



CHAPTER 28

Configuring Advanced Firewall Protection

This chapter describes how to prevent network attacks by configuring protection features, and includes the following sections:

- [Configuring Threat Detection, page 28-1](#)
- [Configuring Connection Settings, page 28-6](#)
- [Configuring IP Audit, page 28-13](#)
- [Configuring the Fragment Size, page 28-20](#)
- [Configuring Anti-Spoofing, page 28-23](#)
- [Configuring TCP Options, page 28-23](#)
- [Configuring Global Timeouts, page 28-26](#)



Note

For Sun RPC server and encrypted traffic inspection settings, which you configure in the Configuration > Firewall > Advanced area (along with many of the topics in this chapter), see [Chapter 25, “Configuring Application Layer Protocol Inspection.”](#)

Configuring Threat Detection

This section describes how to configure scanning threat detection and basic threat detection. Threat detection is available in single mode only.

This section includes the following topics:

- [Configuring Basic Threat Detection, page 28-1](#)
- [Configuring Scanning Threat Detection, page 28-3](#)
- [Configuring Threat Statistics, page 28-4](#)

To view threat detection statistics, see the Firewall Dashboard Tab.

Configuring Basic Threat Detection

Basic threat detection detects activity that might be related to an attack, such as a DoS attack. Basic threat detection is enabled by default.

This section includes the following topics:

- [Basic Threat Detection Overview, page 28-2](#)
- [Configuring Basic Threat Detection, page 28-2](#)

Basic Threat Detection Overview

Using basic threat detection, the security appliance monitors the rate of dropped packets and security events due to the following reasons:

- Denial by access lists
- Bad packet format (such as invalid-ip-header or invalid-tcp-hdr-length)
- Connection limits exceeded (both system-wide resource limits, and limits set in the configuration)
- DoS attack detected (such as an invalid SPI, Stateful Firewall check failure)
- Basic firewall checks failed (This option is a combined rate that includes all firewall-related packet drops in this bulleted list. It does not include non-firewall-related drops such as interface overload, packets failed at application inspection, and scanning attack detected.)
- Suspicious ICMP packets detected
- Packets failed application inspection
- Interface overload
- Scanning attack detected (This option monitors scanning attacks; for example, the first TCP packet is not a SYN packet, or the TCP connection failed the 3-way handshake. Full scanning threat detection (see the [“Configuring Scanning Threat Detection”](#) section on page 28-3) takes this scanning attack rate information and acts on it by classifying hosts as attackers and automatically shunning them, for example.)
- Incomplete session detection such as TCP SYN attack detected or no data UDP session attack detected

When the security appliance detects a threat, it immediately sends a system log message (730100).

Basic threat detection affects performance only when there are drops or potential threats; even in this scenario, the performance impact is insignificant.

Configuring Basic Threat Detection

To enable or disable basic threat detection, on the Configuration > Firewall > Threat Detection pane, click the **Enable Basic Threat Detection** check box.

By default, this option enables detection for certain types of security events, including packet drops and incomplete session detections. You can override the default settings for each type of event if desired.

If an event rate is exceeded, then the security appliance sends a system message. The security appliance tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst rate interval is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each received event, the security appliance checks the average and burst rate limits; if both rates are exceeded, then the security appliance sends two separate system messages, with a maximum of one message for each rate type per burst period.

[Table 28-1](#) lists the default settings.

Table 28-1 Basic Threat Detection Default Settings

Packet Drop Reason	Trigger Settings	
	Average Rate	Burst Rate
<ul style="list-style-type: none"> DoS attack detected Bad packet format Connection limits exceeded Suspicious ICMP packets detected 	100 drops/sec over the last 600 seconds.	400 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	320 drops/sec over the last 120 second period.
Scanning attack detected	5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
	4 drops/sec over the last 3600 seconds.	8 drops/sec over the last 120 second period.
Incomplete session detected such as TCP SYN attack detected or no data UDP session attack detected (combined)	100 drops/sec over the last 600 seconds.	200 drops/sec over the last 20 second period.
	80 drops/sec over the last 3600 seconds.	160 drops/sec over the last 120 second period.
Denial by access lists	400 drops/sec over the last 600 seconds.	800 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	640 drops/sec over the last 120 second period.
<ul style="list-style-type: none"> Basic firewall checks failed Packets failed application inspection 	400 drops/sec over the last 600 seconds.	1600 drops/sec over the last 20 second period.
	320 drops/sec over the last 3600 seconds.	1280 drops/sec over the last 120 second period.
Interface overload	2000 drops/sec over the last 600 seconds.	8000 drops/sec over the last 20 second period.
	1600 drops/sec over the last 3600 seconds.	6400 drops/sec over the last 120 second period.

