

Health

The following topics describe how to use health monitoring:

- Get started with health monitoring, on page 1
- Configure health monitoring, on page 7
- Apply and verify health monitoring, on page 12
- Monitor health, on page 13
- Device exclusion in health monitoring, on page 26
- Sending device health metrics to OpenConfig, on page 30
- Reference, on page 35
- system auditing, on page 48

Get started with health monitoring

Learn how health policies, modules, metrics, events, and alerts work together to monitor Cloud-Delivered Firewall Management Center and its managed appliances.

Health monitoring in Cloud-Delivered Firewall Management Center

Health monitoring tracks hardware, software, and critical system functions in Cloud-Delivered Firewall Management Center and its managed appliances. Cloud-Delivered Firewall Management Center runs health modules at configured intervals, evaluates the results, and presents health information for your deployment.

Health monitoring provides two related functions: collecting time series data for metric analysis, and monitoring system health and generating alerts. Use health monitoring to analyze health and performance trends, review health events, identify conditions that require attention, and send selected health notifications to external destinations.

These categories explain various elements of the Cloud-Delivered Firewall Management Center health monitoring model and how they work together.

Configure health monitoring

- Health policy—Defines what Cloud-Delivered Firewall Management Center monitors for an assigned appliance. A health policy contains enabled health modules and the settings that control how each module operates, such as module-specific criteria, thresholds, and run or collection intervals where applicable. Select modules according to the system conditions and operational outcomes that you need to monitor,

and review module applicability before applying the policy. For more information, refer to Health policies, on page 5.

- Health module—Tests a specific system function or collects information about it. Alert modules report health status. Metrics modules collect time-series data, and some metrics modules also generate alerts when values cross configured thresholds. For more information, refer to Health modules, on page 35.

Health information produced by monitoring

- Health metrics—Time series data collected over time for a monitored system function. Cloud-Delivered Firewall Management Center displays metrics in health monitor dashboards so that you can analyze changes, trends, and relationships between system conditions.
- Health status—Indicates the result of a health module test or the compiled condition of a monitored appliance. Health monitor status categories include **Error**, **Critical**, **Warning**, **Normal**, **Recovered**, and **Disabled** as applicable.
- Health event—A record that Cloud-Delivered Firewall Management Center logs when a health module result meets the configured event conditions. The event contains information about the health module, monitored appliance, reported status, and condition. For more information, refer to Health event views, on page 22.
- Health monitor dashboards—Displays the compiled health status for Cloud-Delivered Firewall Management Center and its managed appliances. It also provides an option to create custom dashboards to view specific health information. For more information, refer to About the health monitor, on page 14.

Respond to health information

- Alert response—Defines how Cloud-Delivered Firewall Management Center sends information to an external destination. For health-event alerting, use email or webhook alert responses. SNMP and syslog are not supported as the primary methods for sending health alerts to external systems.
- Health monitor alert—Associates a severity level, one or more health modules, and an alert response. When the selected severity occurs for a selected module, the health monitor alert triggers its associated alert response. For more information, refer to Health monitor alerts, on page 10.

Health monitoring scope and behavior

- Cloud-Delivered Firewall Management Center automatically reports its own status using the modules configured in the default health policy. All appliances automatically report hardware status through the Hardware Alarms health module.

  Some health modules, such as the **Appliance Heartbeat** module, run on Cloud-Delivered Firewall Management Center and report the status of managed devices. To provide managed-device status, health policies must be deployed to the devices.
- If a health policy enables a module that does not apply to the selected appliance, the health monitor reports that module as **Disabled**. Review module applicability before applying a policy.

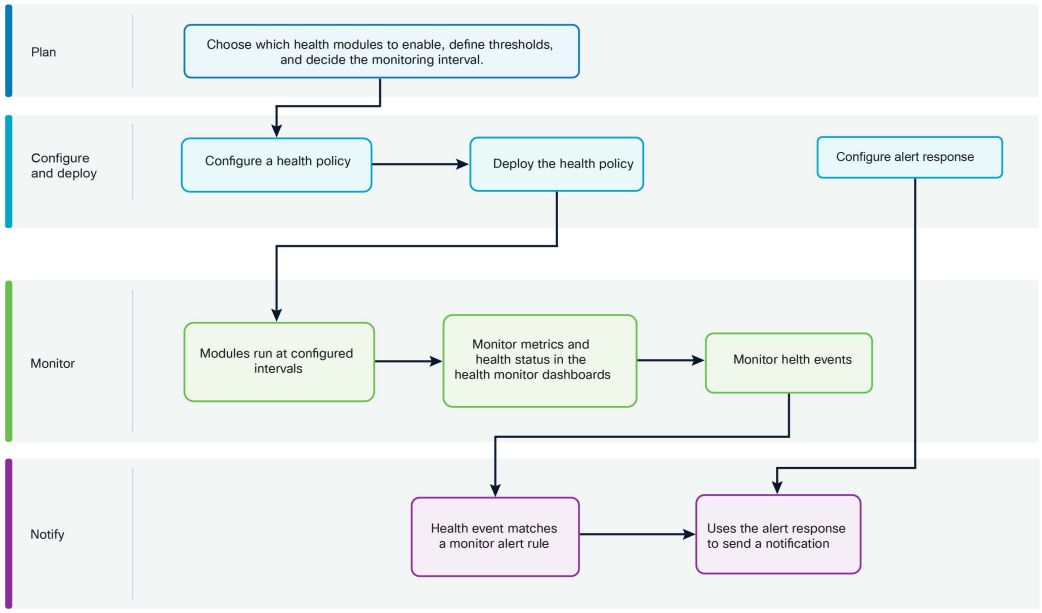
How health monitoring works

Summary

Health monitoring in Cloud-Delivered Firewall Management Center enables you to track appliance health using policies, modules, metrics, events, and notifications, providing visibility and automated alerting for network conditions.

Workflow

Figure 1: Health monitoring workflow



These stages describe how health monitoring information flows from policy configuration to monitoring and notification.

1. Choose the health modules that provide the health metrics and alerts required for your monitoring goal. Check that each module supports your appliance type and review its configuration requirements, dependencies, and available thresholds. For more information, refer to Health modules, on page 35.
2. Create or modify a health policy. Enable the selected health modules and configure their settings, run intervals, and thresholds where needed. For more information, refer to Health policies, on page 5.
3. Deploy the health policy to the appropriate managed appliances. Monitoring begins after you successfully deploy the health policy.
4. The enabled health modules run at the intervals configured in the health policy. A module can collect metric data, report a health alert, or perform both functions, depending on the module.

5. Cloud-Delivered Firewall Management Center displays health statuses and metrics in the health monitor dashboards. Metric data supports analysis over time; it does not necessarily produce a health event for every collected value. For more information, refer to About the health monitor, on page 14.
6. When a module result meets the conditions configured for event generation, Cloud-Delivered Firewall Management Center logs a health event.

   Note that the health event logs the condition in Cloud-Delivered Firewall Management Center and does not, by itself, send an external notification.
7. Cloud-Delivered Firewall Management Center evaluates the health event against configured health monitor alerts. When the event's module and severity match a health monitor alert, Cloud-Delivered Firewall Management Center invokes the alert response associated with that alert.
8. The alert response delivers the notification to its configured destination, such as an email server or webhook endpoint.

Result

During operation, you can run all health modules, or a specific module, on demand to investigate a condition. You can also exclude appliances or health modules if you do not want them to generate health information.

