



Troubleshooting

The following topics describe ways to diagnose problems you may encounter:

- [Best Practices for Troubleshooting, on page 1](#)
- [System Messages, on page 1](#)
- [View Basic System Information, on page 4](#)
- [Manage System Messages, on page 5](#)
- [Memory Usage Thresholds for Health Monitor Alerts, on page 8](#)
- [Disk Usage and Drain of Events Health Monitor Alerts, on page 9](#)
- [Health Monitor Reports for Troubleshooting, on page 13](#)
- [General Troubleshooting, on page 15](#)
- [Connection-Based Troubleshooting, on page 15](#)
- [Advanced Troubleshooting for the Secure Firewall Threat Defense Device, on page 16](#)
- [Feature-Specific Troubleshooting, on page 26](#)




Best Practices for Troubleshooting

- Before you make changes to try to fix a problem, generate a troubleshooting file to capture the original problem. See [Health Monitor Reports for Troubleshooting, on page 13](#) and its subsections.
You may need this troubleshooting file if you need to contact Cisco TAC for support.
- Start your investigation by looking at error and warning messages in the Message Center. See [System Messages, on page 1](#)
- Look for applicable Tech Notes and other troubleshooting resources under the "Troubleshoot and Alerts" heading on the product documentation page for your product.

System Messages

When you need to track down problems occurring in the system, the Message Center is the place to start your investigation. This feature allows you to view the messages that the system continually generates about system activities and status.

To open the Message Center, click on the System Status icon, located next to the Deploy menu in the main menu. This icon can take one of the following forms, depending on the system status:

-  — Indicates one or more errors and any number of warnings are present on the system.
-  — Indicates one or more warnings and no errors are present on the system.
-  — Indicates no warnings or errors are present on the system.

If a number is displayed with the icon, it indicates the total current number of error or warning messages.

To close the Message Center, click anywhere outside of it within the web interface.

In addition to the Message Center, the web interface displays pop-up notifications in immediate response to your activities and ongoing system activities. Some pop-up notifications automatically disappear after five seconds, while others are "sticky," meaning they display until you explicitly dismiss them by clicking **Dismiss** (✕). Click the **Dismiss** link at the top of the notifications list to dismiss all notifications at once.



Tip Hovering your cursor over a non-sticky pop-up notification causes it to be sticky.


The system determines which messages it displays to users in pop-up notifications and the Message Center based on their licenses, domains, and access roles.

Message Types

The Message Center displays messages reporting system activities and status organized into three different tabs:

Deployments

This tab displays current status related to configuration deployment for each appliance in your system, grouped by domain. The system reports the following deployment status values on this tab. You can get additional detail about the deployment jobs by clicking **Show History**.

- **Running (Spinning)** — The configuration is in the process of deploying.
- **Success** — The configuration has successfully been deployed.
- **Warning** () — Warning deployment statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Failure** — The configuration has failed to deploy; see [Configuration Changes that Require Deployment](#). Failed deployments contribute to the message count displayed with the **Error System Status icon**.

Upgrades

This tab displays the current status related to software upgrade tasks for the managed devices. The system reports the following upgrade status values on this tab:

- **In progress**—Indicates that the upgrade task is in progress.
- **Completed**—Indicates that the software upgrade task is completed successful.
- **Failed**—Indicates that the software upgrade task has failed to complete.

Health

This tab displays current health status information for each appliance in your system, grouped by domain. Health status is generated by health modules as described in [About Health Monitoring](#). The system reports the following health status values on this tab:

- **Warning** (⚠) — Indicates that warning limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Yellow Triangle** (⚠). Warning statuses contribute to the message count displayed with the **Warning System Status icon**.
- **Critical** (❗) — Indicates that critical limits have been exceeded for a health module on an appliance and the problem has not been corrected. The Health Monitoring page indicates these conditions with a **Critical** (❗) icon. Critical statuses contribute to the message count displayed with the **Error System Status icon**.
- **Error** (✖) — Indicates that a health monitoring module has failed on an appliance and has not been successfully re-run since the failure occurred. The Health Monitoring page indicates these conditions with a **Error icon**. Error statuses contribute to the message count displayed with the **Error System Status icon**.

You can click on links in the Health tab to view related detailed information on the Health Monitoring page. If there are no current health status conditions, the Health tab displays no messages.

Tasks

Certain tasks (such as configuration backups or update installation) can require some time to complete. This tab displays the status of these long-running tasks, and can include tasks initiated by you or, if you have appropriate access, other users of the system. The tab presents messages in reverse chronological order based on the most recent update time for each message. Some task status messages include links to more detailed information about the task in question. The system reports the following task status values on this tab:

- **Waiting()** — Indicates a task that is waiting to run until another in-progress task is complete. This message type displays an updating progress bar.
- **Running** — Indicates a task that is in-progress. This message type displays an updating progress bar.
- **Retrying** — Indicates a task that is automatically retrying. Note that not all tasks are permitted to try again. This message type displays an updating progress bar.
- **Success** — Indicates a task that has completed successfully.
- **Failure** — Indicates a task that did not complete successfully. Failed tasks contribute to the message count displayed with the **Error System Status icon**.
- **Stopped or Suspended** — Indicates a task that was interrupted due to a system update. Stopped tasks cannot be resumed. After normal operations are restored, start the task again.
- **Skipped** — A process in progress prevented the task from starting. Try again to start the task.

New messages appear in this tab as new tasks are started. As tasks complete (status success, failure, or stopped), this tab continues to display messages with final status indicated until you remove them. Cisco recommends you remove messages to reduce clutter in the Tasks tab as well as the message database.

Message Management

From the Message Center you can:

- Choose to display pop-up notifications.
- Display more task status messages from the system database (if any are available that have not been removed).
- Download a report of all the task manager notifications.
- Remove individual task status messages. (This affects all users who can view the removed messages.)
- Remove task status messages in bulk. (This affects all users who can view the removed messages.)



Tip Cisco recommends that you periodically remove accumulated task status messages from the Task tab to reduce clutter in the display as well the database. When the number of messages in the database approaches 100,000, the system automatically deletes task status messages that you have removed.

