



Health

The following topics describe how to use health monitoring:

- [Requirements and Prerequisites for Health Monitoring, on page 1](#)
- [About Health Monitoring, on page 1](#)
- [Health Policies, on page 15](#)
- [Device Exclusion in Health Monitoring, on page 25](#)
- [Health Monitor Alerts, on page 27](#)
- [About the Health Monitor, on page 30](#)
- [Health Event Views, on page 42](#)
- [History for Health Monitoring, on page 45](#)

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 runs 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 both the leaf domain and its parent domain.

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 1: 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. You can configure this module to generate alerts for selected metrics individually. The module monitors these metrics and detects failures if the difference between ASP drop counter values at any two timestamps exceeds a specified threshold.
Automatic Application Bypass	Alert	Monitors bypassed detection applications.

Module	Type	Description
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. By default, the CPU temperature threshold for the Firepower 1010 is 100 and this module generates a high CPU temperature alert when the CPU temperature of the firepower 1010 device reaches 100 degrees Celsius.
Cluster/HA Failover Status	Alert	For threat defense clusters, alerts when a unit joins, leaves, or is elected primary.
Configuration Resource Utilization	Alert	Alerts if the size of your deployed configurations puts a device at risk of running out of memory. The alert shows you how much memory your configurations require, and by how much this exceeds the available memory. If this happens, reevaluate your configurations. You may be able to reduce the number or complexity of access control rules or intrusion policies.
Connection Statistics	Metrics	Monitors connection statistics and NAT translation counts.
CPU Core Usage	Metrics	Alerts when CPU core use exceeds a configurable threshold. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
Critical Process Statistics	Metrics	Monitors the state of critical processes, their resource consumption, and the restart counts.
Database	Alert	- Monitors the size of the undo log files and raises alerts when the size exceeds threshold limits. Undo logs record the previous state of data before modifications, allowing changes to be reversed. - A warning alert is generated when the undo log file size exceeds 800 MB but is less than 1 GB. - A critical alert is generated when the undo log file size exceeds 1 GB. To change the default threshold values, contact Cisco TAC.
Data Plane CPU Usage	Metrics	Alerts when data plane CPU use exceeds a configurable threshold. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
Data Plane Memory Usage	Metrics	Alerts when data plane memory use exceeds a configurable threshold. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
Deployed Configuration Statistics	Metrics	Monitors statistics about the deployed configuration, such as the number of ACEs and IPS rules.
Disk Status	Alert	Alerts if there is an issue with the hard disk or RAID controller. If this module alerts, contact Cisco TAC. This will prevent upgrade.

Module	Type	Description
Disk Usage	Metrics	This module compares disk usage on the appliance's hard drive to the limits configured for the module and alerts when usage exceeds the thresholds configured for the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds. See Disk Usage and Drain of Events Health Monitor Alerts 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 <code>/</code> and <code>/volume</code> partitions on the appliance and track draining frequency. Although the disk usage module lists the <code>/boot</code> partition as a monitored partition, the size of the partition is static so the module does not alert on the boot partition.
File System Integrity Check	Alert	This module performs a file system integrity check and runs if the system has CC mode or UCAPL mode enabled, or if the system runs an image signed with a DEV key.
Firewall Threat Defense HA	Alert	Alerts if a threat defense high availability pair is split brain.
Firewall Threat Defense Platform Faults	Alert	Monitors Secure Firewall 1000/3100/4200 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 <i>Cisco Firepower 1000/2100 FXOS Faults and Error Messages Guide</i>.
Flow Offload Statistics	Metrics	Monitors hardware flow offload.
Hardware Alarms	Alert	This module determines if hardware needs to be replaced on a physical managed device and alerts based on the hardware status. It also reports on the status of hardware-related daemons.
Inline Link Mismatch Alarms	Alert	Alerts if inline pair interfaces negotiate different speeds.