Your health policy determines what Cloud-Delivered Firewall Management Center monitors and how often it evaluates the condition. The health monitor displays the results. The health event records a qualifying condition. The health monitor alert determines which conditions require notification, and the alert response determines how and where the notification is delivered.

Health event external notifications

Firewall Management Center uses separate configurations to detect a health condition and notify an external system about it.

A health policy defines what Cloud-Delivered Firewall Management Center monitors and which conditions produce a health event. An alert response configuration defines how Cloud-Delivered Firewall Management Center communicates with an external destination, such as an email server. The monitor alert rule configuration connects the health policy and alert response by specifying which health events should trigger a particular alert response.

How the configurations work together

These steps describe how you configure external notifications for health events and how Cloud-Delivered Firewall Management Center processes the health event.

1. Configure and apply the health policy: A health policy enables health modules and defines the criteria used to assess appliance health.

   For example, a health policy can enable the **Memory Usage** module and define **Critical** thresholds.

   Apply the health policy to the appliance, including the Cloud-Delivered Firewall Management Center when you want to monitor Cloud-Delivered Firewall Management Center resources.
2. Configure the alert response: An alert response defines how and where Cloud-Delivered Firewall Management Center sends the external notification. Depending on the response type, the alert response contains the destination address, server or endpoint information, and other communication settings.

For example, an email alert response identifies the recipients, sender, and mail relay host that Cloud-Delivered Firewall Management Center uses to send the message.

For more information about alert responses, refer to Configuring external alerts with alert responses.

3. Configure the monitor alert rule: A monitor alert rule associates a health-event severity and one or more health modules with an alert response.

   For example, a health monitor alert rule can specify:

   - Severity: **Critical**
   - Health module: **Memory Usage**
   - Alert response: **Critical memory email**

   This means that when Cloud-Delivered Firewall Management Center generates a **Critical** event for the **Memory Usage** module, it invokes the associated alert response.

   For more information about health monitor alerts, refer toHealth monitor alerts, on page 10.

4. Cloud-Delivered Firewall Management Center generates a health event: Cloud-Delivered Firewall Management Center evaluates the enabled health modules according to the health policy and its configured run interval. When a health module evaluates an appliance and the result meets a configured condition, Cloud-Delivered Firewall Management Center generates a health event.

   Note that a health event is an event log within Cloud-Delivered Firewall Management Center; it does not send an external notification by itself.

5. Cloud-Delivered Firewall Management Center matches the event and sends the notification: Cloud-Delivered Firewall Management Center evaluates the generated event against the configured health monitor alerts. If the event's module and severity match a monitor alert rule, Cloud-Delivered Firewall Management Center invokes the alert response associated with that rule and sends the notification to the configured external system.

Note Health alerts can take approximately 5–6 minutes to be generated after the corresponding health event occurs.

Health policies

A health policy is a configuration framework that

- enables you to define test criteria for monitoring the health of appliances,
- allows selective enabling of health modules and critical alert settings, and
- provides flexibility to apply, customize, or use default policies for each appliance.

Health policy configuration options

When you configure a health policy, you decide whether to enable each health module for that policy. You also select the criteria that control which health status each enabled module reports each time it assesses the health of a process.

You can 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 Cloud-Delivered Firewall Management Center automatically assigns it the default health policy. To disassociate a health policy from an appliance, you must first associate a different health policy with it. An appliance must have at least one health policy assigned.

Default health policy

The Cloud-Delivered Firewall Management Center setup process creates and applies an initial health policy, in which most—but not all—available health modules are enabled. The system also applies this initial policy to devices added to the Cloud-Delivered Firewall Management Center.

This initial health policy is based on a default health policy. You cannot view or edit the default health policy, but you can copy it when creating a custom health policy.

You can create a custom health policy and set it as your default health policy. The Cloud-Delivered Firewall Management Center applies the default health policy to any managed device when you add the device to the Cloud-Delivered Firewall 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, refer to Set a default health policy, on page 9.

Upgrades and the default health policy

When you upgrade the Cloud-Delivered Firewall Management Center, any new health modules are added to all health policies, including the initial health policy, default health policy, and any other custom health policies. Usually, new health modules are added in an enabled state.



Note For a new health module to begin monitoring and alerting, reapply health policies after upgrade.

Requirements and prerequisites for health monitoring

This reference provides the requirements and prerequisites needed to implement health monitoring functionality.

Model support

Any

Supported domains

Any

Configure health monitoring

Learn how to create, edit, assign, set as default, and delete health policies for managed appliances. Also, learn about health monitor alert severities and how to create, edit, and delete alerts that trigger responses when monitored conditions occur.

Create a health policy

If you want to customize a health policy to use with your appliances, you can create a new policy. The settings in the policy initially populate with the settings from the health policy you choose as a basis for the new policy. You can edit the policy to specify your preferences, such as enable or disable modules within the policy, change the alerting criteria for each module as needed, and specify the run time intervals.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

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 13.
- Edit the policy to specify the module-level policy settings as described in Edit a health policy, on page 7.

Edit a health policy

Edit a health policy when you need to adjust which modules are monitored, change critical or warning thresholds, or reassign policies to different devices.

Before you begin

- Ensure you have access to Cloud-Delivered Firewall Management Center.

- Identify the health policy you need to modify.

Follow these steps to edit a health policy:

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

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—Firewall Threat Defense, Cloud-Delivered Firewall Management Center, or both.
- You cannot choose to include or exclude the individual attributes of CPU and Memory modules.

For information on the modules, refer to Health modules, on page 35.

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, refer to View the device health monitor, on page 16. 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 Firewall Threat Defense devices to an external data collection system which uses the vendor neutral, OpenConfig model. For more information, refer to Configure OpenConfig Streaming Telemetry.

Step 7 To view and modify the devices to which the policy is assigned, do the following:

a) Click **Policy Assignments & Deploy**.

b) From the **Available Devices** list, click the + icon next to the device to which you want to assign the health policy.

c) Click **Apply**.

Note The UI prompts you to save unsaved changes before continuing with assignment and deployment.

Alternatively, you can apply the health policy to your appliance as described in Apply a health policy, on page 13

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 Cloud-Delivered Firewall 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 Cloud-Delivered Firewall Management Center applies the default health policy to any managed device when you add the device to the Cloud-Delivered Firewall 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 **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

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.

Note To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

Step 2 Click **Delete** ( ) next to the policy that you want to delete.

Step 3 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, refer to Set a default health policy, on page 9.

Health monitor alerts

A health monitor alert is a notification mechanism that

- enables association of alert responses with health event levels to trigger notifications based on event severity, and
- allows customization to monitor specific system conditions such as disk space thresholds.

Disk space monitoring alert example

If you are concerned that your appliances may run out of hard disk space, you can automatically send an email to a system administrator when the remaining disk space reaches the warning level. If the hard drive continues to fill, you can send a second email when the hard drive reaches the critical level.

Health monitor alert information

The alerts generated by the health monitor contain specific information that helps identify issues and their severity.

Health monitor alerts include these attributes:

- Severity: Indicates the severity level of the alert.
- Module: Specifies the health module whose test results triggered the alert.
- Description: Includes the health test results that triggered the alert.

The table describes these severity levels.

Table 1: 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.

Create a health monitor alert

When you create a health monitor alert, you create an association between a severity level, a health module, and an alert response. You can use an existing alert or configure a new one specifically to report on system health. When the severity level occurs for the selected module, the alert triggers.

