



Managing IPS Anomaly Detection



Note

From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see [EOL notice](#).

Anomaly detection is designed to recognize network congestion caused by worm traffic that exhibits scanning behavior. Anomaly detection also will identify infected hosts on the network that are scanning for other vulnerable hosts.

Anomaly detection is enabled by default, but there are some configuration settings you should adjust to use it effectively.



Note

The sensor must use IPS software version 6.x or later to configure anomaly detection. In addition, Cisco IOS IPS and the AIP-SSC-5 do not support anomaly detection.

This chapter contains the following topics:

- [Understanding Anomaly Detection, page 41-1](#)
- [Configuring Anomaly Detection, page 41-6](#)

Understanding Anomaly Detection

The anomaly detection component of the sensor detects worm-infected hosts. This enables the sensor to be less dependent on signature updates for protection against worms and scanners, such as Code Red and SQL Slammer and so forth. The anomaly detection component lets the sensor learn normal activity and send alerts or take dynamic response actions for behavior that deviates from what it has learned as normal behavior.



Note

Anomaly detection does not detect email-based worms, such as Nimda.

Anomaly detection detects the following two situations:

- When the network starts on the path of becoming congested by worm traffic.
- When a single worm-infected source enters the network and starts scanning for other vulnerable hosts.

The following topics explain anomaly detection in more detail:

- [Worm Viruses, page 41-2](#)
- [Anomaly Detection Modes, page 41-2](#)
- [Anomaly Detection Zones, page 41-3](#)
- [Knowing When to Turn Off Anomaly Detection, page 41-4](#)
- [Configuring Anomaly Detection Signatures, page 41-4](#)
- [Configuring Anomaly Detection, page 41-6](#)

Worm Viruses

Worm viruses are automated, self-propagating, intrusion agents that make copies of themselves and then facilitate their spread. Worm viruses attack a vulnerable host, infect it, and then use it as a base to attack other vulnerable hosts. They search for other hosts by using a form of network inspection, typically a scan, and then propagate to the next target. A scanning worm virus locates vulnerable hosts by generating a list of IP addresses to probe, and then contacts the hosts. Code Red worm, Sasser worm, Blaster worm, and the Slammer worm are examples of worms that spread in this manner.

Anomaly detection identifies worm-infected hosts by their behavior as a scanner. To spread, a worm virus must find new hosts. It finds them by scanning the Internet using TCP, UDP, and other protocols to generate unsuccessful attempts to access different destination IP addresses. A scanner is defined as a source IP address that generates events on the same destination port (in TCP and UDP) for too many destination IP addresses.

The events that are important for TCP are non-established connections, such as a SYN packet that does not have its SYN-ACK response for a given amount of time. A worm-infected host that scans using TCP generates non-established connections on the same destination port for an anomalous number of IP addresses.

The events that are important for UDP are unidirectional connections, such as a UDP connection where all packets are going in only one direction. A worm-infected host that scans using UDP generates UDP packets but does not receive UDP packets on the same IP address within a time-out period on the same destination port for multiple destination IP addresses.

The events that are important for other protocols, such as ICMP (protocol number 1), are from a source IP address to many different destination IP addresses, that is, packets that are received in only one direction.



Caution

If a worm virus has a list of IP addresses it should infect and does not have to use scanning to spread itself (for example, it uses passive mapping—listening to the network as opposed to active scanning), it will not be detected by anomaly detection worm policies. Worm viruses that receive a mailing list from probing files within the infected host and email this list will not be detected, because no Layer 3 or Layer 4 anomaly is generated.

Anomaly Detection Modes

Anomaly detection initially conducts a “peacetime” learning process when the most normal state of the network is reflected. Anomaly detection then derives a set of policy thresholds that best fit the normal network. This is done in two phases: an initial learning mode phase, followed by the ongoing operational detect mode phase.

Anomaly detection has the following modes:

- Learning accept mode (initial setup)

Although anomaly detection is in detect mode by default, it conducts an initial learning accept mode for the default period of 24 hours. We assume that during this phase no attack is being carried out. Anomaly detection creates an initial baseline, known as a knowledge base, of the network traffic. The default interval value for periodic schedules is 24 hours and the default action is rotate, meaning that a new knowledge base is saved and loaded, and then replaces the initial knowledge base after 24 hours.

