TREND MICRO™

4.1.0

TippingPoint™ Threat Protection System (TPS)

Local Security Manager User Guide

Actionable threat defense against advanced targeted attacks.

TippingPoint

Threat Protection System
Local Security Manager User Guide

Version 4.1.0

May 2016

Legal and notice information

Contents

About this guide

Welcome to the Local Security Manager User Guide.

This section covers the following topics:

Target audience

The intended audience includes technicians and maintenance personnel responsible for installing, configuring, and maintaining TippingPoint security systems and associated devices. Users should be familiar with networking concepts and the following standards and protocols:

- TCP/IP
- UDP
- ICMP
- RADIUS
- LDAP
- Ethernet
- Network Time Protocol (NTP)
- Secure Sockets Layer (SSL)
- Simple Network Time Protocol (SNTP)
- Simple Mail Transport Protocol (SMTP)
- Simple Network Management Protocol (SNMP)

Related documentation

A complete set of product documentation for this product is available online at the Threat Management Center (TMC): *https://tmc.tippingpoint.com*. The product document set generally includes conceptual and deployment information, installation and user guides, CLI command references, safety and compliance information and release notes.

For information about how to access the online product documentation, refer to the *Read Me First* document in your product shipment or on the TMC.

Conventions

This information uses the following conventions.

Typefaces

TippingPoint uses the following typographic conventions for structuring information:

Convention	Element
Bold font	• Key names • Text typed into a GUI element, such as into a box • GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes. Example: Click **OK** to accept.
Italics font	Text emphasis, important terms, variables, and publication titles
`Monospace font`	• File and directory names • System output • Code • Text typed at the command-line
`Monospace, italic font`	• Code variables • Command-line variables
`Monospace, bold font`	Emphasis of file and directory names, system output, code, and text typed at the command line

Messages

Messages are special text that is emphasized by font, format, and icons.

⚠**Warning!** Alerts you to potential danger of bodily harm or other potential harmful consequences.

Caution: Provides information to help minimize risk, for example, when a failure to follow directions could result in damage to equipment or loss of data.

Note: Provides additional information to explain a concept or complete a task.

Important: Provides significant information or specific instructions.

Tip: Provides helpful hints and shortcuts, such as suggestions about how to perform a task more easily or more efficiently.

Support information

Contact the TippingPoint Technical Assistance Center (TAC) by using any of the following options.

Note: Have the following information about your product available:

- Serial number and/or software version for your product
- System logs or event logs if available for your product

Online support

Go to the TippingPoint Threat Management Center (TMC) at:

https://tmc.tippingpoint.com/TMC/

Phone support

North America: +1 866 681 8324

International: see *https://tmc.tippingpoint.com/TMC/*

TPS concepts and the LSM

The TippingPoint Threat Protection System (TPS) protects your network by scanning, detecting, and responding to network traffic according to the filters, action sets, and global settings maintained on each device by a client. Each device provides intrusion prevention for your network according to the amount of network connections and hardware capabilities.

The Local Security Manager (LSM) provides a user-friendly, browser-based GUI for administering the TPS.

This topic describes TPS concepts and the login and navigation procedures of the LSM user interface.

TPS deployment

A single TPS can be installed at the perimeter of your network, at the network core, on your intranet, or in all three locations. The following diagram shows an example of a corporate network with the TPS deployed to a variety of locations.

Figure 1. TPS deployment example

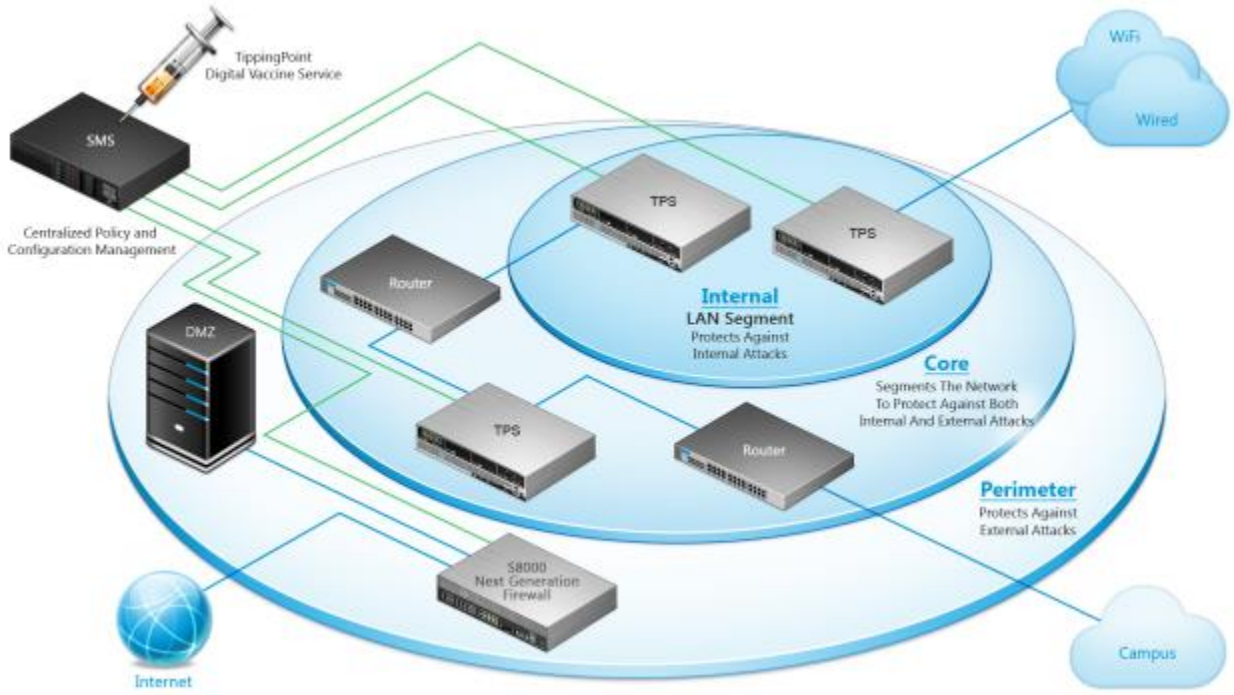

Threat Suppression Engine

The main component of the TPS is the Threat Suppression Engine (TSE), a custom engine that detects and blocks a broad range of attacks at wire speeds. The TSE is a flow-based network security engine, in which each packet is identified as a component of a flow and each flow is tracked in the connection table on the device. A flow is uniquely identified by its packet header information:

- IPv4 or IPv6 protocol (ICMP, TCP, UDP, other)
- source and destination IP addresses
- source and destination ports

The TSE reconstructs and inspects flow payloads by parsing the traffic at the application layer. When a packet matches an IPS filter, the device handles the packets based on the action set configured on the filter. For example, if the action set is **Block**, then the packet is dropped and subsequent packets from the same flow are dropped without inspection. The device provides default actions to block or permit traffic with options to quarantine or rate-limit traffic and to notify users or systems when an action executes. Logging options are also available to review the types of traffic being filtered by the device. You can customize the default action sets, or create your own based on your network requirements.

TPS filtering

The TSE uses Digital Vaccine (DV) filters to police your network and to screen out malicious or unwanted traffic. In addition to the DV filters, the TPS also provides Traffic Management filters, which are custom filters that react to traffic based on source IP address, destination IP address, port, protocol, or other

defined values. Traffic management filters are applied to traffic before DV filters. Depending on how the filters are configured, traffic might or might not require further inspection.

The Digital Vaccine package

DV filters are contained in a Digital Vaccine (DV) package. A DV package is installed and configured to provide out-of-the-box TPS protection for the network. After setting up the TPS, you can customize the filters in the DV through the LSM. To ensure that you have the most up-to-date DV package, use the Update page in the LSM to download the latest package. See *Update the device* on page 118.

The filters within the DV package are developed by TippingPoint's Digital Vaccine Labs to protect the network from specific exploits as well as potential attack permutations to address for Zero-Day threats. These filters include traffic anomaly filters and vulnerability-based filters. Vulnerability-based filters are designed to protect the network from an attack that takes advantage of a weakness in application software. For viruses that are not based on a specific vulnerability in software, the DV provides signature filters. TippingPoint delivers weekly DV updates that can be automatically installed on the TPS device (**System > Update**). If a critical vulnerability or threat is discovered, DV updates are immediately distributed to customers. See *Enable automatic DV updates* on page 120.

Tip: In addition to providing a download location for Digital Vaccine packages, the TMC also provides DV product documentation that includes more detailed information about the filters included in the DV package, filter updates, and other related information.

Additional Digital Vaccine filter subscription services are offered by DVLabs for organizations that experience heavier risk factors for threats that go beyond the scope of the standard Digital Vaccine coverage. These services include the following services:

- Reputation Feed (Rep Feed) — provides reputation filters for suspect IP addresses and domains.
- Malware Filter Package — provides advanced malware protection.

For information about registering for a Digital Vaccine subscription service, contact your TippingPoint customer representative.

Filter components

TPS filters have the following components, which determine the filter type, global and customized settings, and how the system responds when the TSE finds traffic matching the filter:

- **Category** — Defines the type of network protection provided by the filter. The category is also used to locate the filter in the LSM and to control the global filter settings using the Category Setting configuration.
- **Action set** — Defines the actions that execute when the filter is matched.

- **Adaptive Filter Configuration State** — Allows you to override the global Adaptive Filter configuration settings so that the filter is not affected by adaptive filtering. See also *Adaptive filtering* on page 8.
- **State** — Indicates if the filter is enabled or disabled. If the filter is disabled, the TSE does not use the filter to evaluate traffic.

Category settings

Category settings are used to configure global settings for all filters within a specified category group. DV filters are organized into groups based on the type of protection provided:

- **Application Protection Filters** defend against known exploits and exploits that can take advantage of known vulnerabilities targeting applications and operating systems. This filter type includes the subcategories *Exploits*, *Identity Theft*, *Reconnaissance*, *Security Policy*, *Spyware*, *Virus*, and *Vulnerabilities*.
- **Infrastructure Protection Filters** protect network bandwidth and network infrastructure elements, such as routers and firewalls, from attack by using protocols and detecting statistical anomalies. This filter type includes the subcategories *Network Equipment* and *Traffic Normalization*.
- **Performance Protection Filters** block or rate-limit traffic from applications that can consume excessive bandwidth, leaving network resources available for use by key applications. This filter type includes the subcategories *IM*, *P2P*, and *Streaming Media*.

Category Settings are used to assign global configuration settings to filters in a subcategory. For example, if you decide not to use any filters to monitor P2P traffic, you can change the category settings for the Performance Protection P2P filter group to disable these filters. Category settings consist of the following global parameters:

- **State** — Determines whether filters within the subcategory are enabled or disabled. If a category is disabled, all filters in the category are disabled.
- **Action Set** — Determines the action set that filters within a category execute when a filter match occurs. If the *Recommended* action set is configured, filters within the category are configured with the settings recommended by the Digital Vaccine team. If required, you can override the category setting on individual filters by editing the filter to define custom settings.

Filter override settings

For the best system performance, TippingPoint recommends that you use global category settings and the *Recommended* action set for all DV filters. However, in some cases, you might need to override the category settings and recommended action for individual filters due to specific network requirements, or in cases where the recommended settings for a filter interact poorly with your network.

Filter override settings specify custom settings to be applied to the filter in the Security Profile. After a filter has been customized, it is not affected by the global Category Settings that specify the filter State and Action.

Filter limits and exceptions

Limits and exceptions change the way filters are applied based on IP address. For example, you can specify a limit setting so that filters only apply to specific source and destination IP addresses or address ranges. If a filter has both global and filter-level exception settings, the Threat Suppression Engine uses the filter-level settings to determine how to apply the filter. You can configure the following limit and exceptions from the LSM:

- **Filter Exceptions** (specific) — Allow traffic that would normally trigger a filter to pass between specific addresses or address ranges without triggering the filter. Configured from the Filter Edit page, these exceptions apply only to the filter where they were configured.
- **Limit Filter to IP Addresses** (global) — Only apply filters to traffic between specified source and destination IP address pairs. You can configure IP address limits that apply to Application Protection, Traffic Normalization, and Network Equipment Protection filter types. You can configure separate limits that apply only to Performance Protection filters.
- **Exceptions** (global) — Exclude traffic between specified source and destination IP address pairs. You can configure exceptions for the following filter types: Application Protection, Traffic Normalization, Network Equipment Protection, and Performance Protection filters. These exceptions are global for all specified filters.

For more information, see *Edit an IPS profile* on page 59.

Adaptive filtering

With Adaptive Filtering, the Threat Suppression Engine automatically manages filter behavior when the device is under extreme load conditions. This feature protects against the potential adverse effects of a filter that interacts poorly with the network environment by preventing the device from entering High Availability mode. For more information, see *Adaptive filter control* on page 100.

Adaptive filtering works by monitoring each filter to identify any suspected of causing congestion. When it identifies a filter, it manages the filter using one of the following methods, depending on how the global or filter-level Adaptive Filtering is configured:

- **Automatic Mode** — This setting enables the device to automatically disable and generate a system message regarding the defective filter.
- **Manual** — This setting enables the device to generate a system message regarding the defective filter. However, the filter is not disabled.

Jumbo frame support

The TippingPoint Operating System supports inspection of jumbo frames up to 9234 bytes. This includes the 14-byte Ethernet header, 9216 bytes of payload data, and the 4-byte Ethernet checksum.

TippingPoint TPS-specific features

The TippingPoint Threat Protection System (TPS) consists of the TippingPoint 440T and 2200T. These devices provide slightly different support for the features listed in the following sections.

SSL inspection

SSL inspection provides in-line, real-time threat protection for inbound IPv4 traffic that is SSL encrypted. SSL inspection is licensed separately.

Device support: 2200T only

See *SSL inspection profiles and servers* on page 66 for more information.

Throughput upgrades

Throughput upgrades increase the amount of traffic that the device monitors. Throughput upgrades are licensed separately.

Device support: 440T and 2200T

See *Update the license package* on page 121for more information.

VLAN translation

VLAN translation enables the TPS to selectively inspect traffic based on the switch configuration at the aggregation or distribution switch. This feature translates traffic between different VLANs or between VLAN and non-VLAN interfaces.

Device support: 440T and 2200T

See *VLAN translation* on page 53 for more information.

Best effort mode

Best Effort mode protects latency-sensitive applications by not inspecting packets if the latency introduced by inspecting them exceeds the configured threshold. When the latency reaches the specified threshold, permitted traffic is not inspected until latency falls to the user-defined recovery percentage.

Device support: 2200T only

When performing SSL inspection, the latency measure and relief only apply on inspection, and do not apply to the SSL and TCP proxy connections.

Best Effort mode is enabled from the CLI with the `debug np best-effort` command. For detailed information about Best Effort mode, refer to the *Threat Protection System Command Line Interface Reference* .

Inspection bypass rules

Inspection bypass rules describe traffic to be directed through the TPS without inspection. These rules can be applied to traffic according to source or destination IP address, port, or CIDR (Classless Inter-Domain Routing), or to traffic moving through specific ports.

Device support: 440T and 2200T

See *Inspection bypass rules* on page 75 for more information.

IPv6 inspection and management

IPv6 traffic inspection, and IPv6 options are available when configuring the Security Profile options. The majority of existing TippingPoint filters are compatible with both IPv4 and IPv6 traffic. The host management port, default gateway, and management port routes can also be configured with IPv6 addresses. Named network support is not available with IPv6 inspection and management.

Device support: 440T and 2200T

Tip: Named networks, accessible from the LSM through the **System > Named Networks** page, enables you to assign names to specific IPv4 and IPv6 address prefixes.

Inspection of tunneled traffic

Inspection of tunneled traffic includes a wide range of tunneled traffic:

- GRE (Generic Routing Encapsulation)
- GTP (GPRS Tunneling Protocol)
- Mobile IPv4 (IP-in-IP)
- IPv6, including 6-in-4, 4-in-6, and 6-in-6
- Tunnels up to 10 layers of tunneling or a header size of 256 bytes.

Device support: 440T and 2200T

LSM prerequisites

This topic describes the prerequisites for using the LSM. Be sure to first read the *Release Notes* for any late-breaking information that supersedes this document.

Initial setup and installation

Before you can log in to the LSM web interface, complete the initial hardware installation and setup, and connect the appliance to the network. For instructions on installation and setup, see the *Installing your security device* card that shipped with your device.

Note: The device blocks traffic until it has completed the boot sequence.

Browser considerations

You can access the LSM through the following browsers:

- Firefox, V10 or later
- Chrome, V17 or later
- Internet Explorer 8, or later
- Safari, V5.1 or later

Because the LSM manages the TPS through a web browser, take the following security precautions:

- If password caching is on by default in your browser, turn it off.
- Use HTTPS (not HTTP) to ensure secure network communications.

For the latest information on supported browsers, see the *Release Notes*.

Screen resolution

The minimum screen resolution is 1366 x 768. For best results set your screen resolution to 1440 x 900. Lower resolutions might not fully display the contents of some LSM pages.

Logging in to the LSM

The Threat Protection System (TPS) provides simultaneous support for up to 10 web client connections, 10 telnet/SSH (for CLI) connections, and one console connection. Logging in with the CLI is discussed in the *Threat Protection System Command Line Interface Reference*.

After completing the installation steps outlined in the *Installing your security device* card that shipped with your product, log on to the LSM using a supported browser.

1. In the web browser address bar, enter `https://` followed by the IP address or hostname of your TPS.

   Note: The TPS uses a factory-default certificate to secure HTTPS communications from your web browser to the appliance. When you log in to the LSM, you will see an Untrusted Authority warning and a prompt asking if you want to trust the certificate. You can save the certificate to the Trusted Root Certificate store to avoid this warning.

2. The LSM login page is displayed in the browser.
3. Enter your username and password.
4. Click **Log On**. The LSM confirms that your username is valid on the device. If the username is valid, the LSM opens and displays the Dashboard. If the username is not valid, the LSM login page is displayed again.

To exit the LSM, click the Log Off link in the upper-right corner of the page.

Note: When there has been no LSM activity for 15 minutes, connection to the device times out.

Note: Your LSM user role controls what you can see and do within the LSM. User roles have specific capabilities assigned that determine if you have full read and write access or read-only access. For information about user roles, see *User roles* on page 97.

Security notes

Because the LSM manages the device through a web browser, take the following security precautions. Failure to follow these security guidelines can compromise the security of your device.

- Some browser features, such as password caching, are inappropriate for security use and should be turned off.
- The LSM only accepts encrypted HTTPS connections. Unencrypted HTTP connections are not supported.

Caution: TippingPoint recommends that you use the most current version of Internet Explorer or Firefox. For the best user experience, follow these browser recommendations:

- **Internet Explorer**

  Change your cache setting in Internet Explorer for improved browser reliability with TippingPoint devices. Open the Internet Options for your browser (**Tools > Internet Options**). On the General tab, select the **Settings** option for Temporary Internet Files. In the Check for new versions section, select **Every visit to the page**. Save these settings.

  Cookies for previous versions of the LSM might conflict with cookies in the updated version. If the browser receives `404 Page Not Found` errors or displays blank LSM frames, the cookies on the computer might be out of sync. To remedy this, delete the existing cookies and open a new session. On the General tab of the Internet Options dialog, click **Delete Cookies**. Restart Internet Explorer, connect to the LSM, and continue as before.

- **Mozilla Firefox**

  Certificate exceptions cannot be added when managing an IPv6 device on an IPv6 network with Firefox 4 or later. To add a certificate exception in an IPv6 environment, use a different browser or the CLI.

  If your browser receives `404 Page Not Found` errors or displays blank LSM frames, the cookies on the computer might be out of sync. To resolve these issues, clear the cache, delete the cookies, and restart the browser.

- **Pop-Up Blocking**

  If your browser has pop-up blocking enabled, some elements in the LSM might not display correctly. TippingPoint recommends that you enable pop-ups for the LSM by adding the device's URL to the browser pop-up whitelist.

The LSM screen layout

The LSM screen displays information in the following areas:

- **Banner** — Appears along the top of each page and displays basic information, such as the device model number and an online help link.
- **Navigation menus** — Provides access to the different functional areas of the LSM.
- **Workspace** — Displays the pages from which you can monitor the device operation and performance, view current configuration settings, and modify configuration. When you initially log in to the LSM, the workspace automatically displays the Dashboard. When you select a submenu item from the menu bar, the workspace displays the current configuration information for that feature in the form of a table or list.

The following LSM screen capture shows the major areas of the Dashboard, including the banner, navigation menus, and workspace.

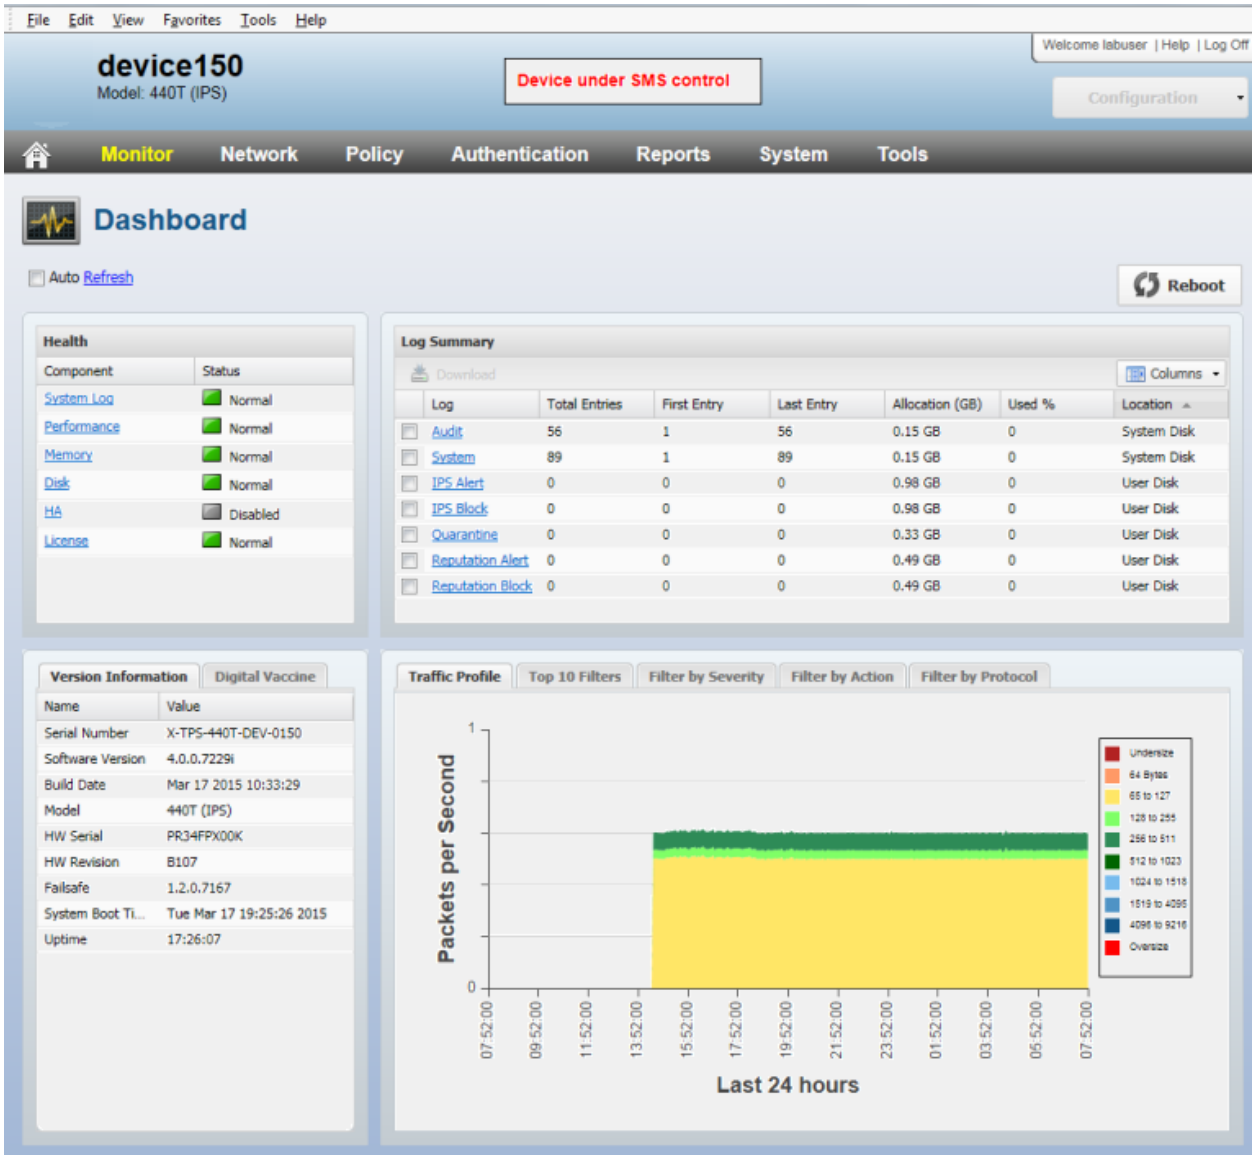

Banner

The banner is displayed along the top of each page and displays the following information:

- The host name
- Appliance model number
- Welcome <username>
- **Help** link
- **Log Off** link
- The Configuration drop-down menu

The Configuration menu enables you to save configuration changes to the TPS. For more information, see *Save configuration changes* on page 16.