If you create or update a threshold in a way that duplicates an existing threshold, you are notified of the conflict. When duplicate thresholds exist, the health monitor uses the threshold that generates the fewest alerts and ignores the others. The timeout value for the threshold must be between 5 and 4,294,967,295 minutes.

Before you begin

- Ensure that you have configured alert responses that govern the Cloud-Delivered Firewall Management Center's communication with the email server or webhook where you send the health alert. For more information, refer to Configuring external alerts with alert responses.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Monitor Alerts**.

Step 2 Click **Add**.

Step 3 In the **Add Health Alert** dialog box, enter a name for the health alert in the **Health Alert Name** field.

Step 4 From the **Severity** drop-down list, choose the severity level you want to use to trigger the alert.

Step 5 From the **Alert** drop-down list, choose the alert response that you want to trigger when the specified severity level is reached. If you have not yet configured the alert responses, click **Alerts** to visit the **Alerts** page and set them.

Step 6 From the **Health Modules** list, choose the health policy modules for which you want the alert to apply.

Step 7 Optionally, in the **Threshold Timeout** field, enter the number of minutes that should elapse before each threshold period ends and the threshold count resets.

Even if the policy run time interval value is less than the threshold timeout value, the interval between two reported health events from a given module is always greater. For example, if you change the threshold timeout to 8 minutes and the policy run time interval is 5 minutes, there is a 10-minute interval (5 x 2) between reported events.

Step 8 Click **Save** to save the health alert.

Edit a health monitor alert

Administrators may need to update the alert parameters for health monitor alerts in order to adjust severity, change modules, or select different responses as requirements evolve.

Before you begin

Follow these steps to edit health monitor alerts:

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Monitor Alerts**.

Step 2 Click the **Edit** (✎) icon that is provided against the required health alert that you want to modify.

Step 3 In the **Edit Health Alert** dialog box, from the **Alert** drop-down list, select the required alert entry, or click **Alerts** link to configure a new alert entry.

Step 4 Click **Save**.

Delete health monitor alerts

Health monitor alerts notify you about system events, issues, or warnings. You may need to delete outdated or irrelevant alerts to streamline monitoring.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Monitor Alerts**.

Step 2 Click **Delete** (🗑) next to the health alert you want to delete.

Step 3 Click **Delete health alert** to delete it.

What to do next

Disable or delete the underlying alert response to ensure that alerting does not continue. For more information, refer to Configuring external alerts with alert responses.

Apply and verify health monitoring

Learn how to apply a health policy so that its enabled modules automatically monitor appliance health, and how to run all health modules on demand for up-to-date device health information.

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 Cloud-Delivered Firewall Management Center.

If you enable a module in a health policy and then apply the policy to an appliance that does not require that health test, the health monitor reports the status for that health module as disabled.

To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored.

When you apply a different policy to an appliance that already has a policy applied, expect some latency in the display of new data based on the newly applied tests.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

Step 2 Click the **Deploy health policy** ( ) next to the policy you want to apply.

Step 3 Choose the appliances where you want to apply the health policy.

Note An appliance must have at least one health policy assigned to it. To stop health monitoring for an appliance, create a health policy with all modules disabled, assign it to the appliance, and remove the appliance's other health policy assignments. The appliance retains one assigned health policy, but no health modules are monitored. To disassociate a health policy from an appliance, you must first associate a different health policy with it.

Step 4 Click **Apply** to apply the policy to the appliances you chose.

The health policy is successfully applied to the selected appliances, and monitoring begins automatically.

What to do next

- Optionally, monitor the task status; see View task messages.

  Monitoring of the appliance starts when the policy is successfully applied.

Monitor health

Use health dashboards, metrics, and event views to assess device and management center status and investigate changes over time. Manage default policies and exclusions to control which conditions affect the health summary.

About the health monitor

The health monitor provides the compiled health status for all devices managed by the Cloud-Delivered Firewall Management Center, plus the Cloud-Delivered Firewall Management Center itself.

The health monitor is composed of:

- The Health Status summary page — Provides you with an at-a-glance view of the health of the Cloud-Delivered Firewall Management Center and all of the devices that the Cloud-Delivered Firewall Management Center manages. In a multi-domain deployment, you can view the health status summary for a device in the domain where the device is located. Devices are listed individually, or grouped according to their geolocation, high availability, or cluster status where applicable.
  - View the health summary of the Cloud-Delivered Firewall Management Center and any device when you hover on the hexagon that represents the device health.
  - The dot to the left of a device indicates its health:
    - Green — No alarms.
    - Orange — At least one health warning.
    - Red — At least one critical health alarm.
- The Monitoring navigation pane — Allows you to navigate the device hierarchy. You can view health monitors for individual devices from the navigation pane.

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Step 2 View the status of the Cloud-Delivered Firewall Management Center and its managed devices in the **Health Status** landing page.

a) Hover your pointer over a hexagon to view the health summary of a device. The popup window shows a truncated summary of the top five health alerts. Click the popup to open a detailed view of the health alert summary.

b) In the device list, click **Expand**(>) and **Collapse** (∨) to expand and collapse the list of health alerts for a device.

When you expand the row, all of the health alerts are listed, including the status, title, and details.

Note Health alerts are sorted by their severity level.

Step 3 Use the Monitoring navigation pane to access device-specific health monitors.

a) In the device list, click **Expand** (>) and **Collapse** (∨) to expand and collapse the list of managed devices.

When you expand the row, all of the devices are listed.

b) Click 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 the graph to pin the metrics statistics box, allowing you to explore the metrics and their values in detail. To close the statistics box and move to another point on the graph, simply click the close button.

You have accessed the health monitor and can view the compiled health status for all managed devices. You can interact with the Health Status landing page and use the Monitoring navigation pane to access device-specific health information.

What to do next

- For information about the compiled health status and metrics for any device managed by the Cloud-Delivered Firewall Management Center, refer to Device health monitors, on page 15.

Device health monitors

The device health monitor provides the compiled health status for any device managed by the Cloud-Delivered Firewall Management Center. The device health monitor collects health metrics for Secure Firewall devices in order to predict and respond to system events.

The device health monitor is comprised of these components:

- System Details: Displays information about the managed device, including the installed Secure Firewall version and other deployment details.
- Troubleshooting & Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Health alerts: A health alert monitor provides an at-a-glance view of the health of the device.
- Time range: An adjustable time window to constrain the information that appears in the various device metrics windows.
- Device metrics: An array of key firewall device health metrics categorized across predefined dashboards, including:
  - CPU: CPU utilization, including the CPU usage by process and by physical cores. The Firewall Threat Defense CPU core allocation dashboard shows core assignments for these categories:
    - **Data Plane**: Handles basic network functions, including core packet forwarding and network data processing.
    - **Snort**: Manages intrusion detection and deep packet inspection features.

Note Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. The Snort core allocation shown in the **CPU** dashboard reflects the total number of Snort process threads assigned to Snort cores.

- **System**: Includes all other system processes. While some processes may have dedicated CPU cores allocated, their usage is combined and displayed under the **System** category in the dashboard.

- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- Disk Usage: Device disk usage, including the disk size and disk utilization per partition.
- Critical Processes: Statistics related to managed processes, including process restarts and other select health monitors such as CPU and memory utilization.

