype	no	no	no	yes
Event Type	no	yes	yes	yes
File Name	yes	yes	no	yes
File Path	no	no	no	yes
File Policy	yes	no	no	no
File Timestamp	no	no	no	yes
HTTP Response Code	yes	yes	no	no
IOC (Indication of Compromise)	no	yes	yes	yes
Message	yes	yes	no	yes
Receiving Continent	yes	yes	yes	no
Receiving 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
Receiving IP	yes	yes	no	yes
Receiving Port	yes	yes	no	no
Security Context	yes	yes	yes	yes
Sending Continent	yes	yes	yes	no
Sending Country	yes	yes	no	no
Sending IP	yes	yes	no	no
Sending Port	yes	yes	no	no
SHA256 / File SHA256	yes	yes	yes	yes
Size (KB) / File Size (KB)	yes	yes	no	yes
SSL Actual Action (search only)	yes	yes	no	no
SSL Certificate Information (search only)	yes	yes	no	no
SSL Failure Reason (search only)	yes	yes	no	no
SSL Status	yes	yes	no	no
SSL Subject/Issuer Country (search only)	yes	yes	no	no
File Storage / Stored (search only)	yes	yes	no	no
Threat Name	no	yes	yes	yes
Threat Score	yes	yes	no	no
Time	yes	yes	yes	yes
Type / File Type	yes	yes	no	yes
URI / File URI	yes	yes	no	no
User	yes	yes	no	yes
Web Application	yes	yes	yes	no
Web Application Category or Tag	yes	yes	yes	no

View Details About Analyzed Files

Tip To see additional options, right-click a file SHA in a table on an event page. For information, see Event Investigation Using Web-Based Resources, on page 620.

File Composition Report

If you configure local malware analysis or dynamic analysis, the system generates a file composition report after analyzing a file. This report allows you to further analyze files and determine whether they may carry embedded malware.

The file composition report lists file properties, any objects embedded in the file, and any detected viruses. The file composition report may also list additional information specific to that file type. When the system prunes stored files, it also prunes the associated file composition report.

To view file composition information, see Using a Network File Trajectory, on page 839.

View File Details in AMP Private Cloud

If you have deployed an AMP private cloud, you can view additional details about analyzed files in your private cloud.

For more information, see the documentation for your private cloud.

Procedure

Sign in directly to your AMP private cloud console.

Threat Scores and Dynamic Analysis Summary Reports

Threat Scores

Table 122: Threat Score Ratings

Threat Score	Numeric Score	Icon
`Low`	0-24	**Low**
`Medium`	25-69	**Medium**
`High`	70-94	**High**
`Very High`	95-100	**Very High**

The 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.

Viewing Dynamic Analysis Results in the Cisco Secure Malware Analytics Cloud

Secure Malware Analytics offers more detailed reporting on analyzed files than is available in the 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 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 the following user roles: Admin, Access Admin, Network Admin

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, see the Secure Malware Analytics documentation.

Using Captured File Workflows

When a managed device captures a file detected in network traffic, it logs an event.

Note If a device captures a file containing malware, the device generates two events: a file event when it detects the file, and a malware event when it identifies malware.

Use this procedure to view a list of captured files in a table and manipulate the event view depending on the information relevant to your analysis. The page you see when you access captured files differs depending on the workflow, which is simply a series of pages you can use to evaluate events by moving from a broad to a more focused view. You can also create a custom workflow that displays only the information that matches your specific needs.

If the system recaptures a file after a configuration change, such as an updated file policy, it updates existing information for that file.

For example, if you configure a file policy to capture files with a **Malware Cloud Lookup** action, the system stores the file disposition and threat score along with the file. Then, if you update your file policy, and the system recaptures the same file due to a new **Detect Files** action, the system updates the file's **Last Changed** value. However, the system does not remove the existing disposition and threat score, even though you did not perform another malware cloud lookup.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Choose **Analysis** > **Files** > **Captured Files**.

Tip
In the table view of events, several fields are hidden by default. To show a hidden field in an event view, expand the search constraints, then click the field name under **Disabled Columns**.

Related Topics

Captured File Fields, on page 832
Predefined Captured File Workflows, on page 650

Configuring Event View Settings, on page 206

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 123: Captured File Fields

Field	Description
Archive Inspection Status	For archive files, the status of archive inspection: • Pending indicates that the system is still inspecting the archive file and its contents. If the file passes through your system again, complete information becomes available. • Extracted indicates that the system was able to extract and inspect the archive's contents. • Failed may, in rare cases, occur if the system is unable to process an extraction. • Depth Exceeded indicates that the archive contains further nested archive files beyond the maximum allowed depth. • Encrypted indicates that the archive file's contents are encrypted and could not be inspected. • Not Inspectable indicates that the system did not extract and inspect the archive's contents. Policy rule actions, policy configuration, and corrupted files are three major reasons for this status. To view the contents of an archive file, right-click on its row in the table to bring up the context menu, then choose **View Archive Contents**.
Category	The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.
Detection Name	The name of the detected malware.

Field	Description
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.
Domain	The domain where the captured file was detected. This field is only present if you have ever configured the management center for multitenancy.
Dynamic Analysis Status	One or more of the following values indicating whether the file was submitted for dynamic analysis: • Analysis Complete — file submitted for dynamic analysis that received a threat score and dynamic analysis summary report • Capacity Handled — file stored because it could not be submitted currently • Capacity Handled (Network Issue) — file stored because it could not be submitted due to a network connectivity issue • Capacity Handled (Rate Limit) — file stored because it could not be submitted due to the maximum number of submissions reached • Device Not Activated — file not submitted because the device is not activated on the on-premises 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 • 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

Field	Description
Dynamic Analysis Status Changed	The last time the file's dynamic analysis status changed.
File Name	The most recently detected file name associated with the file's SHA-256 hash value.
Last Changed	The last time the information associated with this file was updated.
Last Sent	The time the file was most recently submitted to the AMP cloud for dynamic analysis.
Local Malware Analysis Status	One of the following values indicating whether the system performed local malware analysis on a file: • Analysis Complete — the system inspected the file using local malware analysis and pre-classified the file • Analysis Failed — the system attempted to inspect the file using local malware analysis and failed • Manual Request Submitted — a user submitted a file for local malware analysis • Not Analyzed — the system did not inspect the file with local malware analysis
SHA256	The SHA-256 hash value of the file, as well as a network file trajectory icon representing the most recently detected file event and file disposition. To view the network file trajectory, click the trajectory icon.
Storage Status	Indicates whether the file is stored on a managed device: • File Stored • Not Stored (Disposition Was Pending)
Threat Score	The threat score most recently associated with this file. To view the Dynamic Analysis Summary report, click the threat score icon.
Type	The type of file; for example, HTML or MSEXE.

Stored Files Download

Once a device stores a file, as long as the 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.

Manually Submit Files for Analysis

When you manually submit files for analysis, the system runs local analysis, then submits these files to the cloud for dynamic analysis. However, if local analysis is not enabled in a file policy, and you manually submit a file for analysis, the file will only be sent for dynamic analysis.

In addition to executable files, you can also submit file types not eligible for automatic submission, such as .swf, .jar, and others. This allows you to more quickly analyze a broad range of files, regardless of disposition, and pinpoint the exact causes of an incident.

Note The system checks the AMP cloud for updates (no more than once a day) to the list of file types eligible for dynamic analysis and the minimum and maximum file sizes you can submit.

Depending on the situation, there are two ways to submit files for analysis:

Before you begin

In order to manually submit captured files for analysis, one or more file rules must be configured to store files. For information, see the *Network Malware Protection and File Policies* chapter in the Cisco Secure Firewall Management Center Device Configuration Guide.

Procedure

Step 1 To submit a single file for analysis:

a) Select one of the following:

- **Analysis > Files > File Events**
- **Analysis > Files > Malware Events**
- **Analysis > Files > Captured Files**

b) Click **Table View of <Event type or files>**.

c) Right-click a file in the table and select **Analyze File**.

Step 2 To submit multiple captured files for analysis (up to 25 at a time):

a) Select **Analysis > Files > Captured Files**

b) Select the checkbox beside each file to analyze.

c) Click **Analyze**.

Network File Trajectory

The network file trajectory feature maps how hosts transferred files, including malware files, across your network. A trajectory charts file transfer data, the disposition of the file, and if a file transfer was blocked or the file was quarantined. You can determine which hosts and users may have transferred malware, which hosts are at risk, and observe file transfer trends.

You can track the transmission of any file with a AMP cloud-assigned disposition. The system can use information related to detecting and blocking malware from both malware defense and Secure Endpoint to build the a trajectory.

Recently Detected Malware and Analyzed Trajectories

The Network File Trajectory List page displays the malware most recently detected on your network, as well as the files whose trajectory maps you have most recently viewed. From these lists, you can view when each file was most recently seen on the network, the file's SHA-256 hash value, name, type, current file disposition, contents (for archive files), and the number of events associated with the file.

The page also contains a search box that lets you locate files, either based on SHA-256 hash value or file name, or by the IP address of the host that transferred or received a file. After you locate a file, you can click the **File SHA256** value to view the detailed trajectory map.

Network File Trajectory Detailed View

You can trace a file through the network by viewing the detailed network file trajectory. Search for a file's SHA 256 value or click a **File SHA 256** link in the Network File Trajectory list to view details about that file.

The network file trajectory details page has three parts:

- Summary Information — A file's trajectory page displays summary information about the file, including file identification information; when the file was first seen and most recently seen on the network, and by what user; the number of related events and hosts associated with the file; and the file's current disposition. From this section, if the managed device stored the file, you can download it locally, submit the file for dynamic analysis, or add the file to a file list.
- Trajectory Map — A file's trajectory map visually tracks a file from the first detection on your network to the most recent. The map shows when hosts transferred or received the file, how often they transferred the file, and when the file was blocked or quarantined. Vertical lines between data points represent file transfers between hosts. Horizontal lines connecting the data points show a host's file activity over time.

  The map also shows how often file events occurred for the file and when the system assigned the file a disposition or retrospective disposition. You can select a data point in the map and highlight a path that traces back to the first instance the host transferred that file; this path also intersects with every occurrence involving the host as either sender or receiver of the file, and identifies the user involved.
- Related Events — The Events table lists event information for each data point in the map. Using the table and the map, you can pinpoint specific file events, hosts and users on the network that transferred or received this file, related events in the map, and other related events in a table constrained on selected values.

Network File Trajectory Summary Information

The following summary information appears at the top of the details page for a file that appears in the Network File Trajectory list.

Tip To view related file events, click a field value link. The first page in the File Events default workflow opens in a new window, displaying all file events that also contain the selected value.

Table 124: Network File Trajectory Summary Information Fields

Name	Description
Archive Contents	For inspected archive files, the number of files the archive contains.
Current Disposition	One of the following 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.

Name	Description
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.
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 Map and Related Events List

The file trajectory map's y-axis contains a list of all host IP addresses that have interacted with the file. The IP addresses are listed in descending order based on when the system first detected the file on that host. Each row contains all events associated with that IP address, whether a single file event, file transfer, or retrospective event. The x-axis contains the date and time the system detected each event. The timestamps are listed in chronological order. If multiple events occurred within a minute, all are listed within the same column. You can scroll the map horizontally and vertically to view additional events and IP addresses.

The map displays up to 250 events associated with the file SHA-256 hash. If there are more than 250 events, the map displays the first 10, then truncates extra events with an **Arrow**. The map then displays the remaining 240 events.

The first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. If malware events generated by 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 the following 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

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 management center to Secure Endpoint, see Work with Event Data in the Secure Endpoint Console, on page 840. For details about the file trajectory feature in Secure Endpoint, see the Secure Endpoint documentation.

Before you begin

If you are using malware defense, you need the Malware Defense license.

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Files** > **Network File Trajectory**.

Tip
You can also access a file's trajectory from the Context Explorer, dashboard, or event views with file information.

Step 2 Click a **File SHA 256** link in the list.

Step 3 Optionally, enter a complete SHA-256 hash value, the host IP address, or the name of a file you want to track into the search field, and press Enter.

Tip
If only one result matches, the Network File Trajectory page for that file appears.

Step 4 In the Summary Information section, you can:

- Add a file to a file list — To add a file or remove a file from the clean list or custom detection list, click **Edit** ( ).

- Download a file — To download a file, click **Download** ( ), and if prompted, confirm you want to download the file. If the file is unavailable for download, this download file is dimmed.
- Report — Click threat score to view the Dynamic Analysis Summary report.
- Submit for dynamic analysis — Click **AMP Cloud** to submit the file for dynamic analysis. If the file is unavailable for submission or you cannot connect to the AMP cloud, this AMP cloud is dimmed.
- View archive contents — To view information about an archive file's contents, click **View** ( ).
- View file composition — To view a file's composition, click **File List**. If the system has not generated a file composition report, this file list is dimmed.
- View captured files with same threat score — Click the threat score link to view all captured files with that threat score.

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

Step 5 On the trajectory map, you can:

- Locate the first instance — Click an IP address to locate the first time a file event occurred involving an IP address. This highlights a path to that data point, as well as any intervening file events and IP addresses related to the first file event. The corresponding event in the Events table is also highlighted. The map scrolls to that data point if not currently visible.
- Track — Click any data point to highlight a path that includes all data points related to the selected data point, tracking a file's progress through the network.
- View hidden events — Click arrow to view all events not displayed in the File Summary event view.
- View matching file events — Hover your pointer over the **Matching File Event** to view summary information for the event. If you click any event summary information link, the first page of the File Events default workflow appears in a new window with all the extra events constrained based on the file type. The File Summary event view opens in a new window, displaying all file events that match on the criteria value you clicked.

Step 6 In the Events table, you can:

- Highlight — Choose a table row to highlight a data point in the map. The map scrolls to display the selected file event if not currently visible.
- Sort — Click the column headers to sort events in ascending or descending order.

Work with Event Data in the Secure Endpoint Console

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, see the online help in the console or other documentation available from https://www.cisco.com/c/en/us/support/security/fireamp-endpoints/tsd-products-support-series-home.html

To access the Secure Endpoint console from the Secure Firewall Management Center, do one of the following:

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 management center, ensure that the Secure Endpoint contextual cross-launch options are properly enabled. See topics under Event Investigation Using Web-Based Resources, on page 620.

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**.

b) Right-click a file SHA and choose the Secure Endpoint option.

History for File and Malware Events and Network File Trajectory

Feature	Minimum Management Center	Minimum 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**

Feature	Minimum Management Center	Minimum Threat Defense	Details
Unique identifier for connection events in syslogs.	6.4.0.4	Any	The following 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, see Configuration Locations for Syslogs for File and Malware Events, on page 632.

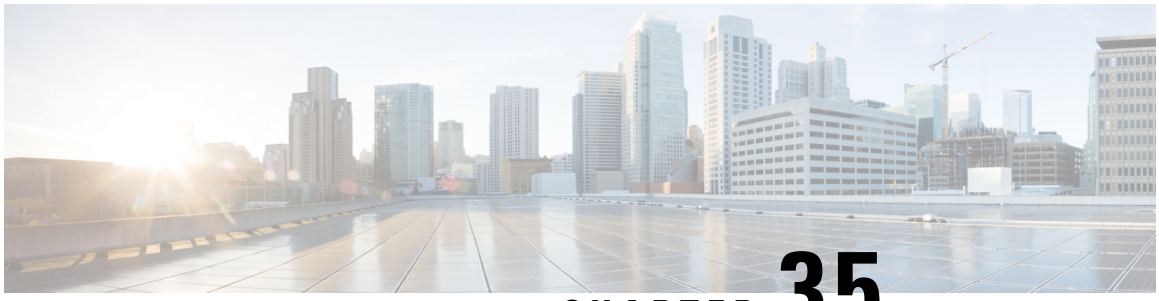

CHAPTER 35

Host Profiles

The following topics describe how to use host profiles:

Requirements and Prerequisites for Host Profiles

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Host Profiles

A host profile provides a complete view of all the information the system has gathered about a single host. To access a host profile:

- navigate from any network map view.
- navigate from any event view that includes the IP addresses of hosts on monitored networks.

Host profiles provide basic information about detected hosts or devices, such as the host name or MAC addresses. Depending on your licenses and system configuration, host profiles can also provide you with the following information:

- the operating system running on a host
- the servers running on a host
- the clients and web applications running on a host
- the protocols running on a host
- the indications of compromise (IOC) tags on a host
- the VLAN tags on a host
- the last twenty-four hours of user activity on your network
- the compliance 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.

Related Topics

Viewing Host Profiles, on page 845

Host Profile Limitations

Unavailable Hosts

A host profile may not be available for every host on your network. Possible reasons include:

- The host was deleted from the network map because it timed out.
- You have reached your host limit.
- The host resides in a network segment that is not monitored by the network discovery policy.

Unavailable Information

The information displayed in a host profile may vary according to the type of host and the information available about the host.

For example:

- If your system detects a host using a non-IP-based protocol like STP, SNAP, or IPX, the host is added to the network map as a MAC host and much less information is available than for an IP host.
- The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data.

(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.

Viewing Host Profiles

Procedure

You have two choices:

- On any network map, drill down to the IP address of the host whose profile you want to view.
- On any event view, click **Host Profile** or **Compromised Host** next to the IP address of the host whose profile you want to view.

Basic Host Information in the Host Profile

Each host profile provides basic information about a detected host or other device.

Descriptions of each of the basic host profile fields follow.

Domain

The domain associated with the host.

IP Addresses

All IP addresses (both IPv4 and IPv6) associated with the host. The system detects IP addresses associated with hosts and, where supported, groups multiple IP addresses used by the same host. IPv6 hosts often have at least two IPv6 addresses (local-only and globally routable), and may also have IPv4 addresses. IPv4-only hosts may have multiple IPv4 addresses.

The host profile lists all detected IP addresses associated with that host. Where available, routable host IP addresses also include a flag icon and country code indicating the geolocation data associated with that address.

Note that only the first three addresses are shown by default. Click **show all** to show all addresses for a host.

Hostname

The fully qualified domain name of the host, if known.

NetBIOS Name

The NetBIOS name of the host, if available. Microsoft Windows hosts, as well as Macintosh, Linux, or other platforms configured to use NetBIOS, can have a NetBIOS name. For example, Linux hosts configured as Samba servers have NetBIOS names.

Device (Hops)

Either:

- the reporting device for the network where the host resides, as defined in the network discovery policy, or
- the device that processed the NetFlow data that added the host to the network map

The number of network hops between the device that detected the host and the host itself follows the device name, in parentheses. If multiple devices can see the host, the reporting device is displayed in bold.

If this field is blank, either:

- the host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy, or
- the host was added using the host input feature and has not also been detected by the 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 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 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 the following values:

- the operating system determined most likely to be running on the host, based on the identity data collected from all sources
- `Pending` if the system has not yet identified an operating system and no other identity data is available

- `unknown` if the system cannot identify the operating system and no other identity data is available for the operating system

Note If the host's operating system is not one the system is capable of detecting, see .

OS Version/Version

The operating system version. If a host is a jailbroken mobile device, `Jailbroken` is indicated in parentheses after the version.

Source

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- Firepower

The system may reconcile data from multiple sources to determine the identity of an operating system.

Viewing Operating System Identities

You can view the specific operating system identities discovered or added for a host. The system uses source prioritization to determine the current identity for the host. In the list of identities, the current identity is highlighted by boldface text.

Note that the **View** is only available if multiple operating system identities exist for the host.

Procedure

Step 1 Click **View** in the **Operating System** or **Operating System Conflicts** section of a host profile.

Step 2 View the information described in Operating Systems in the Host Profile, on page 847.

Step 3 Optionally, click **Delete** ( ) next to any operating system identity.

Note
You cannot delete Cisco-detected operating system identities.

This system removes the identity from the Operating System Identity Information pop-up window and, if applicable, updates the current identity for the operating system in the host profile.

Setting the Current Operating System Identity

You can set the current operating system identity for a host using the web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation. However, if the system detects a conflicting operating system identity for the host after you edit the operating system, an operating system conflict occurs. Both operating systems are then considered current until you resolve the conflict.

Procedure

Step 1 Click **Edit** in the **Operating System** section of a host profile.

Step 2 You have several options:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on the current operating system identity from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and modify the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build**, **Patch**, and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Choose the applicable fixes in the drop-down list, and click **Add**.

Step 8 Optionally, add the relevant patches and extensions using the **Patch** and **Extension** drop-down lists.

Step 9 Click **Finish**.

Related Topics

Operating System Identity Conflicts, on page 850

Operating System Identity Conflicts

An operating system identity conflict occurs when a new identity detected by the system conflicts with the current identity, if that identity was provided by an active source, such as a scanner, application, or user.

The list of operating system identities in conflict displays in bold in the host profile.

You can resolve an identity conflict and set the current operating system identity for a host through the system web interface. Setting the identity through the web interface overrides all other identity sources so that identity is used for vulnerability assessment and impact correlation.

Making a Conflicting Operating System Identity Current

Procedure

Step 1 Navigate to the **Operating System** section of a host profile.

Step 2 You have two choices:

- Click **Make Current** next to the operating system identity you want to set as the operating system for the host.
- If the identity that you *do not* want as the current identity came from an active source, delete the unwanted identity.

Resolving an Operating System Identity Conflict

Procedure

Step 1 Click **Resolve** in the **Operating System Conflicts** section of a host profile.

Step 2 You have the following choices:

- Choose **Current Definition** from the **OS Definition** drop-down list to confirm the current operating system identity through host input, then skip to step 6.
- Choose a variation on one of the conflicting operating system identities from the **OS Definition** drop-down list, then skip to step 6.
- Choose **User-Defined** from the **OS Definition** drop-down list, then continue with step 3.

Step 3 Optionally, choose **Use Custom Display String** and enter the custom strings you want to display in the **Vendor String**, **Product String**, and **Version String** fields.

Step 4 Optionally, to change to an operating system from a different vendor, choose from the **Vendor** and **Product** drop-down lists.