Module	Type	Description
Interface Statistics	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.
Intrusion and File Event Rate	Alert	Alerts if intrusion events per second exceed a configurable threshold. We recommend a warning threshold of 1.5 times your average intrusion event rate, and a critical threshold of 2.5 times. For example, for an average event rate on network segment of 20 events per second, we recommend a warning value of 30 and a critical value of 50. The critical limit must be lower than 1000, and higher than the warning limit. Event rates for your devices are available on System (⚙️) > Monitoring > Statistics. If the rate is zero, the Snort process may be down or the device may not be sending events.
Link State Propagation	Alert	For the ISA 3000, alerts when an interface in a inline set fails.
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 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.

Module	Type	Description
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 existing management center configuration and the out of band configuration changes made.
Process Status	Alert	Alerts when processes on the appliance exit or terminate outside of the process manager. If a process is deliberately exited outside of the process manager, the module status changes to Warning and the health event message indicates which process exited, until the module runs again and the process has restarted. If a process terminates abnormally or crashes outside of the process manager, the module status changes to Critical and the health event message indicates the terminated process, until the module runs again and the process has restarted.
Routing Statistics	Metrics	Monitors the current state of routing table.
SD-WAN Monitoring	Metrics	Monitors the application performance metrics of SD-WAN interfaces.
Snort 3 Statistics	Metrics	Collects Snort 3 statistics for events, flows, and packets.
Snort CPU Usage	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. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
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 <i>Waiting for data</i> . 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.
Snort Memory Usage	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. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
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.

Module	Type	Description
System CPU Usage	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. This module allows you to enable or disable receiving health alerts without interrupting the collection of metrics.
Talos Connectivity Status	Alert	Monitors connectivity with Talos cloud services, required to periodically update the URL filtering database for URL reputation and categorization.
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 2: 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 Core Usage	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. This module is enabled by default. You can enable or disable receiving health alerts for this module without interrupting the collection of metrics.
Critical Process Statistics	Metrics	Monitors the state of critical processes, their resource consumption, and the restart counts.
CSDAC Dynamic Attributes Connector		
Database	Alert	Alerts if the configuration database size is too big. It also monitors the system for database schema or configuration data (sometimes called <i>EO</i>) 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 .

Module	Type	Description
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.
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 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.
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.
Firewall Threat Defense HA	Alert	Alerts if a threat defense high availability pair is split brain.
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.
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. When calculating the memory usage, the management center Memory Usage health module monitors and includes the usage of RAM, swap memory, and cache memory. 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 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. If the management center reaches a critical system memory condition, the system might terminate the processes that use high memory or reboot the management center if memory usage remains high. Complex access control policies and rules can command significant resources and negatively affect performance.
MariaDB Statistics	Metrics	Monitors the status of the MariaDB database, including the database size, number of active connections, and memory use.
Passive Identity Agent Monitor	Alert	Displays connection errors between the management center and the machine on which it's installed. The passive identity agent periodically sends updates to the management center. This health alert is displayed if the management center receives an error from the passive identity agent or if the management center has not received an update or response for five minutes or longer. Try verifying the connection between the two, restarting the passive identity agent software, and checking the health alert again in a few minutes. If issues persist, check the configuration at Integration > Other Integrations > Identity Sources.

Module	Type	Description
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.
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 <i>User Limit</i> 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
Talos Connectivity Status	Alert	Monitors connectivity with Talos, required to download URL filtering and event enrichment data.
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.

Module	Type	Description
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.
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.
Web Server Connection Statistics	Metric	Alerts you when a single IP address exceeds the configurable limit of concurrent HTTP or HTTPS connections to the management center web server. It warns you when maximum number of browser tabs are opened to connect to management center from single IP address, ensuring optimal performance.
Zero-Touch Provisioning	Alert	Alerts if there is a failure when registering a device using the serial number. It also shows errors related to zero-touch provisioning capable management centers in high availability.

Configuring Health Monitoring

Procedure

-
- Step 1** Determine which health modules you want to monitor as discussed in [Health Modules, on page 3](#).
- 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 16](#).
- Step 3** (Optional.) Configure health monitor alerts as discussed in [Creating Health Monitor Alerts, on page 28](#).