Keep the following in mind:

- Anomaly detection does not detect attacks when working with the initial knowledge base, which is empty. After the default of 24 hours, a knowledge base is saved and loaded and now anomaly detection also detects attacks.
- Depending on your network complexity, you may want to have anomaly detection in learning accept mode for longer than the default 24 hours. You configure the mode in the Virtual Sensors policy; see [Defining A Virtual Sensor, page 38-5](#). After your learning period has finished, edit the virtual sensor and change the mode to Detect.

- Detect mode

For ongoing operation, the sensor should remain in detect mode. This is for 24 hours a day, 7 days a week. Once a knowledge base is created and replaces the initial knowledge base, anomaly detection detects attacks based on it. It looks at the network traffic flows that violate thresholds in the knowledge base and sends alerts. As anomaly detection looks for anomalies, it also records gradual changes to the knowledge base that do not violate the thresholds and thus creates a new knowledge base. The new knowledge base is periodically saved and takes the place of the old one thus maintaining an up-to-date knowledge base.

- Inactive mode

You can turn anomaly detection off by putting it in inactive mode. Under certain circumstances, anomaly detection should be in inactive mode, for example, if the sensor is running in an asymmetric environment. Because anomaly detection assumes it gets traffic from both directions, if the sensor is configured to see only one direction of traffic, anomaly detection identifies all traffic as having incomplete connections, that is, as scanners, and sends alerts for all traffic flows.

The following example summarizes the default anomaly detection configuration. If you add a virtual sensor at 11:00 pm and do not change the default anomaly detection configuration, anomaly detection begins working with the initial knowledge base and only performs learning. Although it is in detect mode, it cannot detect attacks until it has gathered information for 24 hours and replaced the initial knowledge base. At the first start time (10:00 am by default), and the first interval (24 hours by default), the learning results are saved to a new knowledge base and this knowledge base is loaded and replaces the initial knowledge base. Because the anomaly detection is in detect mode by default, now that anomaly detection has a new knowledge base, the anomaly detection begins to detect attacks.

Anomaly Detection Zones

By subdividing the network into zones, you can achieve a lower false negative rate. A zone is a set of destination IP addresses. There are three zones, each with its own thresholds: internal, illegal, and external.

The external zone is the default zone with the default Internet range of 0.0.0.0-255.255.255.255. By default, the internal and illegal zones contain no IP addresses. Packets that do not match the set of IP addresses in the internal or illegal zone are handled by the external zone.

We recommend that you configure the internal zone with the IP address range of your internal network. If you configure it in this way, the internal zone is all the traffic that comes to your IP address range, and the external zone is all the traffic that goes to the Internet.

You can configure the illegal zone with IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses or part of your internal IP address range that is unoccupied. An illegal zone can be very helpful for accurate detection, because we do not expect any legal traffic to reach this zone. This allows very low thresholds, which in turn can lead to very quick worm virus detection.

Knowing When to Turn Off Anomaly Detection

Anomaly detection assumes that it gets traffic from both directions. If the sensor is configured to see only one direction of traffic, you should turn off anomaly detection. Otherwise, when anomaly detection is running in an asymmetric environment, it identifies all traffic as having incomplete connections, that is, as scanners, and sends alerts for all traffic flows.

You turn off anomaly detection in the Virtual Sensors policy. Edit the virtual sensor for which you are disabling anomaly detection, and change the Anomaly Detection Mode to Inactive. For information on editing virtual sensors, see [Editing Policies for a Virtual Sensor, page 38-9](#).

Configuring Anomaly Detection Signatures

The Traffic Anomaly engine contains nine anomaly detection signatures covering three protocols (TCP, UDP, and other). Each signature has two subsignatures, one for the scanner and the other for the worm-infected host (or a scanner under worm attack). When anomaly detection discovers an anomaly, it triggers an alert for these signatures. All anomaly detection signatures are enabled by default and the alert severity for each one is set to high.

When a scanner is detected but no histogram anomaly occurred, the scanner signature fires for that attacker (scanner) IP address. If the histogram signature is triggered, the attacker addresses that are doing the scanning each trigger the worm signature (instead of the scanner signature). The alert details state which threshold is being used for the worm detection now that the histogram has been triggered. From that point on, all scanners are detected as worm-infected hosts.