Configuring Scanning Threat Detection

A typical scanning attack consists of a host that tests the accessibility of every IP address in a subnet (by scanning through many hosts in the subnet or sweeping through many ports in a host or subnet). The scanning threat detection feature determines when a host is performing a scan. Unlike IPS scan detection that is based on traffic signatures, the security appliance scanning threat detection feature maintains an extensive database that contains host statistics that can be analyzed for scanning activity.

The host database tracks suspicious activity such as connections with no return activity, access of closed service ports, vulnerable TCP behaviors such as non-random IPID, and many more behaviors.

You can configure the security appliance to send system log messages about an attacker or you can automatically shun the host.

**Caution**

The scanning threat detection feature can affect the security appliance performance and memory significantly while it creates and gathers host- and subnet-based data structure and information.

To configure scanning threat detection, perform the following steps:

- Step 1** To enable scanning threat detection, on the Configuration > Firewall > Threat Detection pane, click the **Enable Scanning Threat Detection** check box.

By default, the system log message 730101 is generated when a host is identified as an attacker.

The security appliance identifies a host as an attacker or as a target if the scanning threat rate is exceeded. The security appliance tracks two types of rates: the average event rate over an interval, and the burst event rate over a shorter burst interval. The burst event rate is 1/30th of the average rate interval or 10 seconds, whichever is higher. For each event detected that is considered to be part of a scanning attack, the security appliance checks the average and burst rate limits. If either rate is exceeded for traffic sent from a host, then that host is considered to be an attacker. If either rate is exceeded for traffic received by a host, then that host is considered to be a target.

Table 28-2 lists the default rate limits for scanning threat detection.

Table 28-2 Default Rate Limits for Scanning Threat Detection

Average Rate	Burst Rate
5 drops/sec over the last 600 seconds.	10 drops/sec over the last 20 second period.
5 drops/sec over the last 3600 seconds.	10 drops/sec over the last 120 second period.

- Step 2** (Optional) To automatically terminate a host connection when the security appliance identifies the host as an attacker, check the **Shun Hosts detected by scanning threat** check box.
- Step 3** (Optional) To except host IP addresses from being shunned, enter an address in the **Networks excluded from shun** field.
- You can enter multiple addresses or subnets separated by commas. To choose a network from the list of IP address objects, click the ... button.
- Step 4** (Optional) To set the duration of a shun for an attacking host, check **Set Shun Duration** and enter a value between 10 and 2592000 seconds. The default length is 3600 seconds (1 hour). To restore the default value, click **Set Default**.

Configuring Threat Statistics

You can configure the security appliance to collect extensive statistics. Threat detection statistics show both allowed and dropped traffic rates. By default, statistics for access lists are enabled.

To view threat detection statistics, see the Firewall Dashboard Tab.

**Caution**

Enabling statistics can affect the security appliance performance, depending on the type of statistics enabled. Enabling statistics for hosts affects performance in a significant way; if you have a high traffic load, you might consider enabling this type of statistics temporarily. Enabling statistics for ports, however, has modest impact.

- To enable *all* statistics, in the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, click the **Enable All Statistics** radio button.
- To disable *all* statistics, on the Configuration > Firewall > Threat Detection pane, click the **Disable All Statistics** radio button.
- To enable only certain statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, click the **Enable Only Following Statistics** radio button, end then check one or more of the following check boxes:

- **Hosts**—Enables host statistics. The host statistics accumulate for as long as the host is active and in the scanning threat host database. The host is deleted from the database (and the statistics cleared) after 10 minutes of inactivity.
- **Access Rules** (enabled by default)—Enables statistics for access rules.
- **Port**—Enables statistics for TCP and UDP ports.
- **Protocol**—Enables statistics for non-TCP/UDP IP protocols.
- **TCP-Intercept**—Enables statistics for attacks intercepted by TCP Intercept (see the “[Enabling Connection Limits, TCP Normalization , and TCP State Bypass](#)” section on page 28-10 to enable TCP Intercept). After you check the TCP-Intercept option, you can set the following options in the Configuration > Firewall > Threat Detection > TCP Intercept Threat Detection area:

Monitoring Window Size—Sets the size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. The security appliance samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

Burst Threshold Rate—Sets the threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, syslog message 733104 is generated.