Notification and information messages are displayed periodically in the banner when the LSM completes an operation or otherwise updates you with information. For example, if you commit configuration changes, the banner displays a notification indicating that the operation was successful.

Navigation menus

The navigation menu bar provides access to the different functional ares of the LSM. Each menu contains submenus that are organized into logical categories that facilitate access to specific feature sets.

Use the menu bar as follows:

- Hover over a menu to display the associated submenus.
- Click a submenu to display configurable options in the workspace.
- Click the Home icon from anywhere in the LSM to return to the Dashboard.

The following screen capture shows the menu bar with the Monitor menu expanded:

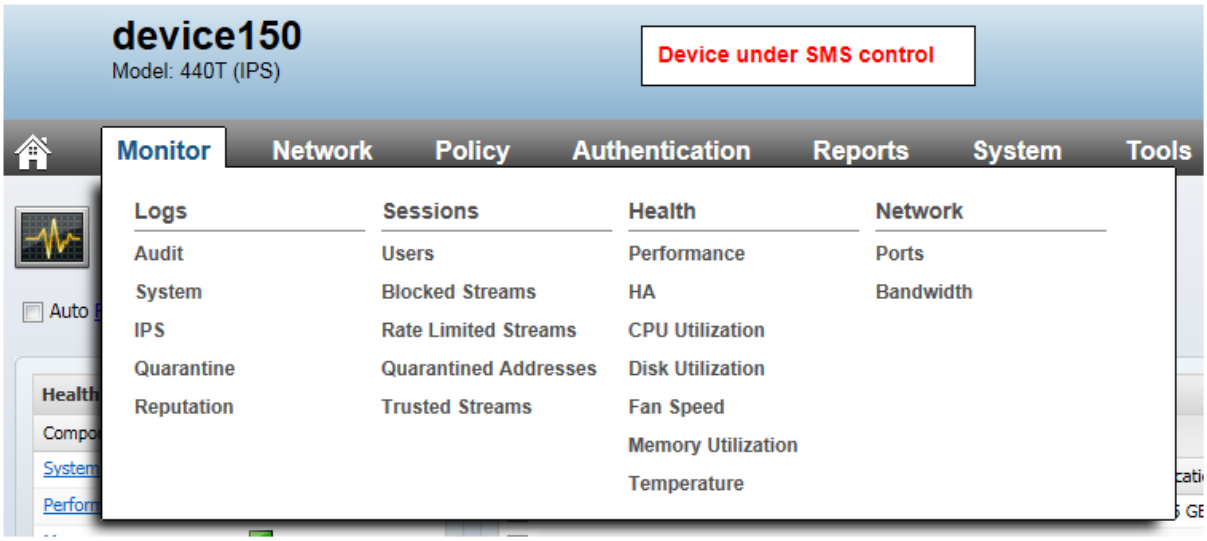

Workspace

The workspace occupies the largest part of the LSM interface. When you initially log in to the LSM, the workspace automatically displays the Dashboard. When you select a submenu item from the menu bar, the workspace displays the current configuration information for that feature in the form of a table or list. It also provides add, edit, and delete options, or feature-specific configurable options.

Common tasks

The LSM provides a consistent presentation on each page so you can easily configure the appliance. There are several areas and tasks in the LSM that you perform the same way, regardless of the specific feature or page you are currently working on. The following topics describes how to perform these common tasks:

Add, Edit, and Delete

To perform an add, edit, or delete operation, click **Add**, **Edit**, or **Delete** located on the left side of the workspace. The following screen capture shows the **Add**, **Edit**, and **Delete** functions.

Add an object

1. Select the menu and submenu for the object you want to add.
2. Click **Add**, just beneath the workspace title or submenus. The LSM displays the appropriate configuration options for the object you want to add.

Edit an existing object

1. Select the menu and submenu for the object you want to edit.
2. Click the checkbox next to the object and then click **Edit**. The LSM displays the appropriate configuration options for the object you want to edit.

Delete an object

1. Select the menu and submenu for the object you want to delete.
2. Click the checkbox next to the object and then click **Delete**.

Save configuration changes

The *Start configuration* is the last saved configuration and maintains the initial configuration of the device. When you reboot the device, the Start configuration is applied to the device.

The *Running configuration*is the Start Configuration plus any committed changes from all users of the device since the last reboot. When you log in to the LSM, it loads the Running configuration. When a user commits

their configuration changes, or when changes are committed automatically, the changes are committed to the Running configuration, and the changes are visible to all users. However, when the device reboots, the Running configuration is reset to the Start configuration. Uncommitted changes and committed changes in the Running configuration are lost.

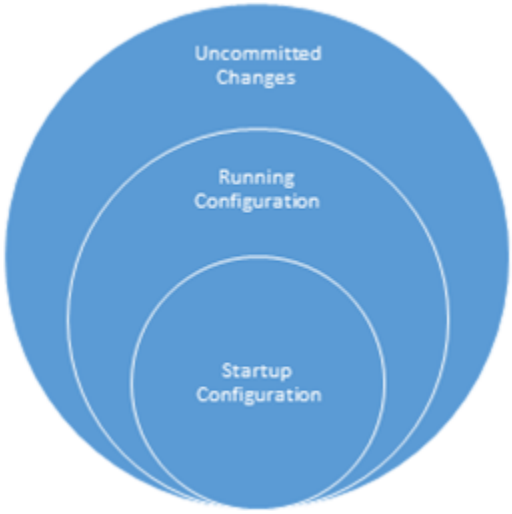

To update the Start configuration, copy the Running configuration to the Start configuration.

1. Select **Configuration** > **Copy Running Configuration to Start**.
2. Select **Configuration** > **Commit pending changes and copy to Start**. This combines the steps for committing and saving changes.

For your convenience, the LSM displays the pending changes count on the Configuration menu. Pending changes that are saved but not committed are lost when you exit the LSM.

Commit inspection profile changes to the device

From the SMS or the LSM, commit inspection profile changes to the appropriate 2200T devices. See the following sections for more information.

Copy the Running configuration to the Start configuration

Copy the Running configuration to the Start configuration to persist any committed configuration changes to the device:

- The *Start configuration* determines the last known configuration of the device and is automatically applied when you reboot the device.
- The *Running configuration* persists any committed configuration changes to the device. When the device reboots, any committed changes in the Running configuration are lost.

Use the Configuration menu to update the Start configuration:

- If you have committed changes to the Running configuration, copy them to the Start configuration by clicking **Copy Running Configuration to Start**.
- If you have uncommitted and committed changes, commit the uncommitted changes to the Running configuration and then copy the Running configuration to the Start configuration by clicking **Commit pending changes and copy to Start**.

View and discard Pending changes

Uncommitted changes are placed into a Pending state until you explicitly commit them to the Running configuration. If you log out of the LSM without committing your Pending changes, the changes are lost.

Use the Configuration menu to manage your Pending changes:

- To view Pending changes, click **View pending changes journal**.
- To discard Pending changes, click **Discard pending changes**.

Refresh the page

You can use one or both of the following page refresh methods:

- **Auto Refresh**

  Click the **Auto Refresh** checkbox to refresh the contents of the page every 60 seconds.

- **Refresh**

  Click the **Refresh** link to perform an instant refresh of the page.

You can force an instant-refresh at any time, even if you enabled **Auto Refresh**.

Note: If no changes are pending in the LSM, you will see any deferred commit changes made by the SMS or CLI when you manually refresh the screen or navigate to a different screen. For more information on deferred commits, see .

Search

When you have a long list of objects in a table or list, the Search mechanism helps you quickly locate a specific object or set of objects with similar settings. The Search feature is located in the upper-right corner of pages on which searching is enabled. For examples of Search and Advanced Search, you can view the **Monitor > Logs > Audit** page.

Perform a search

1. Enter a text string (regular expressions allowed) in the Search field and click **Go**.

A Search Result message is displayed in the banner and any matched records are displayed in the table or list.

2. Click **Clear** to clear the search parameters and return to the default list of objects.

Perform an advanced search

1. Click **Advanced** to display the advanced search options.
2. Enter your search criteria using the pull-down menus and field.
3. Click the plus sign (**+**) to add additional search parameters and further refine your search.
4. Click **Go**.

Sort columns

Some columns of data can be sorted in ascending or descending order. To sort data, click the arrow on the right side of a column, then select **Sort Ascending** or **Sort Descending**.

Show/Hide columns

Use the **Columns** pull-down menu on the right side of a table to hide or show columns in a table. Click the checkbox next to a column name to remove it.

Alternatively, you can click the right side of a column to see the Columns menu, and select or deselect a column heading. Note that the Columns menu is not visible unless you hover over the right side of the column heading.

Dashboard

The Dashboard is the home page for the LSM and is displayed automatically each time you log in. After you have fully configured the system, you can use the Dashboard to quickly assess policy and system performance. You can access the Dashboard at any time by clicking the **Home** icon on the left side of the menu bar.

This topic contains the following information:

Dashboard panels

The Dashboard contains the following four panels. Each contains a specific set of related information.

- Health – Shows the system health status for components of the device, including System Log, Performance, Memory, Disk, HA, and License.
- Log Summary – Lists available logs.
- Performance Graphs – Displays a visual representation of performance patterns of the device, including Temperature, CPU, Disk, and Memory.
- Version Information – Displays build and version information for the hardware, software, Digital Vaccine, and Failsafe. Shows the time of the last system boot and total uptime since the initial boot-up or reboot.

Health

The Health panel includes color health indicators for each of the following components:

- System Log
- Performance
- Memory
- Disk
- HA

- License

For detailed information about each of the health indicators, click on the corresponding links. The colors indicate the current state of each component:

- **Green** — No problems
- **Yellow** — Major warning
- **Red** — Critical warning
- **Grey** — Service is disabled

Click Major and Critical warning indicators to view the error that caused the condition. When you view the error, the indicator is reset and its color changes back to green.

View logs

Use the Log Summary panel of the Dashboard to view and interact with logs.

Review logs in detail by clicking the log link or by viewing the **Monitor > Logs >** `Log Name` page.

For a description of the various logs, see *Log descriptions* on page 21.

Log descriptions

The following table provides a brief description of each log.

Log	Description
Audit	Tracks user activity that might have security implications, including user attempts (successful and unsuccessful) to update the software or change settings, configuration, and user information.
System	Contains information about the software processes that control the device, including startup routines, run levels, and maintenance routines.
IPS Alert	Documents network traffic that triggers IPS filters configured with the following action sets: • Permit + Notify • Permit + Notify + Trace

Log	Description
	• Trust + Notify • Rate Limit + Notify
IPS Block	Documents packets that trigger IPS filters configured with any action that includes a Block + Notify or Block + Notify + Trace action, including Quarantine and TCP Reset action sets.
Quarantine	Records the IP addresses that have been added to and removed from quarantine.
Reputation Alert	Contains messages for network traffic that triggered a reputation filter configured with the Permit + Notify action-set.
Reputation Block	Contains messages for network traffic that triggered a reputation filter configured with the Block + Notify action-set.
SSL Inspection	Contains information about SSL servers that have been configured to log information.

Logs

Using the Log Summary, you can view an entire log or a portion of a log. You can view the log in a new browser window or save a copy of the log to your local Downloads directory.

Download a log

1. Click the **Home** icon or select **System** > **Log Configuration** > **Summary**.
2. Select the log you want to download by clicking its checkbox.
3. Click **Download**.

   The Download System Log dialog is displayed.
4. Choose one of the following Log Entries options:
   - All – Downloads all log entries.
   - Time Range – Downloads entries based on the specified time range.
   - ID Range – Downloads entries based on the line number (or ID) of the log entry.

5. Choose either the Tab delimited format (txt) or Comma delimited format (csv).
6. Click **Open in Browser** to display the log in a new browser window or **OK** to save the log to your local Downloads directory.

Clear log entries

Clearing log entries permanently deletes them from the device.

1. Click the **Home** icon or select **System > Log Configuration > Summary**.
2. Select the log you want to clear by clicking its checkbox.
3. Click **Clear log entries**.

   A confirmation dialog is displayed.

4. Click **OK** to confirm or Cancel to cancel the operation.

   If you click **OK**, all entries are cleared from the log and a success message is displayed in the banner.

Performance graphs

The Performance Graphs panel displays a visual representation of performance aspects of the device. You can monitor system performance at-a-glance.

The following table provides a description of the Performance Graphs:

Graph	Description
Traffic Profile	Displays packet traffic flow per second data over the last 24 hours.
Top 10 Filters	Displays a bar graph of the top 10 attack filters and the number of hits of each.
Filter by Severity	Displays a bar graph of attacks categorized as Low, Minor, Major, and Critical. The graph also lists the percentage of attacks. The severity levels are assigned by the TippingPoint Digital Vaccine team and are included as part of the filter definition.
Filter by Action	Displays the actions taken on filtered traffic. The graph also lists the percentage of packets processed with each action. The following actions are listed:

Graph	Description
	• Trust • Rate limit • Permit • Block
Filter by Protocol	Displays attack traffic categorized by protocol (ARP, Ethernet - Other, ICMP, ICMPv6, UDP, TCP, IPv4 - Other, and IPv6 - Other). The graph also lists the percentage of filtered packets for each protocol.

Version Information

The Version Information panel provides hardware and software version information for the TPS, and version information for the Digital Vaccine package.

Monitor the device

The Monitor page provides complete visibility into system health, traffic flows, and other analytics imperative to system and network administration. It provides administrative control of user sessions, to view or clear all or specific sessions. The XML-based APIs provided at the back end retrieve all the data for all the sessions.

The Monitor page includes reports of various parameters in which the data used to produce the graphic is collected. Monitor requests can be initiated by specifying attributes such as IP addresses, family, port numbers, or protocol. Some requests do not require specifying any attributes.

The columns in the table vary according to the type of report. You can click a heading to sort the table by the column. You can cycle through two sort orders by clicking the column heading: ascending (down arrow) and descending (up arrow). You can click on the **Columns** list to check or uncheck the rows to be included or excluded in the table menu.

The Monitor page provides access to the following areas:

- Logs
- Sessions
- Health
- Network

Monitor logs

In addition to viewing logs, you can also search for logs, sort logs by the newest or oldest entry, download a local copy, and clear log entries.

The Monitor page provides surveillance of the following logs:

- *Audit logs* on page 26
- *System logs* on page 27
- *IPS Block and Alert logs* on page 28
- *Quarantine logs* on page 30
- *Reputation Block and Alert logs* on page 31
- *SSL inspection logs* on page 33

Working with logs

To find, download or clear logs:

Select **Monitor > Logs >** *Log Name*. Substitute the name of the log you want to view for *Log Name*. After the logs are displayed, you can perform any of the following actions:

- Choose **Show newest entries first** or **Show oldest entries first**.
- Click **Download** to download a copy of the log report.
- Click **Clear log entries** to delete all the log entries.
- Search for a specific log:
  - Enter the log that you want to search.
  - Click **Show Advanced** to refine your search criteria.
  - Click **Go** to view the generated report or click **Clear** to clear the search panel.

Audit logs

The Audit log tracks user activity that might have security implications, including user attempts (successful and unsuccessful), to do the following:

- Change user information
- Change routing or network configuration
- Gain access to controlled areas (including the audit log)
- Update system software and attack protection filter packages
- Change filter settings

To maintain a complete history of entries and provide a backup, you can configure the device to send log entries to a remote syslog server from the **System > Log Configuration > Notification Contacts** page.

Note: Only users with Super-user access level can view, print, reset, and download the Audit log.

The following information is displayed in the Audit Logs table:

Column	Description
Log ID	Displays the system-assigned log ID number.
Log Entry Time	Displays the time the log was entered in the format *YYYY-MM-DD HH:MM:SS*.

Column	Description
User	Displays the login name of the user who performed the audited action. The user listed for an event can include SMS, SYS, and CLI.
Access	Displays the access level of the user performing the action.
IP Address	Displays the IP address from which the user performed the action.
Interface	Displays the interface with which the user logged in: WEB for the LSM, CLI for the command line interface. For system-initiated actions, SYS displays in this field.
Component	Displays the area in which the user performed an action (LOGIN, LOGOUT, and Launch Bar Tabs).
Result	Displays the action performed or the result of a LOGIN or LOGOUT attempt.
Action	The action performed as a result. For example, Log Files Reset.

System logs

The System Log records contains information about the software processes that control the device, including startup routines, run levels, and maintenance routines. System log entries can provide useful troubleshooting information if you encounter problems with your device.

To maintain a complete history of entries and provide a backup, you can configure the device to send log entries to a remote syslog server from the **System > Log Configuration > Notification Contacts** page.

Note: Users with any access level can view and print the system log, but only Administrator and Superuser level users can reset this log. System log entries are sent to the syslog server only after the device has fully booted. During the boot sequence, entries cannot be sent because network ports are not yet enabled. When the boot sequence completes, the device sends a startup message to the syslog server.

The following information is displayed in the System Logs table:

Column	Description
Log ID	Displays the system-assigned log ID number.
Log Entry Time	Displays the time the log was entered in the format *YYYY-MM-DD HH:MM:SS*.
Severity Level	Indicates whether the log entry is informational (INFO) or whether it indicates an error or critical condition (ERR or CRIT).
Component	Indicates which software component sent the message to the log.
Message	Text of the log entry.

IPS Block and Alert logs

The IPS Block and Alert logs contain log messages for the network traffic that triggers a reputation filter configured with the action set created by the user.

The following action sets are included for Alert logs:

- Permit + Notify
- Permit + Notify + Trace
- Trust + Notify
- Rate Limit + Notify

The following action sets are included for Block logs:

- Block + Notify
- Block + Notify + Trace

The logs contain IP and Layer 4 information, along with the matching filter.

To maintain a complete history of entries and provide a backup, you can configure the device to send log entries to a remote syslog server from the **System > Log Configuration > Notification Contacts** page.

Note: Any user can view the log, but only administrator and super-user level users can reset the log.

IPS Logs contain the following information:

Column	Description
Log ID	Displays the system-assigned log ID number.
Log Entry Time	Displays the time the log was entered in the format *YYYY-MM-DD HH:MM:SS*.
Severity	Indicates the severity of the triggered filter: • 4: Critical • 3: Major • 2: Minor • 1: Low When the log is downloaded, the severity is indicated with a number.
Action	Indicates the action that triggered the alert.
Filter Name	Displays the name of the triggered filter.
Rate Limit	(Alert Log only) Displays the rate limit. If applicable, the rate limiter rate that was defined in the triggered action set and a link to the action set on which the log entry was generated. This field is blank for Permit and Trust action log entries.
Protocol	Displays the name of the protocol that the action affects.
Interface In	Displays the network interface on which the traffic arrived.
Src Addr	Displays the source address of the triggering traffic.
Src Port	Displays the source port number of the triggering traffic.
Interface Out	Displays the network interface from which the triggering traffic departed.

Column	Description
Dst Addr	Displays the destination address of the triggering traffic.
Dst Port	Displays the destination port number of the triggering traffic.
Virtual Segment	Displays the virtual segment on which the alert or block occurred (such as 1A-1B).
VLAN ID	Displays the identification number of the VLAN.
Hit Count	Displays the number of packets that have been detected if packet tracing is enabled.
Packet Trace	Indicates whether packet tracing is enabled.

Quarantine logs

The Quarantine log records the IP addresses that have been added to and removed from quarantine. Quarantine logging operates independently of a policy's notification contacts. Quarantine events are always recorded in a log file and on the remote syslog server if configured to do so.

Note: Any user can view the log, but only administrator and super-user level users can reset the log.

The following information is displayed in the Quarantine Logs table:

Column	Description
Log ID	Displays the system-assigned log ID number.
Log Entry Time	Displays the time the log was entered in the format *YYYY-MM-DD HH:MM:SS*.
Severity	Indicates the severity of the triggered filter: • 4: Critical • 3: Major

Column	Description
	• 2: Minor • 1: Low When the log is downloaded, the severity is indicated with a number.
Interface In	Displays the network interface on which the traffic arrived.
Src Addr	Displays the source address of the triggering traffic.
Action	Indicates whether the IP address was added or removed to the Quarantine logs.
Filter Name	Displays the name of the triggered filter.

Reputation Block and Alert logs

The Reputation log contains log messages for the network traffic that triggers a reputation filter configured with the action set created by the user. Alert messages are displayed for network traffic that triggered a reputation filter configured with the Permit + Notify action-set. Block messages are displayed for network traffic that triggered a reputation filter configured with the Block + Notify action-set.

The following information is displayed in the Reputation Logs Block table:

Column	Description
Log ID	Displays the system-assigned log ID number.
Log Entry Time	Displays the time the log was entered in the format *YYYY-MM-DD HH:MM:SS*.
Severity	Indicates the severity of the triggered filter: • 4: Critical • 3: Major • 2: Minor

Column	Description
	• 1: Low When the log is downloaded, the severity is indicated with a number.
Action	Indicates whether the IP address was added or removed to the reputation logs.
Filter Name	Displays the name of the triggered filter.
Rate Limit	(Alert Log only) Displays the rate limit. If applicable, the rate limiter rate that was defined in the triggered action set and a link to the action set on which the log entry was generated. This field is blank for Permit and Trust action log entries.
Protocol	Displays the name of the protocol that the action affects.
Interface In	Displays the network interface on which the traffic arrived.
Src Addr	Displays the source address of the triggering traffic.
Src Port	Displays the source port number of the triggering traffic.
Interface Out	Displays the network interface from which the triggering traffic departed.
Dst Addr	Displays the destination address of the triggering traffic.
Dst Port	Displays the destination port number of the triggering traffic.
Virtual Segment	Displays the virtual segment on which the alert or block occurred (such as 1A-1B).
VLAN ID	Displays the identification number of the VLAN.