Step 5 Optionally, to configure the operating system product release level, choose from the **Major**, **Minor**, **Revision**, **Build, Patch,** and **Extension** drop-down lists.

Step 6 Optionally, if you want to indicate that fixes for the operating system have been applied, click **Configure Fixes**.

Step 7 Add the fixes you have applied to the fixes list.

Step 8 Click **Finish**.

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 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 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 management center does not limit the number of server identities from active sources such as user input, scanners, or other applications.

The 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 management center configuration, but user input through the 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 the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or other scanner)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for applications detected by the 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.

Viewing Server Details

Procedure

In a host profile, click **View** ( ) next to a server in the **Servers** section.

Editing Server Identities

You can manually update the identity settings for a server on a host and configure any fixes that you have applied to the host to remove the vulnerabilities addressed by the fixes. You can also delete server identities.

Deleting an identity does not delete the server, even if you delete the only identity. Deleting an identity does remove the identity from the Server Detail pop-up window and, if applicable, updates the current identity for the server in the host profile.

You cannot edit or delete server identities added by a Cisco-managed device.

Procedure

Step 1 Navigate to the **Servers** section of a host profile.

Step 2 Click **View** to open the Server Detail pop-up window.

Step 3 To delete a server identity, click **Delete** ( ) next to the server identity you want to remove.

Step 4 To modify a server identity, click **Edit** ( ) next to the server in the servers list.

Step 5 You have two choices:

- Choose the current definition from the **Select Server Type** drop-down list.
- Choose the type of server from the **Select Server Type** drop-down list.

Step 6 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 7 Optionally, to customize the name and version of the server, choose the **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 8 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 9 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 10 Click **Finish**.

Resolving Server Identity Conflicts

A server identity conflict occurs when an active source, such as an application or scanner, adds identity data for a server to a host, after which the system detects traffic for that port that indicates a conflicting server identity.

Procedure

Step 1 In a host profile, navigate to the **Servers** section.

Step 2 Click resolve next to a server.

Step 3 Choose the type of server from the **Select Server Type** drop-down list.

Step 4 Optionally, to only list vendors and products for that server type, choose the **Restrict by Server Type** check box.

Step 5 Optionally, to customize the name and version of the server, choose **Use Custom Display String**, and enter a **Vendor String** and **Version String**.

Step 6 In the **Product Mappings** section, choose the operating system, product, and versions you want to use.

Example:

For example, if you want the server to map to Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 7 If you want to indicate that fixes for the server have been applied, click **Configure Fixes**, and add the patches you want to apply for that server to the fixes list.

Step 8 Click **Finish**.

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 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.

Deleting Web Applications from the Host Profile

You can delete an application from a host profile to remove applications that you know are not running on the host. Note that deleting an application from a host may bring the host into compliance with a compliance allow list.

Note If the system detects the application again, it re-adds it to the network map and the host profile.

Procedure

Step 1 In a host profile, navigate to the **Applications** section.

Step 2 Click **Delete** ( ) next to the application you want to delete.

Host Protocols in the Host Profile

Each host profile contains information about the protocols detected in the network traffic associated with the host. This information includes:

Protocol

The name of a protocol used by the host.

Layer

The network layer where the protocol runs (`Network` or `Transport`).

If a protocol displaying in the host profile violates a compliance allow list in an activated correlation policy, the 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 re-adds it to the network map and the host profile.

Deleting a Protocol From the Host Profile

Procedure

Step 1 Navigate to the **Protocols** section of a host profile.

Step 2 Click **Delete** ( ) next to the protocol you want to delete.

Indications of Compromise in the Host Profile

The 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 892 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 following information for each:

- **VLAN ID** identifies the VLAN where the host is a member. This can be any integer between zero and 4095 for 802.1q VLANs.
- **Type** identifies the encapsulated packet containing the VLAN tag, which can be either Ethernet or Token Ring.
- **Priority** identifies the priority in the VLAN tag, which can be any integer from zero to 7, where 7 is the highest priority.

If VLAN tags are nested within the packet, the system processes and the management center displays the innermost VLAN tag. The system collects and displays VLAN tag information only for MAC addresses that it identifies through ARP and DHCP traffic.

VLAN tag information can be useful, for example, if you have a VLAN composed entirely of printers and the system detects a Microsoft Windows 2000 operating system in that VLAN. VLAN information also helps the system generate more accurate network maps.

User History in the Host Profile

The user history portion of the host profile provides a graphic representation of the last twenty-four hours of user activity. A typical user logs off in the evening and may share the host resource with another user. Periodic login requests, such as those made to check email, are indicated by short regular bars. A list of user identities is provided with bar graphs to indicate when the user login was detected. Note that for non-authoritative logins, the bar graph is gray.

Note that the system does associate a non-authoritative user login to a host with an IP address of that host, so the user does appear in the host's user history. However, if an authoritative user login is detected for the same host, the user associated with the authoritative user login takes over the association with the host IP address, and new non-authoritative user logins do not disrupt that user association with the host IP address. If you configure capture of failed logins in the network discovery policy, the list includes users that failed to log into the host.

Host Attributes in the Host Profile

You can use *host attributes* to classify hosts in ways that are important to your network environment. Three types of attributes are present in the system:

- *predefined host attributes*
- *compliance allow list host attributes*
- *user-defined host attributes*

After you set a predefined host attribute or create a user-defined host attribute, you must assign a host attribute value.

Note Host attributes can be defined at any domain level. You can assign host attributes created in current and ancestor domains.

Predefined Host Attributes

The 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

Each compliance allow list that you create automatically creates a host attribute with the same name as the allow list. 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.

Creating Text- or URL-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**.
Step 5 Choose the **Type** of attribute that you want to create as described in User-Defined Host Attributes, on page 860
Step 6 Click **Save**.

Creating Integer-Based Host Attributes

When you define an integer-based host attribute, you must specify the range of numbers that the attribute accepts.

Procedure

Step 1 Choose **Analysis** > **Hosts 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 860.
Step 6 In the **Min** field, enter the minimum integer value that can be assigned to a host.
Step 7 In the **Max** field, enter the maximum integer value that can be assigned to a host.
Step 8 Click **Save**.

Creating List-Based Host Attributes

When you define a list-based host attribute, you must supply each of the values for the list. These values can contain alphanumeric characters, spaces, and symbols.

Procedure

Step 1	Choose **Analysis** > **Hosts 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 860.
Step 6	To add a value to the list, click **Add Value**.
Step 7	In the **Name** field, enter the first value you want to add.
Step 8	Optionally, to auto-assign the attribute value you just added to your hosts, click **Add Networks**.
Step 9	Choose the value you added from the **Value** drop-down list.
Step 10	In the **IP Address** and **Netmask** fields, enter the IP address and network mask (IPv4) that represent the IP address block where you want to auto-assign this value.
Step 11	Repeat steps 6 through 10 to add additional values to the list and assign them automatically to new hosts that fall within an IP address block.
Step 12	Click **Save**.

Setting Host Attribute Values

You can set values for predefined and user-defined host attributes. You cannot set values for compliance 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 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.

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.

Second, in the sections associated with operating systems, applications, protocols, and servers, the 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.

Note You can use a host's profile to create a shared host profile for compliance allow lists.

Creating 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.

Related Topics

Building Allow List Host Profiles, on page 946

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.

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, see the *Firepower System Host Input API Guide*.

Descriptions of the columns in the Vulnerabilities sections of the host profile follow.

Name

The name of the vulnerability.

Remote

Indicates whether the vulnerability can be remotely exploited. If this column is blank, the vulnerability definition does not include this information.

Component

The name of the operating system, application protocol, or client associated with the vulnerability.

Port

A port number, if the vulnerability is associated with an application protocol running on a specific port.

Related Topics

Vulnerability Data Fields, on page 904
Vulnerability Deactivation, on page 905

Downloading Patches for Vulnerabilities

You can download patches to mitigate the vulnerabilities discovered on the hosts on your network.

Procedure

Step 1 Access the host profile of a host for which you want to download a patch.

Step 2 Expand the **Vulnerabilities** section.

Step 3 Click the name of the vulnerability you want to patch.

Step 4 Expand the **Fixes** section to display the list of patches for the vulnerability.

Step 5 Click **Download** next to the patch you want to download.

Step 6 Download the patch and apply it to your affected systems.

Deactivating Vulnerabilities for Individual Hosts

You can use the host vulnerability editor to deactivate vulnerabilities on a host-by-host basis. When you deactivate a vulnerability for a host, it is still used for impact correlations for that host, but the impact level is automatically reduced one level.

Procedure

Step 1 Navigate to the **Vulnerabilities** section of a host profile.

Step 2 Click **Edit Vulnerabilities**.

Step 3 Choose the vulnerability from the **Valid Vulnerabilities** list, and click the down arrow to move it to the **Invalid Vulnerabilities** list.

Tip
You can click and drag to choose multiple adjacent vulnerabilities; you can also double-click any vulnerability to move it from list to list.

Step 4 Click **Save**.

What to do next

- Optionally, activate the vulnerability for the host by moving it from the **Invalid Vulnerabilities** list to the **Valid Vulnerabilities** list.

Related Topics

Deactivating Individual Vulnerabilities, on page 866
Deactivating Multiple Vulnerabilities, on page 907

Deactivating Individual Vulnerabilities

If you deactivate a vulnerability in a host profile, it deactivates it for all hosts in your network map. However, you can reactivate it at any time.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the vulnerability detail:

- In an affected host profile, expand the **Vulnerabilities** section, and click the name of the vulnerability you want to enable or disable.
- In the predefined workflow, choose **Analysis** > **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.

Related Topics

Deactivating Vulnerabilities for Individual Hosts, on page 866
Deactivating Multiple Vulnerabilities, on page 907
Operating System Identity Conflicts, on page 850

Scan Results in the Host Profile

When you scan a host using Nmap, or when you import results from an Nmap scan, those results appear in the host profile for any hosts included in the scan.

The information that Nmap collects about the host operating system and any servers running on open unfiltered ports is added directly into the Operating System and Servers sections of the host profile, respectively. In addition, Nmap adds a list of the scan results for that host in the Scan Results section. Note that the scan must find open ports on the host for Scan Results section to appear in the profile.

Each result indicates the source of the information, the number and type of the scanned port, the name of the server running on the port, and any additional information detected by Nmap, such as the state of the port or the vendor name for the server. If you scan for UDP ports, servers detected on those ports only appear in the Scan Results section.

Note that you can run an Nmap scan from the host profile.

Scanning a Host from the Host Profile

You can perform a Nmap scan against a host from the host profile. After the scan completes, server and operating system information for that host are updated in the host profile. Any additional scan results are added to the Scan Results section of the host profile.

Caution Nmap-supplied server and operating system data remains static until you run another Nmap scan or override it with higher priority host input. If you plan to scan a host using Nmap, regularly schedule scans.

Before you begin

- Add an Nmap scan instance; see 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.

Related Topics

Nmap Scan Automation, on page 493

History for Host Profiles

Feature	Minimum Management Center	Minimum 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: management center

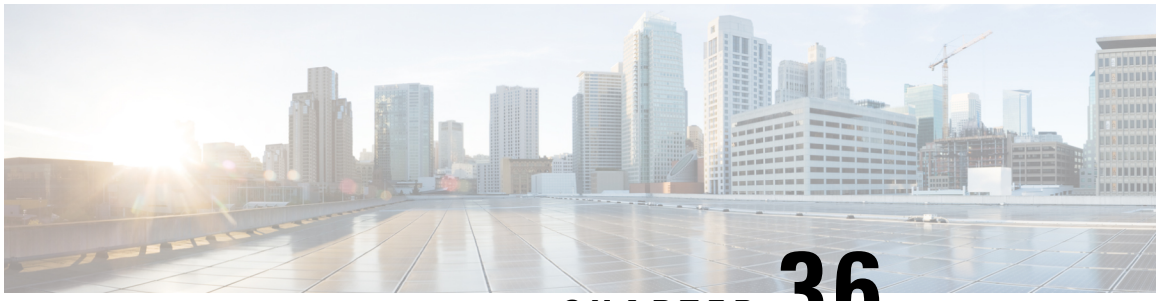

CHAPTER 36

Discovery Events

The following topics describe how to work with discovery events:

Requirements and Prerequisites for Discovery Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Discovery and Identity Data in Discovery Events

The system generates tables of events that represent the changes detected in your monitored network. You can use these tables to review the user activity on your network and determine how to respond. The *network discovery* and *identity* policies specify the kinds of data you want to collect, the network segments you want to monitor, and the specific hardware interfaces you want to use to do it.

You can use discovery and identity event tables to identify threats associated with hosts, applications, and users on your network. The system provides a set of predefined workflows that you can use to analyze the

events that your system generates. You can also create custom workflows that display only the information that matches your specific needs.

To collect and store network discovery and identity data for analysis, you must configure network discovery and identity policies. After you configure an identity policy, you must invoke it in your access control policy and deploy it to the devices you want to use to monitor traffic.

Your network discovery policy provides host, application, and non-authoritative user data. Your identity policy provides authoritative user data.

The following discovery event tables are located under the Analysis > Hosts and Analysis > Users menus.

Discovery Event Table	Populated With Discovery Data?	Populated With Identity Data?
Hosts	Yes	No
Host Indications of Compromise	Yes	No
Applications	Yes	No
Application Details	Yes	No
Servers	Yes	No
Host Attributes	Yes	No
Discovery Events	Yes	Yes
User Indications of Compromise	Yes	Yes
Active Sessions	Yes	Yes
User Activity	Yes	Yes
Users	Yes	Yes
Vulnerabilities	Yes	No
Third-Party Vulnerabilities	Yes	No

Viewing Discovery Event Statistics

The Discovery Statistics page displays a summary of the hosts, events, protocols, application protocols, and operating systems detected by the system.

The page lists statistics for the last hour and the total accumulated statistics. You can choose to view statistics for a particular device, or all devices. You can also view events that match the entries on the page by clicking the event, server, operating system, or operating system vendor listed within the summary.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Statistics**.

Step 2 From the **Select Device** list, choose the device whose statistics you want to view. Optionally, choose **All** to view statistics for all devices managed by the management center.

Step 3 You have the following options:

- In the Statistics Summary, view general statistics as described in The Statistics Summary Section, on page 871.
- In the Event Breakdown, click the type of event you want to view. If no events appear, you may need to adjust the time range as described in Changing the Time Window, on page 675.
- 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 872
The Protocol Breakdown Section, on page 872
The Application Protocol Breakdown Section, on page 873
The OS Breakdown Section, on page 873

The Statistics Summary Section

Descriptions of the rows of the Statistics Summary section follow.

Total Events

Total number of discovery events stored on the 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 management center. Note that the host limit usage only appears if you are viewing statistics for all managed devices.

Note If the host limit is reached and a host is deleted, the host will not reappear on the network map you purge discovery data.

Last Event Received

The date and time that the most recent discovery event occurred.

Last Connection Received

The date and time that the most recent connection was completed.

The Event Breakdown Section

The Event Breakdown section lists a count of each type of discovery event and host input event that occurred within the last hour, as well as a count of the total number of each event type stored in the database.

You can also use the Event Breakdown section to view details on discovery and host input events.

Related Topics

Discovery and Host Input Events, on page 876

The Protocol Breakdown Section

The Protocol Breakdown section lists the protocols currently in use by detected hosts. It displays each detected protocol name, its “layer” in the protocol stack, and the total number of hosts that communicate using the protocol.

The Application Protocol Breakdown Section

The Application Protocol Breakdown section lists the application protocols that are currently in use by detected hosts. It lists the protocol name, the total number of hosts running the application protocol in the past hour, and the total number of hosts that have been detected running the protocol at any point.

You can also use the Application Protocol Breakdown section to view details on servers using the detected protocols.

Related Topics

Server Data, on page 896

The OS Breakdown Section

The OS Breakdown section lists the operating systems currently running on the monitored network, along with their vendors and the total number of hosts running each operating system.

A value of `unknown` for the operating system name or version means that the operating system or its version does not match any of the system's fingerprints. A value of `pending` means that the system has not yet gathered enough information to identify the operating system or its version.

You can use the OS Breakdown section to view details on the detected operating systems.

Related Topics

Host Data, on page 884

Viewing Discovery Performance Graphs

You can generate graphs that display performance statistics for managed devices with discovery events.

New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

Edit the applicable network discovery policy to include applications, hosts, and users. (This may impact system performance.) See Configuring Network Discovery Rules and Actions and Discovered Assets.

You must be an Admin or Maintenance user to perform this task.

Procedure

Step 1 Choose **Overview** > **Summary** > **Discovery Performance**.

Step 2 From the **Select Device** list, choose the 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 874.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph** to graph the selected statistics.

Discovery Performance Graph Types

Descriptions of the available graph types follow.

Processed Events/Sec

Displays a graph that represents the number of events that the Data Correlator processes per second

Processed Connections/Sec

Displays a graph that represents the number of connections that the Data Correlator processes per second

Generated Events/Sec

Displays a graph that represents the number of events that the system generates per second

Mbits/Sec

Displays a graph that represents the number of megabits of traffic that are analyzed by the discovery process per second

Avg Bytes/Packet

Displays a graph that represents the average number of bytes included in each packet analyzed by the discovery process

K Packets/Sec

Displays a graph that represents the number of packets analyzed by the discovery process per second, in thousands

Using Discovery and Identity Workflows

The management center provides a set of event workflows that you can use to analyze the discovery and identity data that is generated for your network. The workflows are, along with the network map, a key source of information about your network assets.

The management center provides predefined workflows for discovery and identity data, detected hosts and their host attributes, servers, applications, application details, vulnerabilities, user activities, and users. You can also create custom workflows.

Procedure

Step 1 To access a predefined workflow:

- Discovery and Host Input Data — See Viewing Discovery and Host Input Events, on page 882.

- Host Data — See Viewing Host Data, on page 884.
- Host Attributes Data — See Viewing Host Attributes, on page 890.
- Host or User Indications of Compromise Data — See View and Work with Indications of Compromise Data, on page 892.
- Server Data — See Viewing Server Data, on page 896.
- Application Data — See Viewing Application Data, on page 899.
- Application Detail Data — See Viewing Application Detail Data, on page 901.
- Active Session Data — See Viewing Active Session Data, on page 917.
- User Data — See Viewing User Data, on page 919.
- User Activity Data — See Viewing User Activity Data, on page 922.
- Network Map — See Viewing Network Maps, on page 608.

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 > Active Sessions page, see 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, see Using Drill-Down Pages, on page 664.

- Navigate Current Page — To navigate within the current workflow page, see Workflow Page Navigation Tools, on page 661.
- Navigate within a Workflow — To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- Navigate to Other Workflows — To navigate to other event views to examine associated events, see Inter-Workflow Navigation, on page 680.
- Sort Data — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- View Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- View User Profile — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.

Related Topics

Using Workflows, on page 656
Purging Data from the Management Center Database, on page 516

Discovery and Host Input Events

The system generates discovery events that communicate the details of changes in your monitored network segments. *New* events are generated for newly discovered network features, and change events are generated for any change in previously identified network assets.

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers discovered running on each host. Optionally, you can configure the system to use exported NetFlow records to generate these new host and server events.

In addition, the system generates new events for each network, transport, and application protocol running on each discovered host. You can disable detection of application protocols in discovery rules configured to monitor NetFlow exporters, but not in discovery rules configured to monitor 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 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 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 management center is reached and a new host is dropped. Compare this with the previous event where old hosts are deleted from the network map when the host limit is reached.

To drop new hosts when the host limit is reached, go to **Policies** > **Network Discovery** > **Advanced** and set **When Host Limit Reached** to **Drop hosts**.

Host IOC Set

This event is generated when an IOC (Indications of Compromise) is set for a host and generates an alert.

Host Timeout

This event is generated when a host is dropped from the network map because the host has not produced traffic within the interval defined in the network discovery policy. Note that individual host IP addresses and MAC addresses time out individually; a host does not disappear from the network map unless all of its associated addresses have timed out.

If you change the networks you want to monitor in your network discovery policy, you may want to manually delete old hosts from the network map so that they do not count against your host limit.

Host Type Changed to Network Device

This event is generated when the system detects that a detected host is actually a network device.

Identity Conflict

This event is generated when the system detects a new server or operating system identity that conflicts with a current active identity for that server or operating system.

If you want to resolve identity conflicts by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

Identity Timeout

This event is generated when server or operating system identity data from an active source times out.

If you want to refresh identity data by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

MAC Information Change

This event is generated when the system detects a change in the information associated with a specific MAC address or TTL value.

This event often occurs when the system detects hosts passing traffic through a router. While each host 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 880

Host Input Event Types

When you view a table of discovery events, the event type is listed in the **Event** column.

Contrast host input events, which are generated when a user takes a specific action (such as manually adding a host), with discovery events, which are generated when the system itself detects a change in your monitored network (such as detecting traffic from a previously undetected host).

You can configure the types of host input events that the system logs by modifying your network discovery policy.

If you understand the information the different types of host input events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the names of the event types can help you craft more effective event searches. Descriptions of the different types of host input events follow.

Add Client

This event is generated when a user adds a client.

Add Host

This event is generated when a user adds a host.

Add Protocol

This event is generated when a user adds a protocol.

Add Scan Result

This event is generated when the system adds the results of an Nmap scan to a host.

Add Port

This event is generated when a user adds a server port.

Delete Client

This event is generated when a user deletes a client from the system.

Delete Host/Network

This event is generated when a user deletes an IP address or subnet from the system.

Delete Protocol

This event is generated when a user deletes a protocol from the system.

Delete Port

This event is generated when a user deletes a server port or group of server ports from the system.

Host Attribute Add

This event is generated when a user creates a new host attribute.

Host Attribute Delete

This event is generated when a user deletes a user-defined host attribute.

Host Attribute Delete Value

This event is generated when a user deletes a value assigned to a host attribute.

