

System Management

The following topics explain how to perform system management tasks such as updating system databases and backing up and restoring the system.

- Installing Software Updates, on page 1
- Backing Up and Restoring the System, on page 10
- Auditing and Change Management, on page 16
- Exporting the Device Configuration, on page 22
- Managing Firewall Device Manager and Firewall Threat Defense User Access, on page 22
- Rebooting or Shutting Down the System, on page 29
- Troubleshooting the System, on page 30
- Hardware Management Tasks, on page 41

Installing Software Updates

You can install updates to the system databases and to the system software. The following topics explain how to install these updates.

Updating System Databases and Feeds

The system uses several databases and Security Intelligence feeds to provide advanced services. Cisco provides updates to these databases and feeds so that your security policies use the latest information available.

Overview of System Database and Feed Updates

Firewall Threat Defense uses the following databases and feeds to provide advanced services.

Intrusion rules

As new vulnerabilities become known, the Cisco Talos Intelligence Group (Talos) releases intrusion rule updates that you can import. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules.

Intrusion rule updates provide new and updated intrusion rules and preprocessor rules, modified states for existing rules, and modified default intrusion policy settings. Rule updates may also delete rules, provide new rule categories and default variables, and modify default variable values.

For changes made by an intrusion rule update to take effect, you must redeploy the configuration.

Intrusion rule updates may be large, so import rules during periods of low network use. On slow networks, an update attempt might fail, and you will need to retry.

Geolocation database (GeoDB)

The Cisco Geolocation Database (GeoDB) is a database of geographical data (such as country, city, coordinates) associated with routable IP addresses.

GeoDB updates provide updated information on physical locations that your system can associate with detected routable IP addresses. You can use geolocation data as a condition in access control rules.

The time needed to update the GeoDB depends on your appliance; the installation usually takes 30 to 40 minutes. Although a GeoDB update does not interrupt any other system functions (including the ongoing collection of geolocation information), the update does consume system resources while it completes. Consider this when planning your updates.

Vulnerability database (VDB)

The Cisco Vulnerability Database (VDB) is a database of known vulnerabilities to which hosts may be susceptible, as well as fingerprints for operating systems, clients, and applications. The firewall system correlates the fingerprints with the vulnerabilities to help you determine whether a particular host increases your risk of network compromise. The Cisco Talos Intelligence Group (Talos) issues periodic updates to the VDB.

The time it takes to update vulnerability mappings depends on the number of hosts in your network map. You may want to schedule the update during low system usage times to minimize the impact of any system downtime. As a rule of thumb, divide the number of hosts on your network by 1000 to determine the approximate number of minutes to perform the update.

After you update the VDB, you must redeploy configurations before updated application detectors and operating system fingerprints can take effect.

Cisco Talos Intelligence Group (Talos) Security Intelligence Feeds

Talos provides access to regularly updated intelligence feeds for use in Security Intelligence policies. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations. These feeds contain addresses and URLs for known threats. When the system updates a feed, you do not have to redeploy. The new lists are used for evaluating subsequent connections.

URL Category/Reputation Database

The system obtains the URL category and reputation database from Cisco Collective Security Intelligence (CSI). If you configure URL filtering access control rules that filter on category and reputation, requested URLs are matched against the database. You can configure database updates and some other URL filtering preferences on **System Settings** > **URL Filtering Preferences**. You cannot manage URL category/reputation database updates the same way you manage updates for the other system databases.

Updating System Databases

You can manually retrieve and apply system database updates at your convenience. Updates are retrieved from the Cisco support site. Thus, there must be a path to the internet from the system's management address.

Alternatively, you can retrieve the update packages from the internet yourself, then upload them from your workstation. This method is primarily meant for air-gapped networks, where there is no path to the internet for retrieving the updates from Cisco. Download the updates from software.cisco.com from the same folders where you would download system software upgrades.

Note In May 2022 we split the GeoDB into two packages: a country code package that maps IP addresses to countries/continents, and an IP package that contains additional contextual data associated with routable IP addresses. The Firewall Device Manager does not and has never used the information in the IP package. This split saves significant disk space in locally managed Firewall Threat Defense deployments. If you are getting the GeoDB from Cisco yourself, make sure you get the country code package, which has the same file name as the old all-in-one package: Cisco_GEODB_Update-*date-build*.

You can also set up a regular schedule to retrieve and apply database updates. Because these updates can be large, schedule them for times of low network activity.

Note While a database update is in progress, you might find that the user interface is sluggish to respond to your actions.

Before you begin

To avoid any potential impact to pending changes, deploy the configuration to the device before manually updating these databases.

Please be aware that VDB and URL category updates can remove applications or categories. You need to update any access control or SSL decryption rules that use these deprecated items before you can deploy changes.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Updates summary.

This opens the Updates page. Information on the page shows the current version for each database and the last date and time each database was updated.

Step 2 To manually update a database, click one of the following options in the section for that database:

- **Update from Cloud**—To have the Firewall Device Manager retrieve the update package from Cisco. This is the easiest and most reliable method, but there must be a path to the internet to use it.
- **(down arrow)** > ***option***—To select the update package from your workstation or a drive connected to your workstation. The option will be one of the following:
  - **Select File**—Select a VDB or Geolocation package. You must select a newer version than the one installed. You cannot downgrade these packages.
  - **Update to Newer Version**—Select an Intrusion Rule package that is newer than the one that is currently installed.
  - **Downgrade to Older Version**—Select an Intrusion Rule package that is older than the one that is currently installed.

Rule and VDB updates require a configuration deployment to make them active. When you update from the cloud, you are asked whether you want to deploy now; click **Yes**. If you click **No**, remember to initiate a deployment job at your earliest convenience.

If you upload your own file, you must always deploy the changes manually.

Note
When manually uploading an intrusion rule package, make sure you upload the right package type for your Snort version, SRU for Snort 2, LSP for Snort 3. You can upload a package for the non-active Snort version, but it will not be activated unless you switch versions. For information on switching Snort versions, see Switching Between Snort 2 and Snort 3.

Step 3 (Optional) To set up a regular database update schedule:

a) Click the **Configure** link in the section for the desired database. If there is already a schedule, click **Edit**.

The update schedules for the databases are separate. You must define the schedules separately.

b) Set the update start time:

- The frequency of the update (Daily, Weekly, or Monthly).
- For weekly or monthly, the days of the week or month you want the update to occur.
- The time you want the update to start. The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

c) For Rule or VDB updates, select the **Automatically Deploy the Update** checkbox if you want the system to deploy the configuration whenever the database is updated.

The update is not effective until it is deployed. The automatic deployment also deploys any other configuration changes that are not yet deployed.

d) Click **Save**.

Note
If you want to remove a recurring schedule, click the **Edit** link to open the scheduling dialog box, then click the **Remove** button.

Updating Cisco Security Intelligence Feeds

Cisco Talos Intelligence Group (Talos) provides access to regularly updated Security Intelligence feeds. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations. When the system updates a feed, you do not have to redeploy. The new lists are used for evaluating subsequent connections.

If you want strict control over when the system updates a feed from the Internet, you can disable automatic updates for that feed. However, automatic updates ensure the most up-to-date, relevant data.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Updates summary.

This opens the Updates page. Information on the page shows the current version for the **Security Intelligence Feeds** and the last date and time the feeds were updated.

Step 2 To manually update the feeds, click **Update Now** in the **Security Intelligence Feeds** group.

If you manually update the feeds on one unit in a high availability group, you need to also manually update them on the other unit to ensure consistency.

Step 3 (Optional.) To configure a regular update frequency:

a) Click the **Configure** link in the section for Cisco Feeds. If there is already a schedule, click **Edit**.

b) Select the desired frequency.

The default is **Hourly**. You can also set a **Daily** update (specify the time of day) or **Weekly** update (select the days of the week and time of day). The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

Click **Delete** to prevent automatic updates.

c) Click **OK**.

Upgrading Firewall Threat Defense

Use this procedure to upgrade a standalone Firewall Threat Defense device. If you need to update FXOS, do that first. To upgrade high availability threat defense, see Upgrading High Availability Firewall Threat Defense.



Caution Traffic is dropped while you upgrade. Even if the system appears inactive or unresponsive, do not manually reboot or shut down during upgrade; you could place the system in an unusable state and require a reimage. You can manually cancel failed or in-progress major and maintenance upgrades, and retry failed upgrades. If you continue to have issues, contact Cisco TAC.

For details on these and other issues you may encounter during upgrade, see Troubleshooting Threat Defense Upgrades, on page 9.

Before you begin

Complete upgrade planning. Make sure your deployment is healthy and successfully communicating.



Note Upgrade planning starts with reading the Secure Firewall Threat Defense release notes. It then includes taking backups, obtaining upgrade packages, and performing associated upgrades (such as FXOS for the Firepower 4100/9300). It also includes checks for necessary configuration changes, readiness checks, disk space checks, and checks for both running and scheduled tasks. For details, see the Secure Firewall Threat Defense upgrade guides for Firewall Device Manager for your version.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the Updates panel.
The System Upgrade panel shows the currently running software version and any upgrade package that you have already uploaded.

Step 2 Upload the upgrade package.

You can upload one package only. If you upload a new package, it replaces the old one. Make sure you have the correct package for your target version and device model. Click **Browse** or **Replace File** to begin the upload.

When the upload completes, the system displays a confirmation dialog box. Before you click **OK**, optionally select **Run Upgrade Immediately** to choose rollback options and upgrade now. If you upgrade now, it is especially important to have completed as much of the pre-upgrade checklist as possible (see the next step).

Step 3 Perform final pre-upgrade checks, including the readiness check.

Revisit the pre-upgrade checklist. Make sure you have completed all relevant tasks, especially the final checks. If you do not run the readiness check manually, it runs when you initiate the upgrade. If the readiness check fails, the upgrade is canceled. For more information, see Running an Upgrade Readiness Check, on page 7.

Step 4 Click **Upgrade Now** to start the upgrade.

a) Choose rollback options.

You can **Automatically cancel on upgrade failure and roll back to the previous version**. With this option enabled, the device automatically returns to its pre-upgrade state upon major or maintenance upgrade failure. Disable this option if you want to be able to manually cancel or retry a failed upgrade.

b) Click **Continue** to upgrade and reboot the device.