Column	Description
Hit Count	Displays the number of packets that have been detected if packet trace is enabled.
Packet Trace	Indicates whether packet tracing is enabled.

SSL inspection logs

The SSL Inspection Log records contains information about SSL sessions. For details, such as connection resets, click **Columns > Details**.

Monitor user sessions

The User Sessions page lists all the currently logged users, locked users, and the IP addresses. If the number of login attempts from a specific user or the IP address exceeds the maximum login attempts, the user or IP address gets locked out.

View active user sessions

Select **Monitor** > **Users > Active Users**. By default, the Active Users table is displayed with the following columns:

Column	Description
User Name	Displays the name that identifies the user.
Idle	Displays the amount of time the user has been active on the device.
Interface	Displays the type of device from which the user logged in.
Logged In	Displays the date and time when the user logged in.
IP Address	Displays the IP address of the device.
Type	Displays the authentication protocol type (LOCAL, LDAP, RADIUS)

Note: Multiple LSM user sessions from the same IP address are not tracked if a user logs in several times from the same IP address.

Log off active users sessions

Select the checkbox next to the Username and click **Log Off** .

View locked users or IP addresses sessions

Select **Monitor > Users > Locked Users/IP Addresses**.

The locked users table has the following columns:

- User Name – Displays the name that identifies the user.
- IP Address – Displays the IP address of the device.
- Time of Lock – Displays the time the user was locked.

Unlock locked users and locked IP addresses

Select the checkbox next to the Username or IP Address and click **Unlock**.

Monitor managed streams

The Managed Streams area enables you to monitor security events, providing visibility into inspection results and traffic flows. You can monitor the following sessions:

- *Blocked streams* on page 34
- *Rate-limited streams* on page 36
- *Quarantined addresses* on page 37
- *Trusted streams* on page 38

Blocked streams

When traffic triggers a filter that has been configured with a Block or Block + Notify action, traffic from the source IP address and port is blocked and an entry is added to the Blocked Streams table, based on the contact configuration in the action set. Only the IPS blocks and IP reputation (not DNS) can create a block entry.

From the Blocked Streams page, you can:

- View and search for information on blocked streams

- Manually clear all or selected blocked stream connections

The Blocked Streams table displays up to 50 entries. Entries are added when the block event occurs. Entries are automatically removed when the connection times out based on the **Connection Table** timeout setting configured from the **Policy > Profiles > Settings** page. You can manually remove an entry from the table with the Flush function, which unblocks the stream.

View blocked streams

1. Select **Monitor > Blocked Streams**. The following information is displayed in the Blocked Streams table:

Field	Description
Protocol	Displays the type of protocol used by the blocked connection.
Source Address	Displays the source IP address of the connection.
Source Port	Displays the source port number of the connection.
Destination Address	Displays the destination IP address of the connection.
Destination Port	Displays the destination port number of the connection.
Source Interface	Displays the network interface on which the traffic arrived.
Destination Interface	Displays the destination network interface.
Virtual Segment	Indicates the virtual segment where traffic was blocked.
Reason	Displays the filter link that details why the traffic connection stream was blocked. Click the link to display and manage the filter.

2. To block the stream, select the checkbox next to the stream and click **Flush Selected** or **Flush All**. On the confirmation dialog, click **OK**.

To search for a specific blocked stream(s):

- Select a protocol (**All, TCP, UDP, ICMP, ICMPv6**) from the list.
- (Optional) Enter either the source or destination IP address.
- (Optional) Enter the port number.
- Click **Go** to view the generated report or click **Clear** to clear the search panel.

Rate-limited streams

When traffic triggers a filter configured with a rate-limit action set, traffic from the source IP and port is limited based on the rate-limit settings. Traffic from the source IP address and port to the destination IP address and port remains rate-limited until the connection time-out period expires, or until the connection is manually terminated.

From the Rate Limited Streams page, you can:

- View and search for information on rate-limited streams
- Manually terminate all or selected rate-limited stream connections

The Rate Limited Streams table displays up to 50 entries. Entries are added when the rate-limit event occurs. Entries are automatically removed when the connection times out based on the **Connection Table** setting configured from the **Policy > Profiles > Settings** page. You can manually remove an entry with the **Flush** functions, which removes the rate limit from the stream.

View rate-limited streams

1. Select **Monitor > Rate Limited Streams**. The following information is displayed in the Rate Limited Streams table:

Column	Definition
Protocol	Displays the type of protocol used by the rate-limited connection.
Source Address	Displays the source IP address of the connection.
Source Port	Displays the source port number of the connection.
Destination Address	Displays the destination IP address of the connection.

Column	Definition
Destination Port	Displays the destination port number of the connection.
Source Interface	Displays the network interface on which the traffic arrived.
Destination Interface	Displays the destination network interface.
Virtual Segment	Indicates the virtual segment where traffic was rate-limited.
Reason	Displays the filter link that details why the traffic connection stream was rate-limited. Click the link to display and manage the filter.

2. To rate-limit the stream, select the checkbox next to the stream and click **Flush Selected** or **Flush All**. On the confirmation dialog, click **OK**.

Search for specific rate-limited streams

1. Select a protocol (**All, TCP, UDP, ICMP, ICMPv6**) from the list.
2. (Optional) Enter the IP address.
3. (Optional) Enter the port number.
4. Click **Go** to view the generated report or click **Clear** to clear the search panel.

Quarantined addresses

When traffic triggers a filter that has been configured with a quarantine action, the host is quarantined and an entry is added to the Quarantined Addresses table, based on the contact configuration in the action set. From the Quarantined Addresses page, you can:

- Manually force an address into quarantine
- Search for quarantined addresses
- Manually release all or selected quarantined hosts

The Quarantined Addresses table displays up to 50 entries.

View quarantined addresses

1. Select **Monitor > Quarantined Addresses**. The following information is displayed in the Quarantined Addresses table:

Column	Definition
IP Address	Displays the IP address under quarantine.
Source Address	Displays the packet's source IP address.
Destination Interface	Displays the destination network interface.
Reason	Indicates the reason the IP address is under quarantine.

2. To remove the IP address from the quarantine, select the checkbox next to the IP address and click **Flush Selected** or **Flush All**. On the confirmation dialog, click **OK**.

Manually force an IP address into quarantine

1. Enter the IP address in the IP Address to Quarantine field.
2. Select the action set from the Action list.
3. Click **Quarantine** to add the IP address to the quarantined addresses table.

   To search for a specific IP address:

   - Enter the IP address in the Search IP Address field.
   - Click **Go** to view the generated report or **Clear** to clear the search panel.

Trusted streams

When traffic triggers a filter configured with a Trust action set, traffic from the source IP and port is recorded in the Trusted Streams table. From the Trusted Streams page, you can:

- View and search for information on trusted streams
- Manually clear all or selected trusted stream connections

The Trusted Streams table displays up to 50 entries. Entries are added when the trust action occurs. Entries are automatically removed when the connection times out based on the **Trusted Streams** flush setting

configured from the **Policy > Profiles > Settings** page. You can manually remove an entry with the **Flush** functions, which removes the trusted stream from the table.

View trusted streams

1. Select **Monitor > Trusted Streams**. The following information is displayed in the Trusted Streams table:

Column	Definition
Protocol	Displays the type of protocol used by the trusted connection.
Source Address	Displays the source IP address of the connection.
Source Port	Displays the source port number of the connection.
Destination Address	Displays the destination IP address of the connection.
Destination Port	Displays the destination port number of the connection.
Source Interface	Displays the network interface on which the traffic arrived.
Destination Interface	Displays the destination network interface.
Virtual Segment	The virtual segment where the stream is trusted.
Reason	The filter link that details why the traffic connection stream was trusted. Click the link to display and manage the filter.

2. To stop trusting the stream, select the checkbox next to the Rule ID and click **Flush Selected** or **Flush All**. On the confirmation dialog, click **OK**.

Search for specific trusted streams

1. Select a protocol (**All, TCP, UDP, ICMP, ICMPv6**) from the list.

2. (Optional) Enter the IP address.
3. (Optional) Enter the port number.
4. Click **Go** to view the generated report or click **Clear** to clear the search panel.

Monitor health

This page displays the current status and network performance of the appliance. It allows you to monitor key network and device metrics and to quickly detect and resolve the device malfunctions and bottlenecks in the network. Health statistics such as performance, port settings, and usage thresholds indicate the state of system components and help you to maintain optimal performance and continued operation of the appliance. It contains the following topics:

Performance

To view the current throughput performance of the device, click **Monitor > Health > Performance**.

The Performance/Throughput table displays the following information:

Column	Description

Component	Congestion, performance, or port being monitored:
	• Congestion – Indicates the traffic congestion impact on the engine.
	  • Green – Engine usage below 10%. This reflects a normal operating state.
	  Note: Usage could be below 10% and still show a yellow Warning state depending on the Performance Protection values set.

Column	Description
	• Yellow – Engine usage between 10% – 25%. This warns that congestion is causing a higher than normal strain on the engine. • Red – Engine usage above 25%. The strain of traffic congestion on the engine has reached a critical level. • Performance – Indicates the total and used bandwidth of the device. A yellow Warning state indicates that Performance Protection has been triggered based on user-configured settings. Configure Performance Protection on the **System > Log Configuration > Performance Protection** page. • Ports – Indicates the port bandwidth used.
Description	Describes the component.
State	The current performance status of the device or the operating status of each port. • **Normal** — Green. Device performance is normal. The port is active without errors. • **Warning** — Yellow. Performance loss or congestion has put undue stress on the engine or has triggered Performance Protection mode. • **Critical** — Red. The port is waiting for traffic or usage in a stand-by mode, or the device has entered Layer 2 Fallback because of Performance Protection. • **Inactive** — Grey. The port is not in use or is disabled.
Throughput	A bar graph depicting the performance of the device and current usage level of the ports. • **Performance** – Indicates the total and used bandwidth of the device. The color of the bar changes according to the status of Performance Protection. • *<Port name>* • Green – Less than 80% of port bandwidth used.

Column	Description
	• Yellow – Between 80% and 99% of port bandwidth used. • Red – Over 99% of port bandwidth used.
Details	Percentage of throughput used.

High availability

High availability (HA) is a system configuration setting that ensures that your network traffic always flows at wire speeds in the event of any internal hardware or software failure on the device. High availability is critical in maintaining network protection from an attack, even in the event of a device failure.

HA allows the user to install two devices in a redundant network configuration. HA keeps the devices in sync with each other; if one experiences a system failure, the network flow can be routed to the other without any interruption.

The HA page displays identifying information for your device and its HA partner device. The State Synchronization table displays each subsystem and its current state. You can force a subsystem state resync by selecting the checkbox next to the **Subsystem** and clicking **Force State Re-Sync**.

The High Availability page lists the current high availability status for the following High Availability features:

- Intrinsic Network High Availability
- Transparent High Availability

For information on configuring HA, see *High Availability settings* on page 103.

CPU utilization

The CPU Utilization page displays a graph with a set of management cores and data cores. The available trend intervals are 24 hours, seven days, and 30 days. The data cores rise when you run heavy traffic through the device, which intensifies CPU usage.

1. Select **Monitor > CPU Utilization**.
2. Select the core from the All Cores list.
3. Select the time interval from the Time Period list.

Note: The warning threshold is 50 percent and the critical threshold is 98 percent. Thresholds do not apply to the data cores.

Disk utilization

The Disk Utilization page displays a high-level view of system disk and user disk usage metrics. Available trend intervals are 24 hours, seven days, and 30 days.

1. Select **Monitor > Disk Utilization**.
2. Select the time interval from the Time Period list.

Note: The warning threshold is 90 percent and the critical threshold is 95 percent.

Fan speed

The Fan speed page graphically displays the fan speeds (in RPM). Available trend intervals are 24 hours, seven days, and one month.

1. Select **Monitor > Fan Speed**.
2. Select the time interval from the **Time Period** list.

Memory utilization

The Memory Utilization page displays a graph for the amount of memory used (in Gigabytes). Available trend intervals are 24 hours, seven days, and 30 days.

1. Select **Monitor > Memory Utilization**.
2. Select the time interval from the Time Period list.

Note: The warning threshold is 90 percent and the critical threshold is 95 percent.

Temperature

The Temperature section displays a graph for the temperature range of the device (in degrees Celsius). Available trend intervals for temperature sensors are 24 hours, seven days, and 30 days.

1. Select **Monitor > Temperature**.
2. Select the time interval from the Time Period list.

Note: The warning threshold is 62 degrees Celsius and the critical threshold is 68 degrees Celsius.

Monitor network

The Monitor Network provides information of the traffic (in bps) for the ports, and shows a graphical representation of the network bandwidth. Ports refer to the physical ports on the device such as 1A, 1B, and so on. The available trend intervals are 24 hours, seven days, and 30 days.

Monitor port health

To monitor port information for each port on the device, click **Monitor > Network > Ports**. The Ports table displays the following information:

Column	Description
Port	Identifies the port number.
Description	Indicates the segment of the port.
Speed	Indicates the port speed.
Duplex	Indicates if the port is set to full or half for duplex.
Media	Indicates the port medium, which can be copper or fiber.
Status	Indicates if the link is down or up.
Received	Indicates the total number of discards, packets, and bytes received on the port.
Transmitted	Indicates the total number of discards, packets, and bytes transmitted on the port.

Monitor network bandwidth

The Network Bandwidth page displays the graph of the traffic (in bps) for the physical ports, such as 1A, 1B, and so on. The available trend intervals are 24 hours, seven days, and 30 days.

1. Select **Monitor > Network Bandwidth**.
2. Click **Ports**.
3. Refine the search using the drop-down lists on the right side of the page.

Monitor SSL bandwidth

The SSL Bandwidth page displays the graph of the decrypted traffic (in bps). The available trend intervals are 24 hours, seven days, and 30 days. Select **Monitor > SSL Bandwidth**.

Network

The **Network** menu pages in the LSM enable you to view and modify the setup of the device so that it can work within your network environment. The following menu options are available:

- **Ports** — Disable, enable, or restart a port, and manage port configuration (auto-negotiation and line speed).
- **Segments** — View and manage segment configuration for Layer-2 Fallback (high availability) and link down synchronization.
- **Virtual Segments** — Create and manage virtual segments to further refine the network traffic classifications.
- **VLAN Translation** — Enable translation of traffic between different VLANs or between VLAN and non-VLAN interfaces.
- **DNS** — Specify domain names and IPv4 or IPv6 server addresses.

This topic discusses the following subjects:

Network ports

Use the Network Ports pages to perform the following tasks:

- View a list of network I/O modules and their ports
- View and edit current port configuration
- Disable/enable Auto Negotiation
- Disable/enable ports
- Restart a port

By default, the device sets all ports to auto-negotiate. With this setting, the device port negotiates the highest line speed supported by both the device port and its link partner. TippingPoint recommends configuring both the device ports and the link partners to auto-negotiate because it is the best, most reliable way to

establish and maintain links. However, if the device cannot establish or maintain a link when auto-negotiate is set, you might need to disable auto-negotiation and configure the line speed and duplex settings.

When configuring the ports, remember that both link partners must be configured with identical settings. If one port is configured to auto-negotiate, the other port must also be configured to auto-negotiate. If only one port is configured to auto-negotiate, the link might come up, but one or both partners may experience poor performance or RX errors.

The following table describes the port configuration parameters.

Column	Description
Port name	Displays the interface ID.
Administrative State	Indicates whether the port is currently enabled or disabled.
Type	Displays the type of port. For example, data or management.
Port	Displays the physical port number assigned to the interface.
MAC Address	Displays the uniquely assigned media access control address for communicating on the physical network segment.
Auto Negotiation	Indicates whether the port auto-negotiates line speed or uses the line speed and duplex settings as a forced port configuration. By default, Auto Negotiation is enabled. • If Auto Negotiation is enabled, the device automatically negotiates the highest common speed and duplex that the device and the link partner both support. • If Auto Negotiation is disabled, the manually configured Line Speed and Duplex settings are used. You might want to disable auto-negotiation on some older networks if the device is unable to establish or sustain the link with its partner.

Column	Description
Port Speed	Displays the port speed of throughput on the port.
Port Duplex	Indicates whether the port is set to full- or half-duplex.
Media	Indicates whether the port is copper or fiber.
Status	Indicates whether connectivity is currently up or down.

Note: When auto-negotiation is disabled, some port options might be limited. For example, 1000 Mbps line speed is only available for copper ports when auto-negotiation is enabled. When auto-negotiation is disabled, the line speed can only be set to 100 Mbps or 10 Mbps.

Edit port settings

1. Select **Network > Ports > Settings**.
2. Select an interface and click **Edit**. The Edit Port Settings dialog is displayed.
3. Select **Enabled** to enable the port settings.
4. Select one of the following options: **Use Auto-Negotiation** or **Manually set Port Speed and Duplex**. For manual settings, select the speed and mode from the drop-down list.

   When auto-negotiation is disabled, some port options might be limited. For example, 1000 Mbps line speed is only available for copper ports when auto-negotiation is enabled. When autonegotiation is disabled, the line speed can only be set to 100 Mbps or 10 Mbps. To run 10G fiber ports in 10G mode, ensure auto-negotiation is disabled (the default). When auto-negotiation is enabled for 10G ports, the speed drops to 1G. When the auto-negotiate feature is turned off, users can configure all fiber ports (SFP, SFP+, QSFP+) only to their default settings even though the hardware might list other optional values. The 12 fixed RJ-45 copper ports can be configured to either 10 Mbps or 100 Mbps using the LSM, or to 10 Mbps, 100 Mbps, or 1 Gbps using the command-line interface.

   The set of speeds available will be tailored to capabilities of specific ports.

5. Click **OK**.

   If you use a copper-fiber translator, disable auto-negotiation on the device before performing a restart. Some translators do not support auto-negotiation. When the copper cable is pulled, these translators do not attempt to auto-negotiate with the device. The device driver attempts to re-initialize the port several times before timing out and placing the port in an Disabled mode.

6. After making port configuration changes, restart the port to ensure proper functioning of the device. See *Restart an interface* on page 49.

Restart an interface

1. Select **Network > Ports > Settings**.
2. Select an interface and click **Restart**, on confirmation click **OK**. This operation restarts the selected interfaces and the network connectivity may be interrupted. Disabled interfaces will not be restarted.

Segments

Use the Segments page to view segments on the device.

Each segment is composed of a pair of ports on the device; for example, ports 1A and 1B compose one segment. These two ports integrate the device into the network.

From the Segments page you can access the Segment Editor page for each segment, where you can complete the following tasks:

- View current high availability and link-down synchronization for each network segment
- Edit HA settings for Layer-2 fallback and link-down synchronization

Segment configuration defines how the device handles traffic and port status. You can specify settings for the following options:

- **Intrinsic Network HA Layer-2 Fallback Action** determines if the device permits all traffic or blocks all packet transfers on that segment if the device goes into high availability.
- **Link Down Synchronization** allows you to configure the device to force both ports down on a segment when the device detects a link state of down on one of the ports. When link-down synchronization is enabled, the device monitors the link state for both ports on a segment. If the link goes down on either port, both ports on the segment are disabled. This functionality propagates the link state across the device.

When link-down synchronization is enabled, segment monitoring begins only after link-up is detected on both ports. When link-down synchronization disables the ports on a segment, two audit log messages are generated. The first message in the audit log corresponds to the port with the link down. The second message corresponds to the segment partner. Additionally, an error message is added to the system log indicating which port was detected with the link-down, activating link-down synchronization for that segment.

Note: In addition to the physical segments on the device, physical segments also have predefined virtual segments that allow you to classify and filter traffic on the network by both physical port and VLAN ID.

A network segment is created by joining an Ethernet pair of interfaces on the device in a transparent, bump-in-the-wire architecture to allow traffic flow and inspection between the two network ports. Segments can be configured between vertical port pairs only.

Segment Name	Displays the name of the segment.
Port Pair	Displays the paired Ethernet interfaces. (For example: 1A + 1B)
Intrinsic HA	Displays Permit if traffic is allowed to flow or Block if traffic is blocked.
Link Down Synchronization	Indicates the action the segment takes when a link goes down (Hub, Breaker, or Wire) and the wait time before the device reflects the port status change (if value ranging from 0–240 seconds is set).
Operating Mode	Displays the current operating mode.

Edit segment, enable L2FB and segment bypass

1. Select **Network > Segments**.
2. Click **Edit**. The Segment Editor dialog is displayed.
3. For Intrinsic HA - L2FB Action, select **Permit** or **Block**.
   - Permit – Allows traffic flow without inspection to continue during fallback.
   - Block – Stops and discards all traffic during fallback on all ports and during a system reboot or power cycle on copper ports. When the device returns to normal operating conditions, traffic flows and is inspected regardless of the block setting.
4. Configure Link Down Synchronization (LDS) settings so that any port state changes are reflected to the partner port in the segment. This ensures that the segment appears as a bump in the wire and does not become a black hole. Select one of the following options:
   - Hub – Take no action when a link goes down.
   - Breaker – When a link goes down, take partner link down until both member ports are restarted or the segment itself is restarted.
   - Wire – When a link goes down, take partner link down until original link restored.

   Enter a wait time (0 to 240 seconds) for Breaker and Wire options. This determines how long the device waits before reflecting the port status change on the partner port.
5. Click **OK**.

Restart a segment

1. Select **Network > Segments**.
2. Select a segment and click **Restart**. When prompted to confirm, click **OK**. This operation restarts the selected segments, including port status, and can interrupt the network connectivity.

Virtual segments

Physical segments have predefined virtual segments. CIDRs and profiles are applied to the virtual segment. Virtual segments enable further management of VLAN traffic.

Virtual segments are saved on the device in a prioritized table, and security profiles and traffic management profiles are applied in order of priority. For example, if port 1A is assigned to two different virtual segments, the profiles that are assigned to the higher-priority segment are applied to the traffic on that port before the profiles assigned to the lower-priority segment. If no physical ports are defined on a virtual segment, the virtual segment will apply to all physical ports.

You can configure a maximum of 64 virtual segments.

The Virtual Segment table has the following configuration parameters:

Parameter	Description
Order	The order of priority. Position values must remain contiguous across all defined virtual segments, so there should never be a gap in the sequence.
Name	The name of the virtual segment. Each name must be unique.
VLAN ID	The ID of the incoming VLAN. If no VLAN IDs are defined on a virtual segment, all VLAN IDs are included. Each VLAN ID in a range counts separately. For example, `vlan-id range 1 5` counts as 5 IDs. You can configure up to 512 VLAN IDs per virtual segment.
Source Address	The source CIDR. If no source addresses are defined, all source addresses are included. Each CIDR counts as a single address. For example, `192.168.1.0/24` counts as 1 address. No more than 512 addresses may be specified.