Host Attribute Set Value

This event is generated when a user sets a host attribute value for a host.

Host Attribute Update

This event is generated when a user changes the definition of a user-defined host attribute.

Set Host Criticality

This event is generated when a user sets or modifies the host criticality value for a host.

Set Operating System Definition

This event is generated when a user sets the operating system for a host.

Set Server Definition

This event is generated when a user sets the vendor and version definitions for a server.

Set Vulnerability Impact Qualification

This event is generated when a vulnerability impact qualification is set.

When a vulnerability is disabled at a global level from being used for impact qualifications, or when a vulnerability is enabled at a global level, this event is generated.

Vulnerability Set Invalid

This event is generated when a user invalidates (or reviews) a vulnerability or vulnerabilities.

Vulnerability Set Valid

This event is generated when a user validates a vulnerability that was previously marked as invalid.

Related Topics

Discovery Event Types, on page 877

Viewing Discovery and Host Input Events

Discovery events workflows allow you to view data from both discovery events and host input events. You can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of discovery events and a terminating host view page. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Choose **Analysis** > **Hosts heading** > **Discovery Events**.

Step 2 You have the following options:

- Adjust the time range as described in Changing the Time Window, on page 675.

  Note
  Events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.

- Learn more about the contents of the columns in the table; see Discovery Event Fields, on page 883.

Related Topics

Using Discovery and Identity Workflows, on page 874

Discovery Event Fields

Descriptions of the fields that can be viewed and searched in the discovery events table follow.

Time

The time that the system generated the event.

Event

The discovery event type or host input event type.

IP Address

The IP address associated with the host involved in the event.

User

The last user to log into the host involved in the event before the event was generated. If only non-authoritative users log in after an authoritative user, the authoritative user remains the current user for the host unless another authoritative user logs in.

MAC Address

The MAC address of the NIC used by the network traffic that triggered the discovery event. This MAC address can be either the actual MAC address of the host involved in the event, or the MAC address of a network device that the traffic passed through.

MAC Vendor

The MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.

When searching this field, enter **`virtual_mac_vendor`** to match events that involve virtual hosts.

Port

The port used by the traffic that triggered the event, if applicable.

Description

The text description of the event.

Domain

The domain of the device that discovered the host. This field is only present if you have ever configured the 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 685

Host Data

The system generates an event when it detects a host and collects information about it to build the host profile. You can use the management center web interface to view, search, and delete hosts.

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.

Viewing Host Data

You can use the 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; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Host Data Fields, on page 885.
- Right-click an item in the table to see options. (Not every column offers options.)
- Assign a host attribute to specific hosts; see Setting Host Attributes for Selected Hosts, on page 891.
- Create traffic profiles for specific hosts, see Creating a Traffic Profile for Selected Hosts, on page 888.

- Create a compliance allow list based on specific hosts, see Creating a Compliance Allow List Based on Selected Hosts, on page 889.

Host Data Fields

When the system discovers a host, it collects data about that host. That data can include the host's IP addresses, the operating system it is running, and more. You can view some of that information in the table view of hosts.

Descriptions of the fields that can be viewed and searched in the hosts table follow below.

Last Seen

The date and time any of the host's IP addresses was last detected by the system. The Last Seen value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system generates a new host event for any of the host's IP addresses.

For hosts with operating system data updated using the host input feature, the Last Seen value indicates the date and time when the data was originally added.

IP Address

The IP addresses associated with the host.

MAC Address

The host's detected MAC address of the NIC.

The MAC Address field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Address field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

MAC Vendor

The host's detected MAC hardware vendor of the NIC.

The MAC Vendor field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Vendor field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

When searching this field, enter `virtual_mac_vendor` to match events that involve virtual hosts.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the

current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-specified criticality value assigned to the host.

NetBIOS Name

The NetBIOS name of the host. Only hosts running the NetBIOS protocol will have a NetBIOS name.

VLAN ID

VLAN ID used by the host.

Hops

The number of network hops from the device that detected the host to the host.

Host Type

The type of host. Can be any of the following: host, mobile device, jailbroken mobile device, router, bridge, NAT device, and load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a device is not identified as a network device, it is categorized as a host.

When searching this field, enter `!host` to search for all network devices.

Hardware

The hardware platform for a mobile device.

OS

One of the following:

- The operating system (name, vendor, and version) either detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple identities, it displays those identities in a comma-separated list.

This field appears when you invoke the hosts event view from the Custom Analysis widget on the dashboard. It is also a field option in custom tables based on the Hosts table.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Conflict

This field is search only.

OS Vendor

One of the following:

- The vendor of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple vendors, it displays those vendors in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Name

One of the following:

- The operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple names, it displays those names in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Version

One of the following:

- The version of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple versions, it displays those versions in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

Source Type

The type of source used to establish the host's operating system identity:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)

- `Firepower` for operating systems detected by the system

The system may reconcile data from multiple sources to determine the identity of an operating system.

Confidence

One of the following:

- the percentage of confidence that the system has in the identity of the operating system running on the host, for hosts detected by the system
- 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner
- `unknown`, for hosts for which the system cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

When searching this field, enter `n/a` to include hosts added to the network map based on NetFlow data.

Notes

The user-defined content of the Notes host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the 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 685
Operating System Identity Conflicts, on page 850

Creating a Traffic Profile for Selected Hosts

A traffic profile is a profile of the traffic on your network, based on connection data collected over a timespan that you specify. After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

You can use the Hosts page to create a traffic profile for a group of hosts that you specify. The traffic profile will be based on connections detected where one of the hosts you specify is the initiating host. Use the sort and search features to isolate the hosts for which you want to create a profile.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create a traffic profile.

Step 2 At the bottom of the page, click **Create Traffic Profile**.

Step 3 Modify and save the traffic profile according to your specific needs.

Related Topics

Introduction to Traffic Profiles, on page 991

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 the hosts that you want to use to create an allow list.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create an allow list.

Step 2 At the bottom of the page, click **CreateAllow List**.

Step 3 Modify and save the allow list according to your specific needs.

Related Topics

Introduction to Compliance Allow Lists, on page 939

Host Attribute Data

The system collects information about the hosts it detects and uses that information to build host profiles. However, there may be additional information about the hosts on your network that you want to provide to your analysts. You can add notes to a host profile, set the business criticality of a host, or provide any other information that you choose. Each piece of information is called a *host attribute*.

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also set attribute values in response to a correlation rule.

Related Topics

Viewing Host Attributes, on page 890
Configuring Set Attribute Remediations, on page 1012

Viewing Host Attributes

You can use the 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; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Host Attribute Data Fields, on page 890.
- Assign a host attribute to specific hosts; see Setting Host Attributes for Selected Hosts, on page 891.

Host Attribute Data Fields

Note that the host attributes table does not display hosts identified only by MAC addresses.

Descriptions of the fields that can be viewed and searched in the host attributes table follow.

IP Address

The IP addresses associated with a host.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the

current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-assigned importance of a host to your enterprise. You can use the host criticality in correlation rules and policies to tailor policy violations and their responses to the importance of a host involved in an event. You can assign a host criticality of low, medium, high, or none.

Notes

Information about the host that you want other analysts to view.

Any user-defined host attribute, including those for compliance 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 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 685

Setting Host Attributes for Selected Hosts

You can configure predefined and user-defined host attributes from a host workflow.

Procedure

Step 1 In a host workflow, check the check boxes next to the hosts to which you want to add a host attribute.

Tip
Use the sort and search features to isolate the hosts to which you want to assign particular attributes.

Step 2 At the bottom of the page, click **Set Attributes**.

Step 3 Optionally, set the host criticality for the hosts you selected. You can choose **None**, **Low**, **Medium**, or **High**.

Step 4 Optionally, add notes to the host profiles of the hosts you selected in the text box.

Step 5 Optionally, set any user-defined host attributes you have configured.

Step 6 Click **Save**.

Indications of Compromise Data

The system correlates various types of data (intrusion events, Security Intelligence, connection events, and file or malware events) to determine whether a host on your monitored network is likely to be compromised by malicious means. Certain combinations and frequencies of event data trigger indications of compromise (IOC) tags on affected hosts. The IP addresses of these hosts appear in event views with a **Red Compromised Host icon** .

When a host is identified as potentially compromised, the user associated with that compromise is also tagged. These users appear in event views with a **Red User icon**.

If a file containing malware is seen again within 300 seconds of being tagged as an IOC, another IOC is not generated. If the same file is seen more than 300 seconds later, a new IOC will be generated.

To configure the system to tag events as indications of compromise, see *Enabling Indications of Compromise Rules* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Related Topics

Editing Server Identities, on page 854

View and Work with Indications of Compromise Data

You can use the 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.

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 > 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 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 > Users > Indications of Compromise**. • If you are using a custom workflow that does not include the User IOC table view, click **(switch workflow)**, then choose **User Indications of Compromise**.

Step 3 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Indications of Compromise Data Fields, on page 894.
- On a Host Indications of Compromise page: View the host profile for a compromised host by clicking **Compromised Host** in the **IP Address** column.
- On a User Indications of Compromise page: View the user profile associated with a compromise by clicking **Red User** in the **User** column.
- Mark IOC events resolved so they no longer appear in the list. To do so, check the check boxes next to the IOC events you want to modify, then click **Mark Resolved**.

- View details of events that triggered the IOC by clicking **View** ( ) in the **First Seen** or **Last Seen** columns.
- See more options: Right-click a value in the table.

Indications of Compromise Data Fields

The following are the fields in Host or User IOC (indication of compromise) tables. Not every IOC-related table includes all fields.

IP Address (When viewing Host IOC data)

The IP address associated with the host that triggered the IOC.

User (When viewing User IOC data)

The username, realm, and authentication source of the user associated with the event that triggered the IOC.

Category

Brief description of the type of compromise indicated, such as `Malware Executed` or `Impact 1 Attack`.

Event Type

Identifier associated with a specific IOC, referring to the event that triggered it.

Description

Description of the impact on the potentially compromised host, such as `This host may be under remote control` or `Malware has been executed on this host`.

First Seen/Last Seen

The first/most recent date and time that events triggering the IOC occurred.

Domain

The domain of the host that triggered the IOC. This field is only present if you have ever configured the management center for multitenancy.

Related Topics

Event Searches, on page 685

Editing Indication of Compromise Rule States for a Single Host or User

When enabled in a network discovery policy, indication of compromise rules apply to all hosts in the monitored network and to authoritative users that are associated with IOC events on that network. You can disable a rule for an individual host or user to avoid unhelpful IOC tags (for example, you may not want to see IOC tags for a DNS server.) If a rule is disabled in the applicable network discovery policy, it cannot be enabled for a specific host or user. Disabling a rule for a particular host does not affect tagging for the user involved in the same event, and vice-versa.

Procedure

Step 1 Navigate to the **Indications of Compromise** section of a host or user profile.

Step 2 Click **Edit Rule States**.

Step 3 In the **Enabled** column for a rule, click the slider to enable or disable it.

Step 4 Click **Save**.

Viewing Source Events for Indication of Compromise Tags

You can use the Indications of Compromise section of the host profile and the user profile to navigate quickly to the events that triggered the IOC tags. Analyzing these events can give you the information you need to determine what, and whether, action is required to address threats of compromise.

Clicking **View** ( ) next to the timestamp of an IOC tag navigates to the table view of events for the relevant event type, constrained to show only the event that triggered the IOC tag.

Only the first instance of a User IOC is displayed in the 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

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 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.

Viewing Server Data

You can use the 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; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Server Data Fields, on page 897.

- 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

Descriptions of the fields that can be viewed and searched in the servers table follow below.

Last Used

The date and time the server was last used on the network or the date and time that the server was originally updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects a server information update.

IP Address

The IP address associated with the host running the server.

Port

The port where the server is running.

Protocol

The network or transport protocol used by the server.

Application Protocol

One of the following:

- the name of the application protocol for the server
- `pending` if the system cannot positively or negatively identify the server for one of several reasons
- `unknown` if the system cannot identify the server based on known server fingerprints or if the server was added through host input and did not include the application protocol

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

Vendor

One of the following:

- the server vendor as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its vendor based on known server fingerprints, or if the server was added to the network map using NetFlow data

Version

One of the following:

- the server version as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its version based on known server fingerprints, or if the server was added to the network map using NetFlow data

Web Application

The web application based on the payload content detected by the system in the HTTP traffic. Note that if the system detects an application protocol of `HTTP` but cannot detect a specific web application, the system supplies a generic web browsing designation.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

Hits

The number of times the server was accessed. For servers added using the host input feature, this value is always 0.

Source Type

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- `Firepower`, `Firepower Port Match`, or `Firepower Pattern Match` for servers detected by the 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 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 685

Application and Application Details Data

When a monitored host connects to another host, the system can, in many cases, determine what application was used. The system detects the use of many email, instant messaging, peer-to-peer, web applications, as well as other types of applications.

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

If you know which applications are running on which hosts, you can use that knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You can obtain the latest information about 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.

Viewing Application Data

You can use the 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; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Application Data Fields, on page 900.
- 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, see Event Investigation Using Web-Based Resources, on page 620
- 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, see Event Investigation Using Web-Based Resources, on page 620.

Application Data Fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the applications table follow.

Application

The name of the detected application.

IP Address

The IP address associated with the host using the application.

Type

The type of application:

Application Protocols

Represents communications between hosts.

Client Applications

Represents software running on a host.

Web Applications

Represents the content or requested URL for HTTP traffic.

Category

A general classification for the application that describes its most essential function. Each application belongs to at least one category.

Tag

Additional information about the application. Applications can have any number of tags, including none.

Risk

How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from Very Low to Very High.

Of Application Protocol Risk, Client Risk, and Web Application Risk, the highest of the three detected, when available, in the traffic that triggered the intrusion event.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from Very Low to Very High.

Of Application Protocol Business Relevance, Client Business Relevance, and Web Application Business Relevance, the lowest of the three detected, when available, in the traffic that triggered the intrusion event.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Domain

The domain of the host using the application. This field is only present if you have ever configured the 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 685

Viewing Application Detail Data

You can use the 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 Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Application Detail Data Fields, on page 902.
- 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, see Event Investigation Using Web-Based Resources, on page 620
- 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, see Event Investigation Using Web-Based Resources, on page 620.

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 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 685

Vulnerability Data

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 management center to:

- Track and review the vulnerabilities for each host.
- Deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability.

Vulnerabilities for vendorless and versionless servers are not mapped unless the applications protocols used by the servers are mapped in the management center configuration. Vulnerabilities for vendorless and versionless clients cannot be mapped.

Related Topics

Mapping Vulnerabilities for Servers, on page 112

Vulnerability Data Fields

Except as noted, these fields appear on all pages under **Analysis > Hosts > 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

See "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 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 685

Vulnerability Deactivation

Deactivating a vulnerability prevents the system from using that vulnerability to evaluate intrusion impact correlations. You can deactivate a vulnerability after you patch the hosts on your network or otherwise judge them immune. Note that if the system discovers a new host that is affected by that vulnerability, the vulnerability is considered valid (and is not automatically deactivated) for that host.

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network. You can deactivate vulnerabilities within the vulnerabilities workflow only on:

- the second page of the default vulnerabilities workflow, **Vulnerabilities on the Network**, which shows only the vulnerabilities that apply to the hosts on your network
- a page in a vulnerabilities workflow, custom or predefined, that you constrained based on IP address using a search.

You can deactivate a vulnerability for a single host using the network map, using the host's host profile, or by constraining the vulnerabilities workflow based on the IP addresses of the host or hosts for which you want to deactivate vulnerabilities. For hosts with multiple associated IP addresses, this function applies only to the single, selected IP address of that host.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Related Topics

Deactivating Vulnerabilities for Individual Hosts, on page 866
Deactivating Individual Vulnerabilities, on page 866

Deactivating Multiple Vulnerabilities, on page 907

Viewing Vulnerability Data

You can use the 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 the following options:

- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Deactivate vulnerabilities so they are no longer used for intrusion impact correlation for currently vulnerable hosts; see Deactivating Multiple Vulnerabilities, on page 907.
- View the details for a vulnerability by clicking **View** ( ) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. See options for viewing additional details at Viewing Vulnerability Details, on page 906.
- View the full text of a vulnerability title by right-clicking the title and choosing **Show Full Text**.

Viewing Vulnerability Details

Procedure

You can view vulnerability details in any of the following ways:

- Choose **Analysis** > **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 > Vulnerabilities**, 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.

Deactivating Multiple Vulnerabilities

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices so long as the vulnerability is activated in the ancestor domain.

Procedure

Step 1 Access the table of vulnerabilities:

- If you are using the predefined vulnerabilities workflow, choose **Analysis** > **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

Deactivating Vulnerabilities for Individual Hosts, on page 866
Deactivating Individual Vulnerabilities, on page 866

Third-Party Vulnerability Data

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, see the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map third-party vulnerability information to the operating system and application definitions in the database. You cannot map third-party vulnerability information to client definitions.

Viewing Third-Party Vulnerability Data

After you use the host input feature to import third-party vulnerability data, you can use the 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 the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Third-Party Vulnerability Data Fields, on page 908.
- View the vulnerability details for a third-party vulnerability by clicking **View** ( ) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page.

Third-Party Vulnerability Data Fields

Descriptions of the fields that can be viewed and searched in the third-party vulnerabilities table follow.

Vulnerability Source

The source of the third-party vulnerabilities, for example, QualysGuard or NeXpose.

Vulnerability ID

The ID number associated with the vulnerability for its source.

IP Address

The IP address associated with the host affected by the vulnerability.

Port

A port number, if the vulnerability is associated with a server running on a specific port.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database.

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 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 685

Active Sessions, Users, and User Activity Data

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, see Active Sessions Data, on page 916.
- 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, see User Data, on page 917.
- 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, see User Activity Data, on page 920.

For more information about the user identity sources that populate these workflows, see 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, see the Cisco Secure Firewall Management Center Device Configuration Guide.	Yes	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
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 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 management center but are not sent to managed devices.	No	Yes	No
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 management center cannot correlate the user in the 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

Field	Description	Active Sessions Table	Users Table	User Activity Table
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 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
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 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

Field	Description	Active Sessions Table	Users Table	User Activity Table
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
First Name	The user's first name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The management center cannot correlate the user in the 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

Field	Description	Active Sessions Table	Users Table	User Activity Table
Last Name	The user's last name, as obtained by a realm. This field is blank if: • You have not configured a realm. • The management center cannot correlate the user in the 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
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 management center cannot correlate the user in the 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 management center may store invalid AIM users (for example, if a user misspelled his or her username).	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 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 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 VPN module of Cisco Secure Client 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 Secure 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 Secure 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 Secure 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** > **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, see User Data, on page 917 and User Activity Data, on page 920.

For information about general user-related event troubleshooting and Remote Access VPN Troubleshooting, see the *Troubleshoot Realms and User Downloads* and *VPN Troubleshooting* in the Cisco Secure Firewall Management Center Device Configuration Guide.

Viewing 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 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see Active Sessions Data, on page 916 and User-Related Fields, on page 910.

User Data

When an identity source reports a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type.

The system updates the users database when one of the following occurs:

- A user on the 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)
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 management center queries the servers based on the interval you specified. It may take five to ten minutes for the management center database to update with user metadata after the system detects a new user login. The 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 management center can store in its database depends on your management center model. When a non-authoritative user login to a host is detected, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user.

Note that traffic-based detection of AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent overuse of user count because of duplicate user records from these protocols, configure traffic-based detection to ignore those protocols.

You can search, view, and delete users from the database; you can also purge all users from the database.

For information about general user-related event troubleshooting, see Cisco Secure Firewall Management Center Device Configuration Guide.

Viewing User Data

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

Step 1 Access the users data:

- If you are using the predefined workflow, choose **Analysis** > **Users** > **Users**.
- If you are using a custom workflow that does not include the table view of users, click **(switch workflow)**, then choose **Users**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see User-Related Fields, on page 910.

User Activity Data

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 management center model.

After you reach the user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

User Indications of Compromise Events

The following user IOC changes are logged in the user activity database:

- When indications of compromise are resolved.
- When indication of compromise rules are enabled or disabled for users.

For information about general user-related event troubleshooting, see the Cisco Secure Firewall Management Center Device Configuration Guide.

Viewing User Activity Data

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for. The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Access the user activity data:

- If you are using the predefined workflow, choose **Analysis** > **Users 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; see Changing the Time Window, on page 675.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
- Perform basic workflow actions; see Using Discovery and Identity Workflows, on page 874.
- Learn more about the contents of the columns in the table; see User-Related Fields, on page 910.

User Profile and Host History

You can learn more about a specific user by viewing the User pop-up window. The page that appears, called the "User Profile" in this document, is titled "User Identity" in the web interface.

You can display the window from:

- any event view that associates user data with other kinds of events
- the table view of 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, see:

- *Indications of Compromise* in the Cisco Secure Firewall Management Center Device Configuration Guide
- Indications of Compromise Data Fields, on page 894

- Editing Indication of Compromise Rule States for a Single Host or User, on page 894
- Resolving Indication of Compromise Tags, on page 895
- Viewing Source Events for Indication of Compromise Tags, on page 895

Host History Section

The host history provides a graphic representation of the last twenty-four hours of the user's activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example, periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

If you use traffic-based detection or captive portal to capture failed logins, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit.

Related Topics

User Data Fields

Viewing User Details and Host History

Procedure

You have two options:

- In any event view that lists users, click user that appears next to a user identity **User icon**, or, for users associated with an indication of compromise, **Red User icon**.
- In any users workflow, click the Users terminating page.