You are automatically logged off and taken to a status page where you can monitor the upgrade until the device reboots. The page also includes an option to cancel the in-progress installation. If you disabled automatic rollback and the upgrade fails, you can manually cancel or retry the upgrade.

Traffic is dropped while you upgrade. For the ISA 3000 only, if you configured hardware bypass for power failure, traffic is dropped during the upgrade but is passed without inspection while the device completes its post-upgrade reboot.

Step 5 Log back in when you can and verify upgrade success.

The Device Summary page shows the currently running software version.

Step 6 Complete post-upgrade tasks.

a) Update system databases. If you do not have automatic updates configured for intrusion rules, VDB, and GeoDB, update them now.
b) Complete any other required post-upgrade configuration changes.
c) Deploy.

Running an Upgrade Readiness Check

Before the system installs an upgrade, it runs a readiness check to ensure the upgrade is valid for the system, and to check other items that sometimes prevent a successful upgrade. If the readiness check fails, you should fix the problems before trying the installation again. If the check has failed, you will be prompted about the failure the next time you try the installation, and you are given the option to force the installation if you want to.

You can also manually run the readiness check prior to initiating the upgrade, as described in this procedure.

Before you begin

Upload the upgrade package you want to check.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the Updates summary.

The **System Upgrade** section shows the currently running software version and any update that you have already uploaded.

Step 2 Look at the **Readiness Check** section.

- If the upgrade check has not been performed yet, click the **Run Upgrade Readiness Check** link. The progress of the check is shown in this area. It should take about 20 seconds to complete the process.
- If the upgrade check has already been run, this section indicates whether the check succeeded or failed. For failed checks, click **See Details** to view more information about the readiness check. After fixing problems, run the check again.

Step 3 If the readiness check fails, you should resolve the issues before you install the upgrade. The detailed information includes help on how to fix indicated problems. For a failed script, click the **Show Recovery Message** link to see the information.

Following are some typical problems:

- FXOS version incompatibility—On systems where you install FXOS upgrades separately, such as the Firepower 4100/9300, an upgrade package might require a different minimum FXOS version than the Firewall Threat Defense software version you are currently running. In this case, you must first upgrade FXOS before you can upgrade the Firewall Threat Defense software.
- Unsupported device model—The upgrade package cannot be installed on this device. You might have uploaded the wrong package, or the device is an older model that is simply no longer supported in the new Firewall Threat Defense software version. Please check device compatibility and upload a supported package, if one is available.
- Insufficient disk space—If not enough space is available, try deleting unneeded files, such as system backups. Delete only those files you have created.

Monitoring Firewall Threat Defense Upgrades

When you start the Firewall Threat Defense upgrade, you are automatically logged off and taken to a status page where you can monitor overall upgrade progress. The page also includes an option to cancel the in-progress installation. If you disabled automatic rollback and the upgrade fails, the page allows you to manually cancel or retry the upgrade.

You can also SSH to the device and use the CLI: **show upgrade status**. Add the **continuous** keyword to view log entries as they are made, and **detail** to see detailed information. Add both keywords to get continuous detailed information.

After the upgrade completes, you lose access to the status page and the CLI when the device reboots.

Canceling or Retrying Firewall Threat Defense Upgrades

Use the upgrade status page or the CLI to manually cancel failed or in-progress major or maintenance upgrades, and to retry failed upgrades:

- Upgrade status page: Click **Cancel Upgrade** to cancel an in-process upgrade. If the upgrade fails, you can click **Cancel Upgrade** to stop the job and to return to the state of the device prior to the upgrade, or click **Continue** to retry the upgrade.
- CLI: Use **upgrade cancel** to cancel an in-process upgrade. If the upgrade fails, you can use **upgrade cancel** to stop the job and to return to the state of the device prior to the upgrade, or use **upgrade retry** to retry the upgrade.



Note By default, Firewall Threat Defense automatically reverts to its pre-upgrade state upon upgrade failure ("auto-cancel"). To be able to manually cancel or retry a failed upgrade, disable the auto-cancel option when you initiate the upgrade. In a high availability deployment, auto-cancel applies to each device individually. That is, if the upgrade fails on one device, only that device is reverted.

Cancel and retry are not supported for patches. For information on reverting a successful upgrade, see Reverting Firewall Threat Defense, on page 8.

Reverting Firewall Threat Defense

If a major or maintenance upgrade succeeds but the system does not function to your expectations, you can revert. Reverting Firewall Threat Defense returns the software to its state just before the last major or maintenance upgrade; post-upgrade configuration changes are not retained. Reverting after patching necessarily removes patches as well. Note that you cannot revert individual patches or hotfixes.

The following procedure explains how to revert from Firewall Device Manager. If you cannot get into Firewall Device Manager, you can revert from the Firewall Threat Defense command line in an SSH session using the **upgrade revert** command. You can use the **show upgrade revert-info** command to see what version the system will revert to.

Before you begin

If the unit is part of a high availability pair, you must revert both units. Ideally, initiate the revert on both units at the same time so that the configuration can be reverted without failover issues. Open sessions with both

units and verify that revert will be possible on each, then start the processes. Note that traffic will be interrupted during the revert, so do this at off hours if at all possible.

For the Firepower 4100/9300 chassis, major Firewall Threat Defense versions have a specially qualified and recommended companion FXOS version. This means that after you revert the Firewall Threat Defense software, you might be running a non-recommended version of FXOS (too new). Although newer versions of FXOS are backwards compatible with older the Firewall Threat Defense versions, we do perform enhanced testing for the recommended combinations. You cannot downgrade FXOS, so if you find yourself in this situation, and you want to run a recommended combination, you will need to reimage the device.

Procedure

Step 1 Select **Device**, then click **View Configuration** in the **Updates** summary.

Step 2 In the **System Upgrade** section, click the **Revert Upgrade** link.

You are presented with a confirmation dialog box that shows the current version and the version to which the system will revert. If there is no available version to revert to, there will not be a **Revert Upgrade** link.

Step 3 If you are comfortable with the target version (and one is available), click **Revert**.

After you revert, you must re-register the device with the Smart Software Manager.

Troubleshooting Threat Defense Upgrades

These issues can occur when you are upgrading any device, whether standalone or in a high availability pair. To troubleshoot issues specific to high availability upgrades, see Troubleshooting High Availability Threat Defense Upgrades.

Upgrade package errors.

To find the correct upgrade package, select or search for your model on the Cisco Support & Download site, then browse to the software download page for the appropriate version. Available upgrade packages are listed along with installation packages, hotfixes, and other applicable downloads. Upgrade package file names reflect the platform, package type (upgrade, patch, hotfix), software version, and build.

Upgrade packages from Version 6.2.1+ are signed, and terminate in .sh.REL.tar. If you are upgrading from an older version, download the package that terminates in .sh instead. The Cisco Support & Download site indicates the correct package for your version. Do not untar signed upgrade packages. Do not rename upgrade packages or transfer them by email.

Cannot reach the device at all during upgrade.

Devices stop passing traffic during the upgrade or if the upgrade fails. Before you upgrade, make sure traffic from your location does not have to traverse the device itself to access the device's management interface.

Device appears inactive or is unresponsive during upgrade.

You can manually cancel in-progress major and maintenance upgrades; see Canceling or Retrying Firewall Threat Defense Upgrades, on page 8. If the device is unresponsive, or if you cannot cancel the upgrade, contact Cisco TAC.

Caution Even if the system appears inactive, do *not* manually reboot or shut down during upgrade. You could place the system in an unusable state and require a reimage.

Upgrade is successful but the system does not function to your expectations.

First, make sure that cached information gets refreshed. Do not simply refresh the browser window to log back in. Instead, delete any "extra" path from the URL and reconnect to the home page; for example, http://threat-defense.example.com/.

If you continue to have issues and need to return to an earlier major or maintenance release, you may be able to revert; see Reverting Firewall Threat Defense, on page 8. If you cannot revert, you must reimage.

Upgrade fails.

When you initiate a major or maintenance upgrade, you use the **Automatically cancel on upgrade failure...** (auto-cancel) option to choose what happens if upgrade fails, as follows:

- Auto-cancel enabled (default): If upgrade fails, the upgrade cancels and the device automatically reverts to its pre-upgrade state. Correct any issues and try again later.
- Auto-cancel disabled: If upgrade fails, the device remains as it is. Correct the issues and retry immediately, or manually cancel the upgrade and try again later.

For more information, see Canceling or Retrying Firewall Threat Defense Upgrades, on page 8. If you cannot retry or cancel, or if you continue to have issues, contact Cisco TAC.

Reimaging the Device

Reimaging a device involves wiping out the device configuration and installing a fresh software image. The intention of reimaging is to have a clean installation with a factory default configuration.

You would reimage the device in these circumstances:

- You want to convert the system from ASA Software to Firewall Threat Defense Software. You cannot upgrade a device running an ASA image to one running a Firewall Threat Defense image.
- The device is not functioning correctly and all attempts at fixing the configuration have failed.

For information on how to reimage a device, see *Reimage the Cisco ASA or Threat Defense Device* or the *Threat Defense Quick Start* guide for your device model. These guides are available at http://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-guides-list.html.

Backing Up and Restoring the System

You can back up the system configuration so that you can restore the device if the configuration becomes corrupted due to subsequent miss-configuration or physical mishap.

You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including the build number, not just the same point release). Do not use the backup and restore process to copy configurations between appliances. A backup file contains information that uniquely identifies an appliance, so that it cannot be shared in this manner.

Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment. The backup also does not include licensing or cloud registration information, so whatever license or cloud registration state exists at the time of restore is retained.

Backups include the configuration only, and not the system software. If you need to completely reimage the device, you need to reinstall the software, then you can upload a backup and recover the configuration.

The configuration database is locked during backup. You cannot make configuration changes during a backup, although you can view policies, dashboards, and so forth. During a restore, the system is completely unavailable.

The table on the Backup and Restore page lists all existing backup copies that are available on the system, including the file name of the backup, the date and time it was created, and the file size. The type of backup (manual, scheduled, or recurring) is based on how you directed the system to create that backup copy.

Tip Backup copies are created on the system itself. You must manually download backup copies and store them on secure servers to ensure that you have the backup copies you need for disaster recovery. The system maintains up to 3 backup copies on the device. New backups replace the oldest backup.

The following topics explain how to manage backup and restore operations.

Backing Up the System Immediately

You can start a backup whenever you want.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.

Step 2 Click **Manual Backup** > **Back Up Now**.