Note During a device upgrade or high-availability failover event, the Firewall Threat Defense device may briefly appear as **Offline** in the device's health monitoring dashboard. This happens because health alerts are cleared during the process and are only updated after the process is complete. Wait for the upgrade or failover operation to finish.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

View the device health monitor

The device health monitor compiles device metrics and displays health status and trends using an array of dashboards.

Follow these steps to view the device health monitor.

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (V) to expand and collapse the list of managed devices.

Step 3 View the **Health Alerts** for the device in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the device. The pop-up window shows a truncated summary of the top five health alerts. Click the pop-up to open a detailed view of the health alert summary.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click the **Show the deployment details on top of the graphs** ( ) 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 View the device monitor reports health and performance metrics in several predefined dashboards by default.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

Note

- Multi-threaded Snort processes are allocated across multiple Snort cores to improve device performance. Note that the Snort core allocation shown in the **CPU**dashboard reflects the total number of Snort process threads assigned to Snort cores.
- The **Process Health** widget in Cloud-Delivered Firewall Management Center displays the CPU% usage for each process. The CPU% metric reflects usage relative to the number of cores in the Firewall Threat Defense device, where 100% corresponds to full utilization of one core. For example, on an 8-core device, 200% indicates two fully utilized cores, leaving six available for other processes. To monitor overall system CPU usage, use the **CPU** widget instead of the **Process Health** widget.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 7 Click the **Add New Dashboard**( ) to create a custom correlation dashboard by building your own variable set from the available metric groups.

For more information, refer to Correlating device metrics, on page 18.

Correlating device metrics

The device health monitor includes an array of key Firewall Threat Defense device metrics that serve to predict and respond to system events. The health of any Firewall Threat Defense device can be determined by these reported metrics.

The device monitor reports these metrics in several predefined dashboards by default. These dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics related to the Snort process.
- ASP Drops: Statistics related to the Accelerated Security Path (ASP) performance and behavior.

You can add custom dashboards to correlate metrics that are interrelated. Select from predefined correlation groups, such as CPU and Snort; or create a custom correlation dashboard by building your own variable set from the available metric groups. For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Before you begin

- To view and correlate the time series data (device metrics) in the health monitor dashboard, enable REST API (**Administration** > **Configuration** > **REST API Preferences**).

Note Correlating device metrics is available only for Firewall Threat Defense 6.7 and later versions. Hence, for Firewall Threat Defense versions earlier than 6.7, the health monitor dashboard does not display these metrics even if you enable REST API.

Follow these steps to correlate device metrics:

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Step 2 In the **Devices** list, click **Expand**(>) and **Collapse** (∨) to expand and collapse the list of managed devices, and choose the device for which you want to modify the dashboard.

Use the Monitoring navigation pane to access device-specific health monitors.

Step 3 Click the **Add New Dashboard**(+) icon to add a new dashboard, and specify a name to identify the dashboard.

Step 4 To create a dashboard from a predefined correlation group, click **Add from Predefined Correlations** drop-down, choose the group, and click **Add Dashboard**. 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.

For a comprehensive list of the supported device metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 5 Click **Add Metrics** to add and select metrics from another group, and to remove an individual metric, click the Remove icon on the right side of the item. Click the delete icon to remove the entire group. Click **Add Dashboard** to add the dashboard to the health monitor.

Step 6 You can **Edit** or **Delete** the predefined dashboards and the custom correlation dashboards.

View system details and generate troubleshooting files

Use this task to view system details and generate and download troubleshooting files when you need to monitor device health or initiate troubleshooting procedures on specific devices through the management center.

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Use the **Monitoring** navigation pane to access device-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (V) to expand and collapse the list of managed devices.

Step 3 Click 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 the link expands the collapsed section to see **System Details** and **Troubleshooting & Links** for the device. The system details include:

- **Version:** The Secure Firewall software version.
- **Model:** The device model.
- **Mode:** The firewall mode. The Firewall Threat Defense device supports two firewall modes for regular firewall interfaces: Routed mode and Transparent mode.
- **VDB:** The Cisco vulnerability database (VDB) version.
- **SRU:** The intrusion rule set version.
- **Snort:** The Snort version.

Step 5 You have these choices:

- Generate troubleshooting files; refer to Generate troubleshooting files for specific system functions.
- Generate and download advanced troubleshooting files; refer to Download advanced troubleshooting files.
- Create and modify health policies; refer to Create a health policy, on page 7.

- Create and modify health monitor alerts; refer to Create a health monitor alert, on page 11.

Cluster health monitor

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these 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 Cloud-Delivered Firewall 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.

View the cluster health monitor

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 Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (V) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

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, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click 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 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 View device monitor health and performance metrics in several predefined dashboards.

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. Refer to 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 event views

Health event views enable you to monitor health events logged by the health monitor in the Cloud-Delivered Firewall Management Center. The views are fully customizable, allowing you to analyze health statuses efficiently and access related event data.

You can search event data to easily access other information that may be related to the events you are investigating. If you understand what conditions each health module tests for, you can more effectively configure alerting for health events.

You can perform many of the standard event view functions on the health event view pages.

View health events

Use this task when you need to monitor the health of your appliances and review any incidents or events recorded.

When you access health events from the **Health Monitor** page on your Cloud-Delivered Firewall Management Center, you retrieve all health events for all managed appliances.

Tip You can bookmark this view to allow you to return to the page in the health events workflow containing the **Health Events** table of events. The bookmarked view retrieves events within the time range you are currently viewing, but you can then modify the time range to update the table with more recent information if needed.

Follow these steps to view health events.

Procedure

Choose **Troubleshooting** > **+ Show more** > **Health** > **Events**.

Tip If you are using a custom workflow that does not include the table view of health events, click **(switch workflow)**. On the **Select Workflow** page, click **Health Events**.

Note If no events appear, you may need to adjust the time range.

View the health events table

The Health Events Table provides actionable insight into system health and allows users to organize, modify, and respond to health events.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Events**.

Step 2 You have the following choices:

- Bookmark: To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**, provide a name for the bookmark, and click **Save**.
- Change Workflow: To choose another health events workflow, click **(switch workflow)**.
- Delete Events: To delete health events, check the check box next to the events you want to delete, and click **Delete**. To delete all the events in the current constrained view, click **Delete All**, then confirm you want to delete all the events.
- Generate Reports: Generate a report based on data in the table view — click **Report Designer**.
- Modify: Modify the time and date range for events listed in the Health table view. Note that events that were generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.
- Navigate: Navigate through event view pages.

- Navigate Bookmark: To navigate to the bookmark management page, click **View Bookmarks** from any event view.
- Navigate Other: Navigate to other event tables to view associated events.
- Sort: Sort the events that appear, change what columns display in the table of events, or constrain the events that appear
- View All: To view event details for all events in the view, click **View All**.
- View Details: To view the details associated with a single health event, click the down arrow link on the left side of the event.
- View Multiple: To view event details for multiple health events, choose the check box next to the rows that correspond with the events you want to view details for and then click **View**.
- View Status: To view all events of a particular status, click status in the Status column for an event with that status.
- Search: Search for specific health events or health metrics recorded within your deployment. Customize your search query to analyze health events and other monitored data and save the search constraints for quick access. For more information, refer to Event searches.

Health event fields

The health monitor modules you enable in your health policy run tests to determine appliance health status. When the health status meets specified criteria, a health event is generated. The health event table includes fields that are searchable and viewable for monitoring appliance health.