History for Working with Discovery Events

Table 126:

Feature	Minimum Management Center	Minimum 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: management center

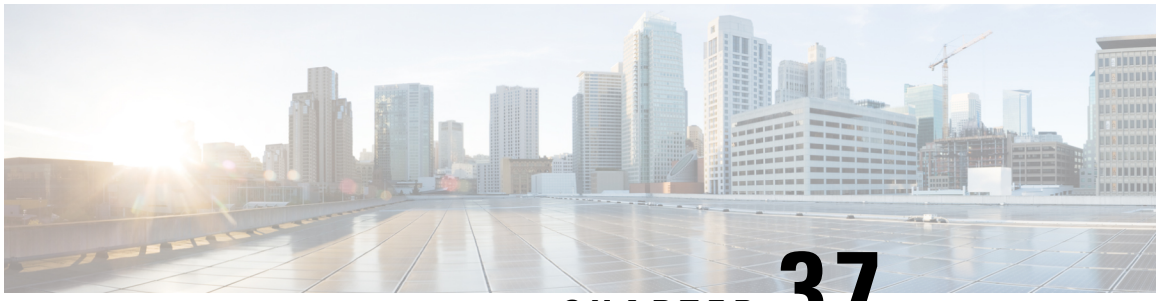

CHAPTER 37

Correlation and Compliance Events

The following topics describe how to view correlation and compliance events.

Viewing Correlation Events

When a correlation rule within an active correlation policy triggers, the system generates a correlation event and logs it to the database.

Note When a compliance allow list within an active correlation policy triggers, the system generates an allow list event.

You can view a table of correlation events, then manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access correlation events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of correlation events. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

You must be an Admin or Security Analyst user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Correlation Events** .

Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Tip
If you are using a custom workflow that does not include the table view of correlation events, click **(switch workflow)**, then choose **Correlation Events**.

Step 2 Optionally, adjust the time range as described in Changing the Time Window, on page 675.

Step 3 Perform any of the following actions:

- To learn more about the columns that appear, see Correlation Event Fields, on page 926.
- To view the host profile for an IP address, click host profile that appears next to the IP address.
- To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- To sort and constrain events or to navigate within the current workflow page, see Using Workflows, on page 656.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To drill down to the next page in the Workflows, constraining on a specific value, see Using Drill-Down Pages, on page 664.
- To delete some or all correlation events, check the check boxes next to the events you want to delete and click **Delete**, or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To navigate to other event views to view associated events, see Inter-Workflow Navigation, on page 680.
- 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, see Event Investigation Using Web-Based Resources, on page 620
- 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, see Event Investigation Using Web-Based Resources, on page 620.

Related Topics

Database Event Limits, on page 59
Workflow Pages, on page 660

Correlation Event Fields

When a correlation rule triggers, the system generates a correlation event. The fields in the correlation events table that can be viewed and searched are described in the following table.

Table 127: Correlation Event Fields

Field	Description
Description	The description of the correlation event. The information in the description depends on how the rule was triggered. For example, if the rule was triggered by an operating system information update event, the new operating system name and confidence level appears.
Device	The name of the device that generated the event that triggered the policy violation.
Domain	The domain of the device whose monitored traffic triggered the policy violation. This field is only present if you have ever configured the 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 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.

Field	Description
Security Intelligence Category	The name of the object that represents or contains the blocked IP address in the event that triggered the policy violation. When searching this field, specify the Security Intelligence category associated with the correlation event that triggered the policy violation. The Security Intelligence category can be the name of a Security Intelligence object, the global Block list, a custom Security Intelligence list or feed, or one of the categories in the Intelligence Feed.
Source Continent or Destination Continent	The continent associated with the source or destination host IP addresses in the event that triggered the policy violation.
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 685

Using Compliance Allow List Workflows

The 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.

Procedure

Step 1 Access an allow list workflow using the **Analysis** > **Correlation** menu.

Step 2 You have the following options:

- Switch Workflow — To use a different workflow, including a custom workflow, click **(switch workflow)**.
- Time Range — To adjust the time range, which is useful if no events appear, see Changing the Time Window, on page 675.
- Host Profile — To view the host profile for an IP address, click **Host Profile**() or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- User Profile (events only) — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.
- Constrain — To constrain the columns that appear, click **Close** (×) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip
  To hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, expand the search constraints, then click the column name under Disabled Columns.

- Drill Down — See Using Drill-Down Pages, on page 664.
- Sort — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- Navigate This Page — See Workflow Page Traversal Tools, on page 661.
- 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 660
Configuring Event View Settings, on page 206

Viewing 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. list events are a special kind of correlation event, and are logged to the management center correlation event database.

You can use the management center to view a table of compliance allow list events. Then, you can manipulate the event view depending on the information you are looking for.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

The page you see when you access allow list events differs depending on the workflow you use. You can use a predefined workflow, which terminates in a table view of events. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

You must be an Admin, Security Analyst, or Discovery Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Allow List** > **Events**.

Step 2 You have the following options:

- To perform basic workflow actions, see Using Compliance Allow List Workflows, on page 928.
- To learn more about the contents of the columns in the table, see Allow List Event Fields, on page 930.
- To see more options, right-click values in the table.

Allow List Event Fields

Allow list events, which you can view and search using workflows, contain the following fields.

Device

The name of the managed device that detected the allow list violation.

Description

A description 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 management center for multitenancy.

Host Criticality

The user-assigned host criticality of 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 that was 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. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Viewing 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.

You can use the management center to view a table of allow list violations for all active allow lists. Then, you can manipulate the event view depending on the information you are looking for.

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.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Allow List** > **Violations**.

Step 2 You have the following options:

- To perform basic workflow actions, see Using Compliance Allow List Workflows, on page 928.
- To learn more about the contents of the columns in the table, see Allow List Violation Fields, on page 932.
- To see more options, right-click values in the table.

Allow List Violation Fields

Allow list violations, which you can view and search using workflows, contain the following fields.

Domain

The domain where the non-compliant host resides. This field is only present if you have ever configured the 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. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

Remediation Status Events

When a remediation triggers, the system logs a remediation status event to the database. These events can be viewed on the Remediation Status page. You can search, view, and delete remediation status events.

Related Topics

Remediation Status Table Fields, on page 934

Viewing Remediation Status Events

The page you see when you access remediation status events differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of remediations. The table view contains a row for each remediation status event. You can also create a custom workflow that displays only the information that matches your specific needs.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Step 2 Optionally, adjust the time range as described in Changing the Time Window, on page 675.

Step 3 Optionally, to use a different workflow, including a custom workflow, click **(switch workflow)** by the workflow title.

Tip
If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then choose **Remediation Status**.

Step 4 You have the following options:

- To learn more about the columns that appear, see Remediation Status Table Fields, on page 934.
- To sort and constrain the events, see Using Workflows, on page 656.
- To navigate to the correlation events view to see associated events, click **Correlation Events**.
- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**. To navigate to the bookmark management page, click **View Bookmarks**.
- To generate a report based on the data in the table view, click **Report Designer** as described in Creating a Report Template from an Event View, on page 529.
- To drill down to the next page in the workflow, see Using Drill-Down Pages, on page 664.
- To delete remediation status events from the system, check the check boxes next to events you want to delete and click **Delete** or click **Delete All** and confirm you want to delete all the events in the current constrained view.
- To search for remediation status events, click **Search**.

Related Topics

Using Workflows, on page 656

Remediation Status Table Fields

The following table describes the fields in the remediation status table that can be viewed and searched.

Table 128: 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 management center for multitenancy.
Policy	The name of the correlation policy that was violated and triggered the remediation.

Field	Description
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 • Unknown/unexpected error If custom remediation modules are installed, you may see additional status messages that are implemented by the custom module.
Rule	The name of the correlation rule that triggered the remediation.
Time	The date and time that the 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 685

Using the Remediation Status Events Table

You can change the layout of the event view or constrain the events in the view by a field value.

When you disable a column, it is disabled for the duration of your session unless you add it back later. If you disable the first column, the Count column is added.

Clicking a value within a row in a table view constrains the table view and does not drill down to the next page.

Tip Table views always include “Table View” in the page name.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Before you begin

You must be an Admin user to perform this task.

Procedure

Step 1 Choose **Analysis** > **Correlation** > **Status**.

Tip
If you are using a custom workflow that does not include the table view of remediations, click **(switch workflow)** menu by the workflow title, then choose **Remediation Status**.

Step 2 You have the following options:

- To learn more about the columns that appear, see Remediation Status Table Fields, on page 934.
- To sort and constrain the events, see Using Workflows, on page 656.

PART IX

Correlation and Compliance

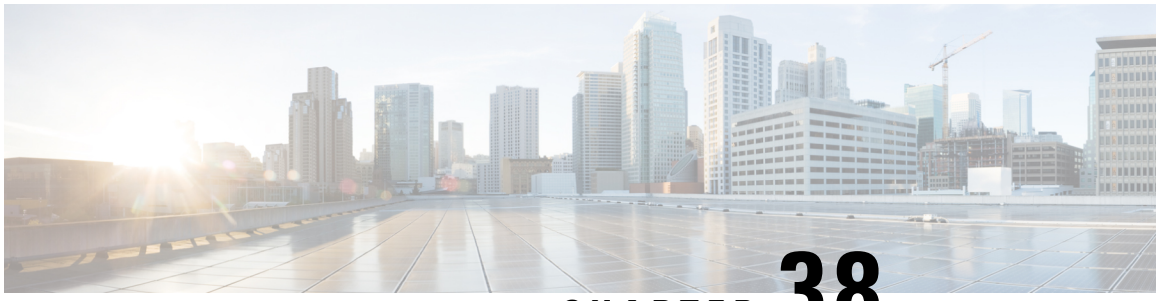

CHAPTER 38

Compliance Lists

The following topics describe how to configure compliance allow lists before you add them to correlation policies.

Introduction to Compliance Allow Lists

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.

A compliance allow list has two main components:

- *Targets* are the hosts you select for compliance evaluation. You can evaluate all or some monitored hosts, constraining by subnet, VLAN, and host attribute. In a multidomain deployment, you can target domains and subnets within or across domains.
- *Host profiles* specify the compliance criteria for the targets. The global host profile is operating system agnostic. You can also configure operating-system specific host profiles, either unique to one 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.

Example: 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

Configuring Correlation Policies, on page 955

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.

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.

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.

Note In addition to targeting IP addresses and ranges in leaf domains, you can also constrain a target network using a higher-level domain. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Compliance 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.

Table 129: 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

Host Profile Type	Appearance	Description
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 not only which operating systems are allowed to run on your network, but also the application protocols, clients, web applications, and protocols that are allowed to run on those operating systems.

For example, you could require that compliant hosts run a particular version of Microsoft Windows. As another example, you could allow SSH to run on Linux hosts on port 22, and further restrict the vendor and version of the SSH client.

Create one host profile for each operating system you want to allow on your network. To disallow an operating system on your network, do not create a host profile for that operating system. For example, to make sure that all the hosts on your network are running Windows, configure the allow list to only contain host profiles for that operating system.

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.

For example, 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.

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, which you can edit. It also displays shared host profiles from ancestor domains, which you cannot edit. To view and edit shared host profiles created in a lower domain, switch to that domain.

Note If you modify a shared host profile (including built-ins), or modify a built-in application protocol, protocol, or client, your change affects every allow list that uses it. If you make unintended changes to or delete these built-in elements, you can reset to factory defaults.

Allow 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 (for example, a port used by SMTP or web servers) active on a host, or a new UDP server running on a host
- detects a change in a discovered TCP or UDP server running on a host, for example, a version change due to an upgrade
- detects a new client or web application running on a host
- drops a client or web application from its database due to inactivity
- detects that a host is communicating with a new network or transport protocol
- detects a new jailbroken mobile device
- detects that a TCP or UDP port has closed or timed out on a host

In addition, you can trigger a compliance change for a host by using the host input feature or the host profile to:

- add a client, protocol, or server to a host
- delete a client, protocol, or server from a host
- set the operating system definition for a host
- change a host attribute for a host so that the host is no longer a valid target

Note To avoid overwhelming you with events, the system does not generate 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 the following 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

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Creating 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 **NewAllow List**.

Step 3 Optionally, enter the **IP Address** and **Netmask** for an initial target network. In a multidomain deployment, choose the **Domain** where the target network resides.

Tip
To survey the entire monitored network, use the default values of 0.0.0.0/0 and ::/0.

Note

After you choose a domain for the target network, you cannot change it. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add—To add the target network without a survey, click **Add**.
- Add and Survey Network—To add and survey the target network, click **Add and Survey Network**.
- Skip—To create an allow list without surveying your network, click **Skip**.

Step 5 Optionally, enter a new **Name** and **Description** for the allow list.

Step 6 Optionally, **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 Setting Target Networks for a Compliance Allow List, on page 945.

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 Building Allow List Host Profiles, on page 946.
- Edit Surveyed Profiles—To edit an existing operating system-specific host profile created by a network survey, click its name and proceed as described in Building Allow List Host Profiles, on page 946.
- 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 Building Allow List Host Profiles, on page 946.
- 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 **SaveAllow List**.

What to do next

- Add the allow list to an active correlation policy as described in Configuring Correlation Policies, on page 955. The system immediately starts evaluating the allow list and generating violations.

Related Topics

Compliance Allow List Target Networks, on page 940
Creating a Compliance Allow List Based on Selected Hosts, on page 889
IP Address Conventions, on page 25

Setting 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. Targeting a subnet in a higher-level domain targets the **same** subnet in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, using literal IP addresses to constrain this configuration can have unexpected results.

Step 4 Add the target network:

- Add — To add the target network without a survey, click **Add**.
- Add and Survey Network — To add and survey the target network, click **Add and Survey Network**.

Step 5 Optionally, click the new target to configure it further:

- Name — Enter a new **Name**.
- Add Networks — To target additional hosts, click **Add** (+), then enter the **IP Address** and **Netmask**. To exclude the network from 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 **SaveAllow List**.

Related Topics

Compliance Allow List Target Networks, on page 940
IP Address Conventions, on page 25

Building 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-ins), 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.

Before you begin

- Create or edit a host profile within an allow list as described in Editing a Compliance Allow List, on page 950, or create or edit a shared host profile as described in Managing Shared Host Profiles, on page 952.

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 Adding an Application Protocol to a Compliance Allow List, on page 948.
- Client — To allow a client, click **Add** (+) and proceed as described in Adding a Client to a Compliance Allow List, on page 948.
- Web Application — To allow a web application, click **Add** (+) and proceed as described in Adding a Web Application to a Compliance Allow List, on page 949.
- Protocol — To allow a protocol, click **Add** (+) and proceed as described in Adding a Protocol to a Compliance Allow List, on page 949.
- Delete — To disallow an item you previously allowed, click **Delete** ( ).
- Edit Properties — To edit the properties of an allowed application protocol, client, or protocol, click its name. The changes you make are reflected in every host profile that uses that element.

Tip
Select the appropriate **Allow all...** check box to allow all application protocols, clients, or web applications for hosts matching this profile.

Step 2 To immediately implement all changes made since the last time you saved, click **SaveAllow List** (or **Save All Profiles** if you are editing a shared host profile).

Adding 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 OpenSSH 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 lists application protocols that have been allowed or are currently allowed by the allow list.
- To allow an application protocol not in the list, select **<New Application Protocol>** and click **OK** to display the application protocol editor. Select the application protocol **Type** and **Protocol** you want to allow. Optionally, restrict the application protocol by **port**, **Vendor**, and **Version**.

  Note
  You must type the vendor and version exactly as they would appear in a table view of applications. If you do not specify a vendor or version, the 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 **SaveAllow List**.

Adding 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 While you are creating or modifying 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 You have two options:

- If the clients you want to allow are listed, select them. The web interface lists clients that have been allowed or are currently allowed by the 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 **SaveAllow List**.

Adding a Web Application to a Compliance Allow List

Using allow list host profiles, you can allow web applications either globally or on specific operating systems.

Procedure

Step 1 While you are creating or modifying a compliance 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 **SaveAllow List**.

Adding a Protocol to a Compliance Allow List

Using allow list host profiles, you can allow protocols either globally or on specific operating systems. ARP, IP, TCP, and UDP are always allowed to run on any host; you cannot disallow them.

Procedure

Step 1 While you are creating or modifying a compliance allow list host profile, click **Add** (+) next to **Allowed Protocols** (or next to **Globally Allowed Protocols** if you are modifying the global host profile).

Step 2 You have two options:

- If the protocols you want to allow are listed, select them. The web interface lists protocols that have been allowed or are currently allowed by the 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 **SaveAllow List**.

Managing 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 **NewAllow List** and proceed as described in Creating a Compliance Allow List, on page 944.
- 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 Editing a Compliance Allow List, on page 950. 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 Managing Shared Host Profiles, on page 952.

Editing 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:

- 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 Building Allow List Host Profiles, on page 946.
- 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.
- Add Target Network — To add a new target network without surveying its hosts, click **Add** ( ) next to Target Networks and proceed as described in Setting Target Networks for a Compliance Allow List, on page 945.
- 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 Building Allow List Host Profiles, on page 946.
- Edit Other Host Profile — To edit a shared or operating-system specific host profile, click the host profile's name and proceed as described in Building Allow List Host Profiles, on page 946.
- Edit Target Network — To edit a target network, click the network's name and proceed as directed in Setting Target Networks for a Compliance Allow List, on page 945.

Step 4 To immediately implement all changes made since the last time you saved, click **SaveAllow List**.

Managing 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-ins), 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 Building Allow List Host Profiles, on page 946.
- Create Shared Host Profile by Survey — To create multiple new shared host profiles by surveying a network, click **Add Target Network** and proceed as described in Setting Target Networks for a Compliance Allow List, on page 945 .
- Delete — To delete a shared host profile, click **Delete** ( ), then confirm your choice.
- Edit — To modify an existing shared host profile (including a built-in shared host profile), click its name and proceed as described in Building Allow List Host Profiles, on page 946.
- 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**.

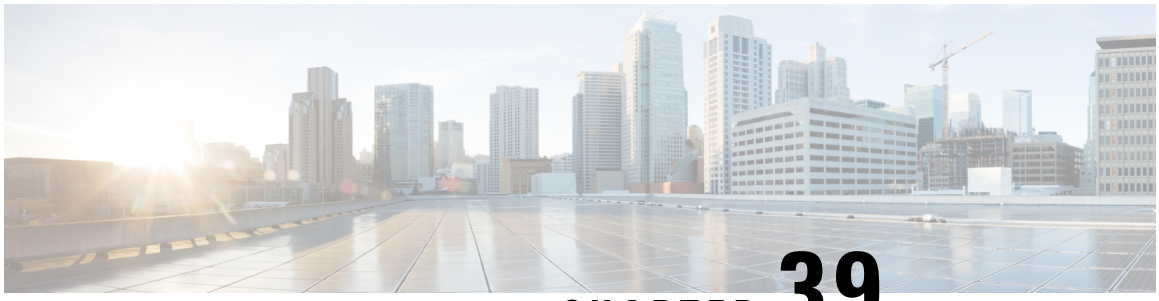

CHAPTER 39

Correlation Policies

The following topics describe how to configure correlation policies and rules.

Introduction to Correlation Policies and Rules

You can use the *correlation* feature to respond in real time to threats to your network, using *correlation policies*.

A correlation *policy violation* occurs when the activity on your network triggers either a *correlation rule* or *compliance allow list* within an active correlation policy.

Correlation Rules

When a correlation rule in an active correlation policy triggers, the system generates a *correlation event*. Correlation rules can trigger when:

- The system generates a specific type of event (connection, intrusion, malware, discovery, user activity, and so on).
- Your network traffic deviates from its normal profile.

You can constrain correlation rules in the following ways:

- Add a *host profile qualification* to constrain the rule using information from the host profile of a host involved in the triggering event.
- Add a *connection tracker* to a correlation rule so that after the rule's initial criteria are met, the system begins tracking certain connections. Then, a correlation event is generated only if the tracked connections meet additional criteria.
- Add a *user qualification* to a correlation rule to track certain users or groups of users. For example, you can constrain a correlation rule so that it triggers only for a particular user's traffic, or traffic from a specific department.

- Add *snooze periods*. When a correlation rule triggers, a snooze period causes that rule not to trigger again for a specified interval. After the snooze period elapses, the rule can trigger again and start a new snooze period.
- Add *inactive periods*. During inactive periods, correlation rules do not trigger.

Although you can configure correlation rules without licensing your deployment, rules that use unlicensed components do not trigger.

Compliance Allow Lists

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*.

Correlation Responses

Responses to correlation policy violations include simple alerts and various remediations (such as scanning a host). You can associate each correlation rule or 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.

Correlation and Multitenancy

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.

Related Topics

Introduction to Compliance Allow Lists, on page 939
Secure Firewall Management Center Alert Responses, on page 551
Introduction to Remediations, on page 1003

Requirements and Prerequisites for Compliance

Model Support

Any

Supported Domains

Any

User Roles

- Admin

Configuring Correlation Policies

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, and 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 Adding Responses to Rules and Allow Lists, on page 955.

Step 8 Click **Save**.

What to do next

- Activate the policy by clicking the slider.

Adding 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. Note that an Nmap remediation does not launch when used as a response to a traffic profile change.

In a multidomain deployment, you can use responses created in the current domain or in ancestor domains.

Procedure

Step 1 In the correlation policy editor, next to a rule or 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 Up Arrow ∧.

Step 3 Click **Update**.

Related Topics

Secure Firewall Management Center Alert Responses, on page 551
Introduction to Remediations, on page 1003

Managing Correlation Policies

Changes made to active correlation policies take effect immediately.

When you activate a correlation policy, the system immediately begins processing events and triggering responses. Note that the system does not generate 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 Configuring Correlation Policies, on page 955.
- Edit — Click **Edit** ( ); see Configuring Correlation Policies, on page 955. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Delete — Click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Configuring Correlation Rules

A simple correlation rule requires only that an event of a certain type occurs. You do not need to provide more specific conditions. For example, correlation rules based on traffic profile changes do not require conditions. You can also create complex correlation rules, with multiple conditions and added constraints.

When you create correlation rule trigger criteria, host profile qualifications, user qualifications, or connection trackers, the syntax varies but the mechanics remain consistent.