Step 3 Enter a name for the backup and optionally a description.

If you decide you want to perform the backup at a future time rather than immediately, you can click **Schedule** instead.

Step 4 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 5 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 6 Click **Back Up Now**.

The system starts the backup process. When the backup is complete, the backup file will appear in the table. You can then download the backup copy to your system and store it elsewhere, if desired.

You can leave the Backup and Restore page after initiating the backup. However, the system will likely be sluggish, and you should consider pausing your work to allow the backup to complete.

In addition, the system will acquire locks on the configuration database during parts or all of the backup, which can prevent you from making changes for the duration of the backup process.

Backing Up the System at a Scheduled Time

You can set up a scheduled backup to back up the system at a specific future date and time. A scheduled backup is a one-time occurrence. If you want to create a backup schedule to regularly create backups, configure a recurring backup instead of a scheduled backup.



Note If you want to delete the schedule for a future backup, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Scheduled Backup** > **Schedule a Backup**.

If you already have a scheduled backup, click **Scheduled Backup** > **Edit** .

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the date and time for the backup.

Step 5 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 6 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 7 Click **Schedule**.

When the selected date and time arrives, the system takes a backup. When completed, the backup copy is listed in the table of backups.

Setting Up a Recurring Backup Schedule

You can set up a recurring backup to back up the system on a regular schedule. For example, you could take a backup every Friday at midnight. A recurring backup schedule helps ensure that you always have a set of recent backups.



Note If you want to delete a recurring schedule, edit the schedule and click **Remove**.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

Step 2 Click **Recurring Backup** > **Configure**.

If you already have a recurring backup configured, click **Recurring Backup** > **Edit**.

Step 3 Enter a name for the backup and optionally a description.

Step 4 Select the **Frequency** and the related schedule:

- **Daily**—Select the time of day. A backup is taken every day at the scheduled time.
- **Weekly**—Select the days of the week and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for every Monday, Wednesday, and Friday at 23:00 hours (11 PM).
- **Monthly**—Select the days of the month and the time of day. A backup is taken on each day you select at the scheduled time. For example, you could schedule backups for the first (1), fifteenth (15), and twenty-eighth (28) days of the month at 23:00 hours (11 PM).

The time you specify is adjusted for Daylight Savings Time, so it will move an hour forward or backward whenever the time is adjusted in your area. You must edit the schedule at the time change if you want to keep this exact time throughout the year.

Step 5 (Optional.) Select the **Encrypt File** option to encrypt the backup file.

If you select the option, you must enter the **Password** (and **Confirm Password**) that will be required to restore the backup file.

Step 6 (ISA 3000 only.) Select the **Location of Backup Files**.

You can create the backup on the **Local Hard Disk** or on the **SD Card**. The benefit of using the SD card is that you can use the card to recover the configuration to a replacement device.

Step 7 Click **Save**.

When the selected dates and times arrive, the system takes a backup. When completed, the backup copy is listed in the table of backups.

The recurring schedule continues to take backups until you change or remove it.

Restoring a Backup

You can restore backups as needed so long as the device is running the same software version (including build number) as it was running when you took the backup. You can restore a backup onto a replacement device only if the two devices are the same model and are running the same version of the software (including build number).

However, you cannot restore a backup if the device is part of a high availability pair. You must first break HA from the **Device** > **High Availability** page, then you can restore the backup. If the backup includes the HA configuration, the device will rejoin the HA group. Do not restore the same backup on both units, because they would then both go active. Instead, restore the backup on the unit you want to go active first, then restore the equivalent backup on the other unit.

If the backup copy you want to restore is not already on the device, you must upload the backup first before restoring it.

During a restore, the system is completely unavailable.



Note The backup does not include the management IP address configuration. Thus, when you recover a backup file, the management address is not replaced from the backup copy. This ensures that any changes you made to the address are preserved, and also makes it possible to restore the configuration on a different device on a different network segment. The backup also does not include licensing or cloud registration information, so whatever license or cloud registration state exists at the time of restore is retained.

Before you begin

If you are restoring the backup on a different system, for example, when replacing a device, best practice is to first register the device and enable any optional licenses required by the features configured in the backup file. The backup file does not include license or cloud services information, so whatever license changes or cloud registrations you make prior to restore are retained.

Procedure

Step 1 Click **Device**, then click **View Configuration** in the Backup and Restore summary.

This opens the Backup and Restore page. The table lists all existing backup copies that are available on the system.

Step 2 If the backup copy that you want to restore is not in the list of available backups, click **Upload** > **Browse** and upload the backup copy.

Step 3 Click the restore icon ( ) for the file.

You are asked to confirm the restore. By default, the backup copy will be deleted after the restore, but you can select **Do not remove the backup after restoring** to keep it before proceeding with the restore.

If the backup file was encrypted, you must enter the **Password** that is required to open the file and unencrypt it.

The system will reboot after restore completes.

Note

After the system reboots, it automatically checks for Vulnerability Database (VDB), Geolocation, and Rules database updates, and downloads them if needed. Because these updates can be large, the initial attempt might fail. Please check the task list, and if a download failed, manually download an update as described in Updating System Databases, on page 2. The system also redeploys policies. Any subsequent deployment will fail until the update is successful.

Step 4 If necessary, click **Device** > **Smart License** > **View Configuration**, re-register the device, and re-enable the required optional licenses.

The backup does not include license or cloud registration information. Thus, if you restore a backup to a new system, for example, when replacing a device, and the system is in evaluation mode, you need to register it and enable any licenses you need. If you registered the device and enabled licenses prior to the restore, no additional changes are required.

If you are simply restoring a previous backup to the same system, you should not need to make any changes to the licenses or cloud registration. However, verify that all optional licenses needed are enabled, as the backup could possible include features that require licenses you disabled after the backup was created.

Managing Backup Files

As you create new backups, the backup files are listed on the Backup and Restore page. Backup copies are not retained indefinitely: as disk space usage on the device reaches the maximum threshold, older backup copies are deleted to make room for newer ones. In addition, when you install any upgrade other than a hot fix, all backup files are deleted. Thus, you should regularly manage the backup files to ensure that you have the specific backup copies you most want to keep.

You can do the following to manage your backup copies:

- Download files to secure storage—To download a backup file to your workstation, click the download icon ( ) for the file. You can then move the file to your secure file storage.
- Upload a backup file to the system—If you want to restore a backup copy that is no longer available on the device, click **Upload** > **Browse File** and upload it from your workstation. You can then restore it.

Note Uploaded files may be renamed to match the original filename. Also, if there are more than 3 backup copies already on the system, the oldest one will be deleted to make room for the uploaded file. You cannot upload files that were created by an older software version.

- Restore a backup—To restore a backup copy, click the restore icon ( ) for the file. The system is unavailable during restore, and will reboot after restore completes. You should deploy the configuration after the system is up and running.
- Delete a backup file—If you no longer want a particular backup, click the delete icon ( ) for the file. You are asked to confirm the deletion. Once deleted, you cannot recover the backup file.

Auditing and Change Management

You can view status information about system events and actions that users have taken. This information can help you audit the system and ensure that the system is being managed properly.

Click **Device** > **Device Administration** > **Audit Log** to see the audit log. In addition, you can find system management information by clicking the **Task List** or **Deployment** icon buttons in the upper right corner.

The following topics cover some of the main concepts and tasks for system auditing and change management.

Audit Events

The audit log can include the following types of event:

Custom Feed Update Event, Custom Feed Update Failed

These events indicate a successfully completed or failed update to a custom Security Intelligence feed. The details include who started the update and information about the feed that was being updated.

Custom Rules File Import Summary Event

These events indicate that you imported a file that contained one or more custom intrusion rule. The event includes a summary of the number of rules added, updated, and deleted, and a differences view that shows details about the imported rules.

Deployment Completed, Deployment Failed: *job name* or *entity name*

These events indicate a successfully completed or failed deployment job. The details include who started the job and information about the job entity. Failed jobs include the error message related to the failure.

The details also include a **Differences View** tab, which shows the changes that were deployed to the device in the job. This is the combination of all the Entity change events for the deployed entities.

To filter on these events, simply click the **Deployment History** pre-defined filter. Note that the event type for these events is Deployment Event, you cannot filter on completed or failed events only.

The event name includes the user-defined job name (if you configure one), or "User (*username*) Triggered Deployment." There are also "Device Setup Automatic Deployment" and "Device Setup Automatic Deployment (Final Step)" jobs that occur during the device setup wizard.

Entity Created, Entity Updated, Entity Deleted: *entity name* (*entity type*)

These events indicate that a change was made to the identified entity or object. The entity details include who made the change, as well as the entity name, type, and ID. You can filter on these items. The details also include a **Differences View** tab, which shows the changes that were made to the object.

HA Action Event

These events relate to actions on the high availability configuration, either actions that you initiated, or actions that the system initiated. HA Action Event is the event type, but the event names are one of the following:

- **HA Suspended**—You intentionally suspended HA on the system.
- **HA Resumed**—You intentionally resumed HA on the system.
- **HA Reset**—You intentionally reset HA on the system.

- **HA Failover: Unit Switched Modes**—Either you intentionally switched modes, or the system failed over due to health metric violations. The message indicates that the active peer became standby, or the standby peer became active.

High Availability Sync Completed

The active unit synchronized the configuration with the standby unit. The event includes the change information for the previous version compared to the synchronized version.

Interface List Scanned

This event indicates that you scanned for changes in the interface inventory.

Pending Changes Discarded

This event indicates that you deleted all pending changes. Any changes indicated in Entity Created, Entity Updated, and Entity Deleted events between this event and the previous Deployment Completed event are removed, and the state of the affected objects is returned to the last deployed version.

Rules Update Event

When running Snort 3, this event from the LSPUpdateServer entity shows detailed information about the intrusion rules that were added, removed, or changed when a new intrusion rules package was downloaded and installed. The event is limited to 100 rules, so if more than 100 are added, removed, or changed, the event will not have complete information.

Task Started, Task Completed, Task Failed

The task events indicate the start and end of a job initiated either by the system or a user. These two events are consolidated into a single task in the task list, which you can see by clicking the **Task List** button in the upper right corner.



Tasks include actions such as deployment jobs and manual or scheduled database updates. Any item in the task list will correspond to two task events in the audit log, an indication of the start of the task, and either a successful completion or a failure.