Average Threshold Rate—Sets the average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, syslog message 733105 is generated.

Click the **Set Default** button to restore the default values.

- To set the number of rate intervals maintained for host statistics, on the Configuration > Firewall > Threat Detection > Scanning Threat Statistics area, choose **1**, **2**, or **3** from the User can specify the number of rate for Threat Detection Host drop-down list. Because host statistics use a lot of memory, reducing the number of rate intervals from the default of 3 reduces the memory usage. By default, the Firewall Dashboard Tab shows information for three rate intervals, for example, for the last 1 hour, 8 hours, and 24 hours. If you set this keyword to **1**, then only the shortest rate interval statistics are maintained. If you set the value to **2**, then the two shortest intervals are maintained.

Configuring Connection Settings

This section describes how to:

- Set maximum TCP and UDP connections
- Set maximum embryonic connections
- Set maximum per-client connections
- Set connection timeouts
- Set dead connection detection
- Disable TCP sequence randomization.
- Configure TCP normalization.
- Configure TCP state bypass.

This section includes the following topics:

- [Information About Connection Limits, page 28-6](#)
- [Information About TCP Normalization, page 28-7](#)
- [Information About TCP State Bypass, page 28-8](#)
- [Enabling Connection Limits, TCP Normalization , and TCP State Bypass, page 28-10](#)

**Note**

You can also configure maximum connections, maximum embryonic connections, and TCP sequence randomization in the NAT configuration. If you configure these settings for the same traffic using both methods, then the security appliance uses the lower limit. For TCP sequence randomization, if it is disabled using either method, then the security appliance disables TCP sequence randomization.

Information About Connection Limits

This section describes why you might want to limit connections, and includes the following topics:

- [Information About TCP Intercept, page 28-6](#)
- [Disabling TCP Intercept for Management Packets for Clientless SSL VPN Compatibility, page 28-7](#)
- [Information About Dead Connection Detection, page 28-7](#)
- [Information About TCP Sequence Randomization, page 28-7](#)

Information About TCP Intercept

Limiting the number of embryonic connections protects you from a DoS attack. The security appliance uses the per-client limits and the embryonic connection limit to trigger TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. TCP Intercept uses the SYN cookies algorithm to prevent TCP SYN-flooding attacks. A SYN-flooding attack consists of a series of SYN packets usually originating from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests. When the embryonic connection threshold of a connection is crossed, the

security appliance acts as a proxy for the server and generates a SYN-ACK response to the client SYN request. When the security appliance receives an ACK back from the client, it can then authenticate the client and allow the connection to the server.

Disabling TCP Intercept for Management Packets for Clientless SSL VPN Compatibility

By default, TCP management connections have TCP Intercept always enabled. When TCP Intercept is enabled, it intercepts the 3-way TCP connection establishment handshake packets and thus deprives the security appliance from processing the packets for Clientless (browser-based) SSL VPN. Clientless SSL VPN requires the ability to process the 3-way handshake packets to provide selective ACK and other TCP options for Clientless SSL VPN connections. To disable TCP Intercept for management traffic, you can set the embryonic connection limit; only after the embryonic connection limit is reached is TCP Intercept enabled.

Information About Dead Connection Detection

Dead connection detection detects a dead connection and allows it to expire, without expiring connections that can still handle traffic. You configure DCD when you want idle, but valid connections to persist.

When you enable DCD, idle timeout behavior changes. With idle timeout, DCD probes are sent to each of the two end-hosts to determine the validity of the connection. If an end-host fails to respond after probes are sent at the configured intervals, the connection is freed, and reset values, if configured, are sent to each of the end-hosts. If both end-hosts response that the connection is valid, the activity timeout is updated to the current time and the idle timeout is rescheduled accordingly.