Note In a multidomain deployment, constraining a correlation rule by an ancestor domain matches events reported by that domain's descendants.

Before you begin

- Confirm that your deployment is collecting the type of information you want to use to trigger correlation events. For example, the information available for any individual connection or connection summary event depends on several factors, including the detection method, the logging method, and event type. The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see Differences between NetFlow and Managed Device Data.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Click **Create Rule**.

Step 3 Enter a **Rule Name** and **Rule Description**.

Step 4 Optionally, choose a **Rule Group** for the rule.

Step 5 Choose a base event type and, optionally, specify additional trigger criteria for the correlation rule. You can choose the following base event types:

- **a VPN troubleshooting event occurs**—See Syntax for VPN Troubleshoot Event Trigger Criteria, on page 958.
- **an intrusion event occurs**—See Syntax for Intrusion Event Trigger Criteria, on page 959.
- **a malware event occurs**—See Syntax for Malware Event Trigger Criteria, on page 961.
- **a discovery event occurs**—See Syntax for Discovery Event Trigger Criteria, on page 963.
- **user activity is detected** —See Syntax for User Activity Event Trigger Criteria, on page 966.
- **a host input event occurs**—See Syntax for Host Input Event Trigger Criteria, on page 966.
- **a connection event occurs**—See Syntax for Connection Event Trigger Criteria, on page 968.
- **a traffic profile changes**—See Syntax for Traffic Profile Changes, on page 971.

Step 6 Optionally, further constrain the correlation rule by adding any or all of the following:

- Host Profile Qualification—Click **Add Host Profile Qualification**; see Syntax for Correlation Host Profile Qualifications, on page 973.
- Connection Tracker—Click **Add Connection Tracker**; see Connection Trackers, on page 976.
- User Qualification—Click **Add User Qualification**; see Syntax for User Qualifications, on page 975.
- Snooze Period—Under Rule Options, use the **Snooze** text field and drop-down list to specify the interval that the system should wait to trigger a correlation rule again, after the rule triggers.
- Inactive Period—Under Rule Options, click **Add Inactive Period**. Using the text field and drop-down lists, specify when and how often you want the system to refrain from evaluating network traffic against the correlation rule.

Tip
To remove a snooze period, specify an interval of `0` (seconds, minutes, or hours).

Step 7 Click **Save Rule**.

Example Simple Correlation Rule

The following simple correlation rule triggers if a new host is detected in a specific subnet. Note that when the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a block of IP addresses, as expressed in special notation such as CIDR.

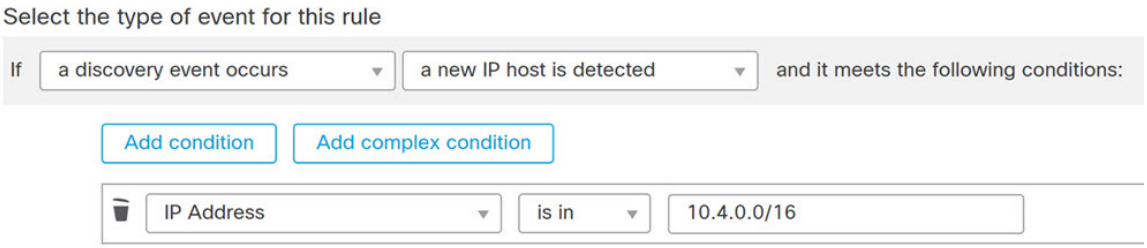

What to do next

- Use the rule in correlation policies as described in Configuring Correlation Policies, on page 955.

Related Topics

Managing Correlation Rules, on page 987
Correlation Rule Building Mechanics, on page 984
Snooze and Inactive Periods, on page 984
Differences between NetFlow and Managed Device Data

Syntax for VPN Troubleshoot Event Trigger Criteria

The following table describes how to build a correlation rule condition when you choose a VPN troubleshooting event as the base event.

Table 130: 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

The following table describes how to build a correlation rule condition when you choose an intrusion event as the base event.

Table 131: 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 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.

Related Topics

Intrusion Event Fields, on page 766
IP Address Conventions, on page 25

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)

The following table describes how to build a correlation rule condition when you choose a malware event as the base event.

Table 132: Syntax for Malware Events

If you specify...	Choose an operator, then...
Application Protocol	Choose one or more application protocols associated with the malware event.
Application Protocol Category	Choose one or more category of application protocol.
Client	Choose one or more clients associated with the malware event.
Client Category	Choose one or more category of client.
Destination Country or Source Country	Choose one or more countries associated with the source or destination IP address in the malware event.
Destination IP, Host IP, or Source IP	Enter a single IP address or address block.
Destination Port/ICMP Code	Enter the port number or ICMP code for destination traffic.
Disposition	Choose either or both **Malware** or **Custom Detection**.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the 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.

If you specify...	Choose an operator, then...
SSL Flow Status	Choose one or more statuses based on the result of the system's attempt to decrypt the traffic.
Source Port/ICMP Type	Enter the port number or ICMP type for source traffic.
Web Application	Choose one or more web applications associated with the malware event.
Web Application Category	Choose one or more category of web application.

Related Topics

File and Malware Event Fields, on page 815
IP Address Conventions, on page 25

Syntax for Discovery Event Trigger Criteria

To base a correlation rule on a discovery event, first specify the type of discovery event you want to use. Your choice determines the set of trigger criteria you can use. The following table lists the discovery event types you can choose.

You cannot trigger a correlation rule on hops changes, or when the system drops a new host due to reaching the host limit. You can, however, choose **there is any type of event** to trigger the rule when any type of discovery event occurs.

Table 133: 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

Choose this option...	To use this discovery event type...
a TCP port closed	TCP Port Closed
a TCP port timed out	TCP Port Timeout
a UDP port closed	UDP Port Closed
a UDP port timed out	UDP Port Timeout
a VLAN tag was updated	VLAN Tag Information Update
an IOC was set	Indication of Compromise
an open TCP port is detected	New TCP Port
an open UDP port is detected	New UDP Port
the OS information for a host has changed	New OS
the OS or server identity for a host has a conflict	Identity Conflict
the OS or server identity for a host has timed out	Identity Timeout
there is any kind of event	any event type
there is new information about a MAC address	MAC Information Change
there is new information about a TCP server	TCP Server Information Update
there is new information about a UDP server	UDP Server Information Update

The following table describes how to build a correlation rule condition when you choose a discovery event as the base event.

Table 134: Syntax for Discovery Events

If you specify...	Choose an operator, then...
Application Protocol	Choose one or more application protocols.
Application Protocol Category	Choose one or more category of application protocol.
Application Port	Enter the application protocol port number.
Client	Choose one or more clients.
Client Category	Choose one or more category of client.
Client Version	Enter the version number of the client.
Device	Choose one or more devices that may have generated the discovery event.

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 management center for multitenancy.
Hardware	Enter the hardware model for the mobile device. For example, to match all Apple iPhones, enter **`iPhone`**.
Host Type	Choose one or more host types. You can choose between a host or one of several types of network device.
IP Address or New IP Address	Enter a single IP address or address block.
Jailbroken	Choose **Yes** to indicate that the host in the event is a jailbroken mobile device or **No** to indicate that it is not.
MAC Address	Enter all or part of the MAC address of the host. For example, if you know that devices from a certain hardware manufacturer have MAC addresses that begin with 0A:12:34, you could choose **begins with** as the operator, then enter **`0A:12:34`** as the value.
MAC Type	Choose whether the MAC address was **ARP/DHCP Detected**. That is, choose whether the system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**), or whether the system is seeing many hosts with that MAC address because, for example, there is a router between the managed device and the host (**is not ARP/DHCP Detected**).
MAC Vendor	Enter all or part of the name of the MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NETBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
OS Name	Choose one or more operating system names.
OS Vendor	Choose one or more operating system vendors.
OS Version	Choose one or more operating system versions.
Protocol or Transport Protocol	Enter the name or number of the transport protocol as listed in http://www.iana.org/assignments/protocol-numbers.
Source	Choose the source of the host input data (for operating system and server identity changes and timeouts).
Source Type	Choose the type of the source for the host input data (for operating system and server identity changes and timeouts).
VLAN ID	Enter the VLAN ID of the host involved in the event.

If you specify...	Choose an operator, then...
Web Application	Choose a web application.

Related Topics

Discovery Event Types, on page 877
Discovery Event Fields, on page 883
IP Address Conventions, on page 25

Syntax for User Activity Event Trigger Criteria

To base a correlation rule on user activity, first choose the type of user activity you want to use. Your choice determines the set of trigger criteria you can use. You can choose:

- **a new user identity is detected**
- **a user logs into a host**

The following table describes how to build a correlation rule condition when you choose user activity as the base event.

Table 135: 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 management center for multitenancy.
IP Address	Enter a single IP address or address block.
Username	Enter a username.

Related Topics

User Activity Data Fields
IP Address Conventions, on page 25

Syntax for Host Input Event Trigger Criteria

To base a correlation rule on a host input event, first specify the type of host input event you want to use. Your choice determines the set of trigger criteria you can use. The following table lists the host input event types you can choose.

You cannot trigger a correlation rule when you add, delete, or change the definition of a user-defined host attribute, or set a vulnerability impact qualification.

Table 136: 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

Choose this option...	To trigger the rule on this event type...
a client is deleted	Delete Client
a host is added	Add Host
a protocol is added	Add Protocol
a protocol is deleted	Delete Protocol
a scan result is added	Add Scan Result
a server definition is set	Set Server Definition
a server is added	Add Port
a server is deleted	Delete Port
a vulnerability is marked invalid	Vulnerability Set Invalid
a vulnerability is marked valid	Vulnerability Set Valid
an address is deleted	Delete Host/Network
an attribute value is deleted	Host Attribute Delete Value
an attribute value is set	Host Attribute Set Value
an OS definition is set	Set Operating System Definition
host criticality is set	Set Host Criticality

The following table describes how to build a correlation rule condition when you choose a host input event as the base event.

Table 137: 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 management center for multitenancy.
IP Address	Enter a single IP address or address block.
Source	Choose the source for the host input data.
Source Type	Choose the type of the source for the host input data.

Related Topics

Host Input Event Types, on page 880
Discovery Event Fields, on page 883
IP Address Conventions, on page 25

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 138: 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.
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).

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 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.
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.

If you specify...	Choose an operator, then...
SSL Certificate Fingerprint	Enter the fingerprint of the certificate used to encrypt the traffic, or choose a subject common name associated with the fingerprint.
SSL Certificate 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.
URL Reputation	Choose one or more URL reputation values for the URL visited in the connection.
Username	Enter the username of the user logged in to either host in the connection.
Web Application	Choose one or more web applications associated with the connection.
Web Application Category	Choose one or more categories of web application.

Related Topics

Connection and Security-Related Connection Event Fields, on page 731

IP Address Conventions, on page 25

Syntax for Traffic Profile Changes

To base a correlation rule on a traffic profile change, first choose the traffic profile you want to use. The rule triggers when network traffic deviates from the pattern characterized by the profile you choose.

You can trigger the rule based on either raw data or on the statistics calculated from the data. For example, you could write a rule that triggers if the amount of data traversing your network (measured in bytes) suddenly spikes, which could indicate an attack or other security policy violation. You could specify that the rule trigger if either:

- the number of bytes traversing your network spikes above a certain number of bytes
- the number of bytes traversing your network spikes above a certain number of standard deviations above or below the mean amount of traffic

  Note that to create a rule that triggers when the number of bytes traversing your network falls outside a certain number of standard deviations (whether above or below), you must specify upper and lower bounds, as shown in the following graphic.

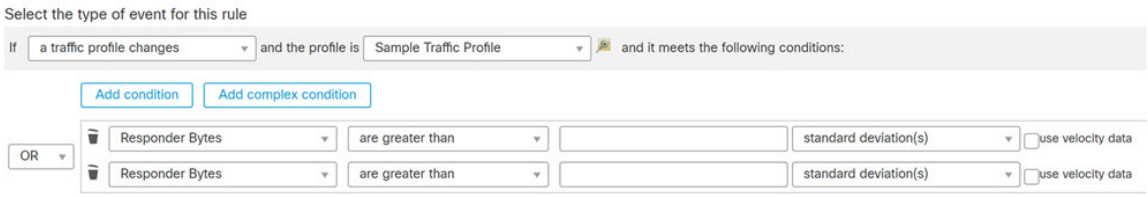

  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *above* the mean, use only the first condition shown in the graphic.

  To create a rule that triggers when the number of bytes traversing is greater than a certain number of standard deviations *below* the mean, use only the second condition.

Check the **use velocity data** check box to trigger the correlation rule based on rates of change between data points. If you wanted to use velocity data in the above example, you could specify that the rule triggers if either:

- the change in the number of bytes traversing your network spikes above or below a certain number of standard deviations above the mean rate of change
- the change in the number of bytes traversing your network spikes above a certain number of bytes

The following table describes how to build a condition in a correlation rule when you choose a traffic profile change as the base event.

Table 139: 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**.

The following table describes how to build a host profile qualification for a correlation rule.

Table 140: 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 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

The following host profile qualification constrains a correlation rule so the rule triggers only if the host involved in the discovery event on which the rule is based is running a version of Microsoft Windows.

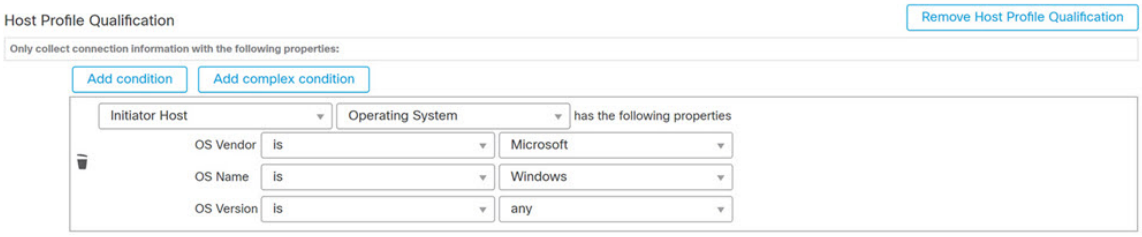

Related Topics

Host Data Fields, on page 885

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 management center-server connection established in an identity realm. This information may not be available for all users in the database.

When you build a user qualification, first specify the identity you want to use to constrain your correlation rule. The identity you can choose depends on the rule's base event type:

- connection event — Choose **Identity on Initiator** or **Identity on Responder**.
- intrusion event — Choose **Identity on Destination** or **Identity on Source**.
- discovery event — Choose **Identity on Host**.
- host input event — Choose **Identity on Host**.

The following table describes how to build a user qualification for a correlation rule.

Table 141: 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.

If you specify...	Choose an operator, then...
Department	Enter a department.
Domain	Choose one or more domains. In a multidomain deployment, constraining by an ancestor domain matches data reported by that domain's descendants. This field is only present if you have ever configured the management center for multitenancy.
Email	Enter an email address.
First Name	Enter a first name.
Last Name	Enter a last name.
Phone	Enter a telephone number.
Username	Enter a username.

Related Topics

User Data Fields

Connection Trackers

A *connection tracker* constrains a correlation rule so that after the rule's initial criteria are met (including host profile and user qualifications), the system begins tracking certain connections. The system generates a correlation event for the rule if the tracked connections meet additional criteria gathered over a time period that you specify.

Tip Connection trackers typically monitor very specific traffic and, when triggered, run only for a finite, specified time. Compare connection trackers with traffic profiles, which typically monitor a broad range of network traffic and run persistently.

There are two ways a connection tracker can generate an event.

Connection Trackers That Fire Immediately When Conditions Are Met

You can configure a connection tracker so that the correlation rule fires as soon as network traffic meets the tracker's conditions. When this happens, the system stops tracking connections for this connection tracker instance, even if the timeout period has not expired. If the same type of policy violation that triggered the correlation rule occurs again, the system creates a new connection tracker.

However, if time expires before network traffic meets the conditions in the connection tracker, the system does not generate a correlation event, and also stops tracking connections for that rule instance.

For example, a connection tracker can serve as a kind of event threshold by generating a correlation event only if a certain type of connection occurs more than a specific number of times within a specific time period. Or, you can generate a correlation event only if the system detects excessive data transfer after an initial connection.

Connection Trackers That Fire at the End of the Timeout Period

You can configure a connection tracker so that it relies on data collected over the entire timeout period, and therefore cannot fire until the end of the timeout period.

For example, if you configure a connection tracker to fire if you detect fewer than a certain number of bytes being transferred during a certain time period, the system waits until that time period passes and then generates an event if network traffic met that condition.

Adding a Connection Tracker

Before you begin

- Create a correlation rule based on a connection, intrusion, discovery, user identity, or host input event. You cannot add a connection tracker to a rule based on a VPN troubleshoot event, malware event or traffic profile change.

Procedure

Step 1 In the correlation rule editor (**Polices** > **Correlation** > **Rule Management**), click Edit, and then click **Add Connection Tracker**.

Step 2 Specify the connections to track; see Syntax for Connection Trackers, on page 977.

Step 3 Based on the tracked connections, specify when you want to generate a correlation event; see Syntax for Connection Tracker Events, on page 980.

Step 4 Specify the interval (in seconds, minutes, or hours) during which the tracker's conditions must be met.

Syntax for Connection Trackers

The following table describes how to build a connection tracker condition that specifies the kind of connections you want to track.

Table 142: Syntax for Connection Trackers

If you specify...	Choose an operator, then...
Access Control Policy	Choose one or more access control policies that handled the connections you want to track.
Access Control Rule Action	Choose one or more access control rule actions associated with the access control rule that logged the connections you want to track. Choose **Monitor** to track connections that match the conditions of any Monitor rule, regardless of the rule or default action that later handles the connections.
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.

If you specify...	Choose an operator, then...
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**)
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.

If you specify...	Choose an operator, then...
NETBIOS Name	Enter the NetBIOS name of the monitored host in the connection.
NetFlow Device	Choose the IP address of the NetFlow exporter you want to track. If you did not add any NetFlow exporters to the network discovery policy, the NetFlow Device drop-down list is blank.
Prefilter Policy	Choose one or more prefilter policies that handled the connections you want to track.
Reason	Choose one or more reasons associated with the connections you want to track.
Security Intelligence Category	Choose one or more Security Intelligence categories associated with the connections you want to track.
TCP Flags	Choose the TCP flag that connections must contain in order to track them. Only connections generated from exported NetFlow records contain TCP flag data.
Transport Protocol	Choose the transport protocol used by the connection.
URL	Enter all or part of the URL visited in the connections you want to track.
URL Category	Choose one or more URL categories for the URL visited in the connections you want to track.
URL Reputation	Choose one or more URL reputation values for the URL visited in the connections you want to track
Username	Enter the username of the user logged into either host in the connections you want to track.
Web Application	Choose one or more web applications.
Web Application Category	Choose one or more web application categories.

Using Event Data to Build a Connection Tracker

You can often use data from the correlation rule's base event when constructing a connection tracker.

For example, assume your correlation rule triggers when the system detects a new client. When you add a connection tracker to this type of correlation rule, the system automatically populates the tracker with constraints that refer to the base event:

- The **Initiator/Responder IP** is set to the **Event IP Address**.
- The **Client** is set to the **Event Client**.

Tip To track connections for a specific IP address or block of IP addresses, click **switch to manual entry** to manually specify the IP. Click **switch to event fields** to go back to using the IP address in the event.

Related Topics

Connection and Security-Related Connection Event Fields, on page 731
IP Address Conventions, on page 25

Syntax for Connection Tracker Events

The following table describes how to how to build a connection tracker condition that specifies when you want to generate a correlation event based on the connections you are tracking.

Table 143: Syntax for Connection Tracker Events

If you specify...	Choose an operator, then enter...
Number of Connections	the total number of connections detected
Number of SSL Encrypted Sessions	the total number of SSL- or TLS-encrypted sessions detected
Total Bytes, Initiator Bytes, or Responder Bytes	one of: • the total bytes transmitted (**Total Bytes**) • the number of bytes transmitted (**Initiator Bytes**) • the number of bytes received (**Responder Bytes**)
Total Packets, Initiator Packets, or Responder Packets	one of: • the total packets transmitted (**Total Packets**) • the number of packets transmitted (**Initiator Packets**) • the number of packets received (**Responder Packets**)
Unique Initiators or Unique Responders	one of: • the number of unique hosts that initiated sessions that were detected (**Unique Initiators**) • the number of unique hosts that responded to connections that were detected (**Unique Responders**)

Sample Configuration for Excessive Connections From External Hosts

Consider a scenario where you archive sensitive files on network 10.1.0.0/16, and where hosts outside this network typically do not initiate connections to hosts inside the network. An occasional connection initiated from outside the network might occur, but you have determined that when four or more connections are initiated within two minutes, there is cause for concern.

The rule shown in the following graphic specifies that when a connection occurs from outside the 10.1.0.0/16 network to inside the network, the system begins tracking connections that meet that criterion. The system then generates a correlation event if the system detects four connections (including the original connection) within two minutes that match that signature.