User Logged In, User Logged Out: *username*

These events show the time and source IP address for the user logging into and out of the Firewall Device Manager. The User Logged Out event occurs for both active log outs and automatic log outs due to idle time being exceeded.

These events do not relate to RA VPN users establishing connections with the device. They also do not include log in/log out to the device CLI.

Viewing and Analyzing the Audit Log

The audit log includes information about system-initiated and user-initiated events such as deployment jobs, database updates, and login/logout of the Firewall Device Manager.

For an explanation of the types of event you can see in the log, see Audit Events, on page 16.

Procedure

Step 1 Click **Device**, then click the **Device Administration** > **View Configuration** link.

Step 2 Click **Audit Log** in the table of contents if it is not already selected.

Events are grouped by date, and within a day, by time, with the most recent date/time at the top of the list. Initially, each event is collapsed, so you see only the time, event name, user who initiated the event, and source IP address of the user. "System" for user and IP address means that the device itself initiated the event.

You can do the following:

- Click > next to the event name to open it and see the event details. Click the icon again to close the event. Many events have a simple list of event attributes, such as event type, user name, source IP address, and so forth. However, Entity and Deployment events have two tabs:
  - **Summary** shows the basic event attributes.
  - **Differences View** shows a comparison of the existing "deployed" configuration with the changes made as part of the event. For deployment jobs, this view can be long and require scrolling. It sums up all differences from the Entity event changes that were part of the deployment job.
- Select a different time range from the drop-down list to the right of the filter field. The default is to view events from the past 2 weeks, but you can change that to the last 24 hours, 7 days, month, or 6 months. Click **Custom** to specify an exact range by entering the start and end date and time.
- Click any link in the log to add a search filter for that item. The list updates so that only those events that include the item are shown. You can also simply click in the **Filter** box and build a filter directly. There are some pre-defined filters beneath the filter box that you can click to load the related filter criteria. For detailed information on filtering the events, see Filtering the Audit Log, on page 18.
- Reload the browser page to refresh the log with the latest events.

Filtering the Audit Log

You can apply a filter to the audit log to narrow your view to certain types of message only. Each element in the filter is an exact, complete match. For example, "User = admin" shows only those events initiated by the user with the name **admin**.

You can use the following techniques, alone or in combination, to build a filter. The list is automatically updated each time you add a filter element.

Click a Predefined Filter

Beneath the **Filter** field are the predefined filters. Simply click a link to load the filter. You are asked for confirmation. If you already have a filter applied, it is replaced; it is not added to.

Clicking Highlighted Items

The easiest way to build a filter is to click on items in the log table or event details that contain the values on which you intend to filter. Clicking an item updates the **Filter** field with a correctly-formulated element

for that value and element combination. However, using this technique requires that the existing list of events contains the desired values.

If you can add a filter element for an item, the item is underlined when you mouse over it and you see the command **Click to Add to Filter**.

Selecting Atomic Elements

You can also build a filter by clicking in the **Filter** field and selecting the desired atomic element from the drop-down list, typing in the match value after the equal sign, then pressing Enter. You can filter on the following elements. Note that not all elements are relevant for every event type.

- **Event Type**—This is usually but not always the same as the event name (without variable qualifiers like entity name or user). For deployment events, the event type is Deployment Event. For an explanation of the event types, see Audit Events, on page 16.
- **User**—The name of the user who initiated the event. The system user is spelled in all capitals: SYSTEM.
- **Source IP**—The IP address from which the user initiated the event. The source IP address for system-initiated events is SYSTEM.
- **Entity ID**—The UUID for the entity or object, which is a long unreadable string such as 8e7021b4-2e1e-11e8-9e5d-0fc002c5f931. Normally, to use this filter you either need to click an entity ID in an event's details, or retrieve the necessary ID through a relevant GET call using the REST API.
- **Entity Name**—The name of the entity or object. For user-created entities, this is typically the name you gave the object, for example, InsideNetwork for a network object. For system-generated entities, or in some cases user-defined entities, this is a predefined but intelligible name, for example, "User (admin) Triggered Deployment" for deployment jobs you do not explicitly name.
- **Entity Type**—The kind of entity or object. These are predefined but intelligible names, such as Network Object. You can find entity types in the API Explorer by looking at the relevant object model for the "type" value. The API types are normally all lower-case with no spaces. If you type them in exactly as shown in the model, the string changes to a more readable format when you press Enter. Typing in either form works. To open API Explorer, click the more options button (⋮) and choose **API Explorer**.

Rules for Complex Audit Log Filters

When building a complex filter that contains more than one atomic element, keep the following rules in mind:

- Elements of the same type have an OR relationship between all values for that type. For example, including "User = admin" and "User = SYSTEM" matches events that were initiated by either user.
- Elements of different types have an AND relationship. For example, including "Event Type = Entity Updated" and "User = SYSTEM" shows only those events where the system updated an entity rather than an active user.
- You cannot use wildcards, regular expressions, partial matches, or simple text string matches.

Examining Deployment and Entity Change History

Deployment and entity events include a **Differences View** tab in the event details. This tab shows a color-coded comparison of the old configuration with the changes.

- For deployment jobs, this is a comparison of the configuration that was running on the device prior to deployment to the changes that were actually deployed.
- For entity events, these are the configuration changes made to the previous version of the object. The previous version might be the version actually on the device, or it might be a change to an object that has not yet been deployed.

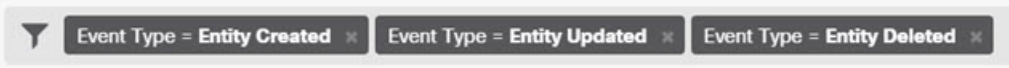
Procedure

Step 1 Click **Device**, then click the **Device Administration** > **View Configuration** link.

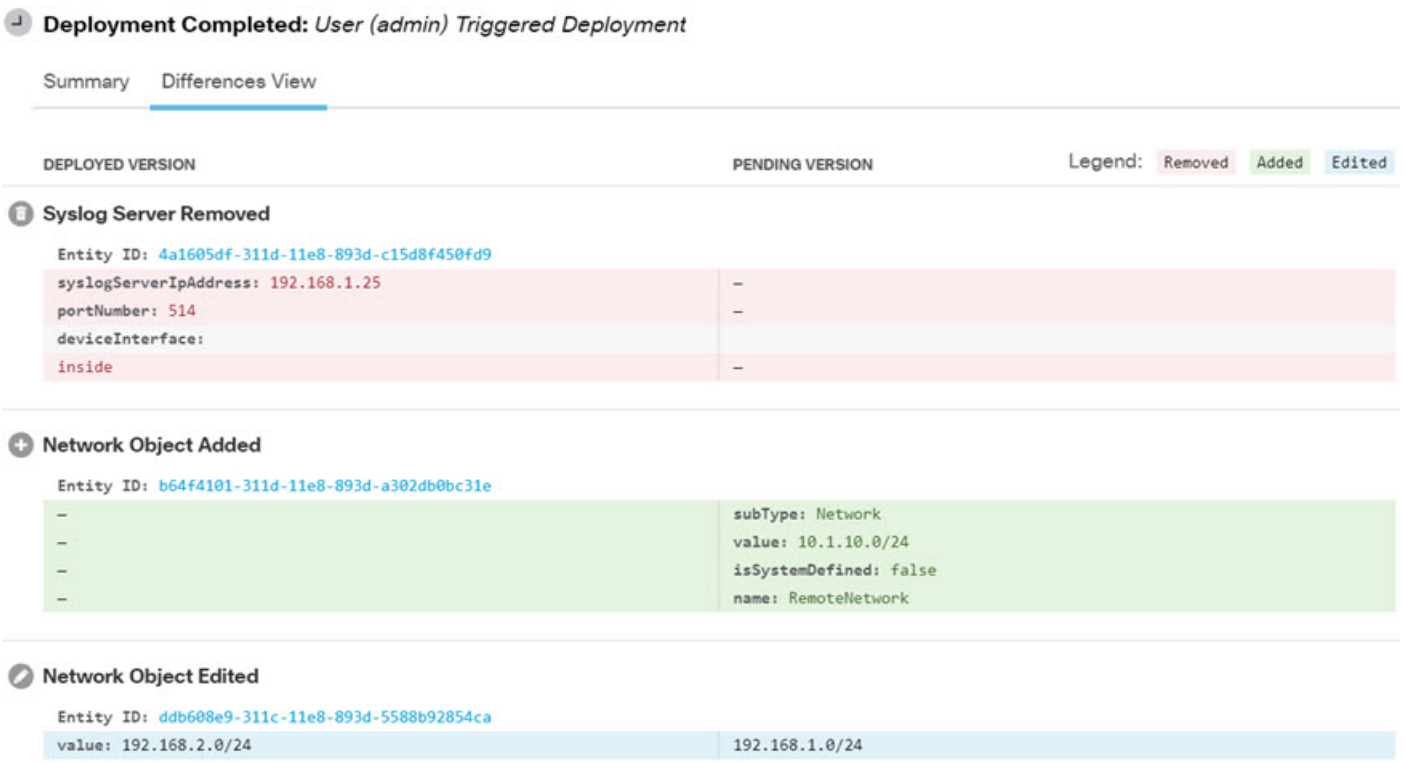
Step 2 Click **Audit Log** in the table of contents if it is not already selected.

Step 3 (Optional.) Filter the messages:

- Deployment events—Click the **Deployment History** predefined filter under the filter box.
- Entity change events—Manually create a filter using the Event Type element for the type of change that interests you. To see all entity changes, include three specifications for Entity Created, Entity Updated, and Entity Deleted. The filter would look like the following:



Step 4 Open the event and click the **Differences View** tab.



The changes are color coded, and the heading indicates the type of object and whether it was Added (Created), Removed (Deleted), or Edited (Updated). Edited objects show only those attributes that were changed or deleted from the object. In deployment jobs, there are separate headings for each entity changed. The heading indicates the entity type for the object.

Discarding All Pending Changes

If you are unsatisfied with a set of configuration changes that have not yet been deployed, you can discard all pending changes. This returns all features to the state that exists on the device. You can then start again on your configuration changes.

Procedure

Step 1 Click the **Deploy Changes** icon in the upper right of the web page.

The icon is highlighted with a dot when there are pending changes.



Step 2 Click **More Options** > **Discard All**.

Step 3 Click **OK** in the confirmation dialog.