Parameter	Description
Destination Address	The destination CIDR. If no destination addresses are defined, all destination addresses are included. Each CIDR counts as a single address. No more than 512 addresses may be specified.
Physical Segments	The physical segment associated with the virtual segment pair.
IPS Profile	The IPS profile assigned to the virtual segment.
Reputation Profile	The Reputation profile assigned to the virtual segment.
Traffic Management Profile	The Traffic Management profile assigned to the virtual segment.
SSL Profile	The SSL Inspection profile assigned to the virtual segment.

Add, insert, or edit a virtual segment

Clicking **Add** adds the new virtual segment after all the other user-created virtual segments. Clicking **Insert** inserts the new virtual segment just before the currently selected virtual segment. Virtual segments that are created by the system can have their profiles modified but are otherwise read-only. All system-created virtual segments always appear at the end of the list.

1. From the LSM menu, click **Network > Virtual Segments**.
2. Click **Add** or **Insert** to create a new virtual segment, or click **Edit** to edit an existing virtual segment.
3. In the Add Virtual Segment dialog or Edit Virtual Segment dialog, specify the following:
    - Name – (Required) Name used to identify the virtual segment. Each virtual segment must have a unique name.
    - Description – An optional parameter to provide more detailed information about the virtual segment.
    - IPS Profile – Security profile that you want to apply to the virtual segment. A virtual segment can have only one IPS profile applied to it.
    - Reputation Profile – Reputation profile that you want to apply to the virtual segment. A virtual segment can have only one Reputation profile applied to it.
    - Traffic Management Profile – Traffic Management profile that you want to apply to the virtual segment. A virtual segment can have only one Traffic Management profile applied to it.

- SSL Inspection Profile – SSL inspection profile that you want to apply to the virtual segment. A virtual segment can have only one SSL inspection profile applied to it.
- Physical Segments – Physical segment associated with the virtual segment. All physical segments are directional.
- Traffic Criteria – (Required) Specify any one or all of the following: VLAN ID, Source IP, and Destination IP. For example, omit VLAN ID and specify Destination IP. When specfying a VLAN ID, specify a value between 1 and 4094 in which the segment is included. There can be no duplicate VLAN IDs or overlapping VLAN ranges. No more than 512 VLAN IDs per virtual segment (a VLAN range of 1–100 counts as 100 IDs). At least one traffic criteria (VLAN ID, source IP address, or destination IP address) must be defined for each virtual segment.
- Source IP Address – Source CIDR associated with the virtual segment. Addresses must be valid and can be IPv4 or IPv6. The host portion of address/mask must be 0 (zero). No more than 250 addresses may be specified.
- Destination IP Address – Destination CIDR associated with the virtual segment. Addresses must be valid and can be IPv4 or IPv6. The host portion of address/mask must be 0 (zero). No more than 250 addresses may be specified.

4. Click **OK**.

Note: Virtual segments must be created with a physically available segment. If creating a virtual segment generates a UDM warning in the system log, ensure you have associated the virtual segment with a valid physical segment.

Move or delete a virtual segment

Only user-created virtual segments can be moved or deleted. Click **Move To** to move a virtual segment to a specific location and priority. Click **Move Up** or **Move Down** to reorder the priority of the virtual segment in the list. Click **Delete** to remove a user-created virtual segment.

VLAN translation

The TPS translates traffic between different VLANs or between VLAN and non-VLAN interfaces. Deploy the TPS on an aggregation or distribution switch to selectively inspect traffic based on the switch configuration.

Note: VLAN translation is supported on both the TippingPoint 440T and 2200T devices.

The following diagram shows a sample TPS deployment where three VLANs connect to a central distribution switch. The traffic is routed from the switch to the TPS, which inspects the traffic, performs the translation tasks, and routes the inspected traffic back onto the network.

Figure 2. Network with VLAN translation

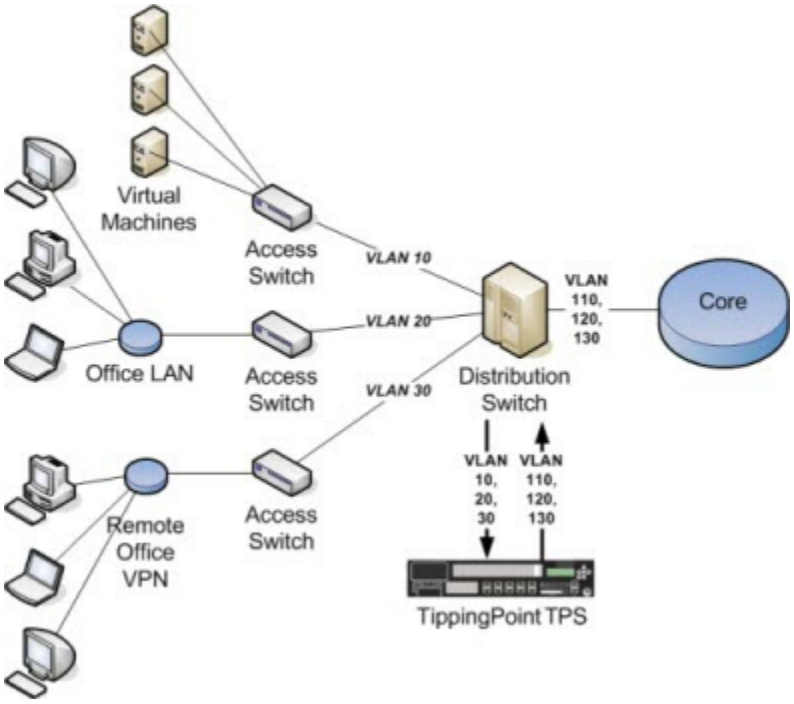

You can configure the aggregation switch to send traffic to the TPS on a selective basis, focusing inspection capacity on the VLANs where the need is greatest.

Note: Security policies are applied to the incoming VLAN ID only. VLAN mappings must be unique, with one incoming VLAN paired with one outgoing VLAN. The TPS does not translate one-to-many VLAN mapping.

The following table describes the VLAN Translations configuration parameters.

Parameter	Description
Incoming Port	The TPS virtual port through which incoming traffic arrives.
Incoming VLAN ID	The ID of the incoming VLAN.
Outgoing VLAN ID	The ID of the outgoing VLAN.
Auto-Reverse	Select this option to enable automatic reverse VLAN translation. This option is disabled by default.

Add or edit a VLAN translation

Add or edit a VLAN translation to selectively inspect traffic based on the switch configuration. This feature translates traffic between different VLANs or between VLAN and non-VLAN interfaces. The TPS creates a separate VLAN translation rule for each port you want to translate. A maximum of 8000 VLAN translation rules can be defined on a 440T or 2200T TPS. If the number of VLAN translation rules you want to commit exceeds the specified limit, the device does not commit your changes.

1. From the LSM menu, click **Network > VLAN Translation**.
2. Click **Add** to create a new VLAN translation, or click **Edit** to edit an existing virtual segment.
3. In the VLAN Translation dialog, specify the following:
   - Incoming Port – The TPS virtual port through which incoming traffic arrives.
   - Incoming VLAN ID – The ID of the incoming VLAN.
   - Outgoing VLAN ID – The ID of the outgoing VLAN.
   - Automatically create reverse translation – Select this option to enable automatic reverse VLAN translation. This option is disabled by default.
4. Click **OK**.

DNS service

You can configure the Domain Name Service (DNS) on the device to resolve DNS names. Additionally, you can configure the domain name and domain search paths used by the device to resolve DNS names.

By default, the DNS service uses the Management Port to send DNS request packets.

To add domain names and server IP addresses:

1. Select **Network > DNS Service**.
2. Under Domain Names:
   a. Enter a valid **Domain Name**.You can also enter optional domain search names.
   b. In the Server IP Addresses panel, enter up to four IPv4 or IPv6 Server addresses and click **OK**.

Manage policies

Policies specify the security features and requirements of your network, such as rules that determine who is allowed to access the network, what applications they can use, what web sites they can visit, and so on. Policies refer to all of the mechanisms available on the device that protect and manage network traffic, including:

- IPS, Reputation, and SSL Inspection profiles
- Action sets
- Notification contacts
- Services

The Policy page provides access to the following areas:

- Profiles
- Objects

Profile configuration

You can monitor and configure the settings for profiles used by the device.

IPS profiles

An IPS profile defines the traffic that the device monitors and the DV filters that the device applies. Traffic monitoring is based on incoming and outgoing port pairs. You can use the default DV filter configuration to protect the virtual segment or customize the configuration as required. The virtual segment specifies both the port and the traffic direction, which allows you to define separate security profiles for traffic in and out of a port.

The default IPS profile is set to ANY incoming ports and ANY outgoing ports, with all IPS filters configured with the default Digital Vaccine settings. With the default profile in place, all incoming and outgoing traffic on any virtual segment configured on the device is monitored according to the filter configuration recommended by TippingPoint. You can edit the default IPS profile to modify the filter settings. You can also create your own IPS profiles, and edit the virtual segments to apply your own IPS profile instead of the default.

Note: Before creating IPS profiles, verify that the network and system configuration on the device is set up correctly for your environment. In particular, configure all required ports before creating the security profiles to protect them.

When an IPS profile is initially created, the recommended settings for all filter categories are enabled.

Use the IPS Profiles page to perform the following tasks:

- View, create, edit, and delete IPS profiles
- Change category settings for a group of filters
- Specify source and destination addresses to limit or exclude
- Override global filter settings and create filter-level settings
- Restore filter to global category settings
- Specify Direct Denial-of-Service (DDoS) filters

The IPS Profile page includes the following information:

Parameter	Description
Profile Name	The name assigned to the IPS profile. The default IPS profile is preconfigured on the device. You can customize this profile to modify global and individual filter settings.
Description	A description of the security profile, if a description has been defined.

To manage the virtual segments associated with IPS profiles, use the Virtual Segments page. See *Virtual segments* on page 51.

Sample IPS profiles

To manage the virtual segments associated with IPS profiles, use the Virtual Segments page. See *Virtual segments* on page 51.

The following table shows a sample port configuration:

Name	Network Port	VLAN
any	any	any
segment1 (A > B)	1A > 1B	any

Name	Network Port	VLAN
segment2 (A > B)	2A > 2B	any
Marketing-A	1A > 1B	6
Marketing-B	2A > 2B	6

The following table lists some IPS security profiles you can create to monitor traffic on a device with the configuration shown in the preceding table.

Name	Virtual Segment(s) (Incoming, Outgoing)	Description
Marketing	Marketing-A ==> Marketing-B Marketing-B ==> Marketing-A	Monitor all VLAN 6 traffic on port 1A > 1B and port 2A > 2B in both directions.
LAN	segment1 (A > B) segment1 (B > A)	Monitor all traffic between ports 1A > 1B and 2A > 2B, except traffic tagged for VLAN 6. VLAN 6 traffic is covered by the Marketing security profile above.

Default IPS profile

The default IPS profile is set to the ANY< ==> ANY virtual segment with all IPS filters configured with the default Digital Vaccine settings. With the default profile in place, all incoming and outgoing traffic in any virtual segment configured on the device is monitored according to the DV filter configuration recommended by TippingPoint.

You can edit the default security profile to customize the virtual segments that it applies to and create custom filter settings, or create your own security profiles as required.

To manage the virtual segments associated with IPS profiles, use the Virtual Segments page. See *Virtual segments* on page 51 for more information.

Applying IPS profiles to traffic

It is possible for a packet to match more than one IPS profile depending on how the virtual segments are configured. Virtual segments compose an ordered list. The device surveys the list beginning at the top and

finds the first virtual segment that matches the traffic. The IPS profile associated with that virtual segment will be applied to the traffic.

Add an IPS profile

1. Select **Policy** > **IPS**.
2. Click **Add**. The Add IPS Profile dialog is displayed.
3. Enter the profile Name.

   Note: You must specify a unique name for each IPS and Reputation Profile that you create in the LSM.

4. (Optional) Enter a description.
5. Select the **IPS Deployment Type** from the list or leave the setting at default. Deployment mode can vary depending on where the device is deployed in your network. The following modes are available:
   - **Default** – Recommended for all deployment scenarios.
   - **Aggressive** – Offers a more aggressive security posture; profiles using this mode might require tuning based upon specific application protocol usage.
   - **Core** – Recommended for deployment in network core.
   - **Edge** – Recommended for deployment in Web Farm/DMZ.
   - **Hyper-Aggressive** – The most aggressive security posture; profiles using this mode might require tuning based upon specific application protocol usage.
   - **Perimeter** – Recommended for deployment in Internet entry point.
6. Click **OK** to add the IPS profile.

Edit an IPS profile

1. Select **Policy > IPS**.
2. Click **Edit**. The **Edit IPS Profile** dialog is displayed.
3. Under the General tab, modify the name, description and the deployment type as required.
4. Under the IPS Category Rules tab, you can optionally select the appropriate values under Application Protection, Infrastructure Protection, and Performance Protection.
5. Under the Limits/Exceptions tab:
   a. (Optional) To limit or exclude the address(es) from application protection and infrastructure protection filters, select a type (limit or exception) from the list, enter the source and destination IP address, and click **Add**.
   b. (Optional) To limit performance protection filters to the IP address(es), enter the source and destination IP address and click **Add**.
6. Under the IPS Filter Overrides tab, you can search for, edit, and delete override filters:

a. Enter a text string in the Keywords field or enter the filter number in the Filter # field.
b. You can select one of the following states: **Any**, **Disabled**, or **Enabled** from the Filter State list.
c. You can select one of the following controls: **Any**, **Category Settings** (defaults), or **Override** (customized) from the Filter Control list.
d. Select a category from the IPS category list. You can choose from all groups under the application protection, infrastructure protection, and performance protection categories.
e. Select an action-set from the Action Set list. You can choose from all the default and custom Action Sets configured on the device.
f. Select any IP protocol under Protocol, and from Severity, select the severity level.
g. Click **Search**. A list of profiles is returned. You can select a profile and click **View** to see the description, or click **Override in this IPS Profile** to override it.
h. Click **Edit** to make any changes, and to delete a row from the table select the row and click **Delete**.

7. Under the DDos tab click **Add** to add a DDoS filter. The Add dialog is displayed.
   a. Enter the name.
   b. Select an action set from the Action Set list.
   c. Enter a IP address under Destination.
   d. Enter the desired values for SYNs per Second (the number of allowed SYN packets per second) under Threshold.
   e. (Optional) Enter a source IP address and click **Add** to exclude a certain IP address from triggering the filter.
   f. Click **OK**.
   g. Select a row and click **Edit** to make changes to a filter.
   h. To delete a row from the table, select the row and click **Delete**.

Reputation profiles and reputation groups

Reputation profiles contain a list of reputation filters. Each filter contains a reputation group and an action set.

Name	Displays the name you have assigned for the profile.
Action when pending	Displays the action to perform when the reputation lookup is pending.
Check Source	Indicates if the source address is checked.

Check Destination	Indicates if the destination address is checked.
IP Exceptions (Source/ Destination)	Displays the source and destination IP exceptions.
DNS Exceptions	Displays the domain names with DNS exceptions.
Reputation Filters	Displays the reputation filters.

For information about reputation groups, see *Reputation groups* on page 86.

Add a reputation profile

1. Select **Policy > Reputation**.

   Note: You must specify a unique name for each Reputation and IPS profile that you create in the LSM.

2. Click **Add**.

   The Add Reputation Proflie dialog is displayed.

3. Enter a unique name and click **OK** or click **OK/Continue** to add another reputation profile.

   The new reputation profile is added with the Instant-Commit feature.

Edit a reputation profile

When specifying an action set for a reputation profile, as a best practice, add quarantine exceptions for hosts that you never want to quarantine, such as the Default Gateway and DNS Server. For example, when a DNS server receives a request from a client and it does not know the answer, it forwards the request to another authoritative DNS server. So, to the IPS the DNS Server can look like an infected host making a DNS request.

1. Select **Policy > Reputation**.
2. Select the checkbox next to the reputation profile you want to edit and click **Edit**.

   The Edit Reputation Profile dialog is displayed.

3. Make the following configurations under the General tab:
   a. Specify how to apply reputation filters to IP addresses by selecting **Source address**, **Destination address**, or **Both source and destination addresses** from the list.
   b. To apply an action to the packets when the reputation lookup is pending, select **Permit** or **Drop**.
   c. To add a reputation filter, click **Add**. The **Add Reputation Filter** dialog is displayed.
   d. Enter a reputation group name.

e. Select **Enabled** under State/Action.

f. Select an action from the **Action Set** list. Click **OK**.

g. Select a row and click **Edit** to make any changes and **Delete** to delete a reputation filter from the table.

4. Under the Exceptions tab:

   a. Enter a source and destination address and click **Add** so that the reputation filters are not applied to specific IP addresses.

   b. Enter a domain name and click **Add** so that the reputation filters are not applied to specific domains. Click **OK**.

TippingPoint ThreatDV

The TippingPoint ThreatDV is a licensed service that identifies and delivers suspect IPv4 and IPv6 and DNS addresses to subscribers. The addresses are tagged with reputation, geographic, and other identifiers for ready and easy security policy creation and management. The service provides the addresses and tags multiple times a day in the same fashion as Digital Vaccines.

Note: While any user can manually create reputation groups and filters, the ThreatDV is available only to users who have licensed the service from TippingPoint. For more information about this service, ask your TippingPoint representative.

Traffic management profiles

Use the Traffic Management Profiles page (**Policy > Profiles > Traffic Management**) to view, create, edit, or delete a traffic management profile and apply traffic management profiles to virtual segments. A traffic management profile consists of the following components:

- **Profile Details** — Profile name and description.
- **Traffic Filters** — One or more filters to manage the traffic based on Protocol or IP address and port. Each filter defines the type of traffic to be monitored and the action to be taken when the filter is triggered.

Traffic that triggers the traffic management filter is managed based on the filter action configured, which can be any of the following:

- **Block** — Traffic that triggers the filter is denied.
- **Allow** — Allows traffic that meets the filter criteria.
- **Rate Limit** — Rate limits traffic that meets the filter criteria.
- **Trust** — Allows traffic that meets the filter criteria through the device without being inspected.

Traffic that is allowed or rate-limited based on a traffic management filter goes on to be inspected based on the security profile configuration (DV filtering). In other words, traffic is not allowed through the device based solely on the traffic management filter criteria, unless the filter is configured with the Trust action.

Note: Quarantine actions take priority over traffic management trust filters.

The Traffic Management Profiles page lists all the traffic management profiles currently configured on the device and includes the following information:

Parameter	Description
Profile Name	The name assigned to the traffic management profile.
Description	A description of the traffic management profile, if a description has been defined.

To manage the virtual segments associated with security profiles, use the Virtual Segments page. See *Virtual segments* on page 51.

This topic discusses the following information:

Applying traffic management profiles to traffic

You can use traffic management filters to prioritize traffic or implement security policy. For example, you might define the following IP filters for your Web servers in a lab that denies access to external users:

- Block traffic if the source is on an external subnet that arrives through port 80 and is destined for the IP address of your Web server.
- Block traffic if the source is your Web server, the source port is 80, and the destination is any external subnet.

You can define multiple traffic management rules in each profile. In general, when defining filters for network segments, more specific filters should come first. For example, a more specific IP filter might block traffic with fully qualified source and destination IP addresses and ports. More general ones, like those that apply to subnets, should follow.

The following table lists several examples of traffic management filters:

Source address	Destination address	Protocol	Source port	Destination port	Action
any	any	UDP	any	53	Allow
any	any	UDP	any	any	Block
any	any	ICMP	any	any	20 Mbps rate-limit
any	1.2.3.4	TCP	any	80	Allow
any	any	TCP	any	80	Block
66.94.234.13	any	IP	any	80	Block

These filters perform the following actions:

- Block all UDP traffic except DNS requests. DNS requests are inspected for attacks.
- Limit all ICMP traffic to 20 Mbps.
- Block all HTTP traffic except for server 1.2.3.4.
- Block IP fragments coming from IP address 66.94.234.13 on any port going to port 80.

Configure a traffic management profile

1. From the LSM menu, click **Policy > Profiles > Traffic Management Profiles**.
2. Click **Add**.
3. On the Traffic Management Profile Editor page, enter the **Name**. You can also enter a description of the profile.
4. Click **Add** to add traffic filters. Configure the following parameters:

Parameter	Description
Name	The name assigned to the traffic management filter.

Parameter	Description
Action	Indicates how the device will manage traffic that triggers the filter. The following options are available: • **Block** — Traffic that triggers the filter is denied. • **Allow** — Allows traffic that meets the filter criteria. • **Rate Limit** — Rate limits traffic that meets the filter criteria. • **Trust** — For trusted servers or traffic, allows traffic that meets the filter criteria to pass through the device without being inspected.
Version	Specifies whether traffic is IPv4 or IPv6. **Note:** Entering an IPv4-mapped address in IPv6 notation will only match addresses that actually appear in IPv6 packets on the wire. They will not match IPv4 packets. Similarly, a range entered in IPv4 notation will only match IPv4 packets, and not IPv6 packets that contain the equivalent IPv4-mapped addresses.
Protocol	Specifies which protocol the filter checks for: **Any**, **TCP**, **UDP**, or **ICMP**. To apply the filter only to IP fragments, select **Apply only to IP fragments**.
Source Address	Specifies the source IP address and port for traffic that will be managed by the filter. IP addresses can be specified in CIDR format.
Destination Address	Specifies the destination IP address and port for traffic that will be managed by the filter. IP addresses can be specified in CIDR format.

5. Click **OK** or click **OK/Continue** to add another traffic filter.
6. Reorder the traffic filters as necessary, generally prioritizing the more specific filters first.
7. Click **OK**.

SSL inspection profiles and servers

An *SSL inspection profile* enables the 2200T TPS to decrypt and inspect inbound encrypted traffic. See the following sections for more information.

Overview

The TippingPoint Threat Protection System (TPS) 2200T provides in-line, real-time threat protection for inbound SSL traffic. The 2200T manages its own private keys and certificates from the servers it is securing; these can either be stored on the device itself or accessed at run-time from the Security Management System (SMS).

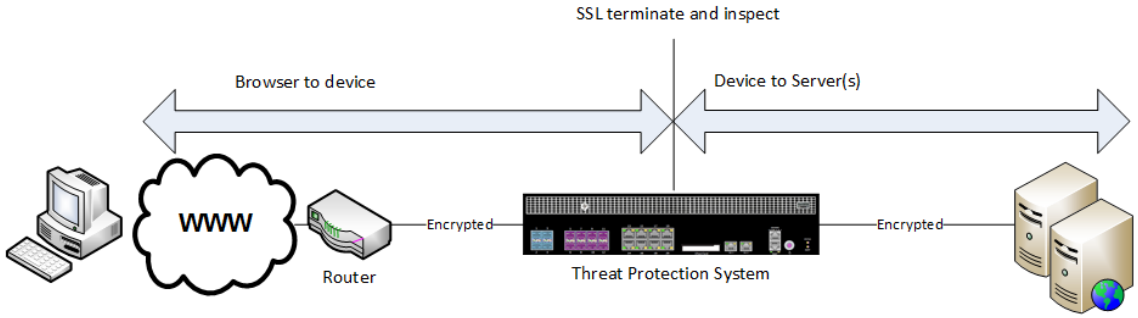

With access to the server certificate and private key, the TPS is a reverse proxy that receives and decrypts SSL data, inspects it using the Threat Suppression Engine, and then encrypts it before sending it to the actual destination.