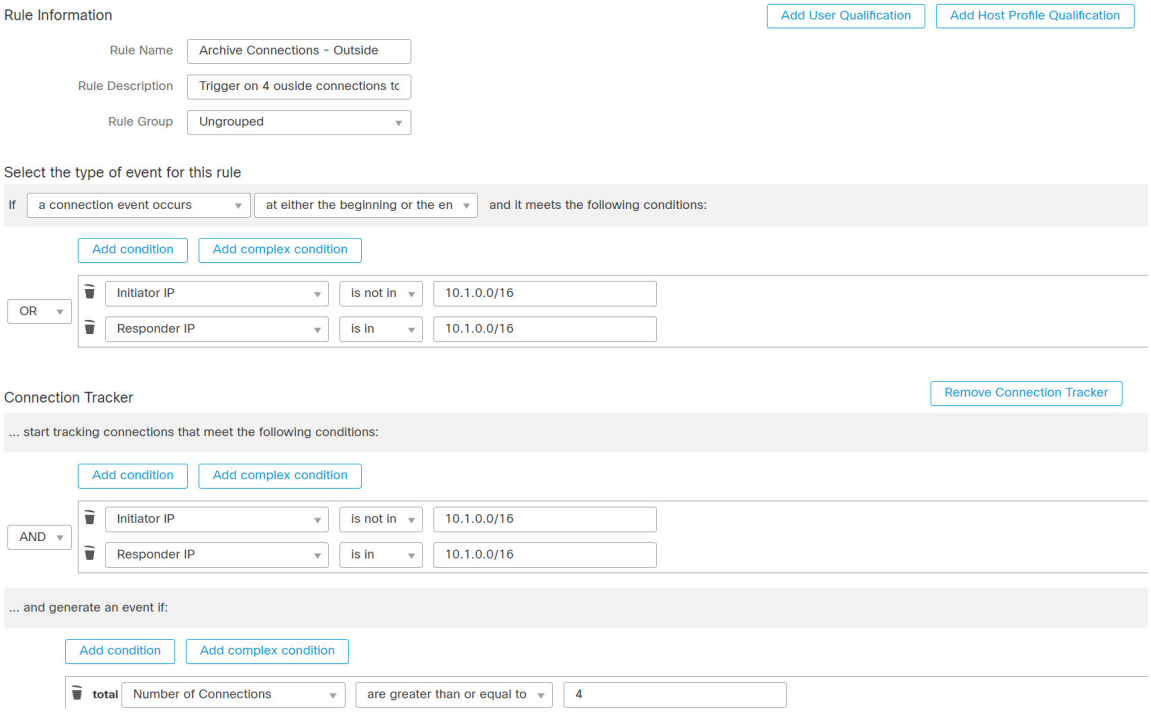

The following diagram shows how network traffic can trigger the above correlation rule.

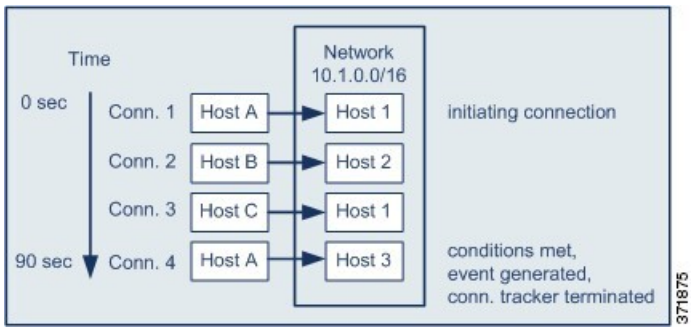

In this example, the system detected a connection that met the basic conditions of the correlation rule, that is, the system detected a connection from a host outside the 10.1.0.0/16 network to a host inside the network. This created a connection tracker.

The connection tracker is processed in the following stages:

- First, the system starts tracking connections when it detects a connection from Host A outside the network to Host 1 inside the network.
- The system detects two more connections that match the connection tracker signature: Host B to Host 2 and Host C to Host 1.
- The system detects a fourth qualifying connection when Host A connects to Host 3 within the two-minute time limit. The rule conditions are met.
- Finally, the system generates a correlation event and the system stops tracking connections.

Sample Configuration for Excessive BitTorrent Data Transfers

Consider a scenario where you want to generate a correlation event if the system detects excessive BitTorrent data transfers after an initial connection to any host on your monitored network.

The following graphic shows a correlation rule that triggers when the system detects the BitTorrent application protocol on your monitored network. The rule has a connection tracker that constrains the rule so that the rule triggers only if hosts on your monitored network (in this example, 10.1.0.0/16) collectively transfer more than 7MB of data (7340032 bytes) via BitTorrent in the five minutes following the initial policy violation.

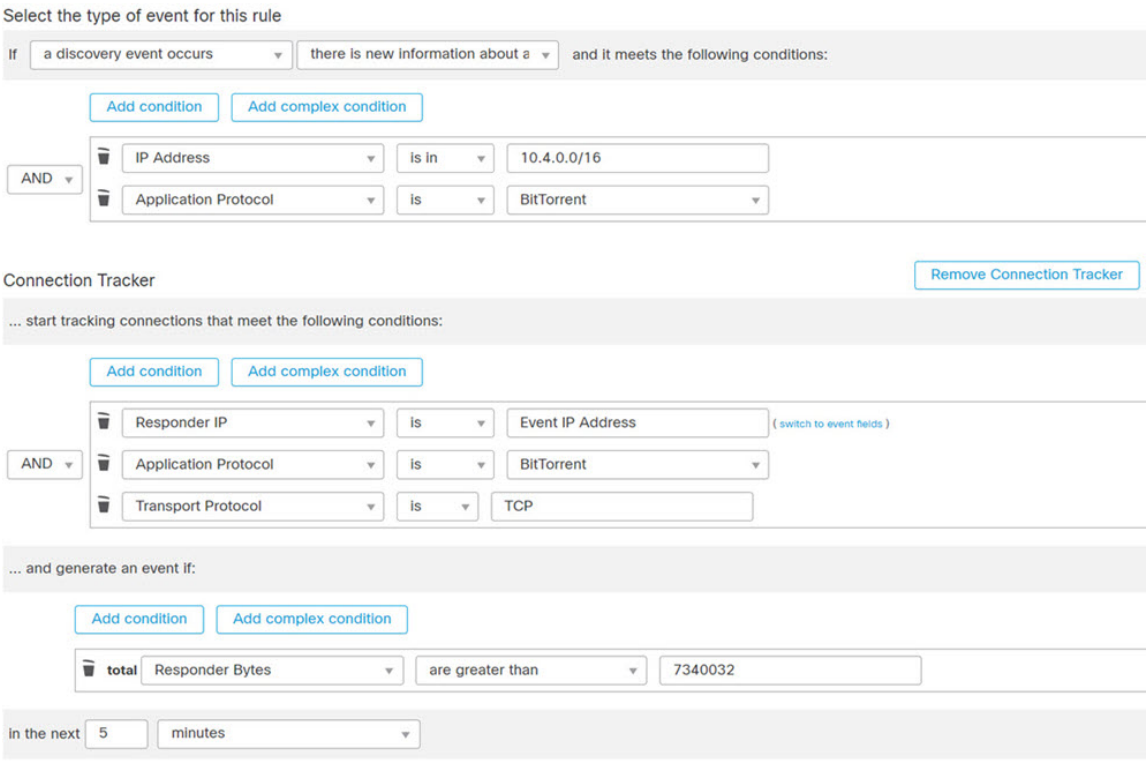

The following diagram shows how network traffic can trigger the above correlation rule.

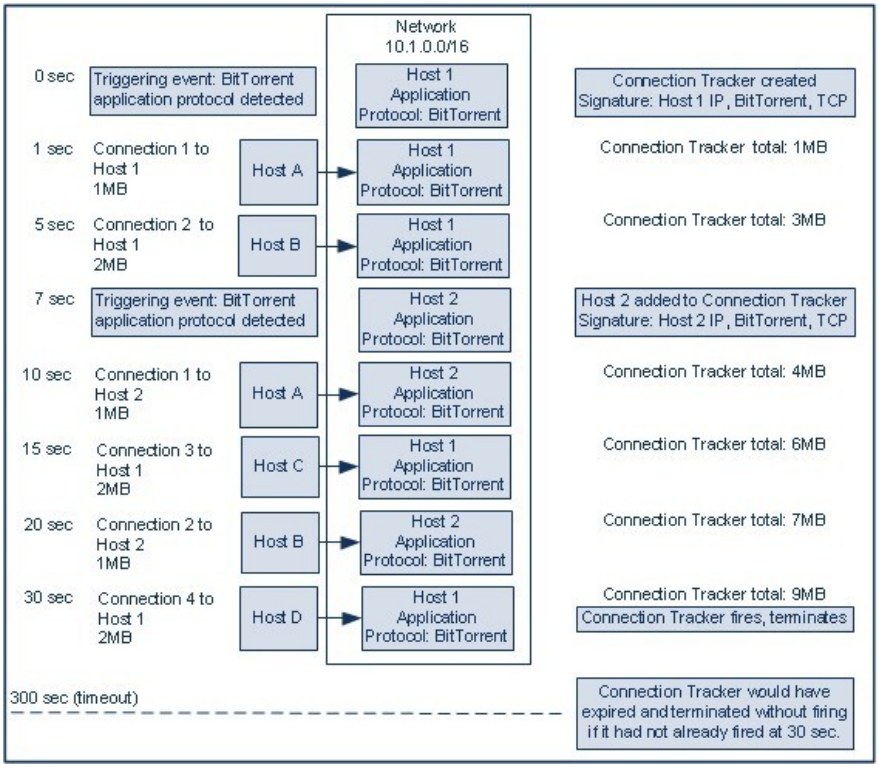

In this example, the system detected the BitTorrent TCP application protocol on two different hosts: Host 1 and Host 2. These two hosts transmitted data via BitTorrent to four other hosts: Host A, Host B, Host C, and Host D.

This connection tracker is processed in the following stages:

- First, the system starts tracking connections at the 0-second marker when the system detects the BitTorrent application protocol on Host 1. Note that the connection tracker will expire if the system does not detect 7MB of BitTorrent TCP data being transmitted in the next 5 minutes (by the 300-second marker).
- At 5 seconds, Host 1 has transmitted 3MB of data that matches the signature:
  - 1MB from Host 1 to Host A, at the 1-second marker (1MB total BitTorrent traffic counted towards fulfilling the connection tracker)
  - 2MB from Host 1 to Host B, at the 5-second marker (3MB total)
- At 7 seconds, the system detects the BitTorrent application protocol on Host 2 and starts tracking BitTorrent connections for that host as well.
- At 20 seconds, the system has detected additional data matching the signature being transmitted from both Host 1 and Host 2:
  - 1MB from Host 2 to Host A, at the 10-second marker (4MB total)
  - 2MB from Host 1 to Host C, at the 15-second marker (6MB total)
  - 1MB from Host 2 to Host B, at the 20-second marker (7MB total)

- Although Host 1 and Host 2 have now transmitted a combined 7MB of BitTorrent data, the rule does not trigger because the total number of bytes transmitted must be **more** than 7MB (**Responder Bytes are greater than 7340032**). At this point, if the system were to detect no additional BitTorrent transfers for the remaining 280 seconds in the tracker's timeout period, the tracker would expire and the system would not generate a correlation event.
- However, at 30 seconds, the system detects another BitTorrent transfer, and the rule conditions are met:
  - 2MB from Host 1 to Host D at the 30-second marker (9MB total)
- Finally, the system generates a correlation event. The system also stops tracking connections for this connection tracker instance, even though the 5-minute period has not expired. If the system detects a new connection using the BitTorrent TCP application protocol at this point, it will create a new connection tracker. Note that the system generates the correlation event *after* Host 1 transmits the entire 2MB to Host D, because it does not tally connection data until the session terminates.

Snooze and Inactive Periods

You can configure *snooze periods* in correlation rules. When a correlation rule triggers, a snooze period instructs the system to stop firing that rule for a specified interval, even if the rule is violated again during the interval. When the snooze period has elapsed, the rule can trigger again (and start a new snooze period).

For example, you may have a host on your network that should never generate traffic. A simple correlation rule that triggers whenever the system detects a connection involving that host may create multiple correlation events in a short period of time, depending on the network traffic to and from the host. To limit the number of correlation events exposing your policy violation, you can add a snooze period so that the system generates a correlation event only for the first connection (within a time period that you specify) that the system detects involving that host.

You can also set up inactive periods in correlation rules. During inactive periods, the correlation rule will not trigger. You can set up inactive periods to recur daily, weekly, or monthly. For example, you might perform a nightly Nmap scan on your internal network to look for host operating system changes. In that case, you could set a daily inactive period on the affected correlation rules for the time and duration of your scan so that those rules do not trigger erroneously.

Correlation Rule Building Mechanics

You build a correlation rule by specifying the conditions under which it triggers. The syntax you can use within conditions varies depending on the element you are creating, but the mechanics are the same.

Most conditions have three parts: a *category*, an *operator*, and a *value*:

- The categories you can choose depend on whether you are building correlation rule triggers, a host profile qualification, a connection tracker, or a user qualification. Within correlation rule triggers, the categories further depend on the base event type for the rule. Some conditions may contain several categories, each of which may have their own operators and values.
- A condition's available operators depend on the category.
- The syntax you can use to specify a condition's value depends on the category and operator. Sometimes you type the value in a text field. Other times, you can choose a value (or multiple values) from a drop-down list.

For example, if you want to generate a correlation event every time a new host is detected, you can create a simple rule with no conditions.

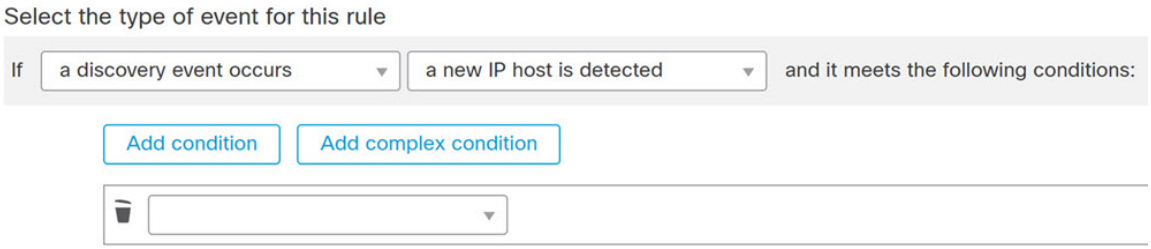

If you want to further constrain the rule and generate an event only if that new host was detected on the 10.4.x.x network, you can add a single condition.

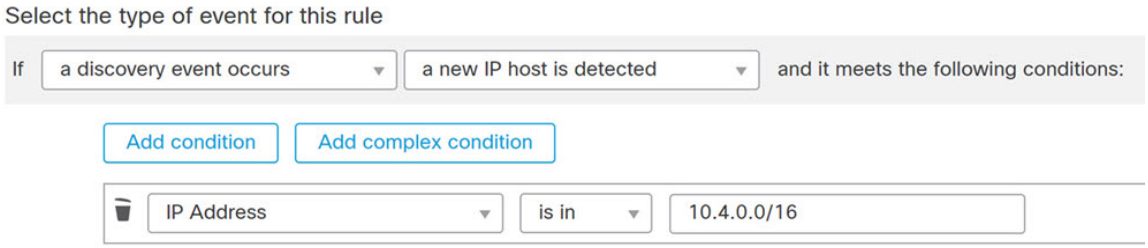

When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

The following rule, which detects SSH activity on a nonstandard port on the 10.4.x.x network and the 192.168.x.x network, has four conditions, with the bottom two constituting a complex condition.

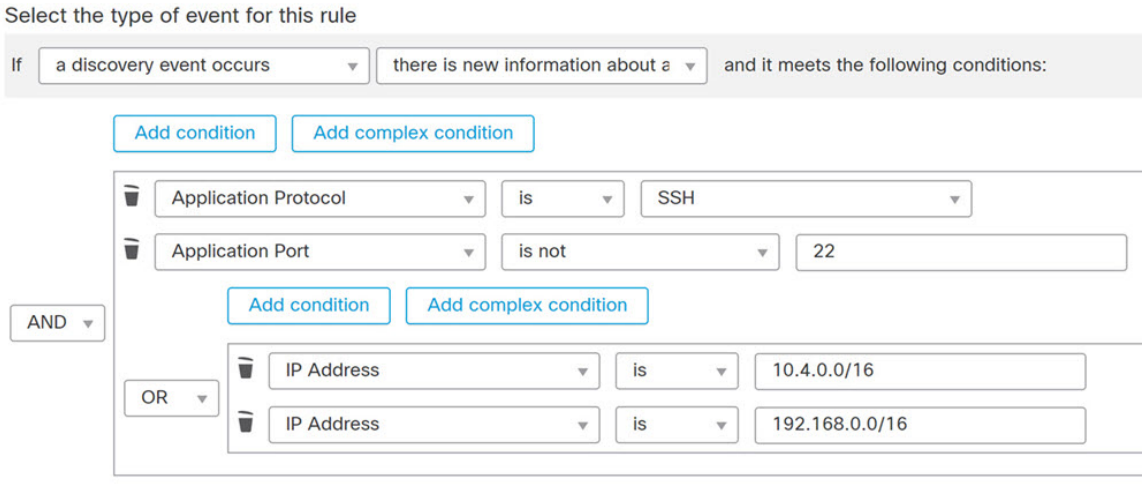

Logically, the rule is evaluated as follows:

```
(A and B and (C or D))
```

Table 144: Rule Evaluation

Where...	Is the condition that states...
A	Application Protocol is SSH
B	Application Port is not 22
C	IP Address is in 10.0.0.0/8
D	IP Address is in 196.168.0.0/16

Caution Evaluating complex correlation rules that trigger on frequently occurring events can degrade system performance. For example, a multicondition rule that the system must evaluate against every logged connection can cause resource overload.

Adding and Linking Conditions in Correlation Rules

Procedure

Step 1 In the correlation rule editor (**Polices** > **Correlation** > **Rule Management**), add a simple or complex condition:

- Simple — Click **Add condition**.
- Complex — Click **Add complex condition**.

Step 2 Link conditions by choosing the **AND** or **OR** operator from the drop-down list to the left of the conditions.

Example: Simple vs Complex Conditions

The following graphic shows a correlation rule with two simple conditions joined by the **OR** operator.

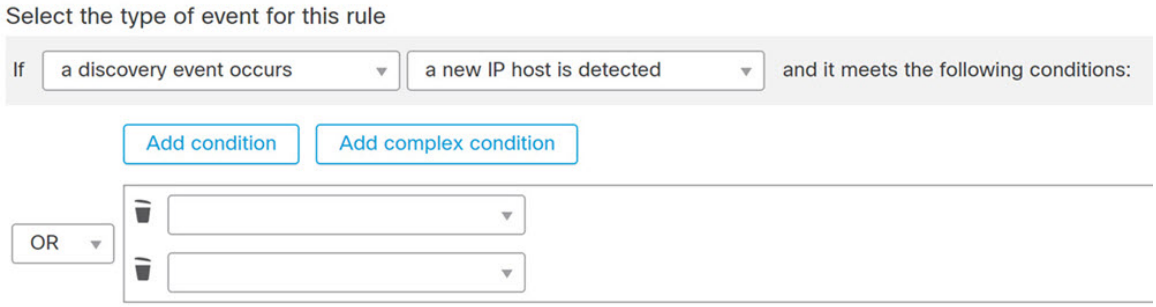

The following graphic shows a correlation rule with one simple condition and one complex condition, joined by the **OR** operator. The complex condition comprises two simple conditions joined by the **AND** operator.

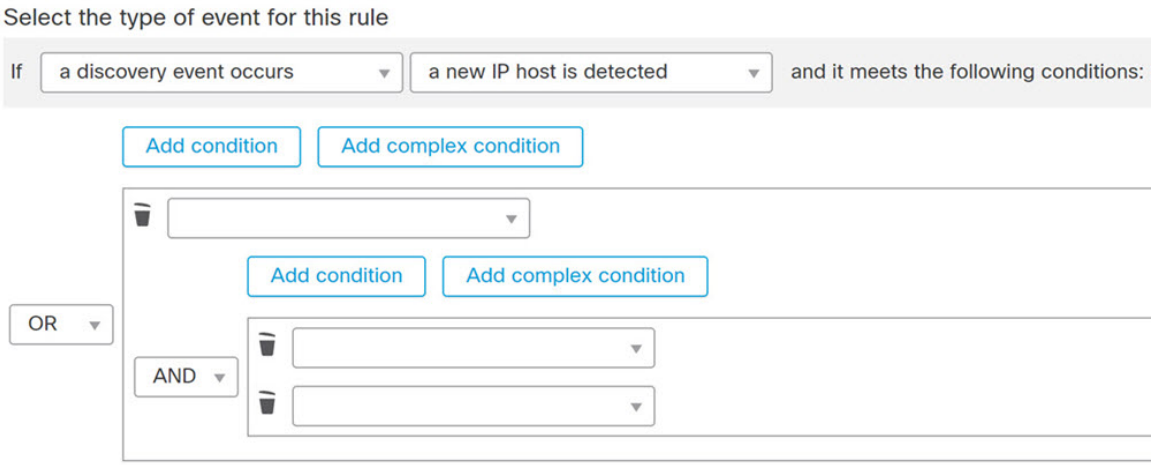

Using Multiple Values in Correlation Rule Conditions

When you are building a correlation condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list.

Procedure

Step 1 In the correlation rule editor, build a condition, choosing **is in** or **is not in** as the operator.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values. You can also click and drag to choose multiple adjacent values.

Step 4 Click the (Right Arrow >) to move the selected entries to **Selected**.

Step 5 Click **OK**.

Managing Correlation Rules

In a multidomain deployment, the system displays correlation rules and groups created in the current domain, which you can edit. It also displays selected correlation rules and groups from ancestor domains, which you cannot edit. To view and edit correlation rules and groups created in a lower domain, switch to that domain.

Note The system does not display configurations from ancestor domains if the configurations expose information about unrelated domains, including names, managed devices, and so on.

Changes made to rules in active correlation policies take effect immediately.

Before you begin

- If you want to delete a rule, delete it from all correlation policies, as described in Managing Correlation Policies, on page 956.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Rule Management**.

Step 2 Manage your rules:

- Create — Click **Create Rule**; see Configuring Correlation Rules, on page 957.
- Create Group — Click **Create Group**, enter a name for the group, and click **Save**. To add a rule to a group, edit the rule.
- Edit — Click **Edit** ( ); see Configuring Correlation Rules, on page 957. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete Rule or Rule Group— Click **Delete** ( ). Deleting a rule group ungroups the rules. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Configuring Correlation Response Groups

You can create a *correlation response group* of alerts and remediations, then activate and assign the group to a correlation rule within an active correlation policy. The system launches all the grouped responses when network traffic matches the correlation rule.