The system discards the changes, and you will see a message that there are no pending changes when the process completes. The system adds a Pending Changes Discarded event to the audit log.

Exporting the Device Configuration

You can export a copy of the currently-deployed configuration in JSON format. You can use the file for archival or record-keeping purposes. Any sensitive data, such as passwords and secret keys, is masked.

You cannot import the file into this or another device. This ability is not a replacement for backing up the system.

You must have completed at least one successful deployment job before you can download the configuration.

Procedure

Step 1 Choose **Device**, then click **View Configuration** in the **Device Administration** group.

Step 2 Click **Download Configuration** in the table of contents.

Step 3 Click **Get Device Configuration** to start a job that creates the file.

If you have previously created a file, you will see a download button and the message **File is ready to download**, with the creation date for the file.

Depending on the size of the configuration, it can take several minutes to generate the file. Check the task list or audit log, or return to this page periodically, until the Export Config job completes and the file is generated.

Step 4 When the file is generated, return to this page and click the **Download the Configuration File** button ( ) to save the file to your workstation.

Managing Firewall Device Manager and Firewall Threat Defense User Access

You can configure an external authentication and authorization source for users to log into Firewall Threat Defense (HTTPS access). You can use an external server in addition to, or instead of, the local user database and the system-defined **admin** user. Note that you cannot create additional local user accounts for Firewall Device Manager access.

Although you can have multiple external Firewall Device Manager user accounts that can change the configuration, these changes are not tracked by user. When one user deploys changes, changes made by all users are deployed. There is no locking: that is, more than one user might attempt to update the same object at the same time, which will result in only one user successfully saving the change. You also cannot discard changes based on user.

You can have 5 concurrent user sessions. If a sixth user logs in, the oldest user session is automatically logged off. There is also an idle timeout, which logs inactive users out after 20 minutes.

You can also configure external authentication and authorization for SSH access to the Firewall Threat Defense CLI. The local database is always checked before using the external source, so you can create additional local users for failsafe access. Do not create duplicate users in both the local and external source. Except for the **admin** user, there is no crossover between the CLI and Firewall Device Manager users: the user accounts are completely separate.



Note When using external servers, you can control access by user to subsets of your devices by either setting up separate AAA server groups, or by creating authentication/authorization policies within the AAA servers that allow the user access to certain Firewall Threat Defense device IP addresses only.

The following topics explain how to configure and manage Firewall Device Manager user access and CLI user access.

Configuring External Authorization (AAA) for the Firewall Device Manager (HTTPS) Users

You can provide HTTPS access to the Firewall Device Manager from an external AAA server. By enabling AAA authentication and authorization, you can provide different levels of access rights, and not have every user log in through the local **admin** account.

These external users are also authorized for the Firewall Threat Defense API and the API Explorer.

You can provide role-based access control (RBAC) by setting up authorization for management users in the AAA server. The levels differ depending on server type. When a user logs into Firewall Device Manager, the username and role are shown in the upper right of the page. After you set up the accounts correctly on the AAA server, you can enable it for administrative access using this procedure.

RADIUS User Authorization

To provide role-based access control (RBAC), update the user accounts on your RADIUS server to define the **cisco-av-pair** attribute (in ISE, but the attribute is spelled Cisco-AVPair in Free RADIUS; check your system for correct spelling). This attribute must be defined correctly on a user account, or the user is denied access to the Firewall Device Manager. Following are the supported values for the **cisco-av-pair** attribute:

- **fdm.userrole.authority.admin** provides full Administrator access. These users can do all actions that the local **admin** user can do.
- **fdm.userrole.authority.rw** provides read-write access. These users can do everything a read-only user can do, and also edit and deploy the configuration. The only restrictions are for system-critical actions, which include installing upgrades, creating and restoring backups, and ending the sessions of the Firewall Device Manager users.
- **fdm.userrole.authority.ro** provides read-only access. These users can view dashboards and the configuration, but cannot make any changes. If the user tries to make a change, the error message explains that this is due to lack of permission.

SAML User Authorization

You identify the field that contains authorization levels when you configure the SAML server identity source. You can configure external users with the following types of authorization access: Administrator, Audit Admin, Cryptographic Admin, Read-Write User, Read-Only User. See Configure SAML Servers.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **AAA Configuration** tab if it is not already selected.

Step 3 Configure the **HTTPS Connection** options:

- **Server Group for Management/REST API**—Select the server group (for external authentication/authorization) or local user database (LocalIdentitySource), that you want to use as a primary authentication source.

  If the server group does not yet exist, click the link to create a new one and create it now. For RADIUS, you will also need to create server objects for each server, to add them to the group, but you can do this while defining the server group. For more information on RADIUS, see RADIUS Servers and Groups. For information on SAML, see Configure SAML Servers.

- **Authentication with LOCAL** (RADIUS.)—If you select an external RADIUS server group, you can specify how to use the local identity source, which contains the local **admin** user account. Select one of the following:
  - **Before External Server**—The system checks the username and password against the local source first.
  - **After External Server**—The local source is checked only if the external source is unavailable.
  - **Never**—(Not recommended.) The local source is never used, so you cannot log in as the admin user.

    Caution
    If you select **Never**, you will not be able to log into the Firewall Device Manager using the **admin** account. You will be locked out of the system if the AAA server becomes unavailable, or if you miss-configure the accounts in the AAA server.

Note
Authentication with LOCAL does not apply when using SAML. With SAML, you can always log in using the local database by entering a local username and password, because you have to explicitly click the **Single-Sign On (SSO)** link to enter SAML credentials.

Step 4 Click **Save**.

Configuring External Authorization (AAA) for the Firewall Threat Defense CLI (SSH) Users

You can provide SSH access to the Firewall Threat Defense CLI from an external RADIUS server. By enabling RADIUS authentication and authorization, you can provide different levels of access rights from a single authentication source, rather than define separate local user accounts on each device.

These SSH external users are **not** authorized for the Firewall Threat Defense API and the API Explorer. The mechanism you use to define authorization for SSH is different from the one required for HTTPS access. However, you can configure the same RADIUS user with both SSH and HTTPS authorization criteria, so that a given user can access the system through both protocols.

To provide role-based access control (RBAC) for SSH access, update the user accounts on your RADIUS server to define the **Service-Type** attribute. This attribute must be defined on a user account, or the user is denied SSH access to the device. Following are the supported values for the **Service-Type** attribute:

- **Administrative (6)**—Provides **config** access authorization to the CLI. These users can use all commands in the CLI.
- **NAS Prompt (7)** or any level other than 6—Provides **basic** access authorization to the CLI. These users can use read-only commands, such as **show** commands, for monitoring and troubleshooting purposes.

After you set up the accounts correctly on the RADIUS server, you can enable it for SSH administrative access using this procedure.

Note Do not create duplicate users in both the local and external source. If you do create duplicate usernames, ensure that they have the same authorization rights. You cannot log in using the password of the external version of the user account when the authorization rights differ in the local user account; you can log in using the local password only. If the rights are the same, the password you use determines if you are logged in as the external or the local user, assuming the passwords are different. Even though the local database is checked first, if a username exists in the local database but the password is incorrect, the external server is checked and if the password is correct for the external source, the login will succeed.

Before you begin

Please inform externally-defined users of the following behavior to set their expectations appropriately:

- The first time an external user logs in, the Firewall Threat Defense creates the required structures but cannot simultaneously create the user session. The user simply needs to authenticate again to start the session. The user will see a message similar to the following: "New external username identified. Please log in again to start a session."
- Similarly, if the user's authorization as defined in the Service-Type changed since the last login, the user will need to re-authenticate. The user will see a message similar to the following: "Your authorization privilege has changed. Please log in again to start a session."

If external users can successfully log into the GUI but not the CLI, there might be a problem with the secret key for the AAA servers. Ensure that the keys do not include & or \ characters. Although these might work with GUI logins, they do not work with SSH logins.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Management Access** link.

If you are already on the System Settings page, simply click **Management Access** in the table of contents.

Step 2 Click the **AAA Configuration** tab if it is not already selected.

Step 3 Configure the **SSH Connection** options:

- **Server Group**—Select the RADIUS server group or local user database (LocalIdentitySource), that you want to use as a primary authentication source. You must select a RADIUS server group to use external authorization.

  If the server group does not yet exist, click the **Create New RADIUS Server Group** link and create it now. You will also need to create RADIUS server objects for each server, to add them to the group, but you can do this while defining the server group. For more information on RADIUS, see RADIUS Servers and Groups.

  Note that SSH connections use the first 2 servers in the group only. If you use a group with 3 or more servers, the additional servers are never tried. In addition, the **Dead Time** and **Maximum Failed Attempts** group attributes are not used.

- **Authentication with LOCAL**—If you select an external server group, you can specify how to use the local identity source. For SSH access, the local database is always checked before the external server.

Step 4 Click **Save**.

Managing the Firewall Device Manager User Sessions

Choose **Monitoring** > **Sessions** to see a list of users who are currently logged into the Firewall Device Manager. The list shows how long each user has been logged in for the current session.

If the same username appears more than once, it means that the user has opened sessions from different source addresses. Sessions are tracked separately based on username and source address, each session with its own unique time stamp.

The system allows 5 concurrent user sessions. If a sixth user logs in, the oldest current session is automatically logged out. In addition, idle users are automatically logged out after 20 minutes of inactivity.

If the Firewall Device Manager user types in the wrong password and fails to log in on 3 consecutive attempts, the user's account is locked for 5 minutes. The user must wait before trying to log in again. There is no way to unlock the Firewall Device Manager user account, nor can you adjust the retry count or lock timeout. (Note that for SSH users, you can adjust these settings and unlock the account.)

If necessary, you can end a user session by clicking the delete icon ( ) for the session. If you delete your own session, you are also logged out. There is no lockout period if you end a session: the user can immediately log back in.

Enabling the Firewall Device Manager Access on a Standby HA Unit for External Users

If you configure external authorization for the Firewall Device Manager users, those users can log into both the active and standby unit of a high availability pair. However, to successfully log into the standby unit for the first time requires a few extra steps compared to logging into the active unit.

After an external user logs into the active unit for the first time, the system creates an object that defines the user and the user’s access rights. An admin or read-write user must then deploy the configuration from the active unit for the user object to appear on the standby unit.

Only after the deployment and subsequent configuration synchronization completes successfully can the external user log into the standby unit.