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 enable or disable health alerts for individual attributes within a health module. By adjusting the settings for health alerts at the attribute level, you can reduce the number of health alerts and streamline your focus on the most critical health alerts without interrupting data collection. You can continue to monitor the attribute from your **Health Monitor** dashboard. When you enable a health module in the health policy, the health alert configuration for all the attributes within that health module gets enabled by default.

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.

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.

You can create a custom health policy and set it as your default health policy. The management center applies the default health policy to any managed device when you add the device to the management center. Note that you cannot delete the health policy that you have set as the default. For detailed instructions for setting a default health policy, see [Set a Default Health Policy, on page 18](#).

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 16](#).
- Edit the policy to specify the module-level policy settings as described in [Edit a Health Policy, on page 17](#).

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](#).
- 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
 - After you enable a health module and when available, use the check box next to an individual attribute within a health module to enable or disable health alerts for that attribute. Note that disabling health alerts

for an attribute does not stop the collection of metrics for that attribute. You can continue to monitor these attributes from the **Health Monitor** dashboard.

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 3](#).

Step 5 Where appropriate, set the **Critical** and **Warning** threshold percentages.

Step 6 In the **Settings** tab, enter the relevant values in the fields:

- **Health Module Run Time 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](#). The metric data is collected for analysis and hence no alerting is associated with it.
- **OpenConfig Streaming Telemetry**—Configure a health metrics telemetry stream from the threat defense devices to an external data collection system which uses the vendor neutral, OpenConfig model. See [Configure OpenConfig Streaming Telemetry](#) for details.

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 16](#)

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**.

Set a Default Health Policy

You can set a user-created health policy as your default health policy. The management center applies the default health policy to any managed device when you add the device to the management center.



Note

Setting a new default health policy does not affect the health policy assigned to devices that are already registered.

Procedure

-
- Step 1** Choose **System** (⚙️) > **Health** > **Policy**.
- Step 2** Click the **More Actions** (⋮) icon next to the health policy that you want to set as the default, and click **Set as Default**.
- Step 3** Click **Proceed**.
-

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.
- Note that you cannot delete a health policy that is set as the default policy. Set another health policy as the default before attempting to delete the current default policy. For more information, see [Set a Default Health Policy, on page 18](#).
-

Send Vendor-Neutral Telemetry Streams Using OpenConfig

OpenConfig is a vendor-independent software layer that provides a single way of streaming network telemetry data to multiple vendors to manage and monitor networks. The OpenConfig streaming telemetry option in the secure firewall uses gNMI (gRPC Network Management Interface) protocol and allows you to control and generate telemetry streams from your threat defense devices to a data collection system.

The firewall threat defense health policy contains all the configurations to support and enable the OpenConfig streaming telemetry functionality. When you deploy the health policy to the device, the OpenConfig streaming telemetry configuration activates a gNMI server and starts listening to Remote Procedure Call (RPC) messages from the data collectors.

Subscription Model of OpenConfig Streaming Telemetry

OpenConfig uses a subscription-based model where the data collectors query the threat defense devices for telemetry data or act as collector for the streamed telemetry data. When a data collector wishes to receive updates and metrics from the threat defense device, it sends a `subscribeRequest` RPC message to the threat defense gNMI server. The subscription request includes details of one or more paths to which the data collector wishes to subscribe. The message also includes subscription mode which describes the longevity of the subscription. The threat defense server supports the following subscription modes:

- *Once subscription*—The threat defense device sends requested data to the gNMI paths only once.
- *Stream subscription*—The threat defense continuously streams telemetry data according to the triggers specified in the `SubscribeRequest` RPC message.
 - *Sampled subscription*—The threat defense server streams the requested data as per the interval specified in the subscription message. The minimum interval that the threat defense support is one minute.
 - *On-change subscription*—The threat defense sends the data whenever the requested values change.

The threat defense server generates `SubscribeResponse` RPC messages according to the type of subscription that is created, at the frequency requested by the data collectors.