Requirements

Make sure your environment meets the following requirements:

- TippingPoint 2200T with SSL Inspection Upgrade license. SSL inspection is not supported on the TippingPoint 440T.
- Inbound IPv4 traffic only.
  - SSL inspection does not support IPv6, including IPv4 over IPv6 tunneling.
  - SSL inspection of outbound IPv4 traffic and IPv6 traffic is not supported.
- VLAN translation cannot be used in conjunction with SSL inspection.
- SSL certificate and private key from the server that hosts the SSL/TLS compliant application.
- Cipher suite support:
  - RSA-3DES-EDE-CBC-SHA1 (enabled by default)
  - RSA-AES128-CBC-SHA1 (enabled by default)

- RSA-AES256-CBC-SHA1 (enabled by default)
- RSA-RC4128-MD5 (disabled by default)
- RSA-RC4128-SHA1 (disabled by default)
- RSA-DES-CBC-SHA1 (disabled by default)

- SSL protocol support:
  - TLS v1.0 (enabled by default)
  - TLS v1.1 (enabled by default)
  - TLS v1.2 (enabled by default)
  - SSL v3.0 (disabled by default)

  Note: TLS Heartbeat Extension (*https://tools.ietf.org/html/rfc6520*) is not supported.

Additional considerations

When deploying SSL inspection, consider the following:

Consideration	Description
Tunneled traffic	Supported SSL encapsulations: • GRE (Generic Routing Encapsulation) * • IPv4 (IP-in-IP) • One layer of tunneling only for both GRE and IPv4-in-IPv4 SSL inspection does not include support for GTP or IPv6 encapsulations. * GRE support includes the mandatory GRE fields. We also support optional GRE key configuration, but the key needs to be the same value for both directions. We do not support other optional GRE fields, such as GRE sequence number.
Quarantine hosts and redirecting HTTP traffic to another site	The 2200T TPS does not redirect SSL sessions on a quarantined host. When configuring an Action Set to **quarantine hosts**, if you also configure the response to **redirect HTTP traffic to a web site**, the host of HTTPS traffic is quarantined, but not redirected.

Consideration	Description
Advanced Distributed Denial of Service (DDoS) filters	When managed by the SMS, if you have configured Advanced Distributed Denial of Service (DDoS) filters to protect against SYN floods, the TPS supports DDoS protection of the SSL server, but not on the ports where we are performing inspection of SSL traffic. For example, if you have a server at 1.1.1.1, and that server responds to HTTP traffic on port 80 and HTTPS traffic on port 443, and you have configured SSL Inspection of the traffic on 1.1.1.1:443, and you have configured DDoS protection of 1.1.1.1 - then you will get DDoS protection on 1.1.1.1:80, but not on 1.1.1.1:443.
Trust as an Action	If traffic management trust is selected, then SSL inspection is bypassed. If a filter hits a trust, then IPS and reputation inspection stop, but the TPS continues to proxy the SSL session between the client and the server.
Packet trace	Packet Trace as an action includes the decrypted traffic.
Traffic capture	Traffic capture by tcpdump does not include the decrypted contents.
L2FB/ZPHA	When the 2200T TPS enters Layer 2 Fallback (L2FB) or Zero Power High Availability (ZPHA), the proxied SSL sessions are cleared.

Configure SSL inspection

Configure SSL inspection to specify the SSL sessions you want the 2200T to inspect. The 2200T does not inspect SSL sessions that do not match the SSL inspection profile. Configuring SSL inspection is a deferred commit operation. After you complete your configuration, commit your changes. The process for setting up SSL inspection is described below:

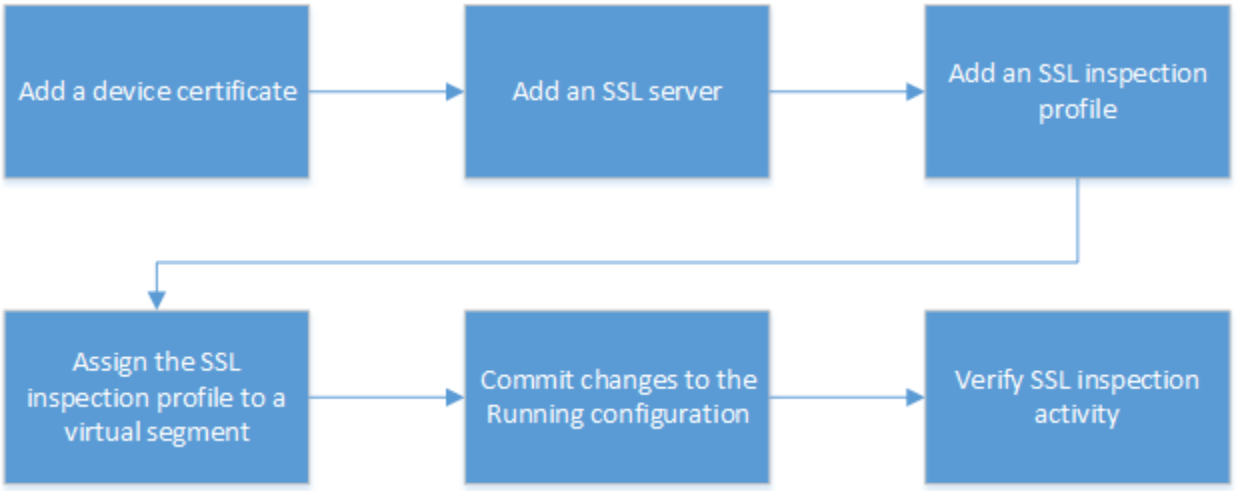

To set up SSL inspection:

1. *Import the SSL server certificate and private key* on page 69.
2. *Add or edit an SSL server* on page 70.
3. *Add or edit an SSL inspection profile* on page 71.
4. *Assign the SSL inspection profile to a virtual segment* on page 72.
5. *Commit changes to the Running configuration* on page 73.
6. *Verify SSL inspection activity* on page 73.

Import the SSL server certificate and private key

From the LSM, add or edit a device certificate to import both the SSL certificate and private key from the server of interest. To commit changes to the TPS, you must import both the SSL certificate and its private key. The TPS does not attempt to validate the status of a device certificate.

1. Select **Authentication > X.509 Certificates > Device Certificates**.
2. In the Device Certificate panel, click **Import** to import a new SSL certificate.

   To update an existing SSL certificate, select the certificate from the list, then click **Import**.
3. Enter the certificate name. We recommend using a naming convention that you can easily and reliably assign the correct certificate to an SSL server.
4. Click **Browse** to locate the file.
5. Select the certificate format, either **Base64 Encoded Certificate (PEM)** or **Encrypted Private Key and Certificate (PKCS12)**.
6. When selecting:

- **PEM** format, the private key must be imported in a separate file. Be sure to select the **Include a Private Key** checkbox, then browse to the private key file. If the private key is encrypted, you must also enter the appropriate password in the Password box.
- **PKCS12** format, you must enter the appropriate password in the Password box. Note that only one certificate/private key pair will be imported, along with all of the CA certificates contained in the file.

7. Click **OK**.

   The appliance imports the certificate and associated private key, and the certificate is displayed in the Device Certificates table.

Add or edit an SSL server

From the LSM, add an SSL server to specify the SSL server configuration to proxy, including the SSL service that is accepted on the SSL detection port. When configuring SSL inspection and Advanced Distributed Denial of Service (DDoS) SYN Proxy filters are configured, if a DDoS filter IP address matches an SSL server address, then DDoS protection will be disabled on any of the configured SSL server ports.

For secure HTTP, IMAP, and POP3 traffic, create a separate SSL server to enable DV filtering on the decrypted SSL service. For example, if the web server accepts POP3S traffic on port 2000, add an SSL server with a Detection Port of 2000 and a Decrypted Service of POP3 to enable DV filters for POP3.

For other SSL services, such as SMTPS, create an SSL server with a Detection Port that identifies the secure traffic, and a Decrypted Service of Other. DV filters are applied to the incoming traffic, but are not applied to the decrypted SSL service.

To inspect more than one decrypted service on a particular SSL server, define the same server IP for each service you want. For example, you can define a server with IP 1.1.1.1 and port 443 (HTTPS), and another server with IP 1.1.1.1 and port 995 (POP3S), and associate them with the same SSL inspection profile.

1. Select **Policy > SSL Inspection > Servers**.
2. In the SSL Servers panel, click **Add** or **Edit** and specify the following settings.
   - **Name -** Enter the server name, for example, `myapp_pop3`. We recommend naming the server so that you easily associate it with your web server.
   - **Server Addresses:** Specify the server IPv4 address or CIDR range.
   - **SSL Detection Ports:** Specify the port range of the encrypted application traffic. For example, if the web server accepts POP3S traffic on port 2000, specify `2000`.
   - **Server Certificate:** Select the SSL certificate for your web server. **Note:** The LSM does not validate the server certificate.
   - **Decrypted Service:** Choose the SSL service that is accepted on the SSL Detection Port to enable filtering for that particular service. If the SSL service you want is not listed, choose **Other**.

- **SSL Cipher Suites:** Select the cipher suites that are offered by the SSL server. If the TPS cannot negotiate a cypher suite with the SSL client, the connection request is refused and the SSL inspection log is updated.
- **Protocols:** Select the cryptographic protocols that are offered by the SSL server.
- **Rekey Interval:** Specify the interval, in seconds, that your web server forces renegotiation of the shared SSL key. If your web server does not offer renegotiation of the shared SSL key, leave this blank.
- **Enable logging:** Select this option to enable the TPS to write log information about SSL inspection to the external log disk (external Cfast). This option collects detailed logging information and should only be enabled for troubleshooting purposes. For example, enable this option if, after you set up SSL inspection, the TPS does not see SSL session activity. By default, this option is disabled. For information about viewing log information, see *Verify SSL inspection activity* on page 73.
- **Allow compression:** Select this option to allow the SSL compression algorithm to be negotiated during the SSL handshake. If your web server does not offer negotiation of SSL compression, disable this option. By default, this option is disabled. If you select this option, and your web server does not offer SSL compression, this setting is ignored.
- **Send TCP reset to server for blocked sessions:** Select this option to always send a TCP reset to the server whenever the TPS blocks an SSL session. This option overrides the TCP reset action set, if enabled, on a DV filter. We recommend that you enable this option so that protected servers can release network resources quickly if flows are blocked. When this option is disabled, the TCP reset action, if enabled on a DV filter, still applies.

3. Click **OK**. You are now ready to assign the SSL server to an SSL inspection profile.

Add or edit an SSL inspection profile

From the LSM, add or edit an SSL inspection profile to specify the SSL traffic that you want to protect. An SSL inspection profile is a set of server policies, each of which specifies an SSL server and a list of client IP exceptions. Assign the SSL inspection profile to the virtual segment that carries the traffic of interest.

1. Select **Policy > SSL Inspection > Profiles**.
2. In the SSL Inspection panel, click **Add** or **Edit**.

   The SSL Profile Editor opens.
3. Enter the SSL profile name, for example, myapp_SSLprofile.
4. Under Server Policies, click **Add**.

   The Add SSL Server Policy dialog box opens.
5. Specify the following settings:
   - **Enable:** Deselect the checkbox to exclude this SSL Server Policy from the SSL inspection profile. By default, this option is selected.

- **Name:** Specify a policy name, for example, that corresponds to the SSL server configuration.
- **SSL Server:** Choose a server to include in SSL inspection.
- **Source Address Exception:** Specify any client IP addresses to exclude from SSL inspection.

6. Click **OK**.

   You are now ready to assign the SSL inspection profile to a virtual segment.

Assign the SSL inspection profile to a virtual segment

From the LSM, assign the SSL inspection profile to the appropriate virtual segments on the 2200T device.

1. From the LSM menu, click **Network > Virtual Segments**.
2. Click **Add** or **Insert** to create a new virtual segment, or click **Edit** to edit an existing virtual segment. Virtual segments that are created by the system can have their profiles modified but are otherwise read-only.
   - Clicking **Add** adds the new virtual segment after all the other user-created virtual segments.
   - Clicking **Insert** inserts the new virtual segment just before the currently selected virtual segment.
   - All system-created virtual segments always appear at the end of the list.
3. In the Add Virtual Segment dialog or Edit Virtual Segment dialog, specify the following:
   - Name – (Required) Name used to identify the virtual segment. Each virtual segment must have a unique name.
   - Description – An optional parameter to provide more detailed information about the virtual segment.
   - IPS Profile – Security profile that you want to apply to the virtual segment. A virtual segment can have only one IPS profile applied to it.
   - Reputation Profile – Reputation profile that you want to apply to the virtual segment. A virtual segment can have only one Reputation profile applied to it.
   - Traffic Management Profile – Traffic Management profile that you want to apply to the virtual segment. A virtual segment can have only one Traffic Management profile applied to it.
   - SSL Inspection Profile – SSL inspection profile that you want to apply to the virtual segment. A virtual segment can have only one SSL inspection profile applied to it.
   - Physical Segments – Physical segment associated with the virtual segment. All physical segments are directional.
   - Traffic Criteria – (Required) Specify any one or all of the following: VLAN ID, Source IP, and Destination IP. For example, omit VLAN ID and specify Destination IP. When specfying a VLAN ID, specify a value between 1 and 4094 in which the segment is included. There can be no duplicate VLAN IDs or overlapping VLAN ranges. No more than 512 VLAN IDs per virtual segment (a

VLAN range of 1–100 counts as 100 IDs). At least one traffic criteria (VLAN ID, source IP address, or destination IP address) must be defined for each virtual segment.

- Source IP Address – Source CIDR associated with the virtual segment. Addresses must be valid IPv4 format. The host portion of address/mask must be 0 (zero). No more than 250 addresses may be specified.
- Destination IP Address – Destination CIDR associated with the virtual segment. Addresses must be valid IPv4 format. The host portion of address/mask must be 0 (zero). No more than 250 addresses may be specified.

4. Click **OK**.

Note: Virtual segments must be created with a physically available segment. If creating a virtual segment generates a UDM warning in the system log, ensure you have associated the virtual segment with a valid physical segment.

Commit changes to the Running configuration

From the LSM, commit your changes to the Running configuration.

Depending on the type of configuration change, the device commits changes to the Running configuration:

- Automatically. An *instant commit* is one that is applied immediately to the Running configuration. Only some items, including Action Sets and Notification Contacts, are instant-commit features. A bright yellow notice is displayed on all features that use instant commit.
- Manually. A *deferred commit* is one that is not immediately committed to the Running configuration. Uncommitted changes are placed into a pending state until you explicitly commit them to the Running configuration. When you log out of the LSM, pending changes are lost.

  Defer your commit until you have completed the necessary configuration changes, then commit all of the changes at once. For example, when creating an SSL server, you must also import a device certificate and assign to the server before you can commit your changes.

To commit your pending changes to the Running configuration, in the Configuration menu, click **Commit pending changes**.

Verify SSL inspection activity

From the LSM, monitor SSL inspection activity.

View information about SSL inspection activity by choosing from the following:

- **Monitor > Sessions > SSL Sessions** displays active session count information for up to 50 SSL sessions. Filter the list to view details for the sessions you want.
- **Monitor > Network > SSL Bandwidth** displays overall SSL traffic seen and amount inspected.
- **Reports > Activity > SSL > Connections** displays the total number of new SSL connections that were created during the 1-minute reporting interval.

- **Reports > Activity > SSL > Connection Rate** displays the average number of new SSL connections created per second during the 1-minute reporting interval.

To view logging information about SSL inspection, choose **Monitor > Logs > SSL Inspection**. The SSL Inspection log displays SSL session information for the SSL servers with logging enabled, including information about SSL sessions that failed to negotiate SSL parameters. Note that by default, when you add an SSL server, logging is disabled. The SSL inspection log does not contain SSL system errors; check the System log.

To display sessions details, such as connection resets, click **Columns > Details**. If you do not see SSL sessions for a particular server, enable logging on that server and then review this log for useful troubleshooting information. When you finish troubleshooting, disable logging on the server. Note that you can also configure notification contacts and thresholds for SSL inspection logs.

The SSL Inspection log does not log SSL sessions that are Blocked or Quarantined:

- Both the IPS Block and Alert logs (**Monitor > IPS**) and the Quarantine log (**Monitor > Quarantine**) have an "SSL Inspected" (y/n) column to report on SSL sessions.
- The Reputation Block and Alert logs (**Monitor > Reputation**) do not report on SSL sessions because Reputation is analyzed prior to SSL Inspection.

Best Practices

Use this checklist to verify your SSL inspection configuration conforms to the recommended best practices.

❑	To help avoid assigning the wrong certificate and private key to a server, use a naming convention for the certificate, private key, and SSL server. The 2200T does not validate the certificate and private key.
❑	Set role-based access controls to limit access to SSL inspection.
❑	Check the System log for errors.
❑	Keep your certificates up-to-date. Whenever you update a certificate on your server, be sure to also update the certificate on the 2200T. If a certificate expires, the System log generates an error.
❑	When managed by the SMS, you can choose to persist private keys on the SMS instead of on the 2200T device. For more information, see the *SMS User Guide*.

Inspection bypass rules

The TippingPoint 440T and 2200T devices enable users to configure inspection bypass rules. Traffic that matches inspection bypass rules is directed through the IPS without inspection. These rules can be applied to traffic according to source or destination IP address, port, or CIDR (Classless Inter-Domain Routing), or to traffic moving through specific ports.

You can now define up to 32 inspection bypass rules on a TippingPoint TPS. Rule configurations that bypass IPv6 traffic or VLAN ranges require additional hardware resources. For example, a single inspection bypass rule for IPv6 or VLAN traffic can result in multiple port-VLAN rule combinations.

Inspection bypass rule	Resulting number of port-VLAN rule combinations
IPv4 traffic on TCP 1556 with untagged traffic or a particular VLAN ID	1
IPv6 traffic on TCP 1556 with untagged traffic or a particular VLAN ID	2
IPv4 traffic on TCP 1556 with VLAN 10 – 100	90
IPv6 traffic on TCP 1556 with VLAN 10 – 100	180

Each TPS supports a maximum number of port-VLAN rule combinations. If the number of configured port-VLAN rule combinations exceeds the maximum threshold for the device, you cannot commit the changes.

For a	Maximum (approximate) number of port-VLAN rule combinations
440T	1200 when bypassing IPv4 or IPv6 traffic
2200T	5500 when bypassing IPv4 traffic 3000 when bypassing IPv6 traffic

From the list of inspection bypass rules, you can reset the **Packet Hit Count** for a particular rule by selecting the rule and clicking **Reset Counts**. To refresh the entire list, click **Refresh** at the top of the page.

Add or edit an inspection bypass rule

Add or edit an inspection bypass rule to enable or disable the rule and to specify the traffic that you do **not** want to inspect.

Inspection bypass rules can also be defined with the `inspection-bypass` context in the Command Line Interface (CLI). Refer to the *Threat Protection System Command Line Interface Reference* for more information.

1. Select **Policy > Inspection Bypass**.
2. In the Inspection Bypass Rules panel, click **Add** or **Edit** and specify the following settings.
   - **Name -** Specify the name of the inspection bypass rule.
   - **Enabled -** Select this option to enable the inspection bypass rule. This option is enabled by default.
   - **Ethernet Type** Choose an option to specify the EtherType or choose Custom to specify the hexadecimal value of the EtherType to bypass. When specifying a hexadecimal value, prepend the value with 0x, for example, 0x0806 for ARP. By default, IP is selected.

     Note: A full list of *EtherType values* can be found at the Internet Assigned Numbers Authority website.
   - **IP Protocol** Choose an option to specify the IP protocol or choose Custom to specify the IP protocol value to bypass.

     Note: A full list of *IP protocol values* can be found at the Internet Assigned Numbers Authority website.
   - **Source address and ports** Specify the source IP address and port range to bypass.
   - **Destination address** Specify the destination IP address and port range to bypass.
   - **VLAN** Choose an option to specify the VLAN traffic to bypass.
     - **ID or Range** Use this option to specify the tagged traffic you do not want to inspect. For example, specify 12-15 to not inspect tagged traffic on VLANs 12 to 15.
     - **None** to bypass all untagged traffic.
     - **Any** to bypass any tagged or untagged traffic. This option is selected by default.
   - **Segments** Bypass a port by choosing the incoming port from the list and clicking Add.

     Note: Inspection bypass applies to incoming traffic only.
3. Click **OK**

Inspection profile settings

Select **Policy > Settings** to configure the following:

Setting	Description
Connection Table	Specifies the global timeout interval for TCP traffic or non-TCP traffic on the connection table. For blocked streams in the connection table, this value determines the time interval that elapses before the blocked connection is cleared from the connection table. Before the timeout occurs, any incoming packets for that stream are blocked at the device. After the connection is cleared (the timeout interval expires), the incoming connection is allowed until traffic matches another blocking filter.
Trust Streams	Specifies the global timeout interval for trusted streams. This value determines the time interval that elapses before the trusted stream is flushed.
Quarantined Addresses	This value determines the time interval that elapses before the quarantined host can be released. After the quarantined host is released (the timeout interval expires), quarantined addresses can be automatically released, if that option is selected.
HTTP Response Processing	Specifies inspection of encoded HTTP responses. • **Accelerated inspection of encoded HTTP responses** — Hardware acceleration is used to detect and decode encoded HTTP responses. • **Inspect encoded HTTP responses** — Enables strict detection and decoding of encoded HTTP responses. • **Ignore encoded HTTP responses** — The device does not detect or decode encoded HTTP responses. **Enable decoding of URL encodings and Numeric Character References (NCR).** This option is enabled by default.
GZIP Decompression	Decompresses files that have been compressed in the gzip file format.

Setting	Description
Asymmetric Network	Specifies whether the device is configured for an asymmetric network. When asymmetric configuration is enabled, the device does not see both sides of a TCP connection. This option is enabled by default. **Note:** DDoS filters and SSL inspection will not be in effect while Asymmetric Network mode is enabled.
DNS Reputation	Allows the device to respond with `NXDOMAIN (name does not exist)` to clients that make DNS requests for hosts that are blocked.
HTTP Mode	Allows all TCP ports to be treated as HTTP ports for inspection purposes. Enable this feature only on devices that primarily handle HTTP traffic so that optimum performance is maintained.
IDS Mode	When IDS mode is enabled, it adjusts the device configuration so that the device operates in a manner suitable for Intrusion Detection System (IDS) scenarios and filter configurations. • Performance protection is disabled. • Adaptive Filtering is set to Manual. • Filters currently set to Block are not switched to Permit, and Block filters can still be set. When IDS Mode settings are changed, reboot the device for the change to take effect. **Important:** Changing IDS Mode does not change Performance Protection mode. For best results, when enabling IDS Mode, go to the **System > Settings > Log Configuration > Performance Protection** page and change Performance Protection to **Always log Alert and Block events** mode.
Reset Security Profile	Removes all user-created security policy configuration changes from the device, including user-created profiles, user-created virtual segments, filter configurations in security profiles, and action sets.