The following anomaly detection event actions are possible:

- Produce alert—Writes the event to the Event Store.
- Deny attacker inline—(Inline only) Does not transmit this packet and future packets originating from the attacker address for a specified period of time.
- Log attacker packets—Starts IP logging for packets that contain the attacker address.
- Deny attacker service pair inline—Blocks the source IP address and the destination port.
- Request SNMP trap—Sends a trap notification to an SNMP trap destination. To use this action, you must configure SNMP trap hosts as described in [Configuring SNMP, page 36-8](#).
- Request block host—Sends a request to ARC to block this host (the attacker). To use this action, you must configure blocking devices as described in [Configuring IPS Blocking and Rate Limiting, page 43-7](#).

You can add actions to the signatures either directly, in the Signatures policy, or to events generated by the signatures based on risk rating in the Event Actions Overrides policy.

The following table lists the anomaly detection worm signatures.

Table 41-1 Anomaly Detection Worm Signatures

Signature ID	Subsignature ID	Name	Description
13000	0	Internal TCP Scanner	Identified a single scanner over a TCP protocol in the internal zone.
13000	1	Internal TCP Scanner	Identified a worm attack over a TCP protocol in the internal zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13001	0	Internal UDP Scanner	Identified a single scanner over a UDP protocol in the internal zone.
13001	1	Internal UDP Scanner	Identified a worm attack over a UDP protocol in the internal zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13002	0	Internal Other Scanner	Identified a single scanner over an Other protocol in the internal zone.
13002	1	Internal Other Scanner	Identified a worm attack over an Other protocol in the internal zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.
13003	0	External TCP Scanner	Identified a single scanner over a TCP protocol in the external zone.
13003	1	External TCP Scanner	Identified a worm attack over a TCP protocol in the external zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13004	0	External UDP Scanner	Identified a single scanner over a UDP protocol in the external zone.
13004	1	External UDP Scanner	Identified a worm attack over a UDP protocol in the external zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13005	0	External Other Scanner	Identified a single scanner over an Other protocol in the external zone.
13005	1	External Other Scanner	Identified a worm attack over an Other protocol in the external zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.
13006	0	Illegal TCP Scanner	Identified a single scanner over a TCP protocol in the illegal zone.

Table 41-1 Anomaly Detection Worm Signatures (continued)

Signature ID	Subsignature ID	Name	Description
13006	1	Illegal TCP Scanner	Identified a worm attack over a TCP protocol in the illegal zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13007	0	Illegal UDP Scanner	Identified a single scanner over a UDP protocol in the illegal zone.
13007	1	Illegal UDP Scanner	Identified a worm attack over a UDP protocol in the illegal zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13008	0	Illegal Other Scanner	Identified a single scanner over an Other protocol in the illegal zone.
13008	1	Illegal Other Scanner	Identified a worm attack over an Other protocol in the illegal zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.

Configuring Anomaly Detection

Use the Anomaly Detection policy to configure anomaly detection settings. The Virtual Sensors policy also contains a setting important for anomaly detection.

This procedure explains the overall configuration of anomaly detection. Before configuring these settings, read the following topics:

- [Understanding Anomaly Detection, page 41-1](#)
- [Worm Viruses, page 41-2](#)
- [Anomaly Detection Modes, page 41-2](#)
- [Anomaly Detection Zones, page 41-3](#)
- [Knowing When to Turn Off Anomaly Detection, page 41-4](#)
- [Configuring Anomaly Detection Signatures, page 41-4](#)

Step 1 Do one of the following to open the Anomaly Detection policy you want to modify:

- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
- (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.

The Anomaly detection policy includes these tabs:

- **Operation Settings**—Defines the worm timeout and identifies any IP addresses that should be ignored by anomaly detection.

- **Learning Accept Mode**—The configuration for learning mode, including how the knowledge base is handled.
- **Internal Zone, Illegal Zone, External Zone**—The zones of your network that you define. You can configure unique settings for each zone. For an explanation of the zones, see [Anomaly Detection Zones, page 41-3](#).

Step 2 Click the **Operation Settings** tab, if necessary, and configure the following:

- **Worm Timeout**—The time in seconds for the worm termination timeout. The range is 120 to 10,000,000 seconds. The default is 600 seconds. For an explanation of how this timeout is used, see [Understanding Anomaly Detection Thresholds and Histograms, page 41-9](#).
- **Enable Ignored Addresses and Source/Destination Addresses to Ignore**—Whether you are configuring a list of addresses that should be ignored while anomaly detection is processing. You can specify a list of source addresses (those that initiate a scan) or destination addresses (the hosts that are scanned).