Information About TCP Sequence Randomization

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

TCP initial sequence number randomization can be disabled if required. For example:

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- You use a WAAS device that requires the security appliance not to randomize the sequence numbers of connections.

Information About TCP Normalization

The TCP normalization feature identifies abnormal packets that the security appliance can act on when they are detected; for example, the security appliance can allow, drop, or clear the packets. TCP normalization helps protect the security appliance from attacks. TCP normalization is always enabled, but you can customize how some features behave.

The TCP normalizer includes non-configurable actions and configurable actions. Typically, non-configurable actions that drop or clear connections apply to packets that are always bad. Configurable actions (as detailed in [“Enabling Connection Limits, TCP Normalization , and TCP State Bypass” section on page 28-10](#)) might need to be customized depending on your network needs.

See the following guidelines for TCP normalization:

- The normalizer does not protect from SYN floods. The security appliance includes SYN flood protection in other ways.
- The normalizer always sees the SYN packet as the first packet in a flow unless the security appliance is in loose mode due to failover.

Information About TCP State Bypass


By default, all traffic that goes through the security appliance is inspected using the Adaptive Security Algorithm and is either allowed through or dropped based on the security policy. The security appliance maximizes the firewall performance by checking the state of each packet (is this a new connection or an established connection?) and assigning it to either the session management path (a new connection SYN packet), the fast path (an established connection), or the control plane path (advanced inspection). See the [“Stateful Inspection Overview” section on page 3-17](#) for more detailed information about the stateful firewall.

TCP packets that match existing connections in the fast path can pass through the security appliance without rechecking every aspect of the security policy. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), can stand in the way of asymmetrical routing solutions: both the outbound and inbound flow of a connection must pass through the same security appliance.

For example, a new connection goes to security appliance 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through security appliance 1, then the packets will match the entry in the fast path, and are passed through. But if subsequent packets go to security appliance 2, where there was not a SYN

packet that went through the session management path, then there is no entry in the fast path for the connection, and the packets are dropped. Figure 28-1 shows an asymmetric routing example where the outbound traffic goes through a different security appliance than the inbound traffic:

Figure 28-1 Asymmetric Routing



If you have asymmetric routing configured on upstream routers, and traffic alternates between two security appliances, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic much as it treats a UDP connection: when a non-SYN packet matching the specified networks enters the security appliance, and there is not an fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Licensing Requirements for TCP State Bypass

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Supported in routed and transparent mode.

Failover Guidelines

Failover is supported.

IPv6 Guidelines

Supports IPv6.

Unsupported Features

The following features are not supported when you use TCP state bypass:

- Application inspection—Application inspection requires both inbound and outbound traffic to go through the same security appliance, so application inspection is not supported with TCP state bypass.
- AAA authenticated sessions—When a user authenticates with one security appliance, traffic returning via the other security appliance will be denied because the user did not authenticate with that security appliance.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The security appliance does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- SSM and SSC functionality—You cannot use TCP state bypass and any application running on an SSM or SSC, such as IPS or CSC.

NAT Guidelines

Because the translation session is established separately for each security appliance, be sure to configure static NAT on both security appliances for TCP state bypass traffic; if you use dynamic NAT, the address chosen for the session on security appliance 1 will differ from the address chosen for the session on security appliance 2.

Default Settings

TCP state bypass is disabled by default.

Enabling Connection Limits, TCP Normalization , and TCP State Bypass

To configure connection limits and TCP normalization, perform the following steps:

-
- Step 1** Configure a service policy on the Configuration > Firewall > Service Policy Rules pane according to [Chapter 23, “Configuring Service Policy Rules.”](#)
- You can configure connection limits as part of a new service policy rule, or you can edit an existing service policy.
- Step 2** On the Rule Actions dialog box, click the **Connection Settings** tab.
- Step 3** To set maximum connections, configure the following values in the Maximum Connections area:
- TCP & UDP Connections—Specifies the maximum number of simultaneous TCP and UDP connections for all clients in the traffic class, up to 65,536. The default is 0 for both protocols, which means the maximum possible connections are allowed.