Table 2: Health event fields

Field	Description
Module Name	Specify the name of the module which generated the health events you want to view. For example, to view events that measure CPU performance, type **CPU**. The search should retrieve applicable CPU Usage and CPU temperature events.
Test Name (Search only)	The name of the health module that generated the event.
Time (Search only)	The timestamp for the health event.
Description	The description of the health module that generated the event. For example, health events generated when a process was unable to execute are labeled **Unable to Execute**.
Value	The value (number of units) of the result obtained by the health test that generated the event. For example, if the Cloud-Delivered Firewall Management Center generates a health event whenever a device it is monitoring is using 80 percent or more of its CPU resources, the value could be a number from 80 to 100.

Field	Description
Units	The units descriptor for the result. You can use the asterisk (*) to create wildcard searches. For example, if the Cloud-Delivered Firewall Management Center generates a health event when a device it is monitoring is using 80 percent or more of its CPU resources, the units descriptor is a percentage sign (%).
Status	The status (Error, Critical, Warning, Normal, Recovered, or Disabled) reported for the appliance.
Device	The appliance where the health event was reported.

Health monitor status categories

The table summarizes the available health monitor status categories, their severity levels, icons, associated colors in pie charts, and descriptions.

Table 3: 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 transitionary state, where, the required data is temporarily unavailable or could not be processed because of changes in the device configuration. Depending on the monitoring cycle, this transitionary state is auto-corrected.
Normal	**Normal** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance.
Recovered	**Recovered** ()	Green	Indicates that all health modules on the appliance are running within the limits configured in the health policy applied to the appliance, including modules that were in a Critical or Warning state.

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.

Device exclusion in health monitoring

A device exclusion is a health monitoring configuration that

- prevents deliberate device or module outages from affecting the overall health summary,
- allows administrators to temporarily disable health monitoring for selected devices, Cloud-Delivered Firewall Management Center where supported, or specific health modules, and
- ensures health events generated during exclusion are marked as disabled and do not affect health calculations.

Device exclusion behavior

In the course of normal network maintenance, you may disable managed devices or make them temporarily unavailable. In Cloud-Delivered Firewall Management Center, you can also exclude Cloud-Delivered Firewall Management Center health monitoring where applicable. Because those outages or expected conditions are deliberate, you may not want their health status to affect the overall health summary.

When you disable health monitoring status, health events are still generated, but they have a disabled status and do not affect the health status for the health monitor. If you remove the device or module from the excluded list, the events that were generated during the exclusion continue to show a status of disabled.

Note If all health modules or a device are added to the exclusion list, the device's health status will be set to disabled. This status refers only to health monitoring and should not be interpreted as the device being offline or disconnected.

You can also disable an individual health module. For example, when you reach the host limit on the Cloud-Delivered Firewall Management Center, you can disable Host Limit status messages. Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Note that on the main Health Monitor page you can distinguish between devices that are excluded if you expand to view the list of devices with a particular status by clicking the arrow in that status row.

Note On Cloud-Delivered Firewall Management Center, Health Monitor exclusion settings are local configuration settings. Therefore, if you exclude a device, then delete it and later re-register it with the Cloud-Delivered Firewall Management Center, the exclusion settings remain persistent. The newly re-registered device remains excluded.

Temporary device exclusion scenario

You can use the health monitor exclude feature to disable health monitoring status reporting for a device, Cloud-Delivered Firewall Management Center where applicable, or specific health modules. For example, if you know that a segment of your network will be unavailable, you can temporarily disable health monitoring for a managed device on that segment to prevent the health status on the Cloud-Delivered Firewall Management Center from displaying a warning or critical state because of the lapsed connection to the device.

To temporarily disable health events from a managed device, or from Cloud-Delivered Firewall Management Center where supported, go to the exclusion configuration page and add the target to the exclusion list. After the setting takes effect, the system no longer considers the excluded device when calculating the overall health status. The Health Monitor Appliance Status Summary lists the device as disabled.

Exclude devices from health monitoring

Exclude devices from health monitoring to disable their events and health status when needed for maintenance or troubleshooting purposes.

You can exclude managed devices individually or by group, model, associated health policy, or other supported scope. In Cloud-Delivered Firewall Management Center exclusion can also apply to Cloud-Delivered Firewall Management Center where supported.

If you need to set the events and health status for an individual managed device to disabled, you can exclude the device from health monitoring. Where supported, you can also exclude Cloud-Delivered Firewall Management Center health monitoring. After the exclusion settings take effect, the device shows as disabled in the Health Monitor Appliance Module Summary, and health events for the device have a status of disabled.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Exclude**.

Step 2 Click **Add Device**.

Step 3 In the **Device Exclusion** dialog box, under **Available Devices**, click **Add** () against the device that you want to exclude from health monitoring.

Step 4 Click **Exclude**.

The selected device is displayed in the exclusion main page.

Step 5 To remove the device from the exclusion list, click **Delete** ().

Step 6 Click **Apply**.

Exclude health policy modules on a device

Excluding health policy modules is useful when certain modules produce irrelevant or temporary health alerts, which you do not want to affect the overall device status. Exclusion settings indicate the number of modules excluded from health monitoring.

Note Excluding health modules for individual interfaces is not supported on devices operating in transparent mode.

Tip Make sure that you keep track of individually excluded modules so you can reactivate them when you need them. You may miss essential warning or critical messages if you accidentally leave a module disabled.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Exclude**.

Step 2 Click **Edit** (✎) next to the Firewall Threat Defense device you want to modify.

Step 3 In the **Exclude Health Modules** dialog box, by default, all the modules of the device are excluded from health monitoring.

Certain modules are applicable to a specific device only; for more information, refer to Health modules, on page 35.

Step 4 To choose modules to be excluded from health monitoring, click the **Enable Module Level Exclusion** link.

The **Exclude Health Modules** dialog box displays all the modules of the device. The modules that are not applicable for the associated health policies are disabled by default. To exclude a module, perform these steps:

a. Click the **Slider** (⬭) button next to the desired module.

b. To specify the duration of the exclusion for the selected modules, from the **Exclude Period** drop-down list, select the duration.

Step 5 (Optional) To stop receiving health status updates for physical interfaces and instead focus on the health status of their subinterfaces, you can disable health monitoring for the physical interfaces while continuing to monitor and receive health alerts for the subinterfaces.

a) If not already enabled, click the toggle button next to the **Interface statistics** module.
b) From the **Exclude Period** drop-down list, choose a duration for which the health status of the physical interfaces is not required.
c) Click the **Exclude specific interfaces** radio button.
d) Check the check boxes next to the physical interfaces that you want to exclude.
e) Selecting a physical interface will automatically exclude both the physical interface and its subinterfaces. Therefore, uncheck the **Subinterfaces** check box next to the corresponding physical interface to continue receiving health alerts for the subinterfaces.

Step 6 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 7 Click **OK**.

Step 8 Click **Apply**.



Note Interface exclusion differs from other module exclusions because you can exclude selected physical interfaces while still monitoring associated subinterfaces, depending on the UI options available in your version.

Exclude interface health monitoring

Use interface exclusion when you want to suppress health status updates for selected physical interfaces without suppressing health monitoring for the entire device. Interface exclusion is useful when interface state is expected during maintenance or when you want to focus on subinterface health instead of physical interface status.

When you configure interface exclusion, select the physical interfaces that you want to exclude from the list of available interfaces. Depending on the available UI options in your version, selecting a physical interface can also exclude its subinterfaces. If subinterface controls are available, you can continue monitoring subinterfaces while excluding only the physical interface.