Deployment Modes for OpenConfig Streaming Telemetry

You can use the following deployment modes for OpenConfig streaming telemetry configuration:

- **DIAL-IN**—In this mode, the gNMI server opens a port on the threat defense and waits for `SubscribeRequest` RPC messages from data collectors. In the device health policy, you can specify the port number to use by the gNMI server and the IP address of the data collector that can connect with the gNMI service. If not specified, the gNMI server uses port number 50051. The Dial-in mode is ideal to use in a trusted network where the endpoint that subscribes to telemetry streams are trusted.
- **DIAL-OUT**—The gNMI service is designed to work in server mode where it accepts subscription requests from gNMI data collectors and serve the telemetry data. If the gNMI data collectors cannot reach the gNMI server, the threat defense uses a tunnel client and establishes a gRPC tunnel with the external server. This tunnel allows exchange of RPC messages between gNMI server and client. The Dial-Out mode is ideal to use when the data collectors are hosted on the cloud or outside the trusted network.

In both dial-in and dial-out mode, all the communication between gNMI server and gNMI client uses TLS encryption and this requires to generate a set of certificates with private keys for the TLS encryption. Dial-out mode requires extra keys for the tunnel infrastructure. See [How to Generate Certificate with Private Key](#) for more information.

Generate Certificates and Private Keys

Generate the CA, server, and client certificate and private key sets required for OpenConfig streaming telemetry configuration.



Note

To ensure that you generate certificates using the same CA, run the following commands together and from the same endpoint. If you want to retry the commands, you must retry all commands.

Before you begin

Procedure

Step 1 Make a folder, for example `keys`, in the endpoint where you want to run the following commands.

Example:

```
mkdir keys
```

Step 2 Create a self-signed CA certificate with a corresponding private key.

Example:

Following sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information:

```
openssl req -x509 -newkey rsa:4096 -days 365 -nodes -keyout keys/ca-key.pem -out
keys/ca-cert.pem -subj "/C=XX
/ST=YY/L=ZZZ/O=Example/OU=EN/CN=gnmi-ca/emailAddress=abc@example.com"
```

The subject information includes the provided Country (C), State (ST), Locality (L), Organization (O), Organizational Unit (OU), Common Name (CN), and email address.

The private key is saved as `ca-key.pem` file, and the certificate is saved as `ca-cert.pem` file in the `keys` folder.

Step 3 Create a self-signed server certificate with the specified Common Name (CN) and Subject Alternative Name (SAN):

Example:

Following sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.200 is the IP address of the threat defense device and 192.168.0.202 is the IP address of the client.

Note

Client IP is not required if you want to use this certificate and key sets in dial-in mode.

```
CN="192.168.0.200"
SAN="IP:192.168.0.200,IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/server-key.pem -out keys/server-req.pem
-subj "/C=XX/ST=YY/L=ZZZ/O=Example/OU=EN/CN=${CN}/emailAddress=abc@example.com"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/server-req.pem -days
60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out keys/server-cert.pem
cat keys/server-key.pem keys/server-cert.pem keys/ca-cert.pem > keys/server-combined.pem
```

The `openssl req` command generates a new RSA private key and a Certificate Signing Request (CSR). The private key is saved as `server-key.pem` file, and the CSR is saved as `server-req.pem` file in the `keys` folder.

The `openssl x509` command processes the CSR and generates a server certificate. The server certificate is saved as `server-cert.pem` file in the `keys` folder.

The `cat` command combines the server key, server certificate, and the CA certificate into a single file named `server-combined.pem` and saves the file in the `keys` folder.

You have to upload the `server-combined.pem` while configuring **OpenConfig Streaming telemetry** from the management center. The gNMI server that runs on the threat defense and the tunnel server (dial-out mode) uses this certificate for TLS communication. If you encrypt the private key with a passphrase, ensure that you specify the passphrase while uploading the certificate to the management center.

Step 4 Create client certificate with the specified Common Name (CN) and Subject Alternative Name (SAN).

Example:

Following sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.202 is the IP address of the client.

```
CN="192.168.0.202"
SAN="IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/client-key.pem -out keys/client-req.pem
-subj "/C=XX/ST=YY/L=ZZZ/O=example/OU=EN/CN=${CN}/emailAddress=abc@example.com"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/client-req.pem -days
60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out keys/client-cert.pem
```

The gNMI client uses the client certificate `client-cert.pem` and the private key for TLS communication.