The addresses can be single host (such as 10.100.10.1), a range of addresses (such as 10.100.10.0-10.100.10.255), or network/host objects that contain single hosts, address ranges, or a combination of hosts and ranges. Click **Select** to select objects from a list or to create new objects.

Step 3 Click the **Learning Accept Mode** tab and define how the knowledge base will be generated and used. For detailed information, see [Configuring Anomaly Detection Learning Accept Mode, page 41-8](#).

Step 4 Configure the internal, illegal, and external zones:

- Define the internal and illegal zones—The internal zones are the IP addresses of your internal network, the network that you manage. The illegal zone should represent IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses or part of your internal IP address range that is unoccupied.

Click the **Internal Zone** and **Illegal Zone** tabs in turn and configure the following on the **General** tab:

- **Enable this zone**—Whether the zone will be processed by anomaly detection.
- **Service Subnets**—The IP addresses that comprise the zone. The default (0.0.0.0) is that no address is included in the zone. Replace 0.0.0.0 to define addresses for the zone.

The addresses can be single host (such as 10.100.10.1), a range of addresses (such as 10.100.10.0-10.100.10.255), or network/host objects that contain single hosts, address ranges, or a combination of hosts and ranges. Click **Select** to select objects from a list or to create new objects.

- Decide whether to enable the external zone—The external zone comprises all IP addresses that are not configured for the internal or illegal zones. You do not explicitly assign addresses to this zone. On the **External Zone** tab, **General** sub-tab, you can enable or disable the zone using the **Enable this zone** checkbox. The external zone is enabled by default.
- Configure scanner thresholds and histograms—Each zone has sub-tabs for **TCP Protocol**, **UDP Protocol**, and **Other Protocols**. On these tabs, you can configure non-default settings for specific services that override the learned histograms. For detailed information about configuring these settings, see [Configuring Anomaly Detection Thresholds and Histograms, page 41-11](#).

At this point, you have finished configuring the basic anomaly detection settings.

Step 5 (Device view only.) Configure the anomaly detection mode. This setting is defined in the **Virtual Sensors** policy. Consider the following tips to select the correct policy:

- If you configured the anomaly detection policy on a virtual sensor (other than vs0, which is represented by the parent IPS device), you must select the parent IPS device, then select the Virtual Sensors policy.

- If you configured the Anomaly Detection policy as a shared policy in Policy view, select the IPS device to which the policy is assigned, or that hosts a virtual sensor to which the policy is assigned.

Then, complete the following steps in the Virtual Sensors policy:

- Select the desired virtual sensor in the table and click the **Edit Row** button.
- In the Modify Virtual Sensors dialog box, select the appropriate option for the Anomaly Detection Mode setting: Detect, Inactive, Learn. The default and normal operational mode is Detect. However, if you are using asymmetric normalizer mode, you might want to set the anomaly detection mode to Inactive. For detailed information about these modes, see [Anomaly Detection Modes, page 41-2](#). For information about the other settings in this dialog box, see [Virtual Sensor Dialog Box, page 38-7](#).
- If you placed anomaly detection in Learning mode, remember to change the mode to Detect after the desired learning period has completed.

Step 6 Add additional actions to the anomaly detection signatures, if desired. For example, you might want to add a deny action so that attacks are dropped. You can alternatively configure event action overrides to add actions based on risk rating. For more information, see [Configuring Anomaly Detection Signatures, page 41-4](#).

Step 7 Manage the knowledge base, if necessary.

If you configured the knowledge base to automatically rotate (on the Learning Accept Mode tab), then the knowledge base is refreshed automatically and manual intervention is not necessary.

If you configured anomaly detection to only save new databases, and not use them, then you need to manually load updated knowledge bases periodically. You cannot do this in Security Manager; use the IPS Device Manager (IDM) instead.

Using IDM (or IME), you can load, delete, and rename knowledge bases, and upload them to or download them from an external server. For more information about what you can do, see the online help for IDM or IME.

Configuring Anomaly Detection Learning Accept Mode

Use the Learning Accept Mode tab of the Anomaly Detection policy to configure whether you want the sensor to create a new knowledge base every so many hours. You can configure whether the knowledge base is created and loaded (Rotate) or saved (Save Only). You can schedule how often and when the knowledge base is loaded or saved.