- **Embryonic Connections**—Specifies the maximum number of embryonic connections per host up to 65,536. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is 0, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.
- **Per Client Connections**—Specifies the maximum number of simultaneous TCP and UDP connections for each client. When a new connection is attempted by a client that already has opened the maximum per-client number of connections, the security appliance rejects the connection and drops the packet.
- **Per Client Embryonic Connections**—Specifies the maximum number of simultaneous TCP embryonic connections for each client. When a new TCP connection is requested by a client that already has the maximum per-client number of embryonic connections open through the security appliance, the security appliance proxies the request to the TCP Intercept feature, which prevents the connection.

Step 4 To configure TCP timeouts, configure the following values in the TCP Timeout area:

- **Connection Timeout**—Specifies the idle time until a connection slot is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
- **Send reset to TCP endpoints before timeout**—Specifies that the security appliance should send a TCP reset message to the endpoints of the connection before freeing the connection slot.
- **Embryonic Connection Timeout**—Specifies the idle time until an embryonic connection slot is freed. Enter 0:0:0 to disable timeout for the connection. The default is 30 seconds.
- **Half Closed Connection Timeout**—Specifies the idle time until a half closed connection slot is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 10 minutes.

Step 5 To disable randomized sequence numbers, uncheck **Randomize Sequence Number**.

TCP initial sequence number randomization can be disabled if another in-line firewall is also randomizing the initial sequence numbers, because there is no need for both firewalls to be performing this action. However, leaving ISN randomization enabled on both firewalls does not affect the traffic.

Each TCP connection has two ISNs: one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN passing in the outbound direction. If the connection is between two interfaces with the same security level, then the ISN will be randomized in the SYN in both directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session.

Step 6 To configure TCP normalization, check **Use TCP Map**.

Choose an existing TCP map from the drop-down list (if available), or add a new one by clicking **New**.

The Add TCP Map dialog box appears.

- a. In the TCP Map Name field, enter a name.
- b. In the Queue Limit field, enter the maximum number of out-of-order packets, between 0 and 250 packets.

The Queue Limit sets the maximum number of out-of-order packets that can be buffered and put in order for a TCP connection. The default is 0, which means this setting is disabled and the default system queue limit is used depending on the type of traffic:

- Connections for application inspection, IPS, and TCP check-retransmission have a queue limit of 3 packets. If the security appliance receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting.
- For other TCP connections, out-of-order packets are passed through untouched.

If you set the Queue Limit command to be 1 or above, then the number of out-of-order packets allowed for all TCP traffic matches this setting. For example, for application inspection, IPS, and TCP check-retransmission traffic, any advertised settings from TCP packets are ignored in favor of the Queue Limit setting. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

- c. In the Timeout field, set the maximum amount of time that out-of-order packets can remain in the buffer, between 1 and 20 seconds.

If they are not put in order and passed on within the timeout period, then they are dropped. The default is 4 seconds. You cannot change the timeout for any traffic if the Queue Limit is set to 0; you need to set the limit to be 1 or above for the Timeout to take effect.

- d. In the Reserved Bits area, click **Clear and allow**, **Allow only**, or **Drop**.

Allow only allows packets with the reserved bits in the TCP header.

Clear and allow clears the reserved bits in the TCP header and allows the packet.

Drop drops the packet with the reserved bits in the TCP header.

- e. Check any of the following options:

- Clear urgent flag—Clears the URG flag through the security appliance. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag, therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable to attacks.
- Drop connection on window variation—Drops a connection that has changed its window size unexpectedly. The window size mechanism allows TCP to advertise a large window and to subsequently advertise a much smaller window without having accepted too much data. From the TCP specification, “shrinking the window” is strongly discouraged. When this condition is detected, the connection can be dropped.
- Drop packets that exceed maximum segment size—Drops packets that exceed MSS set by peer.
- Check if transmitted data is the same as original—Enables the retransmit data checks.
- Drop packets which have past-window sequence—Drops packets that have past-window sequence numbers, namely the sequence number of a received TCP packet is greater than the right edge of the TCP receiving window. If you do not check this option, then the Queue Limit must be set to 0 (disabled).
- Drop SYN Packets with data—Drops SYN packets with data.
- Enable TTL Evasion Protection—Enables the TTL evasion protection offered by the security appliance. Do not enable this option if you want to prevent attacks that attempt to evade security policy.

For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the security appliance and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the security

appliance to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received by the attacker. In this case, an attacker is able to succeed without security preventing the attack.