View Basic System Information

The About page displays information about your appliance, including the model, serial number, and version information for various components of the system. It also includes Cisco copyright information.

Procedure

-
- Step 1** Click **Help** (?) in the toolbar at the top of the page.
- Step 2** Choose **About**.
-

View Appliance Information

Procedure

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

Manage System Messages

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 You have the following choices:

- Click **Deployments** to view messages related to configuration deployments. See [View Deployment Messages, on page 5](#). You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.
- Click **Upgrades** to view messages related to device upgrade tasks. See [Viewing Upgrade Messages](#). See [View Upgrade Messages](#). You must be an Admin user or have **Updates** permission to view these messages.
- Click **Health** to view messages related to the health of your management center and the devices registered to it. See [View Health Messages, on page 7](#). You must be an Admin user or have the **Health** permission to view these messages.

You can navigate to the Health Monitor page by clicking the **Health monitor** link.

- Click **Tasks** to view or manage messages related to long-running tasks. See [View Task Messages, on page 7](#) or [Manage Task Messages, on page 8](#). Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission. You can remove the completed tasks from the notification by clicking the **Remove completed tasks** link.
 - Click the **Download Report** icon to generate a report of all the notifications in the task manager. Choose **Download CSV** or **Download PDF** to download the report.
 - Click **Show Notifications** slider to enable or disable pop-up notification display.
-

View Deployment Messages

You must be an Admin user or have the **Deploy Configuration to Devices** permission to view these messages.

Procedure

Step 1 Click **Notifications** to display the Message Center.

Step 2 Click **Deployments**.

Step 3 You have the following choices:

- Click **total** to view all current deployment statuses.
- Click a status value to view only messages with that deployment status.
- Hover your cursor over the time elapsed indicator for a message (for example, **1m 5s**) to view the elapsed time, and start and stop times for the deployment.

Step 4 Click **show deployment history** to view more detailed information about the deployment jobs.

The Deployment History table lists the deployment jobs in the left column in reverse chronological order.

- a) Select a deployment job.
The table in the right column shows each device that was included in the job, and the deployment status per device.
- b) To view responses from the device, and commands sent to the device during deployment, click download in the **Transcript** column for the device.

The transcript includes the following sections:

- **Snort Apply**—If there are any failures or responses from Snort-related policies, messages appear in this section. Normally, the section is empty.
- **CLI Apply**—This section covers features that are configured using commands sent to the Lina process.
- **Infrastructure Messages**—This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands. Examining these errors can be particularly helpful if you are using FlexConfig policies to configure customized features. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note

There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that the management center sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. The threat defense does not use the security level for anything.

```

===== CLI APPLY =====

FMC >> interface GigabitEthernet0/0
FMC >> nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

View Upgrade Messages

You must be an Admin user or have the **Updates** permission to view these messages.

Procedure

- Step 1** Click **Notifications** to display the Message Center.
- Step 2** Click **Upgrades**.
- Step 3** You can do the following:
 - Click **total** to view all current upgrade tasks.

- Click on a status value to see messages with only that status.
 - Click **Device Management** for more details on the upgrade task.
-

View Health Messages

You must be an Admin user or have the **Health** permission to view these messages.

Procedure

- Step 1** Click **Notifications** to display the Message Center.
- Step 2** Click **Health**.
- Step 3** You have the following choices:
- Click **total** to view all current health statuses. The break-up of the severity, namely, warning, critical, and error is also displayed.
 - Click on a status value to view only messages with that status.
 - Hover your cursor over the relative time indicator for a message (for example, **3 day(s) ago**) to view the time of the most recent update for that message.
 - To view detailed health status information for a particular message, click the message.
 - To view complete health status on the Health Monitoring page, click **Health Monitor**.
-

View Task Messages

Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

- Step 1** Click **Notifications** to display the Message Center.
- Step 2** Click **Tasks**.
- Step 3** You have the following choices:
- Click **total** to view all current task statuses. To view the tasks based on the status, namely, waiting, running, retrying, success, and failures, click on them.
 - Click a status value to view only messages for tasks with the that status.
- Note**
Messages for stopped tasks appear only in the total list of task status messages. You cannot filter on stopped tasks.

- Hover your cursor over the relative time indicator for a message (e.g., **3 day(s) ago**) to view the time of the most recent update for that message.
 - Click any link within a message to view more information about the task.
 - If more task status messages are available for display, click **Fetch more messages** at the bottom of the message list to retrieve them.
-

Manage Task Messages


Everyone can see their own tasks. To see the tasks of other users, you must be an Admin user or have the **View Other Users' Tasks** permission.

Procedure

Step 1 Click System Status to display the Message Center.

Step 2 Click Tasks.

Step 3 You have the following choices:

- If more task status messages are available for display, click on **Fetch more messages** at the bottom of the message list to retrieve them.
 - To remove a single message for a completed task (status stopped, success, or failure), click on **Remove** () next to the message.
 - To remove all messages for all tasks that have completed (status stopped, success, or failure), filter the messages on **total** and click on **Remove all completed tasks**.
 - To remove all messages for all tasks that have completed successfully, filter the messages on **success**, and click on **Remove all successful tasks**.
 - To remove all messages for all tasks that have failed, filter the messages on **failure**, and click on **Remove all failed tasks**.
-

Memory Usage Thresholds for Health Monitor Alerts

The Memory Usage health module compares memory usage on an appliance to the limits configured for the module and alerts when usage exceeds the levels. The module monitors data from managed devices and from the management center itself.

Two configurable thresholds for memory usage, Critical and Warning, can be set as a percentage of memory used. When these thresholds are exceeded, a health alarm is generated with the severity level specified. However, the health alarm system does not calculate these thresholds in an exact manner.

With high memory devices, certain processes are expected to use a larger percentage of total system memory than in a low memory footprint device. The design is to use as much of the physical memory as possible while leaving a small value of memory free for ancillary processes.

Compare two devices, one with 32 GB of memory and one with 4 GB of memory. In the device with 32 GB of memory, 5% of memory (1.6GB) is a much larger value of memory to leave for ancillary processes than in the device with 4 GB of memory (5% of 4GB = 200MB).

To account for the higher percentage use of system memory by certain processes, the management center calculates the total memory to include both total physical memory and total swap memory. Thus the enforced memory threshold for the user-configured threshold input can result in a Health Event where the “Value” column of the event does not match the value that was entered to determine the exceeded threshold.

Beginning with version 7.4.1, the memory usage health module calculates memory usage by taking into account the free memory available, swap memory available, and buffer cache. To avoid premature memory usage health alerts, it is recommended not to exceed the warning and critical alarm thresholds of 88% and 90%, respectively.

The following table shows examples of user-input thresholds and the enforced thresholds, depending on the installed system memory.



Note The values in this table are examples. You can use this information to extrapolate thresholds for devices that do not match the installed RAM shown here, or you can contact Cisco TAC for more precise threshold calculations.

Table 1: Memory Usage Thresholds Based On Installed RAM

User-Input Threshold Value	Enforced Threshold Per Installed Memory (RAM)			
	4 GB	6 GB	32 GB	48 GB
10%	10%	34%	72%	81%
20%	20%	41%	75%	83%
30%	30%	48%	78%	85%
40%	40%	56%	81%	88%
50%	50%	63%	84%	90%
60%	60%	70%	88%	92%
70%	70%	78%	91%	94%
80%	80%	85%	94%	96%
90%	90%	93%	97%	98%
100%	100%	100%	100%	100%

Disk Usage and Drain of Events Health Monitor Alerts

The Disk Usage health module compares disk usage on a managed device’s hard drive and malware storage pack to the limits configured for the module and alerts when usage exceeds the percentages configured for

the module. This module also alerts when the system excessively deletes files in monitored disk usage categories, or when disk usage excluding those categories reaches excessive levels, based on module thresholds.

This topic describes the symptoms and troubleshooting guidelines for two health alerts generated by the Disk Usage health module:

- Frequent Drain of Events
- Drain of Unprocessed Events

The disk manager process manages the disk usage of a device. Each type of file monitored by the disk manager is assigned with a silo. Based on the amount of disk space available on the system the disk manager computes a High Water Mark (HWM) and a Low Water Mark (LWM) for each silo.

To display detailed disk usage information for each part of the system, including silos, LWMs, and HWMs, use the **show disk-manager** command.

Examples

The following is an example of the disk manager information:

```
> show disk-manager
Silo                               Used           Minimum       Maximum
Temporary Files                   0 KB           499.197 MB   1.950 GB
Action Queue Results              0 KB           499.197 MB   1.950 GB
User Identity Events              0 KB           499.197 MB   1.950 GB
UI Caches                          4 KB           1.462 GB     2.925 GB
Backups                           0 KB           3.900 GB     9.750 GB
Updates                           0 KB           5.850 GB     14.625 GB
Other Detection Engine            0 KB           2.925 GB     5.850 GB
Performance Statistics            33 KB          998.395 MB   11.700 GB
Other Events                      0 KB           1.950 GB     3.900 GB
IP Reputation & URL Filtering      0 KB           2.437 GB     4.875 GB
Archives & Cores & File Logs      0 KB           3.900 GB     19.500 GB
Unified Low Priority Events        1.329 MB       4.875 GB     24.375 GB
RNA Events                        0 KB           3.900 GB     15.600 GB
File Capture                      0 KB           9.750 GB     19.500 GB
Unified High Priority Events       0 KB           14.625 GB    34.125 GB
IPS Events                        0 KB           11.700 GB    29.250 GB
```

Health Alert Format

When the Health Monitor process on the management center runs (once every 5 minutes or when a manual run is triggered), the Disk Usage module looks into the diskmanager.log file and, if the correct conditions are met, the health alert is triggered.

The structures of these health alerts are as follows:

- Frequent drain of <SILO NAME>
- Drain of unprocessed events from <SILO NAME>

For example,

- Frequent drain of Low Priority Events
- Drain of unprocessed events from Low Priority Events

It's possible for any silo to generate a *Frequent drain of <SILO NAME>* health alert. However, the most commonly seen are the alerts related to events. Among the event silos, the *Low Priority Events* are often seen because the device generates this type of events frequently.

A *Frequent drain of <SILO NAME>* event has a **Warning** severity level when seen in relation to an event-related silo, because events will be queued to be sent to the management center. For a non-event related silo, such as the *Backups* silo, the alert has a **Critical** severity level because this information is lost.



Important Only event silos generate a *Drain of unprocessed events from <SILO NAME>* health alert. This alert always has a **Critical** severity level.

Additional symptoms besides the alerts can include:

- Slowness on the management center user interface
- Loss of events

Common Troubleshoot Scenarios

A *Frequent drain of <SILO NAME>* event is caused by too much input into the silo for its size. In this case, the disk manager drains (purges) that file at least twice in the last 5-minute interval. In an event type silo, this is typically caused by excessive logging of that event type.

A *Drain of unprocessed events of <SILO NAME>* health alert is caused by a bottleneck in the event processing path.

There are three potential bottlenecks with respect to these Disk Usage alerts:

- Excessive logging — The EventHandler process on threat defense is oversubscribed (it reads slower than what Snort writes).
- Sftunnel bottleneck — The Eventing interface is unstable or oversubscribed.
- SFDataCorrelator bottleneck — The data transmission channel between the management center and the managed device is oversubscribed.

Excessive Logging

One of the most common causes for the health alerts of this type is excessive input. The difference between the Low Water Mark (LWM) and High Water Mark (HWM) gathered from the **show disk-manager** command shows how much space there is available to take on that silo to go from LWM (freshly drained) to the HWM value. If there are frequent drain of events (with or without unprocessed events), review the logging configuration.

- Check for double logging — Double logging scenarios can be identified if you look at the correlator *perfstats* on the management center:

```
admin@FMC:~$ sudo perfstats -Cq < /var/sf/rna/correlator-stats/now
```

- Check logging settings for the ACP — Review the logging settings of the Access Control Policy (ACP). If the logging setting includes both "Beginning" and "End" of connection, modify the setting to log only the end to reduce the number of events.

Communications Bottleneck – Sftunnel

Sftunnel is responsible for encrypted communications between the management center and the managed device. Events are sent over the tunnel to the management center. Connectivity issues and/or instability in the communication channel (sftunnel) between the managed device and the management center can be due to:

- Sftunnel is down or is unstable (flaps).

Ensure that the management center and the managed device have reachability between their management interfaces on TCP port 8305.

The sftunnel process should be stable and should not restart unexpectedly. Verify this by checking the `/var/log/message` file and search for messages that contain the `sftunneld` string.

- Sftunnel is oversubscribed.

Review trend data from the Health Monitor and look for signs of oversubscription of the management center's management interface, which can be a spike in management traffic or a constant oversubscription.

Use as a secondary management interface for eventing. To use this interface, you must configure its IP address and other parameters at the threat defense CLI using the `configure network management-interface` command.

Communications Bottleneck – SFDataCorrelator

The SFDataCorrelator manages data transmission between the management center and the managed device; on the management center, it analyzes binary files created by the system to generate events, connection data, and network maps. The first step is to review the `diskmanager.log` file for important information to be gathered, such as:

- The frequency of the drain.
- The number of files with Unprocessed Events drained.
- The occurrence of the drain with Unprocessed Events.

Each time the disk manager process runs it generates an entry for each of the different silos on its own log file, which is located under `[ngfw]/var/log/diskmanager.log`. Information gathered from the `diskmanager.log` (in CSV format) can be used to help narrow the search for a cause.

Additional troubleshooting steps:

- The command `stats_unified.pl` can help you to determine if the managed device does have some data which must be sent to management center. This condition can happen when the managed device and the management center experience a connectivity issue. The managed device stores the log data on to a hard drive.

```
admin@FMC:~$ sudo stats_unified.pl
```

- The `manage_proc.pl` command can reconfigure the correlator on the management center side.

```
root@FMC:~# manage_procs.pl
```

Before You Contact Cisco TAC

It is highly recommended to collect these items before you contact Cisco TAC:

- Screenshots of the health alert seen.

- Troubleshoot file generated from the management center.
- Troubleshoot file generated from the affected managed device.
- Date and Time when the problem was first seen.
- Information about any recent changes done to the policies (if applicable).
- The output of the stats_unified.pl command as described in [Communications Bottleneck — SFDDataCorrelator, on page 12](#).

Health Monitor Reports for Troubleshooting

In some cases, if you have a problem with your appliance, Support may ask you to supply troubleshooting files to help them diagnose the problem. The system can produce troubleshooting files with information targeted to specific functional areas, as well as advanced troubleshooting files you retrieve in cooperation with Support. You can select any of the options listed in the table below to customize the contents of a troubleshooting file for a specific function.

Note that some options overlap in terms of the data they report, but the troubleshooting files will not contain redundant copies, regardless of what options you select.

Table 2: Selectable Troubleshoot Options

This option...	Reports...
Snort Performance and Configuration	data and configuration settings related to Snort on the appliance
Hardware Performance and Logs	data and logs related to the performance of the appliance hardware
System Configuration, Policy, and Logs	configuration settings, data, and logs related to the current system configuration of the appliance
Detection Configuration, Policy, and Logs	configuration settings, data, and logs related to detection on the appliance
Interface and Network Related Data	configuration settings, data, and logs related to inline sets and network configuration of the appliance
Discovery, Awareness, VDB Data, and Logs	configuration settings, data, and logs related to the current discovery and awareness configuration on the appliance
Upgrade Data and Logs	data and logs related to prior upgrades of the appliance
All Database Data	all database-related data that is included in a troubleshoot report
All Log Data	all logs collected by the appliance database
Network Map Information	current network topology data

Generate Troubleshooting Files for Specific System Functions

You can generate and download customized troubleshooting files that you can send to Support.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, and then click **Generate Troubleshooting Files**.

Note

- When you generate management center troubleshooting files from the Management Center web interface, the file is stored in the management center. Note that only the latest troubleshooting file will be stored in the management center.
- When you generate threat defense troubleshooting files from the Management Center web interface, the file is generated in threat defense and copied to the management center. Note that only the latest threat defense troubleshooting file will be stored in the management center.
- When the troubleshooting files for management center and threat defense are generated from the CLI, all the versions of the troubleshooting files are maintained in management center and threat defense respectively.

Step 2 Choose All Data to generate all possible troubleshooting data, or check individual boxes as described in [View Task Messages, on page 7](#).

Step 3 Click **Generate**.

Step 4 View task messages in the Message Center; see [View Task Messages, on page 7](#).

Step 5 Find the task that corresponds to the troubleshooting files you generated.

Step 6 After the appliance generated the troubleshooting files and the task status changes to `Completed`, click **Click to retrieve generated files**.

Step 7 Follow your browser's prompts to download the file. (The troubleshooting files are downloaded in a single `.tar.gz` file.)

Step 8 Follow the directions from Support to send the troubleshooting files to Cisco.

Download Advanced Troubleshooting Files

You can download troubleshooting files.

Before you begin

You must be an Admin, Maintenance, Security Analyst, or Security Analyst (Read Only) user to perform this task.

Procedure

-
- Step 1** View the health monitor for the appliance.
- Step 2** Choose **System** (⚙) > **Health** > **Monitor**, click the device in the left panel, then **View System & Troubleshoot Details**, and then click **Advanced Troubleshooting**.
- Step 3** In **File Download**, enter the file name supplied by Support.
- Step 4** Click **Download**.
- Step 5** Follow your browser's prompts to download the file.

Note

For managed devices, the system renames the file by prepending the device name to the file name.

- Step 6** Follow the directions from Support to send the troubleshooting files to Cisco.
-

General Troubleshooting

An internal power failure (hardware failure, power surge, and so on) or an external power failure (unplugged cord) can result in an ungraceful shutdown or reboot of the system. This can result in data corruption.

Connection-Based Troubleshooting

Connection-based troubleshooting or debugging provides uniform debugging across modules to collect appropriate logs for a specific connection. It also supports level-based debugging up to seven levels and enables uniform log collection mechanism across modules. Connection-based debugging supports the following:

- A common connection-based debugging subsystem to troubleshoot issues in threat defense
- Uniform format for debug messages across modules
- Persistent debug messages across reboots
- End-to-end debugging across modules based on an existing connection
- Debugging ongoing connections

For more information about the troubleshooting connections, see [Troubleshoot a Connection](#), on page 15.

Troubleshoot a Connection

Procedure

-
- Step 1** Configure a filter to identify a connection using the **debug packet-condition** command.

Example:

```
Debug packet-condition match tcp 192.168.100.177 255.255.255.255 192.168.102.177
255.255.255.255
```

Step 2 Enable debugs for the interested modules and the corresponding levels. Enter the **debug packet** command.

Example:

```
Debug packet acl 5
```

Step 3 Start debugging the packets using the following command:

```
debug packet-start
```

Step 4 Fetch the debug messages from database to analyze the debug messages using the following command:

```
show packet-debugs
```

Step 5 Stop debugging the packets using the following command:

```
debug packet-stop
```

Advanced Troubleshooting for the Secure Firewall Threat Defense Device

You can use Packet Tracer and Packet Capture features to perform an in-depth troubleshooting analysis on a Secure Firewall Threat Defense device. A packet tracer allows a firewall administrator to inject a virtual packet into a security appliance and track the flow from ingress to egress. Along the way, the packet is evaluated against flow and route lookups, ACLs, protocol inspection, NAT, and intrusion detection. This utility is effective because it can simulate real-world traffic by specifying source and destination addresses with protocol and port information. Packet capture is available with the trace option, which provides you with a verdict as to whether the packet is dropped or successful.

For more information about the troubleshooting files, see [Download Advanced Troubleshooting Files, on page 14](#).

Packet Capture Overview

The packet capture feature with trace option allows real packets that are captured on the ingress interface to be traced through the system. The trace information is displayed at a later stage. These packets are not dropped on the egress interface, as they are real data-path traffic. Packet capture for threat defense devices supports troubleshooting and analysis of data packets.

Once the packet is acquired, Snort detects the tracing flag that is enabled in the packet. Snort writes tracer elements, through which the packet traverses. Snort verdict as a result of capturing packets can be one of the following:

Table 3: Snort Verdicts

Verdict	Description
Pass	Allow analyzed packet.

Verdict	Description
Block	Packet not forwarded.
Replace	Packet modified.
AllowFlow	Flow passed without inspection.
BlockFlow	Flow was blocked.
Ignore	Flow was blocked; occurs only for sessions with flows blocked on passive interfaces.
Retry	Flow is stalled, waiting on a enamelware or URL category/reputation query. In the event of a timeout, processing continues with an unknown result: in the case of enamelware, the file is allowed; in the case of URL category/reputation, AC rule lookup continues with an uncategorized and unknown reputation.

Based on the Snort verdict, the packets are dropped or allowed. For example, the packet is dropped if the Snort verdict is **BlockFlow**, and the subsequent packets in the session are dropped before reaching Snort. When the Snort verdict is **Block** or **BlockFlow**, the **Drop Reason** can be one of the following:

Table 4: Drop Reasons

Blocked or Flow Blocked by...	Cause
Snort	Snort is unable to process the packet, erg., snort can't decode packet since it is corrupted or has invalid format.
the App Id preprocessed	App Id module/preprocessed does not block packet by itself; but this may indicate that App Id detection causes other module (erg., firewall) to match a blocking rule.
the SSL preprocessed	There is a block/reset rule in SSL policy to match the traffic.
the firewall	There is a block/reset rule in firewall policy to match the traffic.
the captive portal preprocessed	There is a block/reset rule using the identity policy to match the traffic.
the safe search preprocessed	There is a block/reset rule using the safe-search feature in firewall policy to match the traffic.
the SI preprocessed	There is a block/reset rule a in Security Intelligence tab of AC Policy to block the traffic, erg., DNS or URL SI rule.

Blocked or Flow Blocked by...	Cause
the filterer preprocessed	There is a block/reset rule in filterer tab of AC policy to match the traffic.
the stream preprocessed	There is an intrusion rule blocking/reset stream connection, erg., blocking when TCP normalization error.
the session preprocessed	This session was already blocked earlier by some other module, so session preprocessed is blocking further packets of the same session.
the fragmentation preprocessed	Blocking because earlier fragment of the data is blocked.
the snort response preprocessed	There is a react snort rule, erg., sending a response page on a particular HTTP traffic.
the snort response preprocessed	There is a snort rule to send custom response on packets matching conditions.
the reputation preprocessed	Packet matches a reputation rule, erg., blocking a given IP address.
the x-Link2State preprocessed	Blocking due to buffer overflow vulnerability detected in SMTP.
back orifice preprocessed	Blocking due to detection of back orifice data.
the SMB preprocessed	There is a snort rule to block SMB traffic.
the file process preprocessed	There is file policy that blocks a file, erg., enamelware blocking.
the IPS preprocessed	There is a snort rule using IPS, erg., rate filtering.

The packet capture feature allows you to capture and download packets that are stored in the system memory. However, the buffer size is limited to 32 MB due to memory constraint. Systems capable of handling very high volume of packet captures exceed the maximum buffer size quickly and thereby the necessity of increasing the packet capture limit is required. It is achieved by using the secondary memory (by creating a file to write the capture data). The maximum supported file size is 10 GB.

When the **file-size** is configured, the captured data gets stored to the file and the file name is assigned based on the capture name **recapture**.

The **file-size** option is used when you need to capture packets with the size limit more than 32 MB.

For information, see the [Cisco Secure Firewall Threat Defense Command Reference](#).

Use the Capture Trace

Packet capture is a utility that provides a live snapshot of network traffic passing the specified interface of a device based on a defined criteria. This process continues to capture the packets as long as it has not paused, or the allocated memory has not exhausted.

Packet capture data includes information from Snort and preprocessors about verdicts and actions the system takes while processing a packet. Multiple packet captures are possible at a time. You can configure the system to modify, delete, clear, and save captures.



Note Capturing packet data requires packet copy. This operation may cause delays while processing packets and may also degrade the packet throughput. We recommend that you use packet filters to capture specific traffic data.

Before you begin

To use the packet capture tool on Secure Firewall Threat Defense devices, you must be an Admin or Maintenance user.

Procedure

-
- Step 1** On the management center, choose **Devices > Packet Capture**.
 - Step 2** Select a device.
 - Step 3** Click **Add Capture**.