The default generated filename is YYYY-Mon-dd-hh_mm_ss (that is, year-month-day-hour_minute_second), where Mon is a three-letter abbreviation of the current month.

The knowledge base has a tree structure and contains the following information:

- Knowledge base name
- Zone name
- Protocol
- Service

The knowledge base holds a scanner threshold and a histogram for each service. If you have learning accept mode set to automatic and the action set to rotate, a new knowledge base is created every 24 hours and used in the next 24 hours. If you have learning accept mode set to automatic and the action is set to save only, a new knowledge base is created but not loaded, and the current knowledge base is used. If you do not have learning accept mode set to automatic, no knowledge base is created.

**Tip**

Although you can use Security Manager to configure how knowledge bases are generated, you cannot manage the knowledge bases themselves. Use the IPS Device Manager (IDM), or IPS Manager Express (IME) instead. Using IDM (or IME), you can load, delete, and rename knowledge bases, and upload them to or download them from an external server. For more information about what you can do, see the online help for IDM or IME.

Related Topics

- [Anomaly Detection Modes, page 41-2](#)
- [Configuring Anomaly Detection, page 41-6](#)
- [Understanding Anomaly Detection Thresholds and Histograms, page 41-9](#)

-
- Step 1** Do one of the following to open the Anomaly Detection policy you want to modify:
- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
 - (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.
- Step 2** Click the **Learning Accept Mode** tab and configure the following options:
- **Automatically accept learning knowledge base**—Whether to have the sensor automatically update the knowledge base. If you do not select this option, anomaly detection does not automatically create a new knowledge base, and you cannot configure the other options on this tab.
 - **Action**—Whether to rotate or save the knowledge base when it is created.

If you choose **Rotate** (the default), the new knowledge base is created and loaded according to the schedule you define. If you choose **Save Only**, the new knowledge base is created but not loaded. You can examine it and decide whether to load it into anomaly detection using IDM or IME.
- Step 3** In the **Schedule** field, select the schedule for generating a new knowledge base. The default schedule is periodic starting at 10 AM and running for 24 hours. Options are:
- **Periodic**—Base the schedule on a recurring period. Configure the following options:
 - **Start Time**—The starting time for the learning window in hh:mm:ss format (24-hour clock).
 - **Learning Interval in hours**—How long you want anomaly detection to learn from the network before creating a new knowledge base.
 - **Calendar Schedule**—Base the schedule on specific times of day and days of the week. The dialog box changes to show Time of Day and Days of the Week tables. These times apply to every day selected; you cannot specify different times for different days.
 - To add a time or day, click the **Add Row (+)** button beneath the appropriate table. The time is in hh:mm:ss format (24-hour clock). For day, select the day from the list.
 - To edit an existing time or day, select it and click the **Edit Row (pencil)** button.
 - To delete a time or day, select it and click the **Delete Row (trash can)** button. Ensure that you have at least one time and one day configured.
-

Understanding Anomaly Detection Thresholds and Histograms

Anomaly detection uses thresholds and histograms to determine if scanning behavior is an attack.

During learning mode, anomaly detection develops histograms for each TCP and UDP port, and for other protocols, to create a baseline of the normal behavior of your network (see [Anomaly Detection Modes, page 41-2](#)). For example, the histogram for a TCP port lists the “normal” number of source addresses that make incomplete connections to a certain number of destination addresses during a minute. The histograms contain three buckets: low number of destination addresses (5), medium number (20), and high number (100). (The destination buckets are a fixed number.) Separate histograms are kept for each service and zone (see [Anomaly Detection Zones, page 41-3](#)).

For example, learning mode might develop the following histogram for TCP port 80:

Number of Destination Addresses	Number of Source Addresses
Low (5)	18
Medium (20)	6
High (100)	2

In addition to these learned histograms, the anomaly detection scanner has a threshold setting that you configure. You configure a general scanner threshold, and you can override the threshold (configuring a different value) for any specific service (TCP port, UDP port, or other protocol). Each zone has its own thresholds.