Admin and read-write users can deploy changes after logging into the active unit. However, read-only users cannot deploy the configuration, and must ask a user who has the appropriate rights to deploy the configuration.

Creating Local User Accounts for the Firewall Threat Defense CLI

You can create users for CLI access on Firewall Threat Defense devices. These accounts do not allow access to the management application, but to the CLI only. The CLI is useful for troubleshooting and monitoring purposes.

You cannot create local user accounts on more than one device at a time. Each device has its own set of unique local user CLI accounts.

Procedure

Step 1 Log into the device CLI using an account with config privileges.

The admin user account has the required privileges, but any account with config privileges will work. You can use an SSH session or the Console port.

For certain device models, the Console port puts you into the FXOS CLI. Use the **connect ftd** command to get to the Firewall Threat Defense CLI.

Step 2 Create the user account.

configure user add *username* {**basic** | **config**}

You can define the user with the following privilege levels:

- **config**—Gives the user configuration access. This gives the user full administrator rights to all commands.
- **basic**—Gives the user basic access. This does not allow the user to enter configuration commands. Allowed commands are dig, ping, and traceroute.

Example:

The following example adds a user account named joecool with config access rights. The password is not shown as you type it.

```
> configure user add joecool config
Enter new password for user joecool: newpassword
Confirm new password for user joecool: newpassword
> show user
Login              UID   Auth Access  Enabled Reset    Exp Warn  Str Lock Max
admin             1000  Local Config  Enabled    No  Never  N/A  Dis   No N/A
joecool           1001  Local Config  Enabled    No  Never  N/A  Dis   No   5
```

Note

The user is prompted to change the password on first login. The user can initiate subsequent password changes using the **configure password** command.

Step 3 (Optional.) Adjust the characteristics of the account to meet your security requirements.

You can use the following commands to change the default account behavior.

- **configure user aging** *username max_days warn_days*

  Sets an expiration date for the user's password. Specify the maximum number of days for the password to be valid followed by the number of days before expiration the user will be warned about the upcoming expiration. Both values are 1 to 9999, but the warning days must be less than the maximum days. When you create the account, there is no expiration date for the password.

- **configure user forcereset** *username*

  Forces the user to change the password on the next login.

- **configure user maxfailedlogins** *username number*

  Sets the maximum number of consecutive failed logins you will allow before locking the account, from 1 to 9999. Use the **configure user unlock** command to unlock accounts. The default for new accounts is 5 consecutive failed logins.

- **configure user minpasswdlen** *username number*

  Sets a minimum password length, which can be from 1 to 127.

- **configure user strengthcheck** *username* {**enable** | **disable**}

  Enables or disables password strength checking, which requires a user to meet specific password criteria when changing their password. When a user's password expires or if the **configure user forcereset** command is used, this requirement is automatically enabled the next time the user logs in.

Step 4 Manage user accounts as necessary.

Users can get locked out of their accounts, or you might need to remove accounts or fix other issues. Use the following commands to manage the user accounts on the system.

- **configure user access** *username* {**basic** | **config**}

  Changes the privileges for a user account.

- **configure user delete** *username*

  Deletes the specified account.

- **configure user disable** *username*

  Disables the specified account without deleting it. The user cannot log in until you enable the account.

- **configure user enable** *username*

  Enables the specified account.

- **configure user password** *username*

  Changes the password for the specified user. Users should normally change their own password using the **configure password** command.

- **configure user unlock** *username*

  Unlocks a user account that was locked due to exceeding the maximum number of consecutive failed login attempts.

Rebooting or Shutting Down the System

If necessary, you can reboot or shut down the system.

In addition to the procedure below, you can also perform these tasks through an SSH session or the Firewall Device Manager CLI Console using the **reboot** or **shutdown** commands.

Procedure

Step 1 Click **Device**, then click the **System Settings** > **Reboot/Shutdown** > link.

If you are already on the System Settings page, simply click **Reboot/Shutdown** in the table of contents

Step 2 Click the button that performs the function you need.

- **Reboot**—If you believe the system is not performing correctly and other efforts to resolve the problem have failed, you can reboot the device. In addition, there might be a few procedures that ask you to reboot the device to reload the system software.
- **Shut Down**—Shut down the system to turn off power in a controlled fashion. Use shutdown when you intend to remove the device from the network, for example, to replace it. After shutting down the device, you can turn it back on only from the hardware On/Off switch.

Step 3 Wait until the action completes.

If you have a console connection to the firewall, monitor the system prompts as the firewall shuts down. You will see the following prompt:

```
System is stopped.
It is safe to power off now.
Do you want to reboot instead? [y/N]
```

For the ISA 3000, when shutdown is complete, the System LED will turn off. Wait at least 10 seconds before you remove the power.

You cannot perform other actions in the Firewall Device Manager or the CLI while the system is rebooting or shutting down.

During reboot, the Firewall Device Manager page should refresh when the reboot is complete and bring you to the login page. If you try refreshing the page before the reboot is complete, the web browser might return 503 or 404 errors, based on the operational state of the Firewall Device Manager web server at that point in time.

For shutdown, the system will eventually not be able to respond at all and you will get 404 errors. This is the expected result, because you are completely turning off the system.

Troubleshooting the System

The following topics address some system-level troubleshooting tasks and capabilities. For information on troubleshooting a specific feature, such as access control, see the chapter for the feature.

Pinging Addresses to Test Connectivity

Ping is a simple command that lets you determine if a particular address is alive and responsive. This means that basic connectivity is working. However, other policies running on a device could prevent specific types of traffic from successfully getting through a device. You can use **ping** by opening the CLI console or by logging into the device CLI.



Note Because the system has multiple interfaces, you can control the interface used for pinging an address. You must ensure that you are using the right command, so that you are testing the connectivity that matters. For example, the system must be able to reach the Cisco license server through the Management interface, so you must use the **ping system** command to test the connection. If you use **ping**, you are testing whether an address can be reached through the data interfaces, which might not give you the same result.

The normal ping uses ICMP packets to test the connection. If your network prohibits ICMP, you can use a TCP ping instead (for data interface pings only).

You can ping either an IP address or a fully-qualified hostname (FQDN). For a ping to work on an FQDN, the DNS servers configured for either the management or data interfaces must successfully return an IP address. You must configure DNS servers separately for management and data interfaces. If you do not have DNS servers configured for a specific interface, use the **dig** command to look up the IP address of a given FQDN.

Following are the main options for pinging network addresses.

Pinging an address through the Management interface

Use the **ping system** command.

ping system *host*

The host can be an IP address or fully-qualified domain name (FQDN), such as www.example.com. Unlike pings through the data interfaces, there is no default count for system pings. The ping continues until you stop it using Ctrl+c. For example:

```
> ping system www.cisco.com
PING origin-www.cisco.COM (72.163.4.161) 56(84) bytes of data.
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=1 ttl=242 time=10.6 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=2 ttl=242 time=8.13 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=3 ttl=242 time=8.51 ms
64 bytes from www1.cisco.com (72.163.4.161): icmp_seq=4 ttl=242 time=8.40 ms
^C
--- origin-www.cisco.COM ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3003ms
rtt min/avg/max/mdev = 8.139/8.927/10.650/1.003 ms
>
```

Pinging an address through a data interface using the routing table

Use the **ping** command. Without specifying an interface, you are testing whether the system can generically find a route to the host. Because this is how the system normally routes traffic, this is typically what you want to test.

ping *host*

For example:

```
> ping 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```



Note You can specify the timeout, repeat count, packet size, and even the data pattern to send. Use the help indicator, ?, in the CLI to see the available options.

Pinging an address through a specific data interface

Use the **ping interface** *if_name* command if you want to test connectivity through a specific data interface.

ping interface *if_name host*

For example:

```
> ping interface inside 171.69.38.1
Sending 5, 100-byte ICMP Echos to 171.69.38.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```

Pinging an address through a data interface using TCP ping

Use the **ping tcp** command. A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet.

ping tcp [**interface** *if_name*] *host port*

You must specify the host and TCP port.

You can optionally specify the interface, which is the source interface of the ping, not the interface through which to send the pings. This type of ping always uses the routing table.

A TCP ping sends SYN packets and considers the ping successful if the destination sends a SYN-ACK packet. For example:

```
> ping tcp 10.0.0.1 21
Type escape sequence to abort.
No source specified. Pinging from identity interface.
Sending 5 TCP SYN requests to 10.0.0.1 port 21
from 10.0.0.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Note You can also specify the timeout, repeat count, and the source address for the TCP ping. Use the help indicator, ?, in the CLI to see the available options.

Tracing Routes to Hosts

If you are having problems sending traffic to an IP address, you can trace the route to the host to determine if there is a problem on the network path. A traceroute works by sending UDP packets on an invalid port, or ICMPv6 echoes, to a destination. The routers along the way to the destination respond with an ICMP Time Exceeded Message, and report that error to traceroute. Each node receives three packets, so you get three chances per node to get an informative result. You can use **traceroute** by opening the CLI console or by logging into the device CLI.

Note There are separate commands for tracing a route through a data interface (**traceroute**) or through the virtual management interface (**traceroute system**). Ensure that you use the right command.

The following table describes the possible result per packet as displayed in the output.

Output Symbol	Description
*	No response was received for the probe within the timeout period.
nn msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable.
!H	ICMP host unreachable.
!P	ICMP protocol unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Tracing a route through the virtual management interface

Use the **traceroute system** command.

traceroute system *destination*

The host can be an IPv4/IPv6 address or fully-qualified domain name (FQDN), such as www.example.com. For example:

```
> traceroute system www.example.com
traceroute to www.example.com (172.163.4.161), 30 hops max, 60 byte packets
 1  192.168.0.254 (192.168.0.254)  0.213 ms  0.310 ms  0.328 ms
 2  10.88.127.1 (10.88.127.1)  0.677 ms  0.739 ms  0.899 ms
 3  lab-gw1.example.com (10.89.128.25)  0.638 ms  0.856 ms  0.864 ms
 4  04-bb-gw1.example.com (10.152.240.65)  1.169 ms  1.355 ms  1.409 ms
 5  wan-gw1.example.com (10.152.240.33)  0.712 ms  0.722 ms  0.790 ms
```