Use this option to reduce unnecessary health alerts when traffic loss or interface state changes are expected on a physical interface.

Manage expired health monitor exclusions

When monitoring device or module health, some items may have temporary exclusions. Once these exclusion periods lapse, you need to decide whether to renew the exclusion or clear the device/module from exclusion so alerts resume.

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Exclude**.

The **Warning** (⚠) icon is displayed against the device indicating the expiry of the duration of exclusion of the device or the modules from alerting.

Step 2 To renew the exclusion of the device, click **Edit** (✎) next to the device. In the **Exclude Health Modules** dialog box, click the **Renew** link. The exclusion period of the device is extended with the current value.

Step 3 To clear the device from being excluded, click **Delete** (🗑) next to the device, click **Remove the device from exclusion**, and then click **Apply**.

Step 4 To renew or clear the modules from exclusion, click **Edit** (✎) next to the device. In the **Exclude Health Modules** dialog box, click the **Enable Module Level Exclusion** link, and then click the **Renew** or **Clear** link

against the modules. When you click **Renew**, the exclusion period is extended on the module with the current value.

Sending device health metrics to 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 Firewall 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 telemetry streams

OpenConfig uses a subscription-based model where the data collectors query the Firewall 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 Firewall Threat Defense device, it sends a subscribeRequest RPC message to the Firewall 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 Firewall Threat Defense server supports these subscription modes:

- *Once subscription*: The Firewall Threat Defense device sends requested data to the gNMI paths only once.
- *Stream subscription*: The Firewall Threat Defense continuously streams telemetry data according to the triggers specified in the `SubscribeRequest` RPC message.
    - *Sampled subscription*: The Firewall 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 Firewall Threat Defense sends the data whenever the requested values change.

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

You can use these deployment modes for OpenConfig streaming telemetry configuration:

- **DIAL-IN**: In this mode, the gNMI server opens a port on the Firewall 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 Firewall 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. For more information, refer to How to Generate Certificate with Private Key.

Generate certificates and private keys

Certificates and private keys are required for OpenConfig streaming telemetry. You must use the same certificate authority (CA) for all components to ensure secure and consistent authentication.

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

Procedure

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

Example:

```
mkdir keys
```

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

Example:

This 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:

This 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 Firewall 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 Cloud-Delivered Firewall Management Center. The gNMI server that runs on the Firewall 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 Cloud-Delivered Firewall Management Center.

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

Example:

This 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:

This 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

OpenConfig streaming telemetry allows network monitoring systems to collect real-time data from threat defense devices using the gNMI protocol. This configuration supports both DIAL-IN and DIAL-OUT deployment modes for flexible integration with existing network management infrastructure.

Before you begin

- Make sure the Firewall 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 Firewall 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 Firewall Threat Defense device.

Follow these steps to configure OpenConfig streaming telemetry:

Procedure

Step 1 Choose **Troubleshooting** > **+ Show more** > **Health** > **Policies**.

Step 2 Click the **Edit health policy** icon next to the threat defense health policy that you want to modify.

Step 3 Go to the **Settings** tab, and move the **OpenConfig Streaming Telemetry** slider to enable the configuration. This configuration is disabled by default.

Step 4 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.

Step 5 (Optional) Specify the Passphrase if the PRIVATE KEY files are ENCRYPTED.

Step 6 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 Firewall 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 Firewall Threat Defense device.

b. Click **Add Collector** to add more gNMI collectors. You can add a maximum of five collectors.

Step 7 Specify the username and password to validate the gNMI collector.

The Firewall 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 8 Click **Save**.

What to do next

Deploy the health policy to your Firewall Threat Defense device, for the configuration changes to take effect.

Troubleshoot OpenConfig streaming telemetry

This reference provides troubleshooting procedures for OpenConfig streaming telemetry issues including certificate errors, server response failures, and status verification methods.

Certificate signed by unknown authority

- Ensure that you have uploaded the correct certificate to the Cloud-Delivered Firewall 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 Cloud-Delivered Firewall 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 that the IP SAN is correctly specified in the server certificate. If the configuration applies to more than one Firewall 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. This is the sample response that appears when the telegraf 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 gnmi_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 telegraf

When telegraf is not responding, restart the process using this command on the Firewall 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 this command using the Firewall Threat Defense CLI console:

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

This 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:': {}
```

Reference

Look up health modules and their monitored conditions, and review health event options.

Health modules

Describes health modules that monitor system conditions and provide alerts or metrics for managing device health and performance.

Health modules, or *health tests,* test for the conditions that you specify in a health policy.

Health monitoring includes both health alerts or tests and metric or time series data collection. The two types of health modules are alert modules and metrics modules. Alerts modules (sometimes called *legacy* modules) monitor system infrastructure and report only the health status. When the conditions specified in the health policy for these monitored systems are met, these modules raise health alerts. Metrics modules (sometimes called *telegraf* modules) collect statistics (sometimes called *time series data*) that you can view on the health monitoring dashboard. Some alerts are also generated from metric values when they cross the thresholds specified in the health policy. You can create custom dashboards with your preferred health metrics, allowing you to monitor statistics or troubleshoot appliance health issues.

Note The health alerts generated from the Secure Firewall 200 series device is limited to the essential health modules, to optimize performance and ensure effective resource utilization. For more information about the available health modules, refer to Health alerts for Secure Firewall 200 Series device, on page 46.

Table 4: 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	Alert	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.
Certificate Monitoring	Alert	Alerts when service authentication certificates are near expiration or have expired, based on a configurable threshold (in days). This alert helps you to identify certificates that are about to expire and renew them before a service disruption occurs.
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. For standby Firewall Threat Defense devices in a high availability pair, this widget reflects only the statistics of connections replicated from the active device.

Module	Type	Description
CPU Core Usage CPU Usage (per core)	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 system for database integrity issues related to schema or configuration data (*sometimes called EO*). If you receive an alert about a database integrity issue, contact Cisco TAC. Database integrity issues can prevent upgrades.
		• 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.
		Note that the Secure Firewall 200 series device raises alerts only for database integrity issues related to schema or configuration data.

Data Plane CPU Usage CPU Usage Date Plane	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 Memory Usage Data Plane	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. The Disk Usage module sends a health alert if the size of device configuration history files exceeds the allowed limit. For more information, refer to Disk Usage for Device Configuration History Files Health Monitoring Alert for information about troubleshooting scenarios for the disk usage alerts. 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. Use the **Clear disk space** option to free up disk space by removing the temporary files from your threat defense device. For more information, refer to Clear disk space
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 /6100 platform faults and generate health alerts for the faults. A platform fault represents a failure in the Firewall Threat Defense instance or an alarm threshold that has been raised. During the lifecycle of a platform fault, it can change from one state or severity to another. Each fault includes information about the operational state of the affected object at the time the fault was raised. If the fault is transitional and the failure is resolved, then the object transitions to a functional state. For more information, refer to *Cisco Firepower 1000/2100 FXOS Faults and Error Messages Guide*.
Flow Offload Statistics	Metrics	Monitors hardware flow offload.
FXOS Health	Alert	Alerts when the FXOS https service is not running on the device. This will prevent upgrade.
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.
Identity Limits Monitor (supported only on Secure Firewall 200 Series devices)	Alert	Alerts when the device identity-related mappings and user-to-group mappings exceed the normal limit. Device identity-related mappings include user sessions, SGT Exchange Protocol (SXP) mappings, and dynamic object mappings. For more information, refer to Identity-limits-monitor.
Inline Link Mismatch Alarms	Alert	Alerts if inline pair interfaces negotiate different speeds.

Module	Type	Description
Interface Statistics Interface Status	Alert	Determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces. For physical interfaces, the information includes interface name, link state, and bandwidth. For aggregate interfaces, the information includes interface name, number of active links, and total aggregate bandwidth. **Note** • This module also monitors the high availability standby device traffic flow. Though it is known that the standby device would not be receiving any traffic yet, the Cloud-Delivered Firewall Management Center alerts that the interface is not receiving any traffic. The same alerting principle is applied when traffic is not received by some of the subinterfaces on a port channel. • This module displays the traffic rates according to the values from Lina. However, if you use the **show interface** CLI command to know the interface statistics of your device, the input and output rates in the CLI command result can be different from the traffic rates that appear in the **Interface** widget. The sampling intervals of Lina and the Cloud-Delivered Firewall Management Center interface statistics are different. Due to the difference in sampling interval, throughput values in the Cloud-Delivered Firewall Management Center GUI can be different from the throughput values appears in the device CLI result. • Note that traffic rates in the **Interface Traffic Rate** widget (**Insights & Reports** > **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 **Troubleshooting** > **+ Show more** > **Advanced** > **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, it is recommended that you manually set the **Warning Threshold %** value to `50`. This will further ensure that you receive memory alerts for your appliance in time to address the issue. Complex access control policies and rules can command significant resources and negatively affect performance.

Module	Type	Description
Network Card Reset	Alert	Alerts when a network card restarts due to hardware failure.
NTP Statistics	Metrics	Monitors NTP synchronization status. Disabled by default.
Out of Band Configuration Changes Firewall Management Center Access Configuration Changes	Alert	Monitors configuration changes made on the Cloud-Delivered Firewall Management Center directly using the **configure network management-data-interface** command. This module alerts when there is a conflict between the existing Cloud-Delivered Firewall Management Center configuration and the out of band configuration changes made.
Path Monitoring	Metrics	Monitors data path metrics for the interfaces if path monitoring is enabled on the interfaces.
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. This module also monitors metrics for sending advanced logging events and generates the following alerts: • **Advanced logging events to syslog servers were dropped**: This alert appears when syslog messages are dropped due to memory overflow. • **Advanced Logging events failed to transmit to syslog servers**: This alert appears when the syslog messages failed to transmit due to a connection issue with the syslog server or a configuration error. Check your syslog server status and syslog configuration in Cloud-Delivered Firewall Management Center.
Snort CPU Usage CPU Usage Snort	Metrics	This module checks that the average CPU usage of the Snort processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`.

Module	Type	Description
Snort Identity Memory Usage	Alert	Enables you to set a warning threshold for Snort identity processing and alerts when memory usage exceeds the level configured for the module. The **Critical Threshold %** default value is `80`. This health module specifically keeps track of the total space used for the user identity information in Snort. It displays the current memory usage details, the total number of user-to-IP bindings, and user-group mapping details. Snort records these details in a file. If the memory usage file is not available, the Health Alert for this module displays *Waiting for data*. This could happen during a Snort restart due to a new install or a major update, switch from Snort 2 to Snort 3 or back, or major policy deployment. Depending on the health monitoring cycle, and when the file is available, the warning disappears, and the health monitor displays the details for this module with its status turned Green.
Snort Memory Usage Memory Usage Snort	Metrics	This module checks the percentage of allocated memory used by the Snort process and alerts when memory usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`. 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.
System CPU Usage CPU Usage System	Metrics	This module checks that the average CPU usage of all system processes on the device is not overloaded and alerts when CPU usage exceeds the percentages configured for the module. The **Warning Threshold %** default value is `80`. The **Critical Threshold %** default value is `90`. 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.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the Cloud-Delivered Firewall Management Center from the cloud every 30 minutes. This module alerts you if this information has not been updated on the devices within the time period you have specified. Note that the Secure Firewall 200 series device does not maintain a local URL database and supports Cloud Only lookup. Local URL database related alerts are not available for this device type. 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 **Administration** > **Dynamic Attributes Connector** page • Threat Configuration settings related to expiration of cached URLs, including the Cached URLs Expire setting on the **Integrations** > **Cloud Services** page. (Updates to the URL cache are not monitored by this module.) • Communication issues with the Cisco cloud for sending events. Refer to **Cisco Cloud** box on the **Integration** > **Other Integrations > Cloud Services** page. **Note** Threat Intelligence Director updates are included only if TID is configured on your system and you have feeds. By default, this module sends a warning after 1 hour and a critical alert after 24 hours. If this module indicates failure on the Cloud-Delivered Firewall Management Center or on any devices, verify that the Cloud-Delivered Firewall Management Center can reach the devices.
VPN Statistics	Metrics	Monitors site-to-site and remote access VPN tunnels between Firewall Threat Defense devices.
XTLS Counters	Metrics	Monitors XTLS/SSL flows, memory and cache effectiveness. Disabled by default.

Table 5: Management Center Health Modules

Module	Type	Description
Secure Endpoint Status	Alert	The module alerts if the Cloud-Delivered Firewall 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 Cloud-Delivered Firewall 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 Cloud-Delivered Firewall 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.