When anomaly detection moves to detect mode, where it is actively scanning for worms, the thresholds and histograms are used as follows:

- Histograms are ignored until the threshold for the service is exceeded. For example, consider the above table for TCP/80 traffic. If the threshold is set at 200 (the default), a scanner must scan 200 hosts in a minute to trigger a scanner alert. If 7 source addresses scanned 50 hosts (which is an anomaly in the histogram, which expects no more than 6 hosts scanning 20-99 destinations), but a single scanner scanned only 100 addresses, no alert is generated, and no anomaly is detected.
- When the scanner threshold is exceeded, anomaly detection uses the histogram to determine if the service is under worm attack. In this example, if a source scans more than 200 destinations, anomaly detection evaluates the collected activity in the network. Because 7 hosts have scanned 50 hosts, a worm alert is generated.

When under worm attack, anomaly detection stops learning and clears the current learning information. It also temporarily lowers the thresholds.

- When a worm attack is detected, the worm timeout counter is started. When the timeout is reached, the scanner is reset. If the worm attack continues, new alerts are generated. You configure the worm timeout on the Operation Settings tab of the Anomaly Detection policy.

If you leave the defaults in place, anomaly detection generates histograms based on what it learns about your network from the network’s actual behavior. However, if you understand your network you can fine-tune these histograms to reduce false positives, creating your own definition of expected (or desired or tolerated) behavior for each TCP/UDP port or other protocol, for each zone. You can create your own histograms for only those services that interest you, and leave the defaults for all other ports. Additionally, you can configure the general scanner threshold for each zone, and configure different thresholds for specific services.

For information on configuring thresholds and histograms, see [Configuring Anomaly Detection Thresholds and Histograms, page 41-11](#).

Configuring Anomaly Detection Thresholds and Histograms

Anomaly detection uses thresholds and histograms to determine if scanning behavior is an attack. In most cases, you can use the default thresholds and the histograms that anomaly detection generates during learning mode (see [Anomaly Detection Modes, page 41-2](#)). However, you might want to fine-tune these settings. Changing the thresholds is a more likely change than creating your own histograms.

Before you configure these settings, read [Understanding Anomaly Detection Thresholds and Histograms, page 41-9](#). You must understand how thresholds and histograms are used together to configure them.

-
- Step 1** Do one of the following to open the Anomaly Detection policy you want to modify:
- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
 - (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.
- Step 2** Click the tab for the zone whose thresholds or histograms you want to change. You configure separate values for each zone: **Internal Zone**, **Illegal Zone**, **External Zone**. For an explanation of the zones, see [Anomaly Detection Zones, page 41-3](#).
- The tabs for each zone contain four sub-tabs: General, TCP Protocol, UDP Protocol, and Other Protocols. The General tab defines the IP addresses for the zone and whether the zone is enabled (the External zone includes all IP addresses not specified for the other zones, you do not configure specific addresses for the External zone).
- The other tabs are where you define thresholds and histograms.
- Step 3** Select the tab for the protocol for which you want to modify thresholds or histograms: **TCP Protocol**, **UDP Protocol**, **Other Protocol**.
- On each tab, configure the following options:
- **Enabled**—Whether anomaly detection is enabled for the protocol. You can turn off detection for all of TCP, UDP, or for all non-TCP/UDP protocols with this option. If you deselect the option, any other settings configured on the tab are ignored.
 - **Destination Port Map** or **Protocol Number Map** table—This table lists the TCP/UDP ports, or other protocols, for which you are configuring non-default mappings. By default, all ports and protocols are enabled and use the default scanner threshold.

Add items to this table only if you want to: disable detection for a port or protocol; set a different threshold value for a port or protocol; or configure an explicit histogram for a port or protocol, which will be used instead of the learned histogram.

 - To add a mapping, click the **Add Row (+)** button and fill in the Add Dest or Protocol Map dialog box. For detailed information, see [Dest Port or Protocol Map Dialog Box, page 41-12](#).
 - To edit a mapping, select it and click the **Edit Row (pencil)** button.
 - To delete a mapping, select it and click the **Delete Row (trash can)** button. Deleting a mapping returns the service to the default settings.
 - **Scanner Threshold**—The threshold for all TCP, UDP, or other protocols. This threshold is used for all services except those for which you configured a scanner override in the mapping table. The range is 5 to 1000. The default is 200.
 - **Threshold Histogram**—The default histogram for all TCP, UDP, or other protocols. This histogram is used for all services except those for which you configured a scanner override in the mapping table.

The content of this table is fixed; you cannot add or delete items. However, you can select a row and click **Edit Row (pencil)** to change the number of source addresses configured for a threshold setting. See [Histogram Dialog Box, page 41-13](#).

- Step 4** Repeat the process for each combination of zone and protocol for which you are defining non-default settings.