Step 5 (Optional) For dial-out mode, create the tunnel server certificate with the specified Common Name (CN) and Subject Alternative Name (SAN).

Example:

Following sample command generates a new RSA private key and uses it to create a self-signed X.509 certificate with provided subject information. In this example, 192.168.0.202 is the IP address of the client.

```
CN="192.168.0.202"
SAN="IP:192.168.0.202"
openssl req -newkey rsa:4096 -nodes -keyout keys/tunnel-server-key.pem -out
keys/tunnel-server-req.pem -subj "
/C=XX/ST=YY/L=ZZZ/O=Example/OU=EN/CN=${CN}/emailAddress=abc@example.com"
openssl x509 -req -extfile <(printf "subjectAltName=${SAN}") -in keys/tunnel-server-req.pem
-days 60 -CA keys/ca-cert.pem -CAkey keys/ca-key.pem -CAcreateserial -out
keys/tunnel-server-cert.pem
```

Configure OpenConfig Streaming Telemetry

Before you begin

- Ensure that the threat defense device where you want to deploy the health policy configuration allows installation of the SSL certificate and private key.
- Ensure that you configure a gNMI client that supports the OpenConfig streaming telemetry implementation, from which you can make the gRPC requests to the gNMI server on the threat defense.
- To use dial-out mode and configure OpenConfig streaming telemetry, ensure that you configure a gRPC tunnel server and client on the management system. This tunnel configuration enables communication between the gNMI client and the threat defense device.
- You must be an admin user to perform the following task.

Procedure

- Step 1** Choose **System > Policy**.
- Step 2** Click the **Edit health policy** icon next to the threat defense health policy that you want to modify.
- Step 3** Go to **Settings** tab.

Step 4 Move the **OpenConfig Streaming Telemetry** slider to enable the configuration. This configuration is disabled by default.

Step 5 Upload the **SSL Certificate**. The gNMI server uses this certificate to enable server authentication for the TLS connection and encrypt all communications through the channel.

The OpenConfig streaming telemetry configuration supports only certificate with PEM format. The management center performs the following certificate validations to ensure encrypted communication between the appliance and gNMI collectors without connection failures:

- Verifies that the ASCII text is a valid certificate file.
- Checks the expiration date of the uploaded certificate.
- Verifies the number of expected certificates and private key in the uploaded PEM file. The file must have minimum one certificate and the number of private key in the certificate must always be 1.
- Verify and accept key block types PRIVATE KEY, RSA PRIVATE KEY, ENCRYPTED PRIVATE KEY, or RSA ENCRYPTED PRIVATE KEY.
- For the encrypted PEM file, verify that the Proc-Type: 4,ENCRYPTED? keyword is present.
- Verifies that the passphrase is valid for the encrypted PEM files.

Step 6 (Optional) Specify the Passphrase if the private key files are encrypted.

Step 7 Choose the deployment mode to use for streaming telemetry over gNMI protocol.

For **DIAL-IN** mode:

- a. Assign a port number for the gNMI service.
The gNMI server opens the port and waits for gRPC requests from the collector.
- b. Specify the IPv4/IPv6 address of the gNMI collectors that can connect to the threat defense device.
- c. Click **Add Collector** to add more gNMI collectors. You can add a maximum of five collectors.

For **DIAL-OUT** mode:

- a. Specify the hostname and port number for the gNMI collector, which can subscribe to streaming telemetry from the threat defense device.
- b. Click **Add Collector** to add more gNMI collectors. You can add a maximum of five collectors.

Step 8 Specify the username and password to validate the gNMI collector.

The threat defense server uses this credential to authenticate the gNMI collector when receiving the `SubscribeRequest` RPC message. Each telemetry message is not authenticated using the username and password. The system uses the previously authenticated encrypted streaming channel to carry telemetry messages.

Step 9 Click **Save**.

What to do next

Deploy the health policy to your threat defense device, for the configuration changes to take effect.

Troubleshoot OpenConfig Streaming Telemetry

Certificate signed by unknown authority

- Ensure that you have uploaded the correct certificate to the management center.
- Verify the certificate and key generation steps. Ensure that the IP Subject Alternative Name (SAN) is specified correctly.