Dynamic Attributes Connector Status	Alert	This module monitors the operational status of the Dynamic Attributes Connector (CSDAC) service in Cloud-Delivered Firewall Management Center and raises an alert if the CSDAC service is not available. If the CSDAC service is not being used by a related service, try disabling and re-enabling it in the Cloud-Delivered Firewall Management Center UI, and check the health alert again after some time. For more information on how to enable CSDAC service, refer to Enable the dynamic attributes connector. If a related service such as an Azure AD realm is configured, you cannot disable the CSDAC service directly. First delete the Azure AD realm and then disable and re-enable the CSDAC service. For more information, refer to Configure realms and directories.
Firewall Threat Defense HA Status	Alert	Alerts if a threat defense high availability pair is split brain.
Hardware Statistics	Metrics	Monitors Cloud-Delivered Firewall 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).
ISE Connection Monitor	Alert	This module monitors the status of the server connections between the Cisco Identity Services Engine (ISE) and the Cloud-Delivered Firewall Management Center. ISE provides additional user data, device type data, device location data, SGTs (Security Group Tags), and SXP (Security Exchange Protocol) services.

Module	Type	Description
Local Malware Analysis	Alert	This module monitors ClamAV updates for Local Malware Analysis.
Passive Identity Agent Status Monitor	Alert	Displays connection errors between the Cloud-Delivered Firewall Management Center and the machine on which it's installed. The passive identity agent periodically sends updates to the Cloud-Delivered Firewall Management Center. This health alert is displayed if the Cloud-Delivered Firewall Management Center receives an error from the passive identity agent or if the Cloud-Delivered Firewall 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 **Integrations** > **Identity** > **Identity Sources**.
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.
Security Intelligence	Alert	Alerts if Security Intelligence is in use and the Cloud-Delivered Firewall Management Center cannot update a feed, or feed data is corrupt or contains no recognizable IP addresses. Refer to 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.
Talos Connectivity Status	Alert	Monitors connectivity with Talos, required to download URL filtering and event enrichment data.

Module	Type	Description
Threat Data Updates on Devices	Alert	Certain intelligence data and configurations that devices use to detect threats are updated on the Cloud-Delivered Firewall Management Center from the cloud every 30 minutes. This module alerts you if this information has not been updated on the devices within the time period you have specified. **Note** The Secure Firewall 200 does not maintain a local URL database and supports only cloud-based lookups. Local URL database-related alerts are not available for this device type. 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** > **Security Intelligence** > **Network Lists and Feeds** page. • Dynamic analysis settings configured on the **Administration** > **Dynamic Attributes Connector** page. • Threat Configuration settings related to expiration of cached URLs, including the Cached URLs Expire setting on the **Integrations** > **Cloud Services** page. (Updates to the URL cache are not monitored by this module.) • Communication issues with the Cisco cloud for sending events. Refer to 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 Cloud-Delivered Firewall Management Center or on any devices, verify that the Cloud-Delivered Firewall Management Center can reach the devices.
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.

Module	Type	Description
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.
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 Cloud-Delivered Firewall Management Center s in high availability.

Health alerts for Secure Firewall 200 Series device

To optimize performance and ensure effective resource utilization, the health alerts in the Secure Firewall 200 Series device are limited only to the essential health modules. This table lists the health modules in the Secure Firewall 200 Series device, which generate health alerts. You can view all the metrics in the health monitoring dashboard, which is similar to that of other Threat Defense device models.

Table 6: Health Alerts for Secure Firewall 200 Series Device

Health Module	Health Alert
Certificate Monitoring	Alerts when service authentication certificates are nearing expiration or have expired.
Cluster/HA Failure Status	Provide alerts when a device joins, leaves, or is elected as the primary unit.
Database	Provide alerts on database integrity issues related to schema or configuration data.
Disk Usage	Monitors disk usage in the device's hard drive and alerts when usage exceeds the configured thresholds.
Disk Status	Provide alerts on hard disk or RAID controller issues.
Firewall Threat Defense Platform Faults	Monitors platform faults and generates health alerts for them.
FXOS Health	Alerts when the FXOS HTTPS service is not running in the device.
Identity Process	Monitors the health and operation of identity-related services.

Health Module	Health Alert
Identity Limits Monitor	Monitors the device for identity-related mapping limits and generates alerts when these limits are exceeded. The module generates alerts when these two specific types of limits exceed their normal value: • Number of mappings (user sessions, SXP, dynamic object) limit: Cumulative total of device identity-related mappings that include user sessions, SGT Exchange Protocol (SXP) mappings, and dynamic object mappings. • User/Group mappings limit: Total number of user-to-group mappings downloaded to the device. For more information, refer to Identity-limits-monitor.
Inline Link Mismatch Alarms	Provide alerts if inline pair interfaces negotiate different speeds.
Interface Statistics	Determines if the device currently collects traffic and alerts based on the traffic status of physical interfaces and aggregate interfaces.
Out of band Configuration Changes	Alerts when there is a conflict between the existing Cloud-Delivered Firewall Management Center configuration and the out-of-band configuration changes that are made.
Process Status	Provide alerts when processes in the device are terminated outside of the process manager.
Snort Identity Memory Usage	Enables you to set a warning threshold for Snort identity processing, and alerts when memory usage exceeds the level that is configured for the module.
Snort Reconfiguring Detection	Alerts if a device reconfiguration has failed.
Threat Data Updates on Devices	Monitors updates of threat intelligence data and alerts if this information has not been updated in the devices within the time period you have specified.

Identity limits monitor

Secure Firewall 200 series devices enforce two independent download limits. The device supports a cumulative total of 10,000 user sessions, SXP/SGT mappings, endpoint profiles, and dynamic objects. Separately, it supports up to 10,000 user-group mappings. These limits operate independently; reaching the capacity for one category does not block downloads for the other. The device only stops downloading objects for the specific category that has reached its 10,000-object limit.

Consider some of these remediation measures:

- Create identity mapping filters to limit the the networks to which an identity rule applies. Refer to "Create an identity mapping filter".
- Set an ISE network filter to restrict the data that Cisco ISE reports to the Secure Firewall Management Center. For information, refer to "Activate the pxGrid Cloud identity source".
- Configure realm inclusion and exclusion: When you configure realm to control user groups used in access control policies, you can explicitly select which user groups to download (**Included Groups and Users** field) and which to ignore (**Excluded Groups and Users** field). This selection lets you choose which user groups are downloaded. As a result, it limits the size of the user database on the device and optimizes resource utilization.

  For more information, refer to "Realm directory and synchronize fields".



Caution Generating troubleshooting files in a Firewall Threat Defense device is a CPU-intensive task. Because of limited CPU resources in the Secure Firewall 200 Series device, you may observe higher CPU usage and associated health alerts during this process. To prevent any potential traffic disruption, it's recommended to generate troubleshooting files only when the device is not actively handling network traffic.



Note

- The Secure Firewall 200 Series device raises alerts only for database integrity issues-related configuration data (sometimes called EO).
- The Secure Firewall 200 Series device does not maintain a local URL database, and supports cloud-only URL lookups. Local URL database-related alerts are not available for this device type.

system auditing

The appliances that are part of the system generate an audit record for each user interaction with the web interface.

Audit records

Cloud-Delivered Firewall Management Center logs read-only auditing information for user activity. Audit logs appear in a standard event view that allows you to view, sort, and filter audit log messages based on any item in the audit view. You can easily delete audit information, generate reports, and review detailed records of the changes that users make.

Audit record storage and limitations

The audit log stores a maximum of 100,000 entries. When the number of audit log entries exceeds 100,000, the appliance prunes the oldest records from the database to reduce the number to 100,000.

The audit logs do not display the user or the source IP for login errors:

- When wrong password is used, the source IP is not displayed.

- When the user account does not exist, both source IP and the user are not displayed.
- If the attempt for an LDAP user fails, no audit log is triggered.

Related Topics

Guidelines for SSO configuration on Cloud-Delivered Firewall Management Center

Audit log workflow fields

This table describes the audit log fields that can be viewed and searched.

Table 7: Audit log fields

Field	Description
Time	Time and date that the appliance generated the audit record.
User	User name of the user that triggered the audit event.
Subsystem	In a few cases where a menu path is not relevant, the Subsystem field displays only the event type. For example, **Login** classifies user login attempts.
Message	The action the user performed or the button the user clicked on the page. For example, `Page View` signifies that the user simply viewed the page indicated in the Subsystem, while `Save` means that the user clicked the **Save** button on the page. Changes made to the system appear with a **Compare icon** that you can click to see a summary of the changes.
Source IP	IP address associated with the host used by the user. Note: When searching this field you must type a specific IP address; you cannot use IP ranges when searching audit logs.
Domain	The current domain of the user when the audit event was triggered. This field is only present if you have ever configured the Cloud-Delivered Firewall Management Center for multitenancy.
Configuration Change (search only)	Specifies whether to view audit records of configuration changes in the search results. (`yes` or `no`)
Count	The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

The audit events table view

You can customize the layout of the audit events table view or filter events by specific field value.

Use these options to customize the audit events table view:

- You can change the layout of the event view, or constrain the events in the view by a field value.

- To disable a column, click **Close** (✕) in the desired column heading. Click **Apply** in the pop-up window. When you disable a column, it is disabled for the duration of your session (unless you add it back later).

Note When you disable the first column, the Count column is added.

- To hide or show other columns, or to add a disabled column back to the view, select or clear the appropriate check boxes before you click **Apply**.
- Clicking a value within a row in a table view constrains the table view and does not drill down to the next page in the workflow.

Tip Table views always include "Table View" in the page name.