Dest Port or Protocol Map Dialog Box

Use the Add or Modify Dest Port Map dialog box to add or modify destination port scanner settings for TCP or UDP, and the Add or Modify Protocol Map dialog box to add or modify scanner settings for other protocols.

Before you configure these settings, read the following topics:

- [Understanding Anomaly Detection Thresholds and Histograms, page 41-9](#)
- [Configuring Anomaly Detection Thresholds and Histograms, page 41-11](#)



Tip

You do not need to add a port or protocol to have anomaly detection look for worm attacks against it. By default, all ports and protocols are processed. You need to configure specific settings only if you want to turn off detection on a specific port or protocol, or if you want non-default thresholds or histograms.

Navigation Path

In the Anomaly Detection policy, on the **TCP Protocol**, **UDP Protocol**, or **Other Protocol** sub tabs on the Internal Zone, Illegal Zone, or External Zone tabs, click the **Add Row** button beneath the Destination Port Map or Protocol Number Map tables, or select a row and click the **Edit Row** button. For more information on the steps required to get here, see [Configuring Anomaly Detection Thresholds and Histograms, page 41-11](#).

Field Reference

Table 41-2 Destination Port or Protocol Map Dialog Box

Element	Description
Destination Port Number (Dest Port Map dialog box only.)	The destination port number for which you are defining non-default values. The range is 0 to 65535. Enter a single port number, or the name of a port list object that contains a single port number. Click Select to select an object from a list or to create a new one.
Protocol Number (Protocol Map dialog box only.)	The protocol number for non-TCP/UDP protocols. For a list of protocol numbers, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "Protocol Numbers." Look for a heading (at the time of this writing, the second search hit). The range is 0 to 255. For example, ICMP is protocol 1.
Enabled	Whether to enable this service. If you do not enable the service, the associated port or protocol is not processed by anomaly detection.

Table 41-2 Destination Port or Protocol Map Dialog Box (continued)

Element	Description
Override Scanner Settings	Whether to override the scanner settings for this service or protocol. You must select this option to enable the remaining fields on the dialog box.
Scanner Threshold	The scanner threshold for this port or protocol. The range is 5 to 1000. The default is 200.
Threshold Histogram table	<p>The histograms for this port or protocol. If you leave the table empty, the default histograms are used. You can have up to three rows, for low, medium, and high numbers of destination addresses, with different threshold levels (source addresses) for each.</p> <ul style="list-style-type: none"> To add a threshold, click the Add Row button and fill in the Histogram Dialog Box, page 41-13. The Add button is disabled if you already have three rows. To edit a threshold, select it and click the Edit Row button. You cannot change the destination bucket to one that is already defined in the table. To delete a threshold, select it and click the Delete Row button. Any buckets not included in the table use the default histogram for the bucket.

Histogram Dialog Box

Use the Histogram dialog box to create or modify entries in a histogram. The histograms you create or modify override the default histograms that anomaly detection generates. For detailed information about how these histograms are used, see:

- [Understanding Anomaly Detection Thresholds and Histograms, page 41-9](#)
- [Configuring Anomaly Detection Thresholds and Histograms, page 41-11](#)

Navigation Path

Do one of the following from the Anomaly Detection policy (see [Configuring Anomaly Detection, page 41-6](#)):

- On the TCP Protocol, UDP Protocol, or Other Protocol sub tabs on the Internal Zone, Illegal Zone, or External Zone tabs, select a row in the Threshold Histogram table and click the **Edit Row** button.
- On the Add or Modify Dest or Protocol Map dialog boxes, click the **Add Row** button, or select a row and click the **Edit Row** button. For information on opening the map dialog boxes, see [Dest Port or Protocol Map Dialog Box, page 41-12](#).

Field Reference

Table 41-3 Histogram Dialog Box

Element	Description
Number of Destination IP Addresses	<p>The histogram bucket you are defining. The buckets have a fixed number of destination addresses: Low (5 addresses); Medium (20); High (100).</p> <p>Tip A histogram can have a single entry for each destination bucket (low, medium, high). Thus, you cannot change this value to one that is already defined in the histogram you are editing.</p>
Number of Source IP Addresses	<p>The number of source addresses that are allowed to simultaneously scan the associated number of destination addresses. Enter the desired number.</p> <p>The range is 0 to 4096. If you are editing a histogram, the current value is shown.</p>