- Verify TCP Checksum—Enables checksum verification.
- Drop SYNACK Packets with data—Drops TCP SYNACK packets that contain data.
- Drop packets with invalid ACK—Drops packets with an invalid ACK. You might see invalid ACKs in the following instances:
 - In the TCP connection SYN-ACK-received status, if the ACK number of a received TCP packet is not exactly same as the sequence number of the next TCP packet sending out, it is an invalid ACK.
 - Whenever the ACK number of a received TCP packet is greater than the sequence number of the next TCP packet sending out, it is an invalid ACK.



Note TCP packets with an invalid ACK are automatically allowed for WAAS connections.

- f. To set TCP options, check any of the following options:
 - Clear Selective Ack—Lists whether the selective-ack TCP option is allowed or cleared.
 - Clear TCP Timestamp—Lists whether the TCP timestamp option is allowed or cleared.
 - Clear Window Scale—Lists whether the window scale timestamp option is allowed or cleared.
 - Range—Lists the valid TCP options ranges, which should fall within 6-7 and 9-255. The lower bound should be less than or equal to the upper bound.
- g. Click **OK**.

Step 7 To set the time to live, check **Decrement time to live for a connection**.

Step 8 To enable TCP state bypass, in the Advanced Options area, check **TCP State Bypass**.

Step 9 Click **OK** or **Finish**.

Configuring IP Audit

The IP audit feature provides basic IPS functionality; for advanced IPS functionality on supported platforms, you can install an AIP SSM.

This feature lets you create a named audit policy that identifies the actions to take when a packet matches a predefined attack signature or informational signature. Signatures are activities that match known attack patterns. For example, there are signatures that match DoS attacks. You can configure the security appliance to drop the packet, generate an alarm, or reset the connection.

IP Audit Policy

The IP Audit Policy pane lets you add audit policies and assign them to interfaces. You can assign an attack policy and an informational policy to each interface. The attack policy determines the action to take with packets that match an attack signature; the packet might be part of an attack on your network, such as a DoS attack. The informational policy determines the action to take with packets that match an

informational signature; the packet is not currently attacking your network, but could be part of an information-gathering activity, such as a port sweep. For a complete list of signatures, see the [IP Audit Signature List](#).

Fields

- Name—Shows the names of the defined IP audit policies. Although the default actions for a named policy are listed in this table (“--Default Action--”), they are not named policies that you can assign to an interface. Default actions are used by named policies if you do not set an action for the policy. You can modify the default actions by selecting them and clicking the Edit button.
- Type—Shows the policy type, either Attack or Info.
- Action—Shows the actions taken against packets that match the policy, Alarm, Drop, and/or Reset. Multiple actions can be listed.
- Add—Adds a new IP audit policy.
- Edit—Edits an IP audit policy or the default actions.
- Delete—Deletes an IP audit policy. You cannot delete a default action.
- Policy-to-Interface Mappings—Assigns an attack and informational policy to each interface.
 - Interface—Shows the interface name.
 - Attack Policy—Lists the attack audit policy names available. Assign a policy to an interface by clicking the name in the list.
 - Info Policy—Lists the informational audit policy names available. Assign a policy to an interface by clicking the name in the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Add/Edit IP Audit Policy Configuration

The Add/Edit IP Audit Policy Configuration dialog box lets you add or edit a named IP audit policy that you can assign to interfaces, and lets you modify the default actions for each signature type.

Fields

- Policy Name—Sets the IP audit policy name. You cannot edit the name after you add it.
- Policy Type—Sets the policy type. You cannot edit the policy type after you add it.
 - Attack—Sets the policy type as attack.
 - Information—Sets the policy type as informational.
- Action—Sets one or more actions to take when a packet matches a signature. If you do not choose an action, then the default policy is used.

- Alarm—Generates a system message showing that a packet matched a signature. For a complete list of signatures, see [IP Audit Signature List](#).
- Drop—Drops the packet.
- Reset—Drops the packet and closes the connection.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

IP Audit Signatures

The IP Audit Signatures pane lets you disable audit signatures. You might want to disable a signature if legitimate traffic continually matches a signature, and you are willing to risk disabling the signature to avoid large numbers of alarms.