Certificate is not valid

If the management center displays the error "Request was made for (IP), but the certificate is not valid for (IP)" then verify the server certificate and key generation steps.

- Ensure sure that the IP SAN is correctly specified in the server certificate. If the configuration applies to more than one threat defense device, you must specify all the devices in the IP SAN field.
- If you are using dial-out mode, ensure that the client IP is specified in the server certificate.

Failed to generate response object

If you receive "Failed to generate response object, did not receive any data" error, the gNMI input plug-in is waiting for metric export. Below is the sample response that appears when the telegraph is restarting:

```
root@cronserver:/home/secanup/openconfig-test# gnmic -a $ADDRESS:$PORT --tls-cert $CLIENTCERT
--tls-ca $CACERT --tls-key $CLIENTKEY -u $USER -p $PASS sub --mode once --path
"openconfig-system/system/memory"
rpc error: code = Aborted desc = Error in gnmic_server: failed to generate response object.did
not receive any data
Error: one or more requests failed
```

Wait for the gNMI input plug-in to restart and retry your request.

Restart telegraph

When telegraph is not responding, restart the process using the following command on the threat defense CLI console:

```
pmtool restartbyid hmdaemon
```

Get current status of gNMI server

When OpenConfig streaming telemetry is enabled, to know the status of the gNMI server, run the following command using the threat defense CLI console:

```
curl localhost:9275/OpenConfig/status
```

Below is the sample response to the command:

```
root@firepower:/home/admin# curl localhost:9275/openconfig/status
Mode (Dialin/Dialout): DialIn
Subscription Details:
  Active Subscription Details:
    Stream Mode Subscription Details:
      Total Stream Subscription Request Count: 1
      'Ip of Collector- Subscribe paths:'
        172.16.0.101:45826:
          - /openconfig-system/system/state/hostname
      Sample Subscription Count: 1
```

```

    On Change Subscription Count: 0
  Once Mode Subscription Details:
    Total Subscription Request Count: 0
    Total Subscription Count: 0
    'Ip of Collector- Subscribe paths': {}
Total Subscription Details:
  Stream Mode Subscription Details:
    Total Stream Subscription Request Count: 1
    'Ip of Collector- Subscribe paths':
      172.16.0.101:45826:
        - /openconfig-system/system/state/hostname
Sample Subscription Count: 1
On Change Subscription Count: 0
Once Mode Subscription Details:
  Total Subscription Request Count: 0
  Total Subscription Count: 0
  'Ip of Collector- Subscribe paths': {}

```

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 26](#).

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 3](#).

- 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:
- Click the **Slider** () button next to the desired module.
 - 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 3: 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](#).

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](#).

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 both the leaf domain and its parent domain. 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 34](#) 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 31](#) 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-glance view of the critical processes as well as a tab that lets you see state of all processed, including the CPU and memory usage for each process.
- CPU—The CPU panel lets you toggle between the average CPU usage (default) and the CPU usage of all cores.
- Memory—The Memory panel shows the overall memory usage on the 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 33](#).



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](#) and [Configure Session Timeouts](#) 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](#).
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](#).

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](#)
- Generate and download advanced troubleshooting files; see [Download Advanced Troubleshooting Files](#).
- Create and modify health policies; see [Creating Health Policies, on page 16](#).

- Create and modify health monitor alerts; see [Creating Health Monitor Alerts, on page 28](#).

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 the deployment details on top of the graph** (📅) icon for a deployment overlay on the trend graph, with respect to the selected time range.

The 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 37](#).

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 New Dashboard** (+) icon to add a new dashboard.
- Step 5** Specify a name to identity the dashboard.
- Step 6** To create a dashboard from a predefined correlation group, click **Add from Predefined Correlations** drop-down, choose the group, and click **Add Dashboard**.
- Step 7** To create a custom correlation dashboard, choose a group from the **Select Metric Group** drop-down, then choose 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 Dashboard** to add the dashboard to the health monitor.
- Step 11** You can **Edit** or **Delete** the predefined dashboards and the 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 4: 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.
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 transitional 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 transitional 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.