When used in an active correlation policy, changes to an active group or any of its grouped responses take affect immediately.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Click **Create Group**.

Step 3 Enter a **Name**.

Step 4 Check the **Active** check box if you want to activate the group upon creation.

Deactivated groups do not launch responses.

Step 5 Choose the **Available Responses** to group. then click the Right Arrow to move them to the **Responses in Group**. To move responses the other way, use the Left Arrow .

Step 6 Click **Save**.

What to do next

- If you did not activate the group upon creation and you want to activate it now, click the slider.

Related Topics

Secure Firewall Management Center Alert Responses, on page 551
Introduction to Remediations, on page 1003

Managing Correlation Response Groups

You can delete a response group if it is not used in a correlation policy. Deleting a response group ungroups its responses. You can also temporarily deactivate a response group without deleting it. This leaves the group on the system but does not launch it when policies are violated.

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

Changes made to active, in-use response groups take effect immediately.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Groups**.

Step 2 Manage response groups:

- Activate or Deactivate — Click the slider. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Create — Click **Create Group**; see Configuring Correlation Response Groups, on page 988.
- Edit — Click **Edit** ( ); see Configuring Correlation Response Groups, on page 988. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — Click **Delete** ( ). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

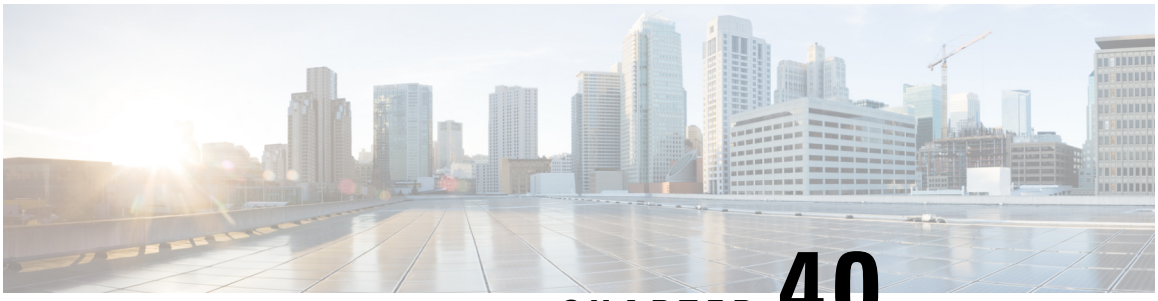

CHAPTER 40

Traffic Profiling

The following topics describe how to configure traffic profiles:

Introduction to Traffic Profiles

A *traffic profile* is a graph of network traffic based on connection data collected over a profiling time window (PTW). This measurement presumably represents normal network traffic. After the learning period, you can detect abnormal network traffic by evaluating new traffic against your profile.

The default PTW is one week, but you can change it to be as short as an hour or as long as several weeks. By default, traffic profiles generate statistics on connection events generated by the system over five-minute intervals. However, you can increase this sampling rate to as long as an hour.

Tip Cisco recommends that the PTW include at least 100 data points. Configure your PTW and sampling rate so that your traffic profiles contain enough data to be statistically meaningful.

The following graphic shows a traffic profile with a PTW of one day and a sampling rate of five minutes.

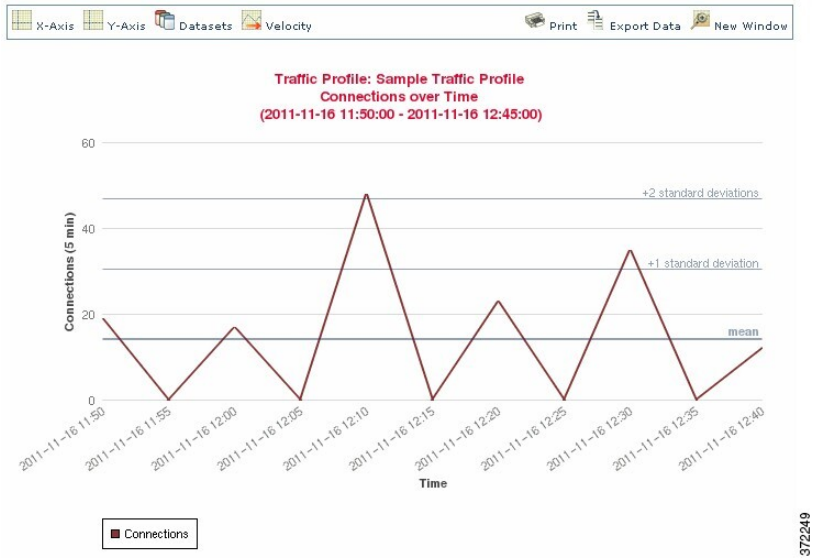

You can also set up inactive periods in traffic profile. Traffic profiles collect data during inactive periods, but do not use that data when calculating profile statistics. Traffic profile graphs plotted over time show inactive periods as a shaded region.

For example, consider a network infrastructure where all the workstations are backed up at midnight every night. The backup takes about 30 minutes and spikes the network traffic. You could configure recurring inactive period for your traffic profile to coincide with the scheduled backups.

Note The system uses end-of-connection data to create connection graphs and traffic profiles. To use traffic profiles, make sure you log end-of-connection events to the management center database.

Implementing Traffic Profiles

When you activate a traffic profile, the system collects and evaluates connection data for the learning period (PTW) you configured. After the learning period, the system evaluates correlation rules written against the traffic profile.

For example, you could write a rule that triggers if the amount of data traversing your network (measured in packets, KBytes, or number of connections) suddenly spikes to three standard deviations above the mean amount of traffic, which could indicate an attack or other security policy violation. Then, you could include that rule in a correlation policy to alert you of the traffic spike or to perform a remediation in response.

Targeting Traffic Profiles

Profile conditions and *host profile qualifications* constrain traffic profiles.

Using profile conditions, you can profile all network traffic, or you can restrict the traffic profile to monitoring a domain, subnets within or across domains, or individual hosts. In a multidomain deployment:

- Leaf-domain administrators can profile network traffic within their leaf domains.
- Higher-level domain administrators can profile traffic within or across domains.

Profile conditions can also constrain traffic profiles using criteria based on connection data. For example, you could set the profile conditions so that the traffic profile only profiles sessions using a specific port, protocol, or application.

Finally, you can also constrain traffic profiles using information about the tracked hosts. This constraint is called a *host profile qualification*. For example, you could collect connection data only for hosts with high criticality.

Note Constraining a traffic profile to a higher-level domain aggregates and profiles the **same** type of traffic in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, profiling traffic across domains can have unexpected results.

Related Topics

Introduction to Correlation Policies and Rules, on page 953

Traffic Profile Conditions

You can create simple traffic profile conditions and host profile qualifications, or you can create more elaborate constructs by combining and nesting conditions.

Conditions have three parts: a category, an operator, and a value:

- The categories you can use depend on whether you are building traffic profile conditions or a host profile qualification.
- The operators you can use depend on the category you choose.
- The syntax you can use to specify a condition's value depends on the category and operator. Sometimes you must enter the value in a text field. Other times, you can pick one or more values from a drop-down list.

For a host profile qualification, you must also specify whether you are constraining the traffic profile using information data about the initiating or responding hosts.

When your construct includes more than one condition, you must link them with an **AND** or an **OR** operator. Conditions on the same level are evaluated together:

- The **AND** operator requires that all conditions on the level it controls must be met.
- The **OR** operator requires that at least one of the conditions on the level it controls must be met.

Unconstrained Traffic Profile

If you want to create a traffic profile that collects data for your entire monitored network segment, you can create a very simple profile with no conditions, as shown in the following graphic.

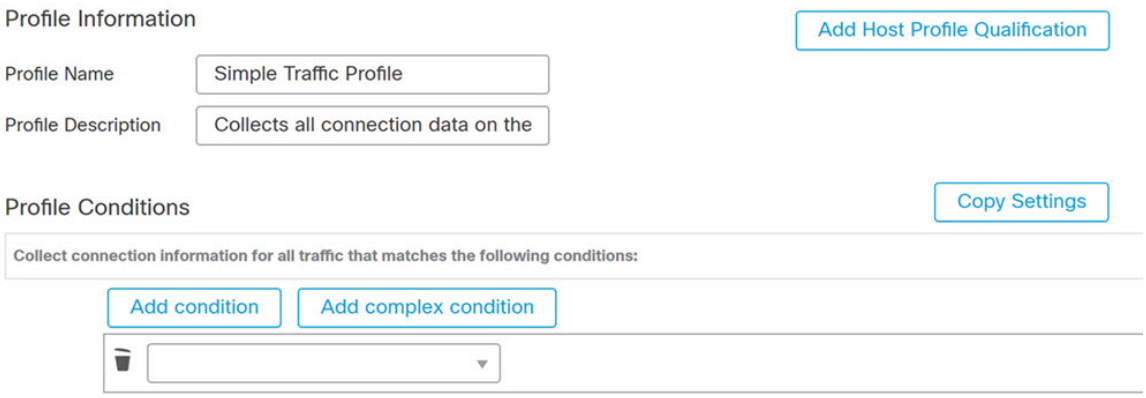

Simple Traffic Profile

If you wanted to constrain the profile and collect data only for a subnet, you can add a single condition, as shown in the following graphic.

Complex Traffic Profile

The following traffic profile contains two conditions linked by **AND**. This means that the traffic profile collects connection data only if both conditions are true. In this example, it collects HTTP connections for all hosts with IP addresses in a specific subnet.

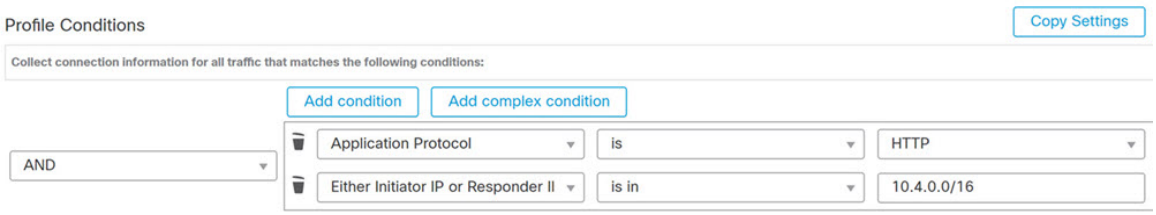

In contrast, the following traffic profile, which collects connection data for HTTP activity in either of two subnets, has three conditions, with the last constituting a complex condition.

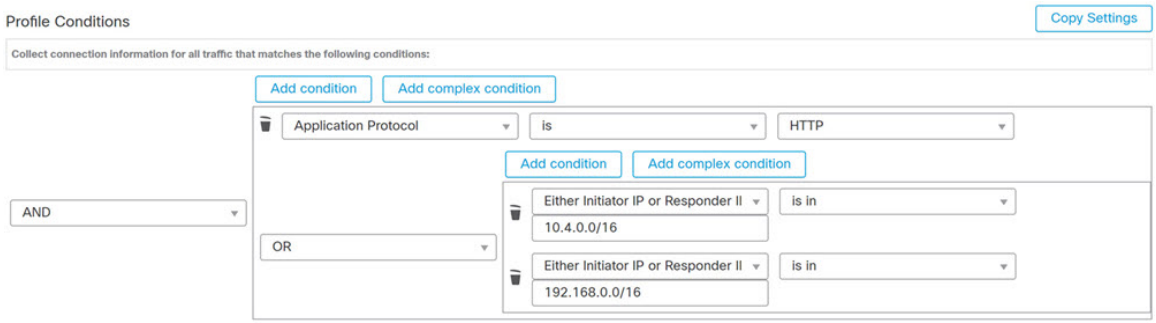

Logically, the above traffic profile is evaluated as follows:

```
(A and (B or C))
```

Where...	Is the condition that states...
A	Application Protocol Name is HTTP
B	IP Address is in 10.4.0.0/16
C	IP Address is in 192.168.0.0/16

Requirements and Prerequisites for Traffic Profiles

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Managing Traffic Profiles

Only rules written against active, complete traffic profiles can trigger a correlation policy violation. A slider next to each traffic profile indicates whether the profile is active and collecting data. A progress bar shows the status of the traffic profile's learning period.

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

Note The system does not display traffic profiles from ancestor domains if the profiles' conditions expose information about unrelated domains, including names, managed devices, and so on.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Manage your traffic profiles:

- Activate/Deactivate — To activate or deactivate a traffic profile, click the slider. Deactivating a traffic profile deletes its associated data. If you reactivate the profile, you must wait the length of its PTW before rules written against it will trigger.
- Create — To create a new traffic profile, click **New Profile** and proceed as described in Configuring Traffic Profiles, on page 996. You can also click **Copy** ( ) to edit a copy of an existing traffic profile.
- Delete — To delete a traffic profile, click **Delete** ( ), then confirm your choice.
- Edit — To modify an existing traffic profile, click **Edit** ( ) and proceed as described in Configuring Traffic Profiles, on page 996. If a traffic profile is active you can only change its name and description.
- Graph — To view the traffic profile as a graph, click **Graph** ( ). In a multidomain deployment, you cannot view the graph for a traffic profile that belongs to an ancestor domain if the graph exposes information about unrelated domains.

Configuring Traffic Profiles

Constraining a traffic profile to a higher-level domain aggregates and profiles the **same** type of traffic in **each** of the descendant leaf domains. The system builds a separate network map for each leaf domain. In a multidomain deployment, profiling traffic across domains can have unexpected results.

Procedure

Step 1 Choose **Policies** > **Correlation**, then click **Traffic Profiles**.

Step 2 Click **New Profile**.

Step 3 Enter a **Profile Name**, and optionally, a **Profile Description**.

Step 4 Optionally, constrain the traffic profile:

- Copy Settings — To copy settings from an existing traffic profile, click **Copy Settings**, choose the traffic profile you want to use, and click **Load**.
- Profile Conditions — To constrain the traffic profile using information from tracked connections, proceed as described in Adding Traffic Profile Conditions, on page 997.
- Host Profile Qualification — To constrain the traffic profile using information from tracked hosts, proceed as described in Adding Host Profile Qualifications to a Traffic Profile, on page 998.

- Profiling Time Window (PTW) — To change the **Profiling Time Window**, enter a time unit, then choose **hour(s)**, **day(s)**, or **week(s)**.
- Sampling Rate — Choose a **Sampling Rate**, in minutes.
- Inactive Period — Click **Add Inactive Period** and use the drop-down lists to specify when and how often you want the traffic profile remain inactive. Inactive traffic profiles do not trigger correlation rules. Traffic profiles do not include data from inactive periods in profile statistics.

Step 5 Save the traffic profile:

- To save the profile and start collecting data immediately, click **Save & Activate**.
- To save the profile without activating it, click **Save**.

Adding Traffic Profile Conditions

Procedure

Step 1 In the traffic profile editor, under Profile Conditions, click **Add condition** or **Add complex condition** for each condition you want to add. Conditions on the same level are evaluated together.

- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 2 Specify a category, operator, and value for each condition as described in Syntax for Traffic Profile Conditions, on page 998 and Traffic Profile Conditions, on page 993.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Using Multiple Values in a Traffic Profile Condition, on page 1001.

When the category represents an IP address, choosing **is in** or **is not in** as the operator allows you to specify whether the IP address *is in* or *is not in* a range of IP addresses.

Example

The following traffic profile collects information on a specific subnet. The category of the condition is **Initiator/Responder IP**, the operator is **is in**, and the value is `10.4.0.0/16`.

Related Topics

IP Address Conventions, on page 25

Adding Host Profile Qualifications to a Traffic Profile

Procedure

Step 1 In the traffic profile editor, click **Add Host Profile Qualification**.

Step 2 Under Host Profile Qualification, click **Add condition** or **Add complex condition** for each condition you want to add. Conditions on the same level are evaluated together.

- To require that all conditions on the level that the operator controls are met, choose **AND**.
- To require that only one of the conditions on the level that the operator controls is met, choose **OR**.

Step 3 Specify a host type, category, operator, and value for each condition as described in Syntax for Host Profile Qualifications in a Traffic Profile, on page 999and Traffic Profile Conditions, on page 993.

If you choose **is in** or **is not in** as the operator, you can select multiple values in a single condition as described in Using Multiple Values in a Traffic Profile Condition, on page 1001.

Example

The following host profile qualification constrains a traffic profile such that it collects connection data only if the responding host in the detected connection is running a version of Microsoft Windows.

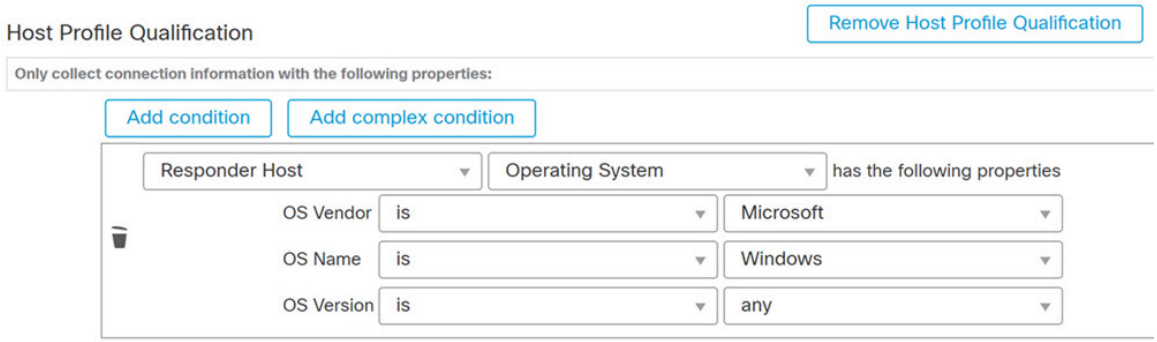

Syntax for Traffic Profile Conditions

The following table describes how to build a traffic profile condition. Keep in mind the connection data available to build a traffic profile depends on several factors, including traffic characteristics and detection method.

Table 145: Syntax for Traffic Profile Conditions

If you choose...	Choose an operator, then...
Application Protocol	Choose one or more application protocols.

If you choose...	Choose an operator, then...
Application Protocol Category	Choose one or more application protocol categories.
Client	Choose one or more clients.
Client Category	Choose one or more client categories.
Connection Type	Choose whether the profile uses connection data from traffic monitored by 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 749
IP Address Conventions, on page 25

Syntax for Host Profile Qualifications in a Traffic Profile

When you build a host profile qualification condition, you must first choose the host you want to use to constrain your traffic profile. You can choose either **Responder Host** or **Initiator Host**. After you choose the host role, continue building your host profile qualification condition.

Although you can add hosts to the network map using NetFlow records, the available information about these hosts is limited. For example, there is no operating system data available for these hosts, unless you provide it using the host input feature. In addition, if your traffic profile uses connection data from exported NetFlow

records, keep in mind that NetFlow records do not contain information about which host in the connection is the initiator and which is the responder. When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known.

To match against *implied* or generic clients, create a host profile qualification based on the application protocol used by the server responding to the client. When the client list on a host that acts as the initiator or source of a connection includes an application protocol name followed by **client**, that client may actually be an implied client. In other words, the system reports that client based on server response traffic that uses the application protocol for that client, not on detected client traffic.

For example, if the system reports **HTTPS client** as a client on a host, create a host profile qualification for **Responder Host** where **Application Protocol** is set to **HTTPS**, because HTTPS client is reported as a generic client based on the HTTPS server response traffic sent by the responder or destination host.

Table 146: 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.

If you choose...	Choose an operator, then...
MAC Address > MAC Type	Choose whether the MAC type is **ARP/DHCP Detected**, that is, whether: • The system positively identified the MAC address as belonging to the host (**is ARP/DHCP Detected**) • The system is seeing many hosts with that MAC address because, for example, there is a router between the device and the host (**is not ARP/DHCP Detected**) • The MAC type is irrelevant (**is any**)
MAC Vendor	Enter all or part of the MAC vendor of hardware used by the host.
Mobile	Choose **Yes** to indicate that the host in the event is a mobile device or **No** to indicate that it is not.
NETBIOS Name	Enter the NetBIOS name of the host.
Network Protocol	Enter the network protocol number as listed in http://www.iana.org/assignments/ethernet-numbers.
Operating System > OS Vendor	Choose one or more operating system vendor names.
Operating System > OS Name	Choose one or more operating system names.
Operating System > OS Version	Choose one or more operating system versions.
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.

Using Multiple Values in a Traffic Profile Condition

When you are building a condition, and the condition syntax allows you to pick a value from a drop-down list, you can often use multiple values from the list.

For example, if you want to add a host profile qualification to a traffic profile that requires that a host be running some flavor of UNIX, instead of constructing multiple conditions linked with the OR operator, use the following procedure.

Procedure

Step 1 While building a traffic profile or host profile qualification condition, choose **is in** or **is not in** as the operator. The drop-down list changes to a text field.

Step 2 Click anywhere in the text field or on the **Edit** link.

Step 3 Under **Available**, choose multiple values.

Step 4 Click the right arrow to move the selected entries to **Selected**.

Step 5 Click **OK**.

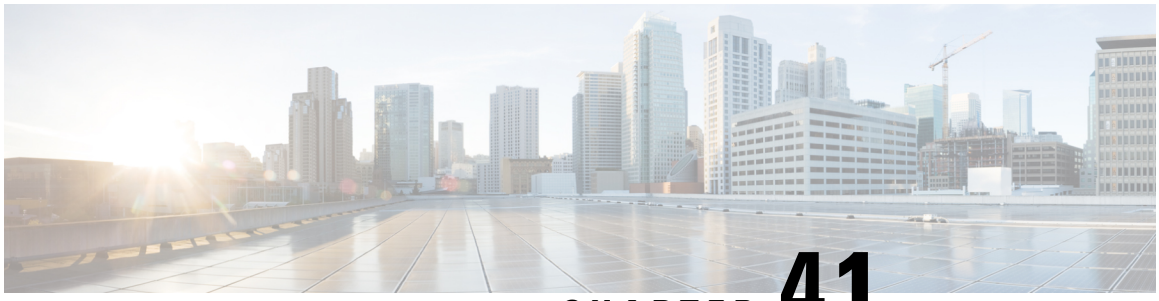

CHAPTER 41

Remediations

The following topics contain information on configuring remediations:

Requirements and Prerequisites for Remediations

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Introduction to Remediations

A *remediation* is a program that the system launches in response to a correlation policy violation.