 - Step 4** Enter the **Name** for capturing the trace.
 - Step 5** Select the **Interface** for the capturing the trace.
 - Step 6** Specify **Match Criteria** details:
 - a) Select the **Protocol**.
 - b) Enter the IP address for the **Source Host**.
 - c) Enter the IP address for the **Destination Host**.
 - d) (Optional) Check **SGT number** check box, and enter a Security Group Tag (SGT).
 - Step 7** Specify **Buffer** details:
 - a) (Optional) Enter a maximum **Packet Size**.
 - b) (Optional) Enter a minimum **Buffer Size**.
 - c) Select either **Continuous Capture** if you want the traffic captured without interruption, or **Stop when full** if you want the capture to stop when the maximum buffer size is reached.

Note
If **Continues Capture** is enabled, and when the allocated memory is full, the oldest captured packets in the memory is overwritten by the new captured packets.

 - d) Check the check box of **Trace**, if you want to capture the details for each packet.
 - e) Enter the value in **Trace Count** field. Default value is 128. You can enter values in the range of 1-1000.
 - Step 8** Click **Save**.

The packet capture screen displays the packet capture details and its status. To have the packet capture page auto refreshed, check the **Enable Auto Refresh** check box and enter the auto refresh interval in seconds.

You can do the following on the packet capture:

- **Edit** (✎) to modify the capture criteria.
- **Delete** (🗑) to delete the packet capture and the captured packets.
- **Clear** (🗑) to erase all the captured packets from a Packet Capture. To erase the captured packets from all of the existing packet captures, click **Clear All Packets**.
- **Pause** (⏸) to temporarily halt capturing packets.
- **Save** (💾) to save a copy of captured packets on a local machine in ASCII or PCAP format. Choose the required format option, and click **Save**. The saved packet capture is downloaded to your local machine.
- To view the details of the packets being captured, click the required capture row.

Packet Tracer Overview

The Packet Tracer tool allows you to test policy configuration by modeling a packet with source and destination addresses, and protocol characteristics. The trace does a policy lookup to validate if the packet will be permitted or denied access based on the configured access rules, NAT, routing, access policies and rate-limiting policies. The packet flow is simulated based on interfaces, source address, destination address, ports, and protocols. This method of testing the packets allows you to verify the effectiveness of your policies and test whether the types of traffic you want to allow or deny are handled as required.

Besides verifying your configuration, you can use the tracer to debug unexpected behavior, such as packets being denied access when they should be allowed. To simulate a packet fully, the packet tracer traces the data path—slow-path and fast-path modules. Initially, processing was transacted on per-session and per-packet basis. The Packet Tracer tool and Capture with Trace feature log the tracing data on per packet basis when the firewall processes packets per session or per packet.

PCAP File

You can initiate a packet tracer using a PCAP file, and that has a complete flow. Currently, only a PCAP with a single TCP/UDP-based flow and a maximum of 100 packets is supported. The packet tracer tool reads the PCAP file, and initializes the state for client and server replay entities. The tool starts replaying the packets in a synchronized manner by collecting and storing the trace output of each packet within the PCAP for subsequent processing and display.

PCAP Replay

Packet replay is executed by the sequence of the packet in the PCAP file, and interferences, if any, to the replay activity terminates it and concludes the replay. The trace output is generated for all the packets in the PCAP on the specified ingress interface and egress interface, thereby providing a complete context for flow evaluation.

PCAP replay is not supported for some features that dynamically modify the packet during replay, such as IPsec, VPN, SSL, HTTPs decryption, NAT, and so on.

Use the Packet Tracer

To use a packet tracer on Secure Firewall Threat Defense devices, you must be an Admin or Maintenance user.

Procedure

Step 1 On the management center, choose **Devices > Packet Tracer**.

Step 2 From the **Select Device** drop-down, choose the device on which you want to run the trace.

Step 3 From the **Ingress Interface** drop-down, choose the ingress interface for the packet trace.

Note

Do not select VTI. VTI as ingress interface is not supported for packet tracer.

Step 4 To use a PCAP replay in the packet-tracer, do the following:

a) Click **Select a PCAP File**.

b) To upload a new PCAP file, click **Upload a PCAP file**. To reuse a recently uploaded file, click the file from the list.

Note

Only .pcap and .pcapng file formats are supported. The PCAP file can contain only a single TCP/UDP based flow with a maximum of 100 packets. The maximum character limit on the PCAP file name (including the file formats) is 64.

c) In the **Upload PCAP** box, you can either drag a PCAP file or click in the box to browse and upload the file. On selecting the file, the upload process starts automatically.

d) Go to this [Step 13](#).

Step 5 To define the trace parameters, from the **Protocol** drop-down menu, select the packet type for the trace, and specify the protocol characteristics:

- **ICMP**—Enter the ICMP type, ICMP code (0-255), and optionally, the ICMP identifier.
- **TCP/UDP/SCTP**—Enter the source and destination port numbers.
- **GRE/IPIP**—Enter the protocol number, 0-255.
- **ESP**—Enter the SPI value for Source, 0-4294967295.
- **RAWIP**—Enter the port number, 0-255.

Step 6 Select the **Source Type** for the packet trace, and enter the source IP address.

Source and destination types include IPv4, IPv6, and fully-qualified domain names (FQDN). You can specify IPv4 or IPv6 addresses and FQDN, if you use Cisco TrustSec.

Step 7 Select the **Source Port** for the packet trace.

Step 8 Select the **Destination** type for the packet trace, and enter the destination IP address.

Destination type options vary depending on the source type that you select.

Step 9 Select the **Destination Port** for the packet trace.

Step 10 Optionally, if you want to trace a packet where the Security Group Tag (SGT) value is embedded in the Layer 2 CMD header (TrustSec), enter a valid **SGT number**.

Step 11 If you want packet tracer to enter a parent interface, which is later redirected to a sub-interface, enter a **VLAN ID**.

This value is optional for non-sub-interfaces only, since all the interface types can be configured on a sub-interface.

Step 12 Specify a **Destination MAC Address** for the packet trace.

If the Secure Firewall Threat Defense device is running in transparent firewall mode, and the ingress interface is VTEP, **Destination MAC Address** is required if you enter a value in **VLAN ID**. Whereas if the interface is a bridge group member, **Destination MAC Address** is optional if you enter a **VLAN ID** value, but required if you do not enter a **VLAN ID** value.