Object configuration

You can monitor and configure the settings for objects used by the device.

Action sets

Note: This is an Instant-Commit feature. Changes take effect immediately.

Action sets determine what the device does when a packet matches a rule or triggers a filter. An action set can contain more than one action, and can contain more than one type of action. The types of action that determine where a packet is sent after it is inspected include the following:

- A permit action allows a packet to reach its intended destination.
- A block action discards a packet. A block action can also be configured to quarantine the host and/or perform a TCP reset.
- A rate limit action enables you to define the maximum bandwidth available for the traffic stream.
- A trust action allows the designated traffic to bypass all inspection; the traffic is transmitted immediately. Trust has lower latency than Permit, and using it can reduce load on the CPU and processors.

Action Name	Description
Recommended	The default action set, as determined by the filter's category settings. When this action set is assigned to a filter, the filter uses the recommended action setting for the default category settings. The recommended action set can enable different configurations for filters within the same category. Under a recommended category setting, some filters are disabled while others are enabled; some might have permit actions assigned while others are set to block.
Block (+TCP Reset)	Blocks a packet from being transferred to the network. TCP Reset is an option for resetting blocked TCP flows.
Block + Notify (+TCP Reset)	Blocks a packet from being transferred. Notifies all selected contacts of the blocked packet. TCP Reset is an option for resetting blocked TCP flows.

Action Name	Description
	When creating an action set with Block + Notify + TCP Reset Destination, when a Reputation filter is hit, the TCP Reset to the Destination IP does not work properly. To resolve this problem, do not use the 'tcp reset' feature or only use 'tcp reset both' when the trigger reason is Reputation.
Block + Notify + Trace (+TCP Reset)	Blocks a packet from being transferred. Notifies all selected contacts of the blocked packet. Logs all information about the packet according to the packet trace settings. TCP Reset is an option for resetting blocked TCP flows.
Permit + Notify	Permits a packet and notifies all selected contacts of the packet.
Permit + Notify + Trace	Permits a packet. Notifies all selected contacts of the packet, and logs all information about the packet according to the packet trace settings.
Trust	Not configured on the device by default; you must create a Trust action set for this action to appear on the table. Allows trusted traffic to pass without inspection. Lower latency than Permit. Cannot be used with DDoS or IP Reputation filters.

The action sets contain the following columns:

Name	Name of the action set.
Action(s)	Actions included in the action set.
Packet Trace	Whether packet tracing is enabled.
Contact(s)	Where the notifications are sent if a notification contact is configured on the action set.
Quarantine	Time taken for the action set to be quarantined.

Add or edit an action set

1. Select **Policy > Objects > Action Sets**.
2. Click **Add** to create a new action set or **Edit** to change an existing one.
3. Under the General tab:
   a. Enter the name of the action set.
   b. (Optional) Select the action from the **Action** list.
   c. Select whether the option to reset a TCP connection is enabled. With **TCP Reset** enabled, the system resets the TCP connection for the source or destination IP when the Block action executes. This option can be configured on Block action sets.
   d. (Optional) Select **Packet Trace**. Packet Trace enables you to capture all or part of a suspicious packet for analysis. You can set the packet trace priority and packet trace verbosity for action sets.
      - Priority sets the relative importance of the information captured. Low priority items are discarded before medium priority items if there is a resource shortage.
      - Verbosity determines how much of a suspicious packet will be logged for analysis. If you choose full verbosity, the whole packet is recorded. If you choose partial verbosity, you can choose how many bytes of the packet (from 64 to 25,618 bytes) the packet trace log records.
4. Under the Notification Contacts tab, configure notification contacts (either human or machine) that get sent messages in response to a traffic-related event. You can configure any of the following notification contacts to be notified when the action is triggered:
   - Remote System Log – Sends messages to a syslog server on your network. This is a default contact available in all action sets.
   - Management Console – Sends messages to the LSM device management application. This default contact is available in all action sets. If this contact is selected, messages are sent to the Alert or IPS Block Log in the LSM, depending on whether a permit or block action has executed.
5. Under the Quarantine tab, assign a quarantine action set to a filter. You can select the following quarantine options for the action set:
   - (Optional) Select **Quarantine hosts that trigger this action** to quarantine the IP addresses that trigger this option.
      - Select **Quarantine hosts after first hit** to quarantine the host after the first hit.
      - Select **Quarantine host after** to activate the quarantine after the specified number of hits (2 – 10,000) during the specified number of minutes (1 – 60).
   - Select **Block non-HTTP traffic sent from quarantined hosts** – To block the non-HTTP requests.
   - Select an action from the **Response to HTTP traffic sent from quarantined hosts** list:

- **Displaying quarantine info** – Select **Event that triggered the quarantine action** to display the events that triggered the quarantine action and select **Text below** to insert custom text.
- **Blocking it** – To block the response to the HTTP traffic.
- **Redirecting to the following site** – To redirect the HTTP requests from the quarantined host to a website.

6. Under the Quarantine Exceptions tab, you can select the following quarantine exceptions for the action set if you enabled the **Quarantine hosts that trigger this action** option in the preceding step:
   - **Only quarantine these hosts** – To quarantine specified hosts, enter the IP address/mask and click **Add**.
   - **Do not quarantine these hosts** – To exclude the specified hosts from quarantine, enter the IP address/mask and click **Add**.
   - **Allow quarantined hosts to access these addresses** – To allow the quarantined hosts to access the specified addresses, enter the IP address/mask and click **Add**.
7. Click **OK** or **OK/Continue** to add another action set.

Notification contacts

Configure notification contacts to send messages to a recipient (either human or machine) in response to a traffic-related event that occurs on the device. The traffic-related event can be the result of triggering an IPS filter configured with an action set that specifies a notification contact. A notification contact can be any of the following:

- **Remote System Log** — Sends messages to a syslog server on your network. This is a default contact available in all IPS action sets.
- **Management Console** — Sends messages to the LSM. This default contact is available in all action sets. If this contact is selected, messages are sent to the Alert or IPS Block Log in the LSM, depending on whether a permit or block action has executed. This notification contact does not require any configuration, although you can change the default name and aggregation period.
- **Email or SNMP** — Sends messages to the email address or specified SNMP. All email or SNMP contacts must be added from the Notification Contacts page. If the default email server is not configured on the device, you are prompted to configure it before adding a contact.

Note: Before creating an Email or notification contact, you must configure Email and SMTP server settings on the device from the **System > Email** page. For details, see *Add an email or SNMP notification contact* on page 85.

After configuring notification contacts, you can select them for IPS filter events when you create or edit the action set assigned to the filter.

Note: Changes to notification settings take effect immediately.

Use the Notification Contacts page to perform the following tasks:

- View existing notification contacts
- Add new contacts

The Notification Contacts page lists the following information:

Parameter	Description
Contact Name	The name assigned to the contact.
Type	The type of contact. The type can be MGMT, SYSLOG, or Email.
Aggregation Period	The aggregation period, in minutes, for the contact.
Other Parameters	Other information about the contact. For example, the Remote System Log contact shows the number of remote syslog servers configured for the device.

This topic discusses the following information:

- *Alert aggregation and the aggregation period* on page 83
- *Configure the management console contact* on page 84
- *Configure the remote system log contact* on page 84
- *Add an email or SNMP notification contact* on page 85

Alert aggregation and the aggregation period

The device uses alert aggregation to prevent system performance problems resulting from an excessive number of notification requests. Because a single packet can trigger an alert, attacks with large numbers of packets could potentially flood the alert mechanism used to send out notifications. Alert aggregation allows you to receive alert notifications at intervals to prevent this flooding. For example, if the aggregation interval is 5 minutes, the system sends an alert at the first IPS filter trigger, collects subsequent alerts and sends them out every five minutes.

On the device, alert aggregation is controlled by the *aggregation period* that you configure when you create a notification contact. This setting is required for all notification contacts.

Caution: Short aggregation periods can significantly affect system performance. The shorter the aggregation period, the higher the system load. In the event of a flood attack, a short aggregation period can lead to system performance problems.

In addition to the user-configured aggregation period, the system also provides alert aggregation services to protect the system from over-active filters that can lower performance.

For email contacts, the aggregation period works in conjunction with the *email threshold* setting configured for the email server. See *Configure email settings* on page 111.

Configure the management console contact

1. On the LSM menu, click **Policy > Notification Contacts**.
2. On the Notification Contacts page, select the **Management Console** checkbox and click **Edit**.
3. Edit the **Contact Name**. By default, it is `Management Console`.
4. Enter the **Aggregation Period** for notification messages in minutes.
5. Click **OK**.

Configure the remote system log contact

Designating a remote system log as the notification contact sends messages to a syslog server on your network. This is a default contact available in all action sets.

Caution: Remote syslog, in adherence to RFC 3164, sends clear text log messages using the UDP protocol with no additional security protections. Use remote syslog only on a secure, trusted network to prevent syslog messages from being intercepted, altered, or spoofed by a third party.

1. On the LSM menu, click **Policy > Notification Contacts**.
2. On the Notification Contacts page, select the checkbox next to **Remote System Log** and click **Edit**.
3. Enter the contact name.

   By default, it is `Remote System Log`.
4. Enter the **Aggregation Period** for notification messages in minutes.
5. Enter the remote system log's host IP address and port number.
6. Select an **Alert Facility** and a **Block Facility**: none or select from a range of 0 to 31.

   The syslog server uses these numbers to identify the message source.
7. Click **Add** to add the remote syslog server.
8. Repeat steps 3–7 to add additional remote system log servers.
9. Click **OK** to save the changes.

Verify that your device can reach the remote system log server on your network. If the remote system log server is on a different subnet than the management port, you might need to configure the routing. For details, see *Add management port routes* on page 109.

Add an email or SNMP notification contact

Note: This is an Instant-Commit feature. Changes take effect immediately.

1. Select **Policy > Notification Contacts**.
2. To add a notification contact, click **Add**.

   The Add Notification Contact dialog is displayed.
3. Select **Email** or **SNMP**.

   Note: SNMP notification contacts require SNMPv2, and do not work when SNMPv2 is disabled.
4. Enter the notification contact name.
5. Enter the **Aggregation Period** for notification messages in minutes (0 – 10,080).
6. Contact:
   - If the contact is an email contact, enter the address where notifications are to be sent in the **To Email Address** field.
   - If the contact is an SNMP contact, enter the community string, host IP address and port number.
7. Click **OK** or **OK/Continue** to enter another contact.

If the email is not sent correctly, ensure that:

- the default email server is configured
- the email server is reachable from the device
- the email allows mail relaying and that you use the account/domain that the email server accepts

Note: You cannot delete the default Remote System Log and Management Console contacts or a Notification Contact if it is currently configured in another Action Set.

Note: Before creating an Email or notification contact, you must configure Email and SMTP server settings on the device from the **System > Email** page. For details, see *Configure email settings* on page 111.

The Action Set and Log Configuration tells the device when to send notifications.

Reputation groups

As a part of IPS profiles, users can create groups of IP addresses and DNS names, known as reputation groups. Reputation filters enable you to apply block, permit, or notify actions across an entire reputation group.

When an IP address or DNS name is added to a reputation group, it is added to the device's reputation database. Incoming traffic is checked against the database, and the appropriate reputation filters are then applied. While the address or name is being looked up, you can choose to have packets from a suspect address dropped or permitted.

Note: Reputation filter hits in the logs appear to report traffic protocol as `ip` instead of `ip6`. These hits are actually showing the matched signature's *`protocolType`* instead of the traffic *`protocolType`*. Traffic protocols can be confirmed by checking the source and destination addresses.

The TippingPoint SMS offers additional reputation features; refer to the *Tipping Point Security Management System User Guide* for more information.

Use the Reputation Groups page to perform the following tasks:

- View existing reputation groups
- Manually create reputation groups
- Delete reputation groups

The Reputation Group feature enables you to create groups of IPv4, IPv6, and DNS addresses and define an action set to apply to all of those addresses.

The Reputation Groups page lists the following information:

Name	Name you have assigned for the profile.
Entries	Number of entries.
Members	Members in the reputation group.
Description	Purpose of this reputation group.

Add or edit a reputation group

1. Select **Policy > Objects > Reputation Groups**.
2. Click **Add** to create a reputation group or **Edit** to edit an existing group.

3. Configure the name and description of the reputation group.
4. Specify whether the members are to be grouped by IP addresses or domains, and then specify the addresses or domains.
5. Click **OK** or click **OK/Continue** to add another reputation group.

Services

The Services page enables you to configure additional ports associated with specific services and protocols to expand the range of traffic scanned by the device or to use in firewall rules. During the inspection process, the device first scans traffic against the standard ports for listed services, and then scans traffic against the list of additional ports you configure. You can configure up to 16 additional ports for each service other than HTTP. For HTTP, only eight additional ports are allowed.

Services contain the following columns:

Name	Name assigned to the service.
Default Ports	The port assigned to the service by default.
Custom Ports	Additional ports assigned to the service by the user.

Configure a service on a custom port

1. Select **Policy** > **Services**.
2. Select the checkbox next to a service and click **Edit**.

   The **Edit Service** dialog is displayed.
3. Enter a port number and click **Add** to add another port for that service.
4. Click **OK**.

Manage authentication

The LSM Authentication pages enable administrators to:

- Create and manage user accounts
- Set user account preferences
- Manage device certificates
- Manage CA certificates
- Log in administratively to the management console
- Configure fine-grained access to the functional areas of the management console using locally defined users, user groups, roles, or an established LDAP or RADIUS server.

Authentication servers

The TPS supports two types of back-end servers for remote authentication:

- RADIUS Server
- LDAP Server

Note: By default, RADIUS and LDAP servers send traffic over the management port.

LDAP groups

Use the LDAP Groups panel to configure up to six v2 or v3 LDAP servers for administrative login authentication and network user authentication.

The TPS device checks the accessibility of each server when it is created or modified. Inaccessible servers get rechecked periodically by the device (approximately once every five minutes). The system log reflects whether the state of the server has changed. To prevent login delays, only accessible servers are contacted in order of priority. If all the servers are inaccessible, the device contacts the highest priority server.

Name	Name of the LDAP group.
Bind DN	Bind Distinguished Name, which identifies the user on the external LDAP server who is permitted to search the LDAP directory. The Bind DN is made up of one or more attribute=value pairs, separated by commas.

TLS	Transport Layer Security, which provides options for encrypting communication to the LDAP server.
Version	LDAP version, either Version 2 or Version 3.
Schema	The specified LDAP Schema: Microsoft Active Directory, Novell Directory, FedoraDS, RFC2798, RFC2307 NIS, Samba SMB, Custom.
Timeout	Time limit on failed connections to the server. The Default value is 2 seconds. This value can be set to 0 – 10 seconds.
Retries	Number of times to retry a search connection after an initial attempt fails. The default value is 1. This value can be set to 0 – 3.
Port	Displays the LDAP Server port.
Servers	Displays the IP address of the LDAP Servers.

Add or edit an LDAP group

1. Select **Authentication > Authentication Servers**.
2. In the LDAP Groups panel, click **Add** to create a new LDAP group or **Edit** to change an existing one.
3. On the General tab of the dialog, enter a name for the group.
4. Select the LDAP Version, either **Version 3** or **Version 2**.
5. Select the LDAP Schema. You can optionally click the ellipses (**...**) button, which opens the Edit LDAP Schema dialog, to configure additional schema information. The default value is **Microsoft Active Directory**.
6. (Optional) Enable or disable options under TLS Configuration. By default, no options are enabled.
   - **TLS Encryption** Enables TLS Encryption.
   - **Start TLS over LDAP** enables TLS security to use both secure and non-secure requests against the LDAP server in a single connection. For example, modifications to the LDAP server are secure, but reading parts of the directory that are open for unauthenticated viewing do not use secure requests.
   - **Valid Server X.509 Certificate** enables the use of an X.509 certificate for secure authentication. Select **Authentication > X.509 Certificates** to import the CA certificate required to validate the server's certificate.
7. (Optional) On the Login tab, configure the following options:

- **Bind DN** – Provides the user permitted to search the LDAP directory.
- **Bind Password** – Provides the password for the user permitted to search the LDAP directory.
- **Base DN for Tree Search** – Indicates the starting point for searches on the LDAP directory.

8. (Optional) On the Server tab, you can modify the default values for the LDAP Server:
   - **Server Port** – Default value port is 636.
   - **Server Retries** – Specifies the number of times to retry a search connection after an initial attempt fails.
   - **Server Timeout** – Specifies a time limit on failed connections.
   - **LDAP Servers** – Specifies the IP address or domain name of the server.
9. Click Add to get the LDAP group you configured added as a member.
10. Click **OK**.

    The LDAP group is added and is displayed in the LDAP Groups table.

RADIUS groups

A RADIUS group is a group of RADIUS servers with a common configuration, including:

- Device user group. If the user group is not defined by the RADIUS server, the default is Administrator
- Authentication protocol and the number of server retries

When the authentication protocol is PEAP/EAP-MSCHAPv2, be sure to also import the CA root certificate. The RADIUS group authenticates against the available CA root certificates on the device.

Add or edit RADIUS group

1. Select **Authentication > Authentication Servers**.
2. In the Radius Groups panel, click **Add** to create a new RADIUS group or **Edit** to change an existing one.
3. In the Add RADIUS Group dialog, enter a name up to 64 characters in length.
4. (Optional) Select the **Default User Group** as None or administrator/operator/superuser or click the ellipses (...), which opens the Add Default User Group dialog, to create a new user group. This is the group a RADIUS user will be assigned if the response contains no Filter-ID attribute.
5. Select the Authentication Protocol from the list.

   To use the PEAP/EAP-MSCHAPv2 protocol, you must first import the CA root certificate for the RADIUS server or servers in the group.
6. (Optional) Specify the number of Server Retries between 0 and 5. The default value is 1.

7. In the RADIUS Servers panel, add a server to the group by specifying the following:

Setting	Description
Server	IP Address of the RADIUS server.
Port	Port on the RADIUS server that listens for authentication requests. The default port is 1812.
Timeout	Timeout, between 1 and 14 seconds, for communication with the RADIUS server. Default is 2 seconds.
NAS ID	(Optional) Value of RADIUS attribute 32, NAS-Identifier. The attribute contains a string identifying the NAS (Network Access Server) used in the RADIUS Access-Request packet.
Password	Case-sensitive string used to encrypt and sign packets between RADIUS clients and the RADIUS server, set in the RADIUS client configuration file. Maximum is 64 characters.

8. Click **Add** and then **OK** or **OK/Continue** to enter another contact.
9. Reorder the RADIUS servers to specify the order in which the IPS communicates with the authentication servers in the group. See the next section for more information.

Reorder RADIUS servers

Reorder RADIUS servers, from top to bottom, to specify the order in which the IPS. Starting from the top-most server in the list, the IPS attempts to communicate with the authentication server. If unsuccessful, the IPS attempts to establish communication with the next server in the list.

1. Select **Authentication > Authentication Servers**.
2. Click the checkbox next to an existing RADIUS Group.
3. Click **Edit**.
4. In the RADIUS Servers panel, select the server and click the up or down arrow.
5. Click **OK**.

Authentication settings

You can configure global authentication settings that apply to local users and groups created on the device. The global authentication settings include options for:

- Password Settings
- Login Settings
- Login Group

Configure authentication settings

1. Select **Authentication > Authentication Settings**.
2. In the Password Settings panel, configure the following settings:
    - Password Security Level – Specifies the level of security required when creating a password. The default value is **Maximum**. Options include:
        - **Maximum** – Basic plus a minimum of two alpha characters, at least one numeric character, and at least one non-alphanumeric character)
        - **Basic** – Must be at least 8 characters in length.
        - **None**
    - Password Expiry Time – Specifies the length of time the password is valid. Default value: **30 days**
    - Password Expiry Action – Specifies the action a user must take if a password expires. Default value: **Force user to change password**
3. In the Login Settings panel, configure the following settings:
    - Maximum Login Attempts – Specifies the number of times a user can attempt to log in. Default value: **4**
    - Failed Login Action – Specifies the action to take if the Maximum Login Attempts is reached. Default value: **Lockout account or IP address**
    - Lockout Time – Specifies the length of time to lock out a user if the Failed Login Action includes a user lockout. Default value: 2 minutes
4. In the Login Group panel, configure the following settings:
    - Administrative Authentication – Specifies the LDAP or RADIUS group to be used for Administrative login to the LSM. The local database of users is always enabled by default.
5. Click **OK**.

Device certificates

The Device Certificates table displays information about certificates that have been imported to the device. It contains the following information for each certificate:

Certificate Name	Displays the name you specified for the certificate.
Common Name	Displays the fully qualified domain name or IP address of the web server.
Expires On	Displays the certificate expiration date.
Status	Displays one of the following certificate statuses: • Valid • Not yet valid – The current date occurs before the certificate "valid from" date. • Expired – The current date occurs after the certificate "expires on" date. • Self-signed – Warning that the certificate is self-signed. • Revoked – Certificate has been revoked by CRL. • Invalid CA – Certificate CA is invalid. • Not verified – Certificate status could not be determined; this results when the CA certificate chain is not complete, or is not marked Trusted.

Add or edit a device certificate

Add or edit a device certificate to import both the SSL certificate and private key from the server of interest. To commit changes to the TPS, you must import both the SSL certificate and its private key. The IPS does not attempt to validate the status of a device certificate.

1. Select **Authentication > X.509 Certificates > Device Certificates**.
2. In the Device Certificate panel, click **Import** to import a new SSL certificate.
    - To update an existing SSL certificate, select the certificate from the list, then click **Import**.

3. Enter the certificate name. We recommend using a naming convention that you can easily and reliably assign the correct certificate to an SSL server.
4. Click **Browse** to locate the file.
5. Select the certificate format, either **Base64 Encoded Certificate (PEM)** or **Encrypted Private Key and Certificate (PKCS12)**.
6. When selecting:
   - **PEM** format, the private key must be imported in a separate file. Be sure to select the **Include a Private Key** checkbox, then browse to the private key file. If the private key is encrypted, you must also enter the appropriate password in the Password box.
   - **PKCS12** format, you must enter the appropriate password in the Password box. Note that only one certificate/private key pair will be imported, along with all of the CA certificates contained in the file.
7. Click **OK**.

   The appliance imports the certificate and associated private key, and the certificate is displayed in the Device Certificates table.

CA certificates

Your device attempts to validate the status of any certificate presented during authentication (such as from an LDAP server). In order to validate a given certificate, you must import a sufficient chain of CA certificates. To import CA certificates, use the **Authentication > X.509 Certificates > CA Certificates** page and add the CA to the Certificate Authority list.

The CA Certificates table contains the following information:

CA Name	Name you specified for the certificate.
Common Name	Name assigned to the CA certificate by the creator. It can be set to any value.
Expires On	Certificate expiration date.
CRL Expiry	Date when the currently loaded Certificate Revocation List (CRL) expires.
Status	One of the following certificate statuses: • Valid

	• Not yet valid – The current date occurs before the certificate "valid from" date. • Expired – The current date occurs after the certificate "expires on" date. • Self-signed – Warning that the certificate is self-signed. • Revoked – Certificate has been revoked by CRL. • Invalid CA – Certificate CA is invalid. • Rejected – Specified purpose of certificate is not acceptable. • Not verified – Certificate status could not be determined; this results when the CA certificate chain is not complete, or not marked Trusted.