For a complete list of signatures, see [IP Audit Signature List](#).

Fields

- Enabled—Lists the enabled signatures.
- Disabled—Lists the disabled signatures.
- Disable—Moves the selected signature to the Disabled pane.
- Enable—Moves the selected signature to the Enabled pane.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

IP Audit Signature List

Table 28-3 lists supported signatures and system message numbers.

Table 28-3 Signature IDs and System Message Numbers

Signature ID	Message Number	Signature Title	Signature Type	Description
1000	400000	IP options-Bad Option List	Informational	Triggers on receipt of an IP datagram where the list of IP options in the IP datagram header is incomplete or malformed. The IP options list contains one or more options that perform various network management or debugging tasks.
1001	400001	IP options-Record Packet Route	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 7 (Record Packet Route).
1002	400002	IP options-Timestamp	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 4 (Timestamp).
1003	400003	IP options-Security	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 2 (Security options).
1004	400004	IP options-Loose Source Route	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 3 (Loose Source Route).
1005	400005	IP options-SATNET ID	Informational	Triggers on receipt of an IP datagram where the IP option list for the datagram includes option 8 (SATNET stream identifier).
1006	400006	IP options-Strict Source Route	Informational	Triggers on receipt of an IP datagram in which the IP option list for the datagram includes option 2 (Strict Source Routing).
1100	400007	IP Fragment Attack	Attack	Triggers when any IP datagram is received with an offset value less than 5 but greater than 0 indicated in the offset field.
1102	400008	IP Impossible Packet	Attack	Triggers when an IP packet arrives with source equal to destination address. This signature will catch the so-called Land Attack.

Table 28-3 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
1103	400009	IP Overlapping Fragments (Teardrop)	Attack	Triggers when two fragments contained within the same IP datagram have offsets that indicate that they share positioning within the datagram. This could mean that fragment A is being completely overwritten by fragment B, or that fragment A is partially being overwritten by fragment B. Some operating systems do not properly handle fragments that overlap in this manner and may throw exceptions or behave in other undesirable ways upon receipt of overlapping fragments, which is how the Teardrop attack works to create a DoS.
2000	400010	ICMP Echo Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 0 (Echo Reply).
2001	400011	ICMP Host Unreachable	Informational	Triggers when an IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 3 (Host Unreachable).
2002	400012	ICMP Source Quench	Informational	Triggers when an IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 4 (Source Quench).
2003	400013	ICMP Redirect	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 5 (Redirect).
2004	400014	ICMP Echo Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 8 (Echo Request).
2005	400015	ICMP Time Exceeded for a Datagram	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 11 (Time Exceeded for a Datagram).
2006	400016	ICMP Parameter Problem on Datagram	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 12 (Parameter Problem on Datagram).
2007	400017	ICMP Timestamp Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 13 (Timestamp Request).

Table 28-3 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
2008	400018	ICMP Timestamp Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 14 (Timestamp Reply).
2009	400019	ICMP Information Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 15 (Information Request).
2010	400020	ICMP Information Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 16 (ICMP Information Reply).
2011	400021	ICMP Address Mask Request	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 17 (Address Mask Request).
2012	400022	ICMP Address Mask Reply	Informational	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the type field in the ICMP header set to 18 (Address Mask Reply).
2150	400023	Fragmented ICMP Traffic	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and either the more fragments flag is set to 1 (ICMP) or there is an offset indicated in the offset field.
2151	400024	Large ICMP Traffic	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP) and the IP length > 1024.
2154	400025	Ping of Death Attack	Attack	Triggers when a IP datagram is received with the protocol field of the IP header set to 1 (ICMP), the Last Fragment bit is set, and $(IP\ offset * 8) + (IP\ data\ length) > 65535$ that is to say, the IP offset (which represents the starting position of this fragment in the original packet, and which is in 8 byte units) plus the rest of the packet is greater than the maximum size for an IP packet.
3040	400026	TCP NULL flags	Attack	Triggers when a single TCP packet with none of the SYN, FIN, ACK, or RST flags set has been sent to a specific host.
3041	400027	TCP SYN+FIN flags	Attack	Triggers when a single TCP packet with the SYN and FIN flags are set and is sent to a specific host.