If the Secure Firewall Threat Defense is running in routed firewall mode, **VLAN ID** and **Destination MAC Address** are optional if the input interface is a bridge group member.

Step 13 (Optional) If you want the packet-tracer to ignore the security checks on the simulated packet, click **Bypass all security checks for simulated packet**. This enables packet-tracer to continue with tracing of packet through the system which, otherwise would have been dropped.

Step 14 (Optional) To allow the packet to be sent out through the egress interface from the device, click **Allow simulated packet to transmit from device**.

Step 15 (Optional) If you want the packet-tracer to consider the injected packet as an IPsec/SSL VPN decrypted packet, click **Treat simulated packet as IPsec/SSL VPN decrypt**.


Step 16 Click **Trace**.

The **Trace Result** displays the results for each phase that the PCAP packets has traveled through the system. Click on the individual packet to view the traces results for the packet. You can do the following:

- Copy (📄) the trace results to clipboard.
- Expand or collapse (⌵) the displayed results.
- Maximize (⌵) the trace result screen.

The time elapsed information that is useful to gauge the processing efforts are displayed for each phase. The total time that is taken for the entire flow of packets flowing from an ingress to an egress interface is also displayed in the results section.

The **Trace History** pane displays the stored trace details for each PCAP trace. It can store up to 100 packet traces. You can select a saved trace and run the packet trace activity again. You can do the following:

- Search for a trace using any of the trace parameters.
- Disable saving of the trace to history using the  button.
- Delete specific trace results.
- Clear all the traces.

CPU Profiler Overview

The CPU profiler gathers data on CPU usage by individual modules or inspectors of Snort 3, while packets are being processed within a specific timeframe. It provides information on the amount of CPU time that each module consumes in relation to the total CPU usage of the Snort 3 process. Using the CPU profiler does not require reloading the configuration or restarting Snort 3, thus minimizing downtime. The profiling result

displays the processing time taken by all the modules during the last profiling session. CPU profiling results are saved in JSON format in the threat defense device and synchronized in the management center.



Caution CPU profiling might cause system performance degrade of approximately 3 percent.

Use the CPU Profiler

Before you begin

You must have Version 7.6 or later devices with Snort 3 to use CPU Profiling.

Procedure

-
- Step 1** From the management center, choose **Devices > Snort 3 Profiling**.
- Step 2** Click the **CPU Profiling** tab.
- Step 3** From the **Select device for CPU Profiling** drop-down list, choose a device for CPU profiling.

Note

You can run multiple profiling sessions simultaneously on different devices.

- Step 4** To start a CPU profiling session, click **Start**. (The session automatically stops after 120 minutes.)
- You can stop the profiling session at any time by clicking **Stop**. However, canceling it before the scheduled 120 minutes might not give precise results.

Note

While the CPU profiling session is in progress, a task is created. Click **Notifications > Tasks** to view the details.

The latest profiling result is automatically displayed in the **CPU Profiling Results** section. The table contains statistics for the processing time taken by all the Snort 3 modules or inspectors during the last profiling session. You can view the CPU profiler output in tabular format:

- **Module:** Name of the module or inspector.
- **% Total of CPU Time:** Percentage of time taken by the module relative to the overall time taken by Snort 3 for processing traffic. If this value is considerably higher than that of other modules, it means the module is contributing more to the unsatisfactory performance of Snort 3.
- **Time (µs):** Total time (in microseconds) taken by each module.
- **Avg/Check:** Average time taken by the module each time the module was invoked.
- **% Caller:** Time taken by the submodule (if configured) relative to the main module. This value is used for debugging purposes.

- Step 5** (Optional) Click **Download Snapshot** to download the profiling result. The downloaded file is in CSV format and contains all the fields from the profiling results page.
- Step 6** (Optional) Click the **Filter by % of Snort time** toggle button to filter out modules whose execution consumed more than *n%* of the profiling time.

Step 7 (Optional) Use the **Search** field to search through the fields in the **CPU Profiling Results** table.

Note

Except for **Module**, you can click any of the other column headers to sort the data.

Step 8 (Optional) Click the **Profiling History** section (collapsible panel on the left) to expand it and view a set of cards representing the previous profiling sessions for a chosen device. When you click a card from history, the details are displayed in the **CPU Profiling Results** section.



Note

If you initiate a deployment while CPU profiling is running, the profiling session is automatically terminated to accommodate the deployment, except for deployments resulting from changes to the access control policy rules and security intelligence. You must run the CPU profiling for the device again.

Rule Profiler Overview

The Snort 3 rule profiler gathers data on the time spent processing a set of Snort 3 intrusion rules, highlights potential issues, and shows the rules with unsatisfactory performance. Within the Snort 3 engine, the rule profiler inspects traffic using the Snort 3 intrusion detection mechanism. The rule profiler displays the top 100 IPS rules that required the maximum time to check. Triggering the rule profiler does not require Snort 3 reload or restart. The rule profiling results are saved in JSON format in the threat defense device and synchronized in the management center.

Use the Rule Profiler

Before you begin

You must have Version 7.6 or later devices with Snort 3 to use Rule Profiling.

Procedure

Step 1 From the management center, choose **Devices > Snort 3 Profiling**.

Step 2 Click the **Rule Profiling** tab.

Step 3 From the **Select device for Rule Profiling** drop-down list, choose a device for rule profiling.

Note

You can run multiple profiling sessions simultaneously on different devices.

Step 4 To start a rule profiling session, click **Start**. (The session automatically stops after 120 minutes.)

You can stop the profiling session at any time by clicking **Stop**. However, canceling it before the scheduled 120 minutes might not give precise results.

Note

While the rule profiling session is in progress, a task is created. Click **Notifications > Tasks** to view the details.