Import a CA certificate

1. Select **Authentication > X.509 Certificates > CA Certificates**.
2. Click **Import**.
3. Enter the CA Certificate Name.

   Any CA certificates bundled with PKCS12 imported certificates will be displayed here with the name `<certificate name>_ca`.
4. Click **Browse** to locate and select the CA Certificate File.
5. Click **OK**.

Users and groups

In the Users and Groups menu, you can access the following sub-menus:

- *User groups* on page 95
- *Local users* on page 96
- *User roles* on page 97

User groups

The TPS provides a predefined set of user groups that each have an assigned role with set access privileges. Each user group can have an associated role that determines the type of administrative functions that are allowed. If a user group does not have any management roles, it can still be used in policy configuration.

Administrative users can create, edit, and delete any user group except the default groups:

- Administrator – Has Read/Write privileges to all TPS capabilities except administering local users, user groups, and roles. Administrator privileges are for an enhanced administrator user who can view, manage, and configure functions and options in the system.
- Operator – Has Read-only privileges to all TPS capabilities. Operator privileges are for a base-level administrator user who monitors the system and network traffic.
- SuperUser – Has Read/Write/Execute privileges to all TPS capabilities. SuperUser privileges include full access to all LSM and CLI functions.

Add or edit a user group

1. Select **Authentication > User Groups.**
2. Click **Add** to create a new user group or **Edit** to change an existing one.
3. Enter a name.
4. (Optional) Specify the Mapped LDAP Group Name. Any LDAP user who is a member of the LDAP group is handled as a member of the group.
5. (Optional) Select an Administrative Role or click the ellipses (**...**) to add a new role. For more information, see *Add or edit user roles* on page 97.
6. Select an available user and click the right arrow to add the user as a member of the group.
7. Click **OK** or **OK/Continue** to add another user group.

Local users

You can create users and add them to a users group on the local device database. A local user can be a member of multiple user groups.

The Local Users table lists all the configured local users, their administrative roles, the user groups to which they belong, and the whether they are currently enabled or disabled.

Add or edit a local user

1. Select **Authentication > Local Users.**
2. Click **Add** to create a new local user or **Edit** to change an existing one.
3. (Optional) Remove the check from **Enabled** to disable this user.
4. Enter a name. User names can contain lowercase letters, uppercase letters, numbers and hyphens. A username cannot be all numbers and cannot start with a hyphen.
5. Enter and password using at least one uppercase letter, one lowercase letter, one number, and one special character, between 8 and 64 characters long.
6. Confirm the password.

7. (Optional) Select a group from the list or click the ellipses (**...**) to add a user group. For information on adding a user group see *Add or edit a user group* on page 96.
8. Click **OK** or **OK/Continue** to add another local user.

User roles

Device administrators with the Superuser role can create custom user roles using the Operator, Administrator, and Superuser roles as templates for each new role. For a description of the privileges associated with each of these default roles, see *User groups* on page 95.

Capabilities can be removed or modified as needed to custom user roles. This enables more granular control over access privileges based on requirements that correlate with a user's tasks and responsibilities. Only custom user roles can be edited; the default user roles cannot be edited.

Hover over each user role in the User Roles table to see all the capabilities available to someone with that assigned role.

You can create up to four custom user roles.

Add or edit user roles

Note: Only custom user roles can be edited; the default user roles cannot be edited.

1. Select **Authentication > User Roles.**
2. Click **Add** to create a user role or **Edit** to change an existing custom user role.
3. Enter a name.
4. (Optional) Enter a description for the user role.
5. Select one of the default roles to use as a template base role for the new role.
6. Check or uncheck each capability for the new role.
7. Select either **Read-only** or **Read/Write** for the state.
8. Click **OK** or **OK/Continue** to add another user role.

   The new role is displayed in the User Roles table.

Reports

Reports enable you to visualize your network activity and measure how current security policies are performing. You can use the reports to analyze traffic patterns and then fine-tune policy as needed.

In addition to the reports available on the Reports page, you can also access reporting information on the Dashboard and Monitor pages. The Dashboard provides information in the form of graphs on device performance. The Monitor page provides additional graphical reports on system health.

Most reports offer several different views of the report data. You can select a different view of the data by selecting an option from drop-down list located on the right side of the page. Not all reports offer the same view options. See the individual report descriptions to see the view options for that report.

You can use one or both of the following refresh methods:

- **Auto Refresh** – Click the **Auto Refresh** checkbox to refresh the contents of the page every 30 seconds.
- **Refresh** – Click the **Refresh** link to perform an instant refresh of the page. You can force an instant-refresh at any time, even if you enabled Auto Refresh.

This topic contains the following information:

Activity reports

Activity reports contain information about network traffic and network activity. The Activity reports include:

Rate Limiters report

When traffic triggers an IPS filter configured with a rate-limit action set, traffic from the source IP and port is limited based on the rate-limit settings. Traffic from the source IP address and port to the destination IP address and port remains rate-limited until the connection timeout period expires based on the Connection Table timeout setting configured from the **Policy > Settings** page. The default timeout setting is 1800 seconds (30 minutes).

This report shows the percentage consumed for each rate limiter pipe. Data will only be displayed in this report if you define at least one rate limit action set and assign it to a profile.

The following view options are available for this report:

- Last minute
- Last hour
- Last 24 hours
- Last 7 days
- Last 30 days

Traffic Profile report

This report shows the number of permitted packets per second, grouped by packet size. Packet size is represented by a color depicted on the legend.

The following view options are available for this report:

- Last 24 hours
- Last 7 days
- Last 30 days

SSL Connections report

The SSL Connections report has two sections:

- **Connections** displays the total number of new SSL connections that were created during the 1-minute reporting interval.
- **Connection Rate** displays the average number of new SSL connections created per second during the 1-minute reporting interval.

The following view options are available for this report:

- Last 24 hours
- Last 7 days
- Last 30 days

Security reports

Security reports contain information about the performance and activity for the device. The Security reports include:

- *Adaptive filter control* on page 100

Adaptive filter control

This report shows Application and IPS Security filters that are being evaluated too frequently. These filters might be causing extra system overhead without ultimately matching any traffic, which can cause performance degradation of the device. This can indicate a defective filter, or maybe a filter that does not perform well in your network environment. By default, the device automatically detects a filter that is not performing correctly and disables it. From this report page you can perform the following actions:

- Modify the filter mode. The filter mode options are:
  - Automatic Mode – Automatically disables the filter and generates a system message regarding the filter.
  - Manual Mode – Generates a system message regarding the filter. Marks the filter as Congested because it is causing device congestion. Does not automatically disable the filter.
- Change the severity level for Adaptive Filter Control log messages.
- Reset the Adaptive Filter Control status.
  - If the filter was disabled in automatic mode, this will re-enable the filter and it will start filtering again.
  - If the filter was disabled in manual mode, this changes the congested state back to false.
- View filter settings
- Download packet capture

Note: Changes to Adaptive Filter Control status take effect immediately.

DDoS

This report shows how often DDoS filters were triggered over a selected time period. This report displays rejected connections over the following view periods:

- Last minute
- Last hour
- Last 24 hours
- Last 35 days

Quarantines

This report provides data on the number of hosts that were quarantined over a selected time period. The following view options are available for this report:

- Packets Blocked
- Total Hosts
- Source Pages
- Redirected pages
- Last minute
- Last hour
- Last 24 hours
- Last 35 days

Top filter matches

The Top Filter Matches report has two sections:

- Top 25 Filters
- Filter Matches

View Top 25 filters

The Top 25 Filter Matches report includes only the IPS filters. It shows the 25 IPS filters with the most hits. The hit counts continue to increment until you reboot the system or click the **Reset Counters** button. The filter numbers are displayed on the y-axis.

Select **Reports > Top Filter matches > Top 25 Filters**, to view the top 25 filter matches.

View filter matches

Select **Reports > Top Filter matches > Filter matches**. The Filter Matches report has three views from which you can select – Severity, Action, and Protocol.

- The Severity report displays the percentage of filter matches that are critical, major, minor, and low severity.
- The Action report displays the percentage of filter matches for different actions (block, permit, rate limit, and trust).
- The Protocol report displays the percentage of filter matches for different protocols (ICMP, UDP, TCP, IPv4 - Other, ARP, Ethernet - Other, ICMPv6, and IPv6 - Other).

Click **Reset Counters** at the top right to set the count back to zero for the report that you are currently monitoring. If you are currently viewing the Filter Matches by Protocol report and click Reset counters, it will affect only the counts for that report. It will not reset the counts for Filter Matches by Action, Filter Matches by Severity, or Top 25 Filters. If you want to reset the counters for all of these reports, reset each of them separately.

Manage the system

The System menu provides access to configuration settings for the device.

This topic contains the following information:

High Availability settings

Use the **System > Settings > High Availability** page to provide High Availability (HA) functionality to minimize network downtime in the event of a device failure. The High Availability feature allows you to view the current HA status and change the configuration as needed. Two types of High Availability mechanisms are available:

- Intrinsic Network HA (INHA) and Zero Power High Availability (ZPHA) for individual device deployment.
- Transparent High Availability (TRHA) for devices deployed in a redundant configuration in which one device takes over for the other in the event of system failure.

Intrinsic Network HA

When the system fails, the device goes into Layer-2 Fallback mode and either permits or blocks all traffic on each segment, depending on the Layer-2 Fallback action setting for the segment. When the device is in Layer-2 Fallback mode, any traffic allowed through the device will not be inspected; it simply passes through the device.

A lack of reported errors or congestion through the TSE does not guarantee that the components receive correct and error-free traffic. The INHA monitors for several points of failure and applies failure detection logic against the system. All components for the INHA are checked for failure.

The device performs the following checks to detect a failed condition and trigger a Layer-2 Fallback:

- Check back-pressure — Presence of back-pressure indicates packets are queued for processing. It indicates a failure if it does not process packets.
- Determine traffic requirements — If the device does not pass traffic, the ability to detect a failure is more difficult. A minimum rate of traffic must pass through the device for best failure detection.
- Handle non-atomic nature of the data path — Packet pass through each component at different times and rates. The status of each component is determined independently of each other. INHA uses sampling to determine health.
- Check and transmit the inbound receive counters — Each component has receive counters incremented by packets received from the previous component. The component transmits these counters incremented as packets to the next component. These counters are the most accurate and most complicated way of detecting health.
- Dropped packets exceeds threshold — If too many packets awaiting deep inspection are queued up, packets are dropped.
- Memory lows — Whether available system memory is too low for proper operations.
- Various chip set errors — Represents possible hardware problems.

Each component also has a specific set of functions for failure checking.

You can view and configure the Layer-2 Fallback behavior for each segment from the Network Segments page (**Network > Segments**). The default setting for each segment is to permit all traffic. This setting is usually preferred by service providers because it prevents a device outage from becoming a network outage. However, for greater security, you might want to change the default Layer-2 Fallback setting to **block all** to guarantee that no uninspected traffic enters the network.

You can view and manually change the current INHA state (normal or Layer-2 Fallback) from the High Availability menu page.

Transparent HA

Transparent High Availability allows the user to install two devices in a redundant network configuration. TRHA keeps the devices in sync with each other so that if one experiences a system failure, the network flow can be routed to the other with no interruption in intrusion prevention services.

Note: Data might not reach the peer machine if "active" machines are under extremely heavy load.

TRHA requires you to configure two devices with the same settings. Before you can configure the TRHA settings on a device, the devices must meet the following network setup and communication requirements:

- Both devices must be from the same product family and running the same TOS version.
- Each device must have a secure connection to the network and to the other device in the TRHA pair. For TPS devices, you must use the HA ports.

- Both devices must be able to communicate on TCP port 9591.
- A device configured with TRHA can only connect and communicate with a partner configured to talk to other devices configured with TRHA. In other words, both machines participating must point to each other. Each device must be configured with the partner serial number and IP address.

Caution: A hijacked device or a rogue device that "steals" the IP address of a TRHA partner can communicate with a legitimate Threat Protection System.

You can configure the TRHA settings and manually enable or disable the feature from the High Availability menu page.

Note: If your system has two devices communicating through TRHA, a change to the global timeout for the connection table at one device will not propagate to the other device. You must make this change on each device accordingly.

Zero-Power HA

Zero-Power High Availability (ZPHA) provides high availability if the device loses power or is powered down. ZPHA can be configured to be in **Normal** or **Bypass** mode. When the current state is **Normal**, traffic does not pass through the device when it is not powered on. When the current state is **Bypass**, traffic passes through an unpowered device as if it was not there.

The following table shows how the traffic would be handled with different states L2FB and ZPHA:

L2FB state	ZPHA state	Traffic status
Normal	Normal	Traffic inspected as per device configuration
L2FB	Normal	Traffic inspected based on segment Layer 2 Fallback action setting
Normal	Bypass	Traffic passed uninspected
L2FB	Bypass	Traffic passed uninspected

Intrinsic Network High Availability

The System > Settings > High Availability page shows the Intrinsic Network High Availability state of the device and allows you to change the current Intrinsic network HA state of the device.

The Intrinsic Network HA panel indicates the Intrinsic Network HA state of your device:

- **Normal** state indicates the device is operating normally and inspecting traffic.
- **Layer-2 Fallback** state indicates that the device is in Layer-2 fallback mode. Traffic on each segment is blocked or permitted according to the Layer-2 fallback action setting (permit or block) for the segment. While in Layer-2 Fallback mode, the IPS does not inspect traffic.

To change the operational Intrinsic Network HA state of device:

1. Click **Change**.
2. Select a network HA option:
    - **Layer-2 Fallback** forces the device to go into Layer-2 Fallback mode. Traffic flowing through each segment on the device is either blocked or permitted based on the Layer-2 fallback Action setting for that segment. Any permitted traffic will not be inspected.
    - **Normal** resumes normal inspection.
3. Click **OK**.

Transparent High Availability

Transparent High Availability (TRHA) is for devices deployed in a redundant configuration in which one device takes over for the other in the event of system failure. For TPS devices, you must use the HA ports.

Use the Transparent HA panel for configuring TRHA:

Setting	Description
Current State	Current TRHA state (Enabled or Disabled).
HA Port Link	Link status for the physical HA ports (Up or Down).
Enabled	Enables Transparent HA. If selected, you must specify the serial number for the other device configured as the HA partner.
Encrypt Traffic	Encrypts traffic between the device and the HA partner. If this option is selected, you must provide a passphrase for the encryption. The passphrase can be no longer than 32 characters and can consist of alphanumeric characters, the hyphen (-), underscore(_), and ampersand (&). **Note:** If the HA port network traffic is physically secure, you do not need to encrypt the traffic, which improves performance.

When an HA partner is configured, the Transparent HA Partner Status panel displays the partner's serial number, model number, and software version.

Zero-Power High Availability

The System > Settings > High Availability page shows the Zero Power High Availability (ZPHA) state of the device and allows you to change the current ZPHA state of the device.

The Zero-Power HA panel indicates the current ZPHA state of the device:

- **Normal** indicates the device is operating normally and inspecting traffic.
- **Bypass** indicates traffic passes through the device without any inspection. When the ZPHA is in Bypass mode, the Layer-2 fallback Action setting for each segment is ignored.

To change the current HA state:

1. Click **Change**.
2. Select a ZPHA option:
    - **Bypass** allows traffic to pass through the device without inspection. If the device is rebooted, the ZPHA operational state returns to Normal.
    - **Normal** resumes normal inspection.
3. Click **OK**.

Configure the management interface

The Management Port is a separate network connection on the TPS that communicates with the device. This allows you to connect the appliance to a dedicated management network, separating the management network from the data networks. However, the management network and the data networks are permitted to overlap with each other. The TPS ships with a default IP address of 192.168.0.1. You can use the **System > Management Port** page to modify the default configuration.

Management interface settings

Use the **System > Management Interface > Settings** page to configure the following options:

- Enable or disable the CLI and Web Interface.
- Specify identification information for the appliance, such a name and location.

Enable the command line and Web interfaces

Enable or disable management access to the TPS through the following:

- Local Security Manager (LSM) — Web-based GUI for managing one device. The LSM provides HTTPS (secure management) access. This access requires access from a supported web browser (Internet Explorer, Mozilla Firefox, and Netscape). Using the LSM, you have a graphical display for reviewing, searching, and modifying settings. The GUI interface also provides reports to monitor the device traffic, triggered filters, and packet statistics.
- Command Line Interface (CLI) — Command line interface for reviewing and modifying settings on the device. The CLI is accessible through SSH (secure access).

1. Select **System > Management Interface**.
2. Check or uncheck the following options:
   - **Enable HTTPS** – Must be enabled to manage the appliance over the network using the LSM.
   - **Enable SSH** – Must be enabled to manage the appliance over the network using the CLI.

Disable TLS versions

Disable older TLS versions to secure the management interface. When deciding which TLS versions to disable, keep in mind that the LSM, SMS, and Captive Portal communicate through the device's management interface.

Tip: If you cannot connect to the LSM after disabling TLS versions, they can be re-enabled using the `{running-gen}tls` CLI command.

1. Select **System > Management Interface**.
2. Check or uncheck the following options:
   - **Enable TLSv1.0**
   - **Enable TLSv1.1**
   - **Enable TLSv1.2**

Modify device details

The Device Details panel provides a method for entering identification information for the appliance, such as a specific name for the appliance, its location in your facility, and a person to contact regarding the appliance.

1. Select **System > Management Interface**.
2. In the Device Details panel, enter a host name, host location, and contact.

Management port settings

The management port lets you connect your device to a dedicated management network. The device separates the management network from the networks connected to the Network Ports.

Use the **System > Management Interface > Management Port** page to configure management port IP addresses and routes.

Change the management port IP address

You can change the IP address of your device to match your corporate IP address standards or security policy for management devices.

1. In the LSM, select **System > Management Interface > Management Port**.
2. Enter an IPv4 or IPv6 address.

   The IP address is used to connect to your TippingPoint device. Must be a valid address on the network segment to which the device is connected, in the form xxx.xxx.xxx.xxx/xx. If the routing prefix size is not specified, the default is 24. TippingPoint recommends configuring the management port on the device to use a non-routed IP address from the RFC 1918 Private Address space. This helps to prevent a direct attack on the management port from the Internet.
3. Click **OK**.

Add management port routes

If you use a separate management network, you might need to configure static IP routes to allow remote network management to the device. The management port uses separate IP routes to those used on the network ports and cannot use dynamic routing.

Routing options enable communication with network subnets other than the subnet on which the management port is located. If your device only communicates with devices on the local network subnet, you do not need to configure a management port route, regardless of whether you are using IPv4 or IPv6.

However, if you are using IPv4 and the device does communicate with devices that are not on the local IPv4 subnet or accessible through the default gateway, you must define the routes to these devices.

Caution: Modifying the management port routes can interrupt the LSM session.

1. Select **System > Management Interface > Management Port**.
2. In the Management Port Routes panel, enter the IP subnet, gateway address and distance.
3. Click **Add** and then **OK**.

Set management port filters

Use the Management Port Filters page to allow or deny specific services from specific IP addresses on the management port. By default, the management port allows management unless you configure management filters.

Note: Modifying the management port filters can interrupt the LSM session. For example, if you deny HTTPS, you can no longer access the LSM through the management port.

1. Select **System > Management Interface > Management Port Filters**.
2. Click **Allow** or **Deny** for the Default Rule.

   The Default Rule allows or denies traffic if there are no Management Port Filters set that apply to the incoming traffic.
3. Select an action, and set whether the action is allowed or denied.
4. Select a service from the list.
5. Enter the IP address that you want to allow or deny access for the service selected in the preceding step.
6. Click **OK**.

Set the date and time

Your device uses the system time in log files. To ensure log file accuracy, facilitate log analysis, and establish predictable scheduling, configure the correct time zone and timekeeping mechanism before using the device in a live environment.

Use the Date/Time page to manage the time setting on the device. You can manually change the current system and time zone or use Network Time Protocol (NTP). If using NTP, you must have access to at least one NTP server.

Set the current time and time zone manually

1. Select **System > Date/Time**.
2. In the Current Device Time panel, click **Change** to change the Current Time. The Change Time dialog is displayed.
3. Click **Set to Browser Time**, or **Specify Time** and enter the time manually, and then click **OK**.
4. Select a region and city in the Time zone panel.

   The default time zone for the device is Greenwich Mean Time. If you change the default, the LSM logs display time data based on the specified time zone. Although the system logs are kept in GMT, the LSM translates these time values into local time values for viewing purposes.
5. Click **OK**.

Synchronize time with NTP

To avoid the man-in-the-middle vulnerability of SNTP servers, users can configure a Network Time Protocol (NTP) server to authenticate NTP messages received from NTP servers and peers. Any attacks of the NTP infrastructure that attempt to inject false time messages must have these messages authenticated (if the **Enabled** option is selected). When authentication is enabled, all time messages are authenticated by the client before they can be accepted as a time source. TippingPoint recommends adding between and four and eight NTP servers.

You can synchronize the device time using the NTP.

1. Select the **System > Date/Time** menu.
2. In the Synchronize Time panel, click **Enabled** to use the NTP server.
3. (Optional) Change the polling period from the default of 32 seconds as necessary.
4. (Optional) Click the **Authentication Keys** button to specify an ID and Key ID on a server.
   a. In the Add Authentication Keys dialog, specify a number between 1 and 65535 for the Key ID on the server.
   b. Specify an Authentication Key value that corresponds to an authentication key on an NTP server.
   c. Click **Add** to add a member for each NTP server.
   d. Click **OK**.
5. Enter the server name or IP address, or click DHCP to get the server IP address.

   At least one NTP Server is required. To add more than one, click the plus (+) icon.
6. Click browse (**...**) to select the authentication key.
7. (Optional) Check **Preferred** to make this the preferred server.
8. Click **OK**.

Configure email

Note: This is an Instant-Commit feature. Changes take effect immediately.

Use the Email page to configure the default email settings to use when sending alerts, notifications, and logs by email. After configuring the email server, you must also configure email contact information from the Notification Contacts page (**Policy > Notification Contacts**). For more information, see *Notification contacts* on page 82.

Configure email settings

1. Select **System > Email**.
2. Enter the email recipient in the **To Email Address** field.
3. Enter the email sender in the **From Email Address** field.

   This address is used as the Reply-To address for messages sent from the device. Consider entering your device name as your company domain, as in `devicename@your_company.com`.
4. Enter your company domain address in the **From Domain Name** field.
5. Enter the IP address of your mail server in the **SMTP Server IP Address** field.

   The device must be able to reach the SMTP server that will be handling the email notifications.

6. Set the **Email Threshold (per minute)** option using the arrow keys.

   By default, the system allows 10 email alerts per minute. On the first email alert, a one-minute timer starts. The systems sends email notifications until it reaches the threshold and then blocks subsequent alerts. After one minute, the system resumes sending email alerts. The system generates a message in the System log whenever email notifications are blocked.

Manage data security

The **System > Data Security** menu provides access to the data security page. You can configure and manage the following items:

- System master key.
- Encrypt the removable disk that stores logs.