```
 6  wag-gw1.example.com (10.152.240.73)  13.868 ms  10.760 ms  11.187 ms
 7  rbb-gw2.example.com (172.30.4.85)  7.202 ms  7.301 ms  7.101 ms
 8  rbb-gw1.example.com (172.30.4.77)  8.162 ms  8.225 ms  8.373 ms
 9  sbb-gw1.example.com (172.16.16.210)  7.396 ms  7.548 ms  7.653 ms
10  corp-gw2.example.com (172.16.16.58)  7.413 ms  7.310 ms  7.431 ms
11  dmzbb-gw2.example.com (172.16.0.78)  7.835 ms  7.705 ms  7.702 ms
12  dmzdcc-gw2.example.com (172.16.0.190)  8.126 ms  8.193 ms  11.559 ms
13  dcz05n-gw1.example.com (172.16.2.106)  11.729 ms  11.728 ms  11.939 ms
14  www1.example.com (172.16.4.161)  11.645 ms  7.958 ms  7.936 ms
```

Tracing a route through a data interface

Use the **traceroute** command.

traceroute *destination*

The host can be an IPv4/IPv6 address or fully-qualified domain name (FQDN), such as www.example.com, if you configure DNS servers for the data interfaces. If you do not have DNS servers configured for a specific interface, use the **dig** command to look up the IP address of a given FQDN. For example:

```
> traceroute 209.165.200.225
Tracing the route to 209.165.200.225
 1  10.83.194.1 0 msec 10 msec 0 msec
 2  10.83.193.65 0 msec 0 msec 0 msec
 3  10.88.193.101 0 msec 10 msec 0 msec
 4  10.88.193.97 0 msec 0 msec 10 msec
 5  10.88.239.9 0 msec 10 msec 0 msec
 6  10.88.238.65 10 msec 10 msec 0 msec
 7 172.16.7.221 70 msec 70 msec 80 msec
 8 209.165.200.225 70 msec 70 msec 70 msec
```

Note You can specify the timeout, time to live, number of packets per node, and even the IP address or interface to use as the source of the traceroute. Use the help indicator, ?, in the CLI to see the available options.

Making the Device Appear on Traceroutes

By default, the Firewall Threat Defense device does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the device, and increase the rate limit on ICMP unreachable messages. To accomplish this, you must create a FlexConfig object that configures the required service policy rule and other options.

For a detailed discussion of service policies and traffic classes, see the *Cisco ASA Series Firewall Configuration Guide* available from https://www.cisco.com/c/en/us/support/security/asa-firepower-services/products-installation-and-configuration-guides-list.html.

Note If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences. Keep these considerations in mind when defining your traffic class.

Procedure

Step 1 Click **View Configuration** in **Device** > **Advanced Configuration**.

Step 2 Click **FlexConfig** > **FlexConfig Objects** in the Advanced Configuration table of contents.

Step 3 Create the object to decrement TTL.

a) Click the **+** button to create a new object.

b) Enter a name for the object. For example, **Decrement_TTL**.

c) In the **Template** editor, enter the following lines, including indentations.

```
icmp unreachable rate-limit 50 burst-size 1
policy-map global_policy
 class class-default
  set connection decrement-ttl
```

d) In the **Negate Template** editor, enter the lines required to undo this configuration.

Just as you need to include the parent commands to enter the correct sub-mode for a command to enable it, you also need to include those commands in the negate template.

The negate template will be applied if you remove this object from the FlexConfig policy (after having deployed it successfully), and also during an unsuccessful deployment (to reset the configuration to its previous condition).

Thus, for this example, the negate template would be the following:

```
no icmp unreachable rate-limit 50 burst-size 1
policy-map global_policy
 class class-default
  no set connection decrement-ttl
```

e) Click **OK** to save the object.

Step 4 Add the objects to the FlexConfig policy.

Only those objects selected in the FlexConfig policy get deployed.

a) Click **FlexConfig Policy** in the table of contents.

b) Click **+** in the Group List.

c) Select the Decrement_TTL object and click **OK**.

The preview should update with the commands in the template. Verify you are seeing the expected commands.

d) Click **Save**.

You can now deploy the policy.

Troubleshooting NTP

The system relies on accurate and consistent time to function correctly and to ensure that events and other data points are handled accurately. You must configure at least one, but ideally three, Network Time Protocol (NTP) servers to ensure the system always has reliable time information.

The device summary connection diagram (click **Device** in the main menu) shows the status of the connection to the NTP server. If the status is yellow or orange, then there is an issue with the connection to the configured servers. If the connection problem persists (it is not just a momentary issue), try the following.

- First, ensure that you have at least three NTP servers configured on **Device** > **System Settings** > **NTP**. Although this is not a requirement, reliability is greatly enhanced if you have at least three NTP servers.
- Ensure that there is a network path between the management interface IP address (defined on **Device** > **System Settings** > **Management Interface**) and the NTP servers.
  - If the management interface gateway is the data interfaces, you can configure static routes to the NTP servers on **Device** > **Routing** if the default route is not adequate.
  - If you set an explicit management interface gateway, log into the device CLI and use the **ping system** command to test whether there is a network path to each NTP server.
- Log into the device CLI and check the status of the NTP servers with the following commands.
  - **show ntp**—This command shows basic information about the NTP servers and their availability. However, the connectivity status in the Firewall Device Manager uses additional information to indicate the status, so there can be inconsistency in what this command shows and what the connectivity status diagram shows. You can also issue this command from the CLI console.
  - **system support ntp**—This command includes the output of **show ntp** plus the output of the standard NTP command **ntpq**, which is documented with the NTP protocol. Use this command if you need to confirm NTP synchronization.

    Look for the section "Results of 'ntpq -pn.' For example, you might see something like the following:

```
Results of 'ntpq -pn'
remote                   : +216.229.0.50
refid                    : 129.7.1.66
st                       : 2
t                        : u
when                     : 704
poll                     : 1024
reach                    : 377
delay                    : 90.455
offset                   : 2.954
jitter                   : 2.473
```

    In this example, the + before the NTP server address indicates that it is a potential candidate. An asterisk here, *, indicates the current time source peer.

    The NTP daemon (NTPD) uses a sliding window of eight samples from each one of the peers and picks out one sample, then the clock selection determines the true chimers and the false tickers. NTPD then determines the round-trip distance (the offset of a candidate must not be over one-half the round trip delay). If connection delays, packet loss, or server issues cause one or all the candidates to be rejected, you would see long delays in the synchronization. The adjustment also occurs over a very long period of time: the clock offset and oscillator errors must be resolved by the clock discipline algorithm and this can take hours.

Note If the refid is .LOCL., this indicates the peer is an undisciplined local clock, that is, it is using its local clock only to set the time. The Firewall Device Manager always marks the NTP connection yellow (not synchronized) if the selected peer is .LOCL. Normally, NTP does not select a .LOCL. candidate if a better one is available, which is why you should configure at least three servers.

Troubleshooting DNS for the Management Interface

You must configure at least one DNS server for use by the Management interface. The server is needed for cloud connections to services such as smart licensing, database updates (such as GeoDB, rules, and VDB), and any other activity that needs domain name resolution.

Configuring a DNS server is rather trivial. You simply enter the IP addresses of the DNS servers you use when you initially configure the device. You can later change them on the **Device** > **System Settings** > **DNS Server** page.

However, the system can fail to resolve fully-qualified domain names (FQDN) due to network connectivity issues or problems with the DNS server itself. If you find the system cannot use your DNS servers, consider the following actions to identify and resolve the problem. Also see Troubleshooting General DNS Problems.

Procedure

Step 1 Determine if you have a problem.

a) Use SSH to log into the device CLI.

b) Enter **ping system www.cisco.com**. If you get an "unknown host" message like the following, then the system could not resolve the domain name. If the ping is successful, then you are done: DNS is working. (Press Ctrl+C to stop the ping.)

```
> ping system www.cisco.com
ping: unknown host www.cisco.com
```

Note
It is critical that you include the **system** keyword in the **ping** command. The **system** keyword sends the ping through the management IP address, which is the only interface that uses the management DNS server. Pinging www.cisco.com is also a good option, because you need a route to that server for smart licensing and updates.

Step 2 Verify the configuration for the management interface.

a) Click **Device** > **Interfaces**, edit the Management interface, and verify the following. If you make changes, the changes are applied immediately on clicking **Save**. If you change the Management address, you will need to reconnect and log back in.

- The gateway IP address is correct for the Management network. If you using the data interfaces as the gateway, subsequent steps will verify that configuration.

- If you are not using the data interfaces as a gateway, verify that the Management IP address/subnet mask and the gateway IP address are on the same subnet.

b) Click **Device** > **System Settings** > **DNS Server** and verify that the right DNS servers are configured.

If you are deploying the device on your network edge, your service provider might have specific requirements about the DNS server you can use.

c) If you are using the data interfaces as the gateway, verify that you have the required routes.

You need a default route for 0.0.0.0. You might need additional routes if the DNS server is not available through the gateway for the default route. There are two basic situations:

- If you are using DHCP to obtain an address for the outside interface, and you selected the **Obtain Default Route using DHCP** option, the default route is not visible in the Firewall Device Manager. From SSH, enter **show route** and verify that there is a route for 0.0.0.0. Because this is the default configuration for the outside interface, this is a likely situation that you might encounter. (Go to **Device** > **Interfaces** to view the configuration of the outside interface.)
- If you are using a static IP address on the outside interface, or you are not obtaining the default route from DHCP, then open **Device** > **Routing**. Verify that the correct gateway is being used for the default route.

If the DNS server cannot be reached through the default route, you must define a static route to it on the **Routing** page. Note that you should not add routes for directly connected networks, that is, networks that are connected directly to any of the system's data interfaces, as the system can route to those networks automatically.

Also verify that there are no static routes that are misdirecting traffic to the server out the wrong interface.

d) If the deployment button indicates that there are undeployed changes, deploy them now and wait for deployment to complete.



e) Retest **ping system www.cisco.com**. If you still have problems, continue with the next step.

Step 3 In the SSH session, enter **dig www.cisco.com**.

- If **dig** indicates that it got a response from the DNS server, but the server could not find the name, it means that DNS is configured correctly, but the DNS server you are using does not have an address for the FQDN. This error is indicated by the NXDOMAIN status. The response would look similar to the following:

```
> dig www.cisco.com

; <<>> DiG 9.11.4 <<>> www.cisco.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 43246
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1280
; COOKIE: 78b1c6b2b3ef5b689fc2f65260db9e9b36a7d9fefb301943 (good)
;; QUESTION SECTION:
;www.cisco.com.                 IN      A
```