The latest profiling result is automatically displayed in the **Rule Profiling Results** section. The table contains statistics for the rules that took the most time to process (sorted in descending order) by the total time (in microseconds [μ s]). You can view the IPS rule profiler output in tabular format:

- **% of Snort Time:** Time spent for processing the rule, relative to the total time of the Snort 3 operation.
- **Rev:** Revision number of the rule.
- **Checks:** Number of times the IPS rule was executed.
- **Matches:** Number of times the IPS rule fully matched in the traffic.
- **Alerts:** Number of times the IPS rule triggered an IPS alert.
- **Time:** Time (in microseconds) spent by Snort in checking the IPS rule.
- **Avg/Check:** Average time spent by Snort on a single check of the rule.
- **Avg/Match:** Average time spent by Snort on a single check, which resulted in a match.
- **Avg/Non-Match:** Average time spent by Snort on a single check, which did not result in a match.
- **Timeouts:** Number of times the rule exceeded the Rule Handling - Threshold configured in the Latency-Based Performance Settings of the access control policy.
- **Suspends:** Number of times the rule was suspended due to consecutive threshold violations.

- Step 5** (Optional) Click **Download Snapshot** to download the profiling result. The downloaded file is in CSV format and contains all the fields from the profiling results page.
- Step 6** (Optional) Click the **Filter by % of Snort time** toggle button to filter out rules whose execution took more than $n\%$ of the profiling time. In general, a rule is considered as performing unsatisfactorily if it consumes 0.2% or more of Snort's overall processing time.
- Step 7** (Optional) Use the **Search** field to search through the fields in the **Rule Profiling Results** table.
- Step 8** (Optional) Click the **Rule Profiling History** section (collapsible panel on the left) to expand it and view a set of cards representing the previous profiling sessions for a chosen device. When you click a card from the history, the details are displayed in the **Rule Profiling Results** section.



Note If you initiate a deployment while rule profiling is running, the profiling session is automatically terminated to accommodate the deployment, except for deployments resulting from changes to the access control policy rules and security intelligence. You must run the rule profiling for the device again.

How to use the Threat Defense Diagnostic CLI from the Web Interface

You can execute the selected threat defense diagnostic CLI commands from the management center. The commands **ping** (except **ping system**), **traceroute**, and select **show** commands run in the diagnostic CLI rather than the regular CLI.

When you run the **show** commands, if the message `Unable to execute the command properly. Please see logs for more details` is displayed, it means that the command is not valid in the diagnostic CLI. For example, **show access-list** works, but this message will be displayed if you enter **show**

access-control-policy. To use non-diagnostic commands, use SSH to log in to a device outside management center.

For more information on the threat defense CLI, see the [Cisco Secure Firewall Threat Defense Command Reference](#).

Before you begin

- You must be an Admin, Maintenance, or Security Analyst to use the diagnostic CLI.
- The purpose of diagnostic CLI is to enable the quick use of a few commands that are useful in troubleshooting a device. For access to the full range of commands, open an SSH session directly with the device.
- In deployments using management center high availability, diagnostic CLI is available only in the active management center.

Procedure

Step 1 Choose **Devices > Threat Defense CLI**.

You can also access the CLI tool through the health monitor for the device (**System** (⚙️) > **Health > Monitor**). From there, you can select the device, click the **View System and Troubleshoot Details** link, click **Advanced Troubleshooting**, then click **Threat Defense CLI** on that page.

Step 2 From the **Device** drop-down list, choose the device on which to execute the diagnostic command.

Step 3 From the **Command** drop-down list, choose the command that you want to execute.

Step 4 Enter the command parameters in the **Parameters** field.

See the [Cisco Secure Firewall Threat Defense Command Reference](#) for the valid parameters.

For example, to execute **show access-list** command, choose **show** from the **Command** drop-down list, then enter **access-list** in the **Parameters** field.

Note

Do not type the full command in the **Parameters** field. Type only the relevant keywords.

Step 5 Click **Execute** to view the command output.

If the message `Unable to execute the command properly. Please see logs for more details.` is displayed, examine the parameters closely. There might be syntax errors.

This message can also mean that the command you are trying to execute is not a valid command within the context of the diagnostic CLI (which you have accessed from the device using the **system support diagnostic-cli** command). Log in to the device using SSH to use these commands.

Feature-Specific Troubleshooting

See the following table for feature-specific troubleshooting tips and techniques.

Table 5: Feature-Specific Troubleshooting Topics

Feature	Relevant Troubleshooting Information
Application control	<i>Best Practices for Application Control</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
LDAP external authentication	Troubleshooting LDAP Authentication Connections
Licensing	Troubleshooting Smart Licensing
User rule conditions	<i>Troubleshoot User Control</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
User identity sources	For troubleshooting information on ISE/ISE-PIC, TS Agent Identity Source, Captive Portal Identity Source, and Remote Access VPN Identity Source, see the corresponding sections in the Cisco Secure Firewall Management Center Device Configuration Guide Troubleshooting LDAP Authentication Connections
URL filtering	<i>Troubleshoot URL Filtering</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Realms and user data downloads	<i>Troubleshoot Realms and User Downloads</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Network discovery	<i>Troubleshooting Your Network Discovery Strategy</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Custom Security Group Tag (SGT) rule conditions	<i>Custom SGT Rule Conditions</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
SSL rules	Chapter on SSL rules in the Cisco Secure Firewall Device Manager Configuration Guide
Cisco Threat Intelligence Director (TID)	<i>Troubleshoot Secure Firewall threat intelligence director</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Secure Firewall Threat Defense syslog	<i>About Configuring Syslog</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Intrusion performance statistics	<i>Intrusion Performance Statistic Logging Configuration</i> in the Cisco Secure Firewall Management Center Device Configuration Guide
Connection-based Troubleshooting	Connection-Based Troubleshooting, on page 15