Set the master key

The system master key encrypts the keystore. If disk encryption is enabled, it also encrypts all data on the removable user disk. If you change the master key while encryption is enabled on the disk, all traffic logs, snapshots, and packet capture data on the disk will be lost.

1. Select **System > Data Security**.
2. Enter the master key in the Master Key box. The master key must meet the following criteria:
   - Between 9 and 32 characters in length
   - Combination of uppercase and lowercase alpha and numeric characters
   - Must contain at least one special character, such as ! @ # $ %
3. (Optional) Check **Show Password** to see the master key.
4. Click **Apply**.

Enable user disk encryption

1. Select **System > Data Security**.
2. On the System page, check **Enable User Disk Encryption**.

   The master key encrypts all data on the removable disk and the system keystore, which stores system keys and certificates.

   Note: Changing the master key while encryption is enabled erases all traffic log, snapshot, and packet capture data on the removable disk.

Configure logs

The **System > Log Configuration** menu provides access to the Log Configuration page.

The logs provide information on system events and traffic-related events triggered by the filters that are configured on the device. Each menu page also provides functions to manage the log files. Logs also indicate the interfaces through which administrators interacted with the system, such as the LSM, the CLI, or an SMS.

You can configure and manage the following items for logs:

- Associate Notification Contacts to System, Audit, and Quarantine logs.
- Manage the alerting threshold for the Alert and Block logs to improve device performance.
- Clear all entries from a log and download logs.

Manage notification contacts

Use the Notification Contacts page to configure notification contacts and thresholds for the following logs:

- System
- Audit
- Quarantine
- SSL Inspection

You can manage the notifications for other logs from the **Policy > Notification Contacts** page.

By default, all notifications are sent to the Management Console. However, you can change this setting for the System and Quarantine logs by editing the default configuration and selecting a different Severity Threshold. The Threshold Severity level cannot be changed for the Audit log.

To edit the default notification contact configuration for the logs:

1. Select **System > Log Configuration**.
2. Click **Notification Contacts**.
3. Click **Edit** for the log you want to modify.
4. In the Edit Log Notification Contacts dialog, select a severity from the Severity Threshold list.

   Note: This can be configured only for System and Quarantine logs.

none	Notifications are not sent under any condition.

Debug	A debug condition occurred.
Info	Informational message.
Notice	Normal, but significant conditions exist.
Warning	A warning condition occurred.
Error	An error occurred.
Critical	A critical condition exists.
Alert	Action must be taken immediately.
Emergency	System is unstable.

5. Click **OK**.

Protect device performance

When traffic congestion on the device significantly impacts performance, use the **System > Log Configuration > Performance Protection** page to temporarily disable logs for Alert and Block events.

The default configuration is to disable Alert and Block events for 600 seconds (10 minutes) when device congestion renders a packet loss value of 1 percent.

You can disable Performance Protection by selecting **Always log Alert and Block events**.

To enable Performance Protection:

1. Select **Disable logging Alert and Block events for a period of time if the device is congested**.
2. Adjust the performance threshold by entering a packet loss percentage value between 0.1 and 99.9 percent.
3. Adjust the period that logging is disabled by entering a value between 60 seconds (1 minute) and 3600 seconds (1 hour).
4. Click **OK** to save your changes.

View and download a log

Use the **System > Log Configuration > Summary** page to view and download logs, get size and location information for each log, and clear log entries.

1. Select **System > Log Configuration**.
2. Click **Summary**.
3. Select a log and click **Download**.
4. In the Download Systems Log dialog, specify **All**, a **Time Range**, or an **ID Range**.

   For **ID Range**, enter the starting and ending entry numbers (line numbers) of the log.
5. (Optional) Change the format from text to a comma-separated value.
6. Check **Open in browser** and then **OK** to view the log in a new browser tab, or click **OK** to download the log to your default Downloads directory.

Configure SMS

The Security Management System (SMS) enables you to remotely monitor and manage your device. When the device is under SMS control, you can use the LSM to do the following:

- View, manage, and edit device configuration.
- Review logs and reports.
- Configure security policy.

When an appliance is under SMS management, the message `DEVICE UNDER SMS CONTROL` is displayed in red at the top of each page in the LSM and the SMS page displays the serial number and the IP address of the controlling SMS. In this state, you can view system configuration and status.

Note: Changes to SMS system settings take effect immediately.

To configure SMS management:

1. Select **System > SMS**.

   The default value is **Any IP Address**, which means that any SMS can manage the appliance.
2. To enter the IP address of a specific SMS, click **Specific IP Address / CIDR** and enter the IP or CIDR address in the box.

   To specify a range of IP addresses, enter an IP address block (10.100.230.0/24, for example). This allows any SMS on the specified IP subnet to manage the appliance.
3. Click **Manage**.
4. To release a device from SMS management, click **Unmanage**.

Configure SNMP

The **System > SNMP** menu provides access to the SNMPv2c and SNMPv3 configuration pages.

Enable SNMP

After you enable SNMP and commit the change, SNMP creates an Engine ID using the Management Port's MAC address. The engine ID uniquely defines an SNMP node (or engine) and associates it to a user.

1. Select **System > SNMP**.
2. On the General page, check **Enable SNMP**.
3. Click **OK.**

   After you commit the change, it might take a couple of seconds to start the SNMP demon. In the unlikely case of a collision with another device, you can change the Engine ID to a different value; however, the new value must be unique. Note that changing the Engine ID regenerates each read-only user, which will affect connectivity.

Add or edit an SNMPv2c community

You can create multiple communities to support NMS, IPs, or subnets. Each community can have multiple rules, although the source IP address must be different. For example, you can create a rule for a Community named Public with a Source IP Address of 1.1.1.1. You can have a second rule for Public with a Source IP Address of 2.2.2.2.

1. Select **System > SNMP > SNMPv2c**.
2. In the Communities table, click **Add** to create a new community or **Edit** to change an existing one.
3. In the Add SNMPv2c Community dialog, enter a community name.
4. Select one of the following options:
   - **Any Source IP Address**
   - **Specific IP Address/DNS Name / CIDR**
5. Click **OK** or **OK/Continue** to add another community.

Add or edit an SNMPv2c trap destination

1. Select **System > SNMP > SNMPv2c**.
2. In the Trap Destinations table, click **Add** to create a trap destination or **Edit** to change an existing one.
3. In the Add SNMPv2c Trap Destination dialog, enter a community name.
4. Enter an IP address or DNS name.

5. (Optional) Enter a port. By default, the SNMP manager receives requests on port 162.
6. (Optional) To send a notification acknowledgement to the SNMP manager, check **Send as Inform Request**.
7. Click **OK** or **OK/Continue** to add another trap destination.

   The new rule is displayed in the Trap Destinations table.

Add or edit an SNMPv3 user

1. Select **System > SNMP > SNMPv3**.
2. In the Users table, click **Add** to create a new user or **Edit** to change an existing one.
3. In the Add SNMPv3 Read-Only User dialog, enter a username.
4. (Optional) To configure authentication, select **MD5** or **SHA** authentication type.
   a. Enter a passphrase at least 8 characters long.
   b. Select **show passphrase** to view the phrase.
5. (Optional) To configure privacy (encryption), select **DES** or **AES**.
   a. Enter a passphrase at least 8 characters long.
   b. Select **show passphrase** to view the phrase.
6. Click **OK** or **OK/Continue** to add another user.

Add an SNMPv3 trap destination

1. Select **System > SNMP > SNMPv3**.
2. In the Trap Destinations table, click **Add** to create a trap destination or **Edit** to change an existing one.
3. In the Add SNMPv3 Trap Destination dialog, enter an IP address or DNS name.
4. (Optional) Enter a port. By default, the SNMP manager receives requests on port 162.
5. Enter a username to specify the assigned user for this trap.
6. (Optional) To send a notification acknowledgement to the SNMP manager, check **Send as Inform Request**.
7. (Optional) To configure authentication, select **MD5** or **SHA** authentication type.
   a. Enter a passphrase at least 8 characters long.
   b. Select **show passphrase** to view the phrase.
8. (Optional) To configure privacy (encryption), select **DES** or **AES**.
   a. Enter a passphrase at least 8 characters long.
   b. Select **show passphrase** to view the phrase.
9. Click **OK** or **OK/Continue** to add another trap destination.

Update the device

Note: This is an Instant-Commit feature. Changes take effect immediately.

The **System > System, DV, Licenses** menu enables you to perform the following tasks:

- Install a new software version
- Roll back to a previous software version
- Install and remove Digital Vaccine (DV) packages
- Install license packages

Upgrade the software to a newer version

Upgrade the software to install a newer TOS version with the latest improvements or additions onto the device. TippingPoint Technical Support releases software updates on the Threat Management Center (TMC). You can download and install updates from this site. Installing a new software package forces a reboot of the device.

Note: You cannot upgrade the software to an earlier TOS version. Instead, use the rollback feature to return to a previously installed image.

1. Log in to the TMC at:

   https://tmc.tippingpoint.com/TMC/

2. After you log in, select **Releases > Software > TPS**.
3. Download the latest release to a thumb drive or your local system.
4. When the download completes, log out of the TMC.
5. In the LSM menu bar, select **System > Update > System, DV, Licenses**.
6. In the Software Versions panel, click **Install**.

   The Install System Software dialog is displayed.

7. Click the **Browse** button to select the package you downloaded from the TMC.
8. Select the package and click **Install**.

   The package installs and the system is rebooted. After the system reboots, the login page is displayed.

   After upgrading the software, create a snapshot to save the configuration. For more information, see *Snapshots* on page 122.

   Note: When you use the rollback feature to return to a previously installed image, or when you restore a snapshot, you must set the master key to the value that was used when that image was in operation or

when the given snapshot was created. Failure to have the correct master key on the system results in the system keystore being inaccessible. This impacts access to x509 private keys.

Roll back to a previous version

A rollback operation reverts the currently running software on your device to a previous working version that you select.

When you perform a TOS rollback, current configuration settings are preserved, but filter settings revert to the settings that were in effect when the rollback version was archived. Any filter changes made after the target rollback version are deactivated, including attack protection filter updates.

Note: If you perform a rollback, read the release notes for both the software version you are rolling back from and the software version you are rolling back to. The release notes contain information that addresses migration issues between versions.

1. From the menu bar, select **System > Update > System, DV, Licenses**.
2. In the Software Version panel, select the version you want to roll back to, and click **Rollback To**.

   The Software Rollback dialog is displayed warning you that any configuration changes made since this version was last run will be lost.
3. Click **OK** to start the rollback operation.

Digital Vaccine packages

When TippingPoint Technical Support discovers new types of network attacks, or when detection methods for existing threats improve, the Digital Vaccine team at the Threat Management Center (TMC) creates and releases new filters to add to your filter database. These filters are released as Digital Vaccine (DV) packages.

Note: When you download and install a DV package, verify that the package you download is not larger than the listed amount of free space. An unpacked package might require more space than anticipated, depending on saved snapshots and rollback versions and the size of the available update. To make sure the appliance has enough disk space, you can delete previously installed software images from the Update page.

When a new DV package is available for download, the TMC team sends notifications to existing customers. You have two options to update the DV on your appliance:

- Configure the Auto DV option on your appliance so that the appliance checks for new DV packages and automatically updates the appliance as necessary.
- Manually download and install the DV package.

Install a Digital Vaccine

1. Log in to the TMC at:

https://tmc.tippingpoint.com/TMC/

2. After you log in, select **Releases > Digital Vaccine> Digital Vaccine**.
3. Download the latest release to a thumb drive or your local system.
4. When the download completes, log out of the TMC.
5. From the LSM menu bar, select **System > Update > System, DV, Licenses**.

   The System Software, Digital Vaccine and Licenses page is displayed.
6. In the Digital Vaccine Packages table, click **Install**.

   The Install Digital Vaccine dialog is displayed.
7. Click the **Browse** button to select the package you downloaded from the TMC.
8. Click **Install**.

   Note: For DV packages, you cannot rollback to a previous version. To use a previous version, download that version from the TMC.

Enable automatic DV updates

1. Select **System > Update > Software, DV, Licenses**.
2. Click **Auto DV**.
3. Check **Enabled** to enable auto configuration.
4. Select the type of schedule that you want to use for the DV update process.
   - **Periodic** – Performs an update every number of days starting from a set day. The option includes a time to perform the update.
   - **Calendar** – Performs an update on a set day and time per week.
5. If you are using a proxy, specify the proxy parameters and a username and password.
6. Click **Commit**.

   When Auto DV is configured, the system automatically checks the DV version when you open the Auto DV Update page. The status is listed on the right side of the page. To perform an update immediately, click **Update Now**.

License packages

If your device is managed by a Security Management System (SMS), licenses are automatically updated. Otherwise, you can access new licenses by logging in to your account on the TMC. For more information see *Install a license package* on page 122.

The License Version panel displays the following information:

Feature	Name of the licensed feature or service: **License** – The default license for TippingPoint Support devices. **Update TOS** – Enables you to update the TOS on the appliance. **Update DV** – Enables you to update to the latest Digital Vaccine. **Auxiliary DV** – Optional license that enables you to update additional Digital Vaccine features. **Reputation DV** – Optional license for updating Reputation Filters. **SSL Inspection** - Optional license to permit SSL inspection (requires reboot). **Throughput Upgrade** - Optional license to upgrade the inspection throughput of the device (requires reboot).
Status	Current status of the license.
Permit	Whether the feature is currently enabled.
Expiration	License expiration date.
Details	Any additional information about the license.

Update the license package

Use the TMC to update the license package and assign a license upgrade to a TPS. License upgrades enable you to increase the inspection throughput of a 440T or 2200T, or enable SSL inspection on a 2200T.

If you purchased an SSL Inspection Upgrade license when you ordered the device, you must also update the license package and assign the license to the device. You are not required to assign the license that you purchased with the device to that device. However, an SSL Inspection license is only applicable to a 2200T device.

SSL Inspection and Throughput Upgrades are licensed separately. To request a license, contact your sales representative.

1. Log in to the TMC at: *https://tmc.tippingpoint.com/TMC/*

2. From the TMC, select **My Account > License Manager**.
3. Assign the license updates to the device.
4. Download the updated license package to a thumb drive or your local system.
5. When the download completes, log out of the TMC.

Install a license package

1. Log in to the TMC at:

   https://tmc.tippingpoint.com/TMC/

2. After you log in, select **My Account > TippingPoint License Package**.
3. Download the necessary license to a thumb drive or your local system.
4. When the download completes, log out of the TMC.
5. From the LSM menu bar, select **System > Update > System, DV, Licenses**.
6. In the License Version panel, click **Install**.
7. In the Install License Package dialog, click **Browse** to select the package you downloaded from the TMC.
8. Click **Install**.

Snapshots

The **System > Update > Snapshots** menu enables you to perform the following tasks:

- Create a snapshot
- Import a local snapshot
- Manage current snapshots

A *snapshot* enables you to restore a device to a previously known working state. Restore a snapshot to the same device or to a different device. You can also export a snapshot and send it to TippingPoint Technical Support for assistance with troubleshooting or debugging the device. All snapshots are stored on the external user-disk, also known as the CFast card.

When restoring a snapshot to a different device, keep in mind:

- Both devices must be the same model, such as the TippingPoint 2200T, and have the same TOS version.
- Before you restore a snapshot, you must set the master key to the value that was used when that image was in operation or when the given snapshot was created. Failure to have the correct master key on the system results in the system keystore being inaccessible. This impacts access to x509 private keys. For example, if private keys are persisted on the original device to perform SSL inspection, the same system master key is required on the new device to access the restored key store. If you set the master key after you restore the snapshot, you must reboot the device again to apply the updated master key.

- The snapshot includes the license package. The license package provides license information for each of your TippingPoint devices. If the license package that was included in the snapshot is outdated, restore the snapshot and then download and install an updated license package from the TMC.
- The SSL inspection license, if it was enabled when the snapshot was taken, is restored with the snapshot.

  Be sure to update the license package on the TMC to assign the SSL inspection license to the correct device. After you install an updated license package and reboot the device, for example, to update DV filters, the license package determines whether to enable or disable SSL inspection.
- If an external ZPHA was configured on the original device, be sure to add an external ZPHA to the target device or update the device configuration to remove ZPHA.

The Current Snapshots table contains the following information:

Name	A user-specified descriptive name for the snapshot.
Date	The date the snapshot was generated.
Software Build	The software version that was running when the snapshot was created.
Digital Vaccine	The version number of the Digital Vaccine package running when the snapshot was created.
Model Type	The model name of the device where the snapshot was created.

Create a snapshot

Note: This is an Instant-Commit feature. Changes take effect immediately.

This procedure describes how to take a snapshot.

Note: The snapshot includes the stored Start configuration only. It does not include the in-memory running configuration. To include this information, select **Configuration > Commit pending changes and Copy to Start** before you take the snapshot. As a best practice, create a snapshot after upgrading the appliance software.

1. Verify that the external user disk (CFast card) has been installed and properly mounted.
2. From the menu bar, select **System > Update > Snapshots**.
3. In the Create Snapshot panel, enter a descriptive name in the Snapshot Name field.

4. (Optional) Click the appropriate checkboxes:
   - **Include DV Reputation Database** – Includes your custom IP and DNS reputation entries.
   - **Include Manual Reputation Database** – Includes the Threat DV package.
   - **Include Management Port, Cluster and HA Configuration** – Includes your configuration settings for the Management Port and HA. For more information about Management Port Settings, see *Configure the management interface* on page 107.
5. Click **Create**.

   The system starts the snapshot creation process. After the snapshot is created, it is displayed in the Current Snapshots table and on the CFast card.
6. In **Current Snapshots**, select **Export** to export the snapshot from the CFast card to another drive.

Restore a snapshot

Restoring a snapshot forces a system reboot. You can only restore a snapshot when running the same software version that was used to create the snapshot. If you have upgraded the software, you must install the previous software version before restoring that snapshot. As a best practice, review the licensing and configuration after restoring the snapshot to a different device, and take a snapshot after upgrading your software. You lose any configuration changes made to the current configuration when you restore a snapshot.

1. From the menu bar, select **System > Update > Snapshots**.

   The Snapshots page is displayed.
2. In the Current Snapshots panel, click the checkbox next to the snapshot you want to restore.
3. Click **Restore**.

   The system loads the snapshot and restores the device to the configuration specified in the snapshot. After the snapshot is loaded, the device reboots and returns you to the login page.

Shut down the device

Use the `shutdown` command to shutdown the TPS. On a 440T TPS, when you remove power from the device or issue the `shutdown` command, you must wait 1 minute before attempting to restore power to the device. This issue is not applicable to the 2200T TPS.

Tools

The Tools menu provides quick and convenient access to common network utilities. It also provides a Tech Support Report feature that creates a report that can be sent to TippingPoint Technical Support for troubleshooting assistance.

This topic discusses the following information to describe how to use the Tools menu:

Issue a ping

Use the Ping utility to find out if a specific host or device is accessible. This utility supports both IPv4 and IPv6.

1. Select **Tools > Ping**.

   The Ping page is displayed.

2. Click **IPv4** or **IPv6** to specify the Internet Protocol version.
3. Enter an IP address or the name of the system you want to ping.
4. (Optional) Enter the IP address of the system sending the ping.
5. (Optional) Change the default settings for repetitions, data size, and time to live.

Repetition	The number of times the ping is attempted. Set between 1 and 20. Default: 4
Data Size	The size of the packet being sent. Set between 1 and 65468. Default: 56
Time to Live	The length of time the packet can be used before it is discarded. Set between 1 and 800.

	Default: 255

6. Click **Start Ping**.

   Results of the ping are displayed in the Results panel.

7. Click **Stop Ping** if you want to stop the ping.

Issue a trace route

Use the Trace Route utility to display the path and transit information for packets being sent across an IP network.

1. Select **Tools > Trace Route**.

   The Trace Route page is displayed.

2. Click **IPv4** or **IPv6** to specify the Internet Protocol version.
3. Enter an IP address or the name of the system you want to trace.
4. (Optional) Enter the local IP address as the source.
5. Click **Start Trace**.

   Results of the trace are displayed in the Results panel.

Tech Support Report

The Tech Support Report collects diagnostic information into a report that TippingPoint Technical Support can use to debug and troubleshoot system issues. It includes diagnostic commands, log files, and optionally a full system snapshot. The Tech Support Report snapshot captures the system's current running configuration.

If you include a snapshot with your Tech Support Report, the snapshot does not contain the following sensitive information:

- User names and passwords
- LDAP and RADIUS server passwords
- SNMPv3 passphrase
- HA passphrase
- Keystore

After the report is created, you can export it to your local system. You can then email the file to TippingPoint Technical Support for assistance.

Do not attempt to restore a Tech Support Report snapshot to your device. All sensitive information including user names and passwords are removed and you will be unable to log in. If you attempt to restore a Tech Support Report and are unable to log in, then phone TippingPoint Technical Support.

To create a snapshot you can use at a later time, use **System > Snapshots**.

Note: Only one report can exist on the device. When you create a new report, the previous report is replaced.

Create a Tech Support Report

1. Select **Tools > Tech Support Report**.
2. (Optional) In the Tech Support Report Options panel, check **Include Snapshot**.
3. (Optional) By default, all IPS and Reputation logs are included in the report. TippingPoint Technical Support uses these logs to troubleshoot issues. However, if you do not want to include the logs, uncheck **Include IPS and Reputation logs**.
4. Click **Create Report**.

   The Tech Support Report is created.
5. Click **Export** to download a tar.zip file of the report to your local Downloads directory.
6. Send the report file in an email to TippingPoint Technical Support.

Traffic capture

The Traffic Capture page provides a listing of captured traffic that you can download for inspection. To capture traffic, you must enable the Packet Trace option when you create an Action Set. For more information on creating an Action Set, see *Action sets* on page 79. The Cfast removable disk stores all captured traffic files.

You can also use the `actionsets` context of the CLI to create an action set to capture traffic. You can use the CLI commands `tcpdump` with the `record` option to create an on-demand packet capture dump. For example `ips{}tcpdump ethernet1 record eth1_capture maxsize 4`. You must specify either a `count` (max number of packets to capture) or a `maxsize` (maximum packet capture file size in millions of bytes) to prevent accidentally filling up the Cfast removable disk.

View captured traffic

1. Select **Tools > Traffic Capture**.

   The Traffic Capture page is displayed. The Traffic Capture table contains the following columns:

Date	The start date of traffic capture.

Time	The start time of traffic capture.
Filename	The file name of the traffic capture.
Type	The type of traffic captured.
Bytes	The size of the file in bytes.

2. Click **Refresh** to refresh the table.

Download captured traffic

1. Select **Tools > Traffic Capture**.

The Traffic Capture page is displayed.

1. Click the checkbox next to the file you want to download.
2. Click **Download**.

The file is downloaded to your local system.

When you are done, you can delete the captured traffic.

Delete captured traffic

1. Select **Tools > Traffic Capture**.

   The Traffic Capture page is displayed.

2. Click the checkbox next to the file you want to delete.
3. Click **Delete**.

TREND MICRO INCORPORATED
225 E. John Carpenter Freeway, Suite 1500
Irving, Texas 75062 U.S.A.
Phone: +1 (817) 569-8900, Toll-free: (888) 762-8736
Email: support@trendmicro.com

www.trendmicro.com
Item Code: APEM47348/160315