```
;; AUTHORITY SECTION:
.                       3600    IN      SOA     a.root-servers.net.
nstld.verisign-grs.com. 2021062901 1800 900 604800 86400

;; Query time: 13 msec
;; SERVER: 10.163.47.11#53(10.163.47.11)
;; WHEN: Tue Jun 29 22:28:43 UTC 2021
;; MSG SIZE  rcvd: 145
```

Resolution: In this case, you need to configure a different DNS server, or get the one you have updated so it can resolve the FQDNs you need resolved. Work with your network administrator or ISP to get the IP address of a DNS server that will work for your network.

- If the command times out, then the system cannot reach your DNS servers, or all of the DNS servers are currently down and not responding (which is less likely). Continue with the next step.

Step 4 Use the **traceroute system** *DNS_server_ip_address* command to trace the route to the DNS server.

For example, if the DNS server is 10.100.10.1, enter:

```
> traceroute system 10.100.10.1
```

Following are the possible results:

- The traceroute completes and reaches the DNS server. In this case, there is in fact a route to the DNS server and the system can reach it. Thus, there is no routing problem. However, for some reason, DNS requests to this server are not getting a response.

  Resolution: There is a possibility that a router or firewall along the path is dropping UDP/53 traffic, which is the port used for DNS. You might try a DNS server along a different network path. This is a difficult problem to resolve, as you will need to determine which node is blocking traffic, and work with the system administrator to get the access rules changed.

- The traceroute cannot reach even one node, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1  * * *
 2  * * *
 3  * * *
(and so forth)
```

  Resolution: In this case, the routing problem is within your system. Try doing a **ping system** for the gateway IP address. Re-verify the configuration of the management interface as mentioned in earlier steps, and ensure that you have the required gateways and routes configured.

- The traceroute makes it through a few nodes before it can no longer resolve the route, which would look like the following:

```
> traceroute system 10.100.10.1
traceroute to 10.100.10.1 (10.100.10.1), 30 hops max, 60 byte packets
 1  192.168.0.254 (192.168.0.254)  0.475 ms  0.532 ms  0.542 ms
 2  10.88.127.1 (10.88.127.1)  0.803 ms  1.434 ms  1.443 ms
 3  site04-lab-gw1.example.com (10.89.128.25)  1.390 ms  1.399 ms  1.435 ms
 4  * * *
 5  * * *
```

```
6  * * *
```

Resolution: In this case, routing breaks down at the last node. You might need to work with the system administrator to get correct routes installed in that node. However, if there is intentionally no route to the DNS server through the node, you need to change your gateway, or create your own static route, to point to a router that can route traffic to the DNS server.

Analyzing CPU and Memory Usage

To view system-level information about CPU and memory usage, select **Monitoring** > **System** and look for the CPU and Memory bar graphs. These graphs show information collected through the CLI using the **show cpu system** and **show memory system** commands.

If you open the CLI console or log into the CLI, you can use additional versions of these commands to view other information. Typically, you would look at this information only if you are having persistent problems with usage, or at the direction of the Cisco Technical Assistance Center (TAC). Much of the detailed information is complex and requires TAC interpretation.

Following are some highlights of what you can examine. You can find more detailed information about these commands in Cisco Firepower Threat Defense Command Reference at http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

- **show cpu** displays data plane CPU utilization.
- **show cpu core** displays usage for each CPU core separately.
- **show cpu detailed** displays additional per-core and overall data plane CPU usage.
- **show memory** displays data plane memory usage.

Note Some of the keywords (not mentioned above) require that you first set up profiling or other features using the **cpu** or **memory** commands. Use these features at the direction of TAC only.

Viewing Logs

The system logs information for a wide variety of actions. You can use the **system support view-files** command to open a system log. Use this command while working with the Cisco Technical Assistance Center (TAC) so that they can help you interpret the output and to select the appropriate log to view.

The command presents a menu for selecting a log. Use the following commands to navigate the wizard:

- To change to a sub-directory, type in the name of the directory and press Enter.
- To select a file to view, enter **s** at the prompt. You are then prompted for a file name. You must type the complete name, and capitalization matters. The file list shows you the size of the log, which you might consider before opening very large logs.
- Press the space bar when you see --More-- to see the next page of log entries; press Enter to see just the next log entry. When you reach the end of the log, you are taken to the main menu. The --More-- line

shows you the size of the log and how much of it you have viewed. **Use Ctrl+C to close the log and exit the command if you do not want to page through the entire log.**

- Type **b** to go up one level in the structure to the menu.

If you want to leave the log open so you can see new messages as they are added, use the **tail-logs** command instead of **system support view-files**.

The following example shows how view the cisco/audit.log file, which tracks attempts to log into the system. The file listing starts with directories at the top, then a list of files in the current directory.

```
> system support view-files

===View Logs===

============================
Directory: /ngfw/var/log
----------sub-dirs----------
cisco
mojo
removed_packages
setup
connector
sf
scripts
packages
removed_scripts
httpd
-----------files------------
2016-10-14 18:12:04.514783 | 5371       | SMART_STATUS_sda.log
2016-10-14 18:12:04.524783 | 353        | SMART_STATUS_sdb.log
2016-10-11 21:32:23.848733 | 326517     | action_queue.log
2016-10-06 16:00:56.620019 | 1018       | br1.down.log

<list abbreviated>


([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: cisco

============================
Directory: /ngfw/var/log/cisco
-----------files------------
2017-02-13 22:44:42.394907 | 472        | audit.log
2017-02-13 23:40:30.858198 | 903615     | ev_stats.log.0
2017-02-09 18:14:26.870361 | 0          | ev_stats.log.0.lck
2017-02-13 05:24:00.682601 | 1024338    | ev_stats.log.1
2017-02-12 08:41:00.478103 | 1024338    | ev_stats.log.2
2017-02-11 11:58:00.260805 | 1024218    | ev_stats.log.3
2017-02-09 18:12:13.828607 | 95848      | firstboot.ngfw-onbox.log
2017-02-13 23:40:00.240359 | 6523160    | ngfw-onbox.log

([b] to go back or [s] to select a file to view, [Ctrl+C] to exit)
Type a sub-dir name to list its contents: s

Type the name of the file to view ([b] to go back, [Ctrl+C] to exit)
> audit.log
2017-02-09 18:59:26 - SubSystem:LOGIN, User:admin, IP:10.24.42.205, Message:Login successful,

2017-02-13 17:59:28 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:36 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login failed,
```

```
2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Login successful,

2017-02-13 22:44:42 - SubSystem:LOGIN, User:admin, IP:10.24.111.72, Message:Unlocked account.,

<remaining log truncated>
```

Creating a Troubleshooting File

Cisco Technical Assistance Center (TAC) personnel might ask you to submit system log information when you submit a problem report. This information assists them with diagnosing the problem. You do not need to submit a diagnostics file unless asked to do so.

The following procedure explains how to create and download the diagnostics file.

Procedure

Step 1 Click **Device**.

Step 2 Under **Troubleshooting**, click **Request File to be Created** or **Re-Request File to be Created** (if you have created one before).

The system starts generating the diagnostic file. You can go to other pages and return here to check status. When the file is ready, the date and time of the file creation is shown along with a download button.

Step 3 When the file is ready, click the download button.

The file is downloaded to your workstation using your browser's standard download method.

Hardware Management Tasks

The following topics explain how to do some hardware maintenance tasks. See the various hardware guides for more complete information and information about other hardware management tasks.

Replacing an ISA 3000 Device

The ISA 3000 has an SD card that you can remove and insert in another ISA 3000 device. If you create system backups on the SD card, you can use this capability to easily replace a device. You simply take the SD card from the failing device and insert it in the new device. The backups are then available for you to restore.

To ensure that you have the necessary backups, configure the backup job to create the backup on the SD card.

Hot swap an SSD

If you have two SSDs, they form a RAID when you boot up. You can perform the following tasks at the Firewall Threat Defense CLI while the firewall is powered up:

- Hot swap one of the SSDs—If an SSD is faulty, you can replace it. Note that if you only have one SSD, you cannot remove it while the firewall is powered on.
- Remove one of the SSDs—If you have two SSDs, you can remove one.
- Add a second SSD—If you have one SSD, you can add a second SSD and form a RAID.

  When you use only one SSD in Secure Firewall devices, the **Drive State** will show as **degraded**.



Caution Do not remove an SSD without first removing it from the RAID using this procedure. You can cause data loss.

Procedure

Step 1 Remove one of the SSDs.

a) Remove the SSD from the RAID.

configure raid remove-secure local-disk {**1** | **2**}

The **remove-secure** keyword removes the SSD from the RAID, disables the self-encrypting disk feature, and performs a secure erase of the SSD. If you only want to remove the SSD from the RAID and want to keep the data intact, you can use the **remove** keyword.

Example:

```
> configure raid remove-secure local-disk 2
```

b) Monitor the RAID status until the SSD no longer shows in the inventory.

show raid

After the SSD is removed from the RAID, the **Operability** and **Drive State** will show as **degraded**. The second drive will no longer be listed as a member disk.

Example:

```
> show raid
Virtual Drive
ID:                     1
Size (MB):              858306
Operability:            operable
Presence:               equipped
Lifecycle:              available
Drive State:            optimal
Type:                   raid
Level:                  raid1
Max Disks:              2
Meta Version:           1.0
Array State:            active
Sync Action:            idle
Sync Completed:         unknown
Degraded:               0
Sync Speed:             none
```

```
RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

Device Name:                nvme1n1
Disk State:                 in-sync
Disk Slot:                  2
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                degraded
Presence:                   equipped
Lifecycle:                  available
Drive State:                degraded
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   1
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:
```

c) Physically remove the SSD from the chassis.

Step 2 Add an SSD.

a) Physically add the SSD to the empty slot.

b) Add the SSD to the RAID.

configure raid add local-disk {**1** | **2**}

It can take several hours to complete syncing the new SSD to the RAID, during which the firewall is completely operational. You can even reboot, and the sync will continue after it powers up. Use the **show raid** command to show the status.

If you install an SSD that was previously used on another system, and is still locked, enter the following command:

configure raid add local-disk {**1** | **2**} *psid*

The *psid* is printed on the label attached to the back of the SSD. Alternatively, you can reboot the system, and the SSD will be reformatted and added to the RAID.