When a remediation runs, the system generates a *remediation status event*. Remediation status events include details such as the remediation name, the correlation policy and rule that triggered it, and the exit status message.

The system supports several remediation modules:

- Cisco ISE Adaptive Network Control (ANC) — applies or clears ISE-configured ANC policies involved in a correlation policy violation

- Cisco IOS Null Route — blocks traffic sent to a host or network involved in a correlation policy violation (requires Cisco IOS Version 12.0 or higher)
- Nmap Scanning — scans hosts to determine running operating systems and servers
- Set Attribute Value — sets a host attribute on a host involved in a correlation policy violation

Tip You can install custom modules that perform other tasks; see the *Firepower System Remediation API Guide*.

Implementing Remediations

To implement a remediation, first create at least one *instance* for the module you choose. You can create multiple instances per module, where each instance is configured differently. For example, to communicate with multiple routers using the Cisco IOS Null Route remediation module, configure multiples instances of that module.

You can then add multiple *remediations* to each instance that describe the actions you want to perform when a policy is violated.

Finally, associate remediations with rules in correlation policies, so that the system launches the remediations in response to correlation policy violations.

Remediations and Multitenancy

In a multidomain deployment, you can install custom remediation modules at any domain level. The system-provided modules belong to the Global domain.

Though you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Related Topics

Secure Firewall Management Center Alert Responses, on page 551
Nmap Scanning
Adding Responses to Rules and Allow Lists, on page 955

Cisco ISE EPS Remediations

If you have Endpoint Protection Service (EPS) enabled and configured in your ISE deployment, you can configure your management center to launch remediations using ISE. When fully configured, ISE EPS remediations run the following **Mitigation Actions** on the source or destination host involved in a correlation policy violation:

- **`quarantine`**—Limits or denies an endpoint's access the network
- **`unquarantine`**—Reverses an endpoint's quarantine status and allows full access to the network
- **`shutdown`**—Deactivates an endpoint's network attached system (NAS) port to disconnect it from the network

You can also exempt specific IP addresses from ISE EPS remediation.

Note Your ISE version and configuration impact how you can use ISE. 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*.

Configuring 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 Adding an ISE EPS Instance, on page 1005.

Step 3 Add one or more ISE EPS remediations as described in Adding ISE EPS Remediations, on page 1006.

What to do next

- Assign remediations as responses to correlation policy violations as described in Adding Responses to Rules and Allow Lists, on page 955.

Adding an ISE EPS Instance

Create ISE EPS instances to group individual remediations by logging type.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **pxGrid Mitigation(v1.0)** as the module type and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Set **Enable Logging** option to enable or disable system logging.

Step 5 Click **Create**.

What to do next

- Create an ISE EPS remediation as described in Adding Set Attribute Value Remediations, on page 1013.

Related Topics

IP Address Conventions, on page 25

Adding ISE EPS Remediations

Create one or more ISE EPS remediations within an instance to run **Mitigation Actions** on the source or destination host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Create an ISE EPS instance as described in Adding an ISE EPS Instance, on page 1005.

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 Adding Responses to Rules and Allow Lists, on page 955.

Cisco IOS Null Route Remediations

The Cisco IOS Null Route remediation module allows you to block an IP address or range of addresses using Cisco's "null route" command. This drops all traffic sent to a host or network by routing it to the router's NULL interface. This does not block traffic sent from the violating host or network.

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Configuring Remediations for Cisco IOS Routers

Note Do not use a destination-based remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

Caution When a Cisco IOS remediation is activated, there is no timeout period. To unblock the IP address or network, you must manually clear the routing change from the router.

Before you begin

- Confirm that your Cisco router is running Cisco IOS 12.0 or higher.
- Confirm that you have level 15 administrative access to the router.

Procedure

Step 1 Enable Telnet on the Cisco router as described in the documentation provided with your Cisco router or IOS software.

Step 2 On the management center, add a Cisco IOS Null Route instance for each Cisco IOS router you plan to use; see Adding a Cisco IOS Instance, on page 1008.

Step 3 Create remediations for each instance, based on the type of response you want to elicit on the router when correlation policies are violated:

- Adding Cisco IOS Block Destination Remediations, on page 1009
- Adding Cisco IOS Block Destination Network Remediations, on page 1009
- Adding Cisco IOS Block Source Remediations, on page 1010
- Adding Cisco IOS Block Source Network Remediations, on page 1011

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Adding a Cisco IOS Instance

If you have multiple routers where you want to send remediations, create a separate instance for each router.

Before you begin

- Configure Telnet access on the Cisco IOS router as described in the documentation provided with the router or IOS software.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Cisco IOS Null Route** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 In the **Router IP** field, enter the IP address of the Cisco IOS router you want to use for the remediation.

Step 5 In the **Username** field, enter the Telnet user name for the router. This user must have level 15 administrative access on the router.

Step 6 In the **Connection Password** fields, enter the Telnet user's user password.

Step 7 In the **Enable Password** fields, enter the Telnet user's enable password. This is the password used to enter privileged mode on the router.

Step 8 In the **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 Adding Cisco IOS Block Destination Remediations, on page 1009, Adding Cisco IOS Block Destination Network Remediations, on page 1009, Adding Cisco IOS Block Source Remediations, on page 1010, and Adding Cisco IOS Block Source Network Remediations, on page 1011.

Related Topics

IP Address Conventions, on page 25

Adding Cisco IOS Block Destination Remediations

The Cisco IOS Block Destination remediation blocks traffic sent from the router to the destination host involved in a correlation policy violation. Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 1008.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Adding Cisco IOS Block Destination Network Remediations

The Cisco IOS Block Destination Network remediation blocks traffic sent from the router to the network of the destination host involved in a correlation policy violation. Do not use this remediation as a response to a correlation rule that is based on a discovery or host input event. These events are associated with source hosts.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 1008.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Destination Network** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or use CIDR notation to describe the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Related Topics

IP Address Conventions, on page 25

Adding Cisco IOS Block Source Remediations

The Cisco IOS Block Source remediation blocks traffic sent from the router to the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 1008.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Adding Cisco IOS Block Source Network Remediations

The Cisco IOS Block Source Network remediation blocks traffic sent from the router to the network of the source host involved in a correlation policy violation.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Add a Cisco IOS instance as described in Adding a Cisco IOS Instance, on page 1008.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Block Source Network** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 In the **Netmask** field, enter the subnet mask or CIDR notation that describes the network that you want to block traffic to.

For example, to block traffic to an entire Class C network when a single host triggered a rule (this is not recommended), use `255.255.255.0` or `24` as the netmask.

As another example, to block traffic to 30 addresses that include the triggering IP address, specify `255.255.255.224` or `27` as the netmask. In this case, if the IP address `10.1.1.15` triggers the remediation, all IP addresses between `10.1.1.1` and `10.1.1.30` are blocked. To block only the triggering IP address, leave the field blank, enter `32`, or enter `255.255.255.255`.

Step 6 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Related Topics

IP Address Conventions, on page 25

Nmap Scan Remediations

The system integrates with Nmap™, an open source active scanner for network exploration and security auditing. You can respond to a correlation policy violation using an Nmap remediation, which triggers an Nmap scan remediation.

For more information about Nmap scanning, see Nmap Scanning.

Set Attribute Value Remediations

You can respond to a correlation policy violation by setting a host attribute value on the host where the triggering event occurred. For text host attributes, you can use the description from the event as the attribute value.

Configuring Set Attribute Remediations

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Create a set attribute instance as described in Adding a Set Attribute Value Instance, on page 1012.

Step 3 Add a set attribute remediation as described in Adding Set Attribute Value Remediations, on page 1013.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Related Topics

Predefined Host Attributes, on page 859
User-Defined Host Attributes, on page 860

Adding a Set Attribute Value Instance

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 From the **Add a New Instance** list, choose **Set Attribute Value** and click **Add**.

Step 3 Enter an **Instance Name** and **Description**.

Step 4 Click **Create**.

What to do next

- Create a set attribute remediation as described in Adding Set Attribute Value Remediations, on page 1013.

Adding Set Attribute Value Remediations

The Set Attribute Value remediation sets a host attribute on a host involved in a correlation policy violation. Create a remediation for each attribute value you want set. For text attributes, you can use the description from the triggering event as the attribute value.

In a multidomain deployment, you cannot add a remediation to an instance created in an ancestor domain.

Before you begin

- Create a set attribute instance as described in Adding a Set Attribute Value Instance, on page 1012.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Next to the instance where you want to add the remediation, click **View** ( ).

Step 3 In the **Configured Remediations** section, choose **Set Attribute Value** and click **Add**.

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

Step 4 Enter a **Remediation Name** and **Description**.

Step 5 To use this remediation in response to an event with source and destination data, choose an **Update Which Host(s) From Event** option.

Step 6 For text attributes, specify whether you want to **Use Description From Event For Attribute Value**:

- To use the description from the event as the attribute value, click **On** and enter the **Attribute Value** you want to set.
- To use the **Attribute Value** setting for the remediation as the attribute value, click **Off**.

Step 7 Click **Create**, then click **Done**.

What to do next

- Assign remediations as responses to correlation policy violations; see Adding Responses to Rules and Allow Lists, on page 955.

Managing Remediation Modules

In a multidomain deployment, the system displays remediation modules installed in the current domain, which you can delete. It also displays modules installed in ancestor domains, which you cannot delete. To manage remediation modules in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Manage your remediation modules:

- Configure — To view the Module Detail page for a module and configure its instances and remediations, click **View** ( ). In a multidomain deployment, you cannot use the Module Detail page to add, delete, or edit instances in the current domain for a module installed in an ancestor domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Managing Remediation Instances, on page 1014.
- Delete — To delete a custom module that is not in use, click **Delete** ( ). You cannot delete system-provided modules.
- Install — To install a custom module, click **Choose File**, browse to the module, and click **Install**. For more information, see the *Firepower System Remediation API Guide*.

Managing Remediation Instances

The Instances page lists all configured instances for all remediation modules.

In a multidomain deployment, the system displays remediation instances created in the current domain, which you can edit. It also displays instances created in ancestor domains, which you cannot edit. To manage remediation instances in a lower domain, switch to that domain.

Though you cannot add a remediation to an instance created in an ancestor domain, you can create a similarly configured instance in the current domain and add remediations to that instance. You can also use remediations created in ancestor domains as correlation responses.

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Manage your remediation instances:

- Add—To add an instance, choose the remediation module for which you want to add an instance and click **Add**. For system-provided modules, see:
  - Adding an ISE EPS Instance, on page 1005
  - Adding a Cisco IOS Instance, on page 1008
  - Cisco Secure Firewall Management Center Device Configuration Guide
  - Adding a Set Attribute Value Instance, on page 1012

  For help adding a custom module, see the documentation for that module, if available.
- Configure—To configure instance details and add remediations to the instance, click **View** ( ).

- Delete—To delete an instance that is not in use, click **Delete** ( ).

Managing Instances for a Single Remediation Module

The Module Detail page displays all of the instances and remediations configured for a particular remediation module.

In a multidomain deployment, you can access the Module Detail page for remediation modules installed in the current domain and in ancestor domains. However, you cannot use the Module Detail page to add, delete, or edit instances in the current domain for a module installed in an ancestor domain. Instead, use the Instances page (**Policies** > **Actions** > **Instances**); see Managing Remediation Instances, on page 1014 .

Procedure

Step 1 Choose **Policies** > **Actions** > **Modules**.

Step 2 Next to the remediation module whose instances you want to manage, click **View** ( ).

Step 3 Manage your remediation instances:

- Add — To add an instance, click **Add**. For system-provided modules, see:
  - Adding an ISE EPS Instance, on page 1005
  - Adding a Cisco IOS Instance, on page 1008
  - Cisco Secure Firewall Management Center Device Configuration Guide
  - Adding a Set Attribute Value Instance, on page 1012

  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

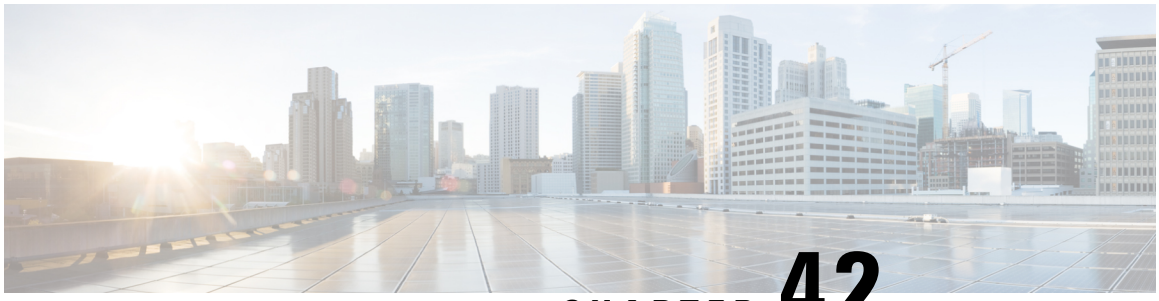

CHAPTER 42

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

When you use SSH to log into the management center, you access the CLI. Although we strongly discourage it, you can then access the Linux shell using the `expert` command .

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, as well as perform limited configuration operations.

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 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

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

Invokes the Linux shell.

Syntax

```
expert
```

Example

```
> expert
```

? (question mark)

Displays context-sensitive help for CLI commands and parameters. Use the question mark (?) command as follows:

- To display help for the commands that are available within the current CLI context, enter a question mark (?) at the command prompt.
- To display a list of the available commands that start with a particular character set, enter the abbreviated command immediately followed by a question mark (?).
- To display help for a command's legal arguments, enter a question mark (?) in place of an argument at the command prompt.

Note that the question mark (?) is not echoed back to the console.

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

Displays the product version and build as well as 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

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.

generate-troubleshoot

Generates troubleshooting data for analysis by Cisco.

Syntax

```
system generate-troubleshoot option1 optionN
```

Where options are one or more of the following, space-separated:

- `ALL`: Run all of the following options.
- `SNT`: Snort Performance and Configuration
- `PER`: Hardware Performance and Logs
- `SYS`: System Configuration, Policy, and Logs
- `DES`: Detection Configuration, Policy, and Logs
- `NET`: Interface and Network Related Data
- `VDB`: Discover, Awareness, VDB Data, and Logs
- `UPG`: Upgrade Data and Logs
- `DBO`: All Database Data
- `LOG`: All Log Data
- `NMP`: Network Map Information

Example

```
> system generate-troubleshoot VDB NMP
starting /usr/local/sf/bin/sf_troubleshoot.pl…
Please, be patient. This may take several minutes.
The troubleshoot options codes specified are VDB,NMP.
Getting filenames from [usr/local/sf/etc/db_updates/index]
Getting filenames from [usr/local/sf/etc/db_updates/base-6.2.3]
Troubleshooting information successfully created at
/var/common/results-06-14-2018—222027.tar.gz
```

lockdown

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

Reboots the appliance.

Syntax

```
system reboot
```

Example

```
> system reboot
```

restart

Restarts the appliance application.

Syntax

```
system restart
```

Example

```
> system restart
```

shutdown

Shuts down the appliance.

Syntax

```
system shutdown
```

Example

```
> system shutdown
```

History for the Secure Firewall Management Center CLI

Feature	Minimum Management Center	Minimum Threat Defense	Details
Automatic CLI access for the management center	6.5	Any	When you use SSH to log into the 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 management center. As a consequence of deprecating this option, the virtual management center no longer displays the **System** > **Configuration** > **Console Configuration** page, which still appears on physical management centers.
Ability to enable and disable CLI access for the management center	6.3	Any	New/Modified screens: New check box available to administrators in management center web interface: **Enable CLI Access** on the **System**(⚙) > **Configuration** > **Console Configuration** page. • Checked: Logging into the management center using SSH accesses the CLI. • Unchecked: Logging into 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: management center

Feature	Minimum Management Center	Minimum Threat Defense	Details
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: management center

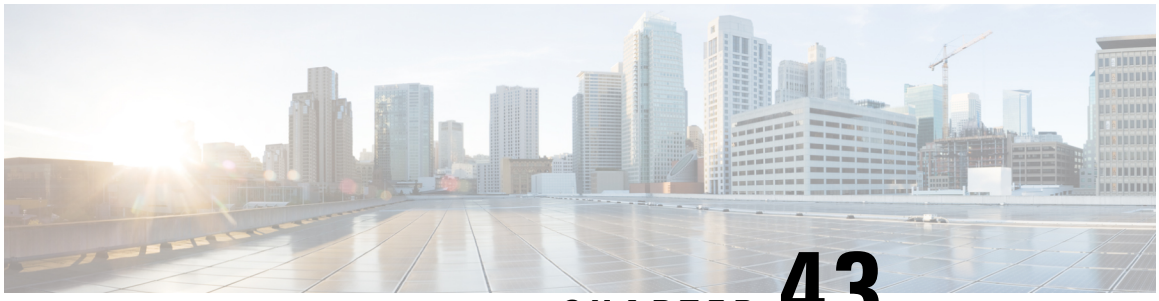

CHAPTER 43

Security, Internet Access, and Communication Ports

The following topics provide information on system security, internet access, and communication ports:

Security and Hardening

To safeguard the management center, you should install it on a protected internal network. Although the 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 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 management center. This allows you to securely control the devices from the management center. You can also configure multiple management interfaces to allow the 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 man-in-the-middle attack.

Communication Ports

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.

Ports for Management Center

The management center uses these ports to communicate.

Table 147: Inbound Ports for Management Center

Inbound Port	Protocol/Feature	Details
22/tcp	SSH	Secure remote connections to the appliance.
161/udp	SNMP	Allow access to MIBs via SNMP polling.
443/tcp	HTTPS	**Required.** Access the management center web interface.
443/tcp	HTTPS	Onboard an on-prem 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.

Table 148: Outbound Ports for 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.
53/tcp 53/udp	DNS	**Required.** DNS

Outbound Port	Protocol/Feature	Details
67/udp 68/udp	DHCP	DHCP
80/tcp	HTTP	Send and receive data from the internet. See Internet Resources Accessed, on page 1031.
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 Resources Accessed, on page 1031.
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 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.

Ports for Managed Devices

Managed devices use these ports to communicate.

Table 149: Inbound Ports for Managed Devices

Inbound Port	Protocol/Feature	Details
22/tcp	SSH	Secure remote connections to the appliance.
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 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.

Table 150: Outbound Ports for Managed Devices

Outbound Port	Protocol/Feature	Details
53/tcp 53/udp	DNS	DNS
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.

Outbound Port	Protocol/Feature	Details
443/tcp	HTTPS	Send and receive data from the internet; see Internet Resources Accessed, on page 1031.
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 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 Resources Accessed

In addition to the system accessing the internet, your browser may contact Google (google.com) or Amplitude (amplitude.com) web analytics servers to provide non-personally-identifiable usage data to Cisco.

Internet Resources Accessed by Management Center

The 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 you access. Some features also require device access; see the next table.

Table 151: Internet Resources Accessed by Management Center

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.	Each peer downloads its own certificates.	cisco.com/security/pki

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
Security intelligence	Download security intelligence feeds.	Active peer downloads, syncs to standby.	intelligence.sourcefire.com
URL filtering	Download URL category and reputation data. Manually query (look up) URL category and reputation data. Query for uncategorized URLs.	Active peer downloads, syncs to standby.	URLs: • regsvc.sco.cisco.com • est.sco.cisco.com • updates-talos.sco.cisco.com • updates-dyn-talos.sco.cisco.com • updates.ironport.com IPv4 blocks: • 146.112.62.0/24 • 146.112.63.0/24 • 146.112.255.0/24 • 146.112.59.0/24 IPv6 blocks: • 2a04:e4c7:ffff::/48 • 2a04:e4c7:fffe::/48
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

Feature	Reason	High Availability	Resource
Cisco Smart Software Manager	Communicate with the Smart Software Manager.	Active peer communicates.	www.cisco.com smartreceiver.cisco.com
Cisco Success Network	Transmit usage information and statistics.	Active peer communicates.	api-sse.cisco.com:8989 dex.sse.itd.cisco.com dex.eu.sse.itd.cisco.com
Cisco Support Diagnostics	Accepts authorized requests and transmits usage information and statistics.	Active peer communicates.	api-sse.cisco.com:8989
Cisco XDR integration	Configure devices to send events to the Cisco Security Cloud.	Active peer communicates.	Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide
Time synchronization	Synchronize time in your deployment. Not supported with a proxy server.	Both peers communicate with the NTP server.	User configured
RSS feeds	Display the Cisco Threat Research Blog on the dashboard.	Both peers communicate.	blog.talosintelligence.com
Upgrades	Download product (management center and device) upgrades.	Upgrade packages do not sync.	support.sourcefire.com cdo-ftd-images3-us-west-2.amazonaws.com
Intrusion rules	Download intrusion rules (SRU/LSP).	Active peer downloads, syncs to standby.	talosintelligence.com
Vulnerability database	Download VDB updates.	Active peer downloads, syncs to standby.	support.sourcefire.com
Geolocation database	Download GeoDB updates.	Active peer downloads, syncs to standby.	support.sourcefire.com
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 Resources Accessed by Managed Devices

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 you access.

Table 152: Internet Resources Accessed by Managed Devices

Feature	Reason	High Availability/Clustering	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.	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.
Cisco XDR integration	Send events to the Cisco Security Cloud.	All units send events.	Cisco Secure Firewall Threat Defense and Cisco XDR Integration Guide