Status Level	Status Icon	Status Color in Pie Chart	Description
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 36](#).
- 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 5: 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 <code>CPU</code> . 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 <code>Unable to Execute</code> .
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 6:

Feature	Minimum Management Center	Minimum Threat Defense	Details
View health status summary for a device in both the leaf domain and its parent domain	7.6.1	Any	In a multi-domain deployment, you can view the health status summary for a device in both the leaf domain and its parent domain.
Set default health policy	7.6.0	Any	You can now set a user-created health policy as your default health policy. The management center applies the default health policy to any managed device when you add the device to the management center. To set a default health policy, choose System (⚙️) > Policy, click the More Actions (⋮) icon next to the health policy that you want to set as the default, and then click Set as Default.
Customize your health policy to minimize health alerts while continuing to collect data.	7.6.0	Any	You can now disable health alerts for individual attributes within the CPU, Memory, and ASP Drop health modules without stopping the data collection. By disabling health alerts for specific attributes, you can minimize the health alert noise and focus on the most critical issues. New/modified screens: Click System (⚙️) > Health > Policy, and then click the Edit icon next to either the Firewall Threat Defense health policy or the Firewall Management Center health policy.
Health alert for Talos connectivity	7.6.0	7.6.0	Management center now alert if the Talos connectivity daemon fails to report its health status within the specified time frame.
Updated the management center memory usage module default thresholds.	7.4.1	Any	The default thresholds for the management center memory usage warning and critical alarms are now set to 88% and 90% respectively. New/modified screens: System(⚙️) > Health > Policy > edit the Firewall Management Center Health Policy > Health Modules > Memory Usage.
Improved memory usage calculation for management center.	7.4.1	Any	The management center memory usage module to consider the amount of swap memory available and the cache memory when calculating the memory usage to accurately determine the memory usage and send health alerts. New/modified screens: System(⚙️) > Health > Monitor > Firewall Management Center > Add New Dashboard.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Health alerts for NTP server sync issues.	7.4.1	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.
CPU and chassis temperature alerts for firepower 1010 and 1100 devices.	7.4.1	7.4.1	Alerts for CPU or chassis temperature warnings or critical levels appear in the management center for Firepower 1010 and 1100 devices.
Stream telemetry to an external server using OpenConfig.	7.4	7.4	You can now send metrics and health monitoring information from your threat defense devices to an external server (gNMI collector) using OpenConfig. You can configure either threat defense or the collector to initiate the connection, which is encrypted by TLS. New/modified screens: System(⚙️) > Health > Policy > Firewall Threat Defense Policies > Settings > OpenConfig Streaming Telemetry.
Health monitor usability enhancements.	7.4	Any	Improved Add New Dashboard dialog box which helps to create the custom dashboards with ease. Included option to edit or delete the predefined device health monitor dashboards. New/modified screens: System(⚙️) > Health > Monitor > Devices > Add New Dashboard.
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.

Feature	Minimum Management Center	Minimum Threat Defense	Details
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
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—<i>Critical</i> alert is displayed for abnormal thermal status, and <i>Normal</i> 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.

Feature	Minimum Management Center	Minimum Threat Defense	Details
Health monitor usability enhancements.	7.1	Any	Following UI page were improved for better usability and presentation of data: • Policy • Exclude • Monitor Alerts New/modified screens: . • System( > Health > Policy • System( > Health > Exclude • System( > Health > Monitor Alerts
Elephant flow detection.	7.1	Any	The health monitor includes the following enhancements: • The Connection statistics includes active elephant flows. • The Connection Group Metrics includes the number of active elephant flows. The Elephant Flow Detection feature is not supported on the Cisco Firepower 2100 series.
Discontinued high unmanaged disk usage alerts.	7.0.6	Any	The Disk Usage health module no longer alerts with high unmanaged disk usage. After upgrade, you may continue to see these alerts until you either deploy health policies to managed devices (stops the display of alerts) or upgrade the devices (stops the sending of alerts). Note Versions 7.0–7.0.5, 7.1.x, 7.2.0–7.2.3, and 7.3.x continue to support these alerts. If your 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.