Table 28-3 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
3042	400028	TCP FIN only flags	Attack	Triggers when a single orphaned TCP FIN packet is sent to a privileged port (having port number less than 1024) on a specific host.
3153	400029	FTP Improper Address Specified	Informational	Triggers if a port command is issued with an address that is not the same as the requesting host.
3154	400030	FTP Improper Port Specified	Informational	Triggers if a port command is issued with a data port specified that is <1024 or >65535.
4050	400031	UDP Bomb attack	Attack	Triggers when the UDP length specified is less than the IP length specified. This malformed packet type is associated with a denial of service attempt.
4051	400032	UDP Snork attack	Attack	Triggers when a UDP packet with a source port of either 135, 7, or 19 and a destination port of 135 is detected.
4052	400033	UDP Chargen DoS attack	Attack	This signature triggers when a UDP packet is detected with a source port of 7 and a destination port of 19.
6050	400034	DNS HINFO Request	Informational	Triggers on an attempt to access HINFO records from a DNS server.
6051	400035	DNS Zone Transfer	Informational	Triggers on normal DNS zone transfers, in which the source port is 53.
6052	400036	DNS Zone Transfer from High Port	Informational	Triggers on an illegitimate DNS zone transfer, in which the source port is not equal to 53.
6053	400037	DNS Request for All Records	Informational	Triggers on a DNS request for all records.
6100	400038	RPC Port Registration	Informational	Triggers when attempts are made to register new RPC services on a target host.
6101	400039	RPC Port Unregistration	Informational	Triggers when attempts are made to unregister existing RPC services on a target host.
6102	400040	RPC Dump	Informational	Triggers when an RPC dump request is issued to a target host.
6103	400041	Proxied RPC Request	Attack	Triggers when a proxied RPC request is sent to the portmapper of a target host.
6150	400042	ypserv (YP server daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP server daemon (ypserv) port.
6151	400043	ypbind (YP bind daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP bind daemon (ypbind) port.

Table 28-3 Signature IDs and System Message Numbers (continued)

Signature ID	Message Number	Signature Title	Signature Type	Description
6152	400044	yppasswdd (YP password daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP password daemon (yppasswdd) port.
6153	400045	ypupdated (YP update daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP update daemon (ypupdated) port.
6154	400046	ypxfrd (YP transfer daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the YP transfer daemon (ypxfrd) port.
6155	400047	mountd (mount daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the mount daemon (mountd) port.
6175	400048	rexid (remote execution daemon) Portmap Request	Informational	Triggers when a request is made to the portmapper for the remote execution daemon (rexid) port.
6180	400049	rexid (remote execution daemon) Attempt	Informational	Triggers when a call to the rexid program is made. The remote execution daemon is the server responsible for remote program execution. This may be indicative of an attempt to gain unauthorized access to system resources.
6190	400050	statd Buffer Overflow	Attack	Triggers when a large statd request is sent. This could be an attempt to overflow a buffer and gain access to system resources.

Configuring the Fragment Size

By default, the security appliance allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the security appliance. Fragmented packets are often used as DoS attacks.

Fields

- Fragment table:
 - Interface—Lists the available interfaces of the security appliance.
 - Size—Sets the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200.
 - Chain Length—Specifies the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.

- Timeout—Specifies the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.
- Edit—Opens the Edit Fragment dialog box.
- Show Fragment—Opens a panel and displays the current IP fragment database statistics for each interface of the security appliance.

Changing Fragment Parameters

To modify the IP fragment database parameters of an interface, perform the following steps:

-
- Step 1** Choose the interface to change in the Fragment table and click **Edit**. The Edit Fragment dialog box appears.
- Step 2** In the Edit Fragment dialog box, change the Size, Chain, and Timeout values as desired, and click **OK**. If you make a mistake, click **Restore Defaults**.
- Step 3** Click **Apply** in the Fragment panel.
-

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Show Fragment

The Show Fragment panel displays the operational data of the IP fragment reassembly module.

Fields

- Size—*Display only*. Displays the number of packets in the IP reassembly database waiting for reassembly. The default is 200.
- Chain—*Display only*. Displays the number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- Timeout—*Display only*. Displays the number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds displayed, all fragments of the packet that were already received will be discarded. The default is 5 seconds.
- Threshold—*Display only*. Displays the IP packet threshold, or the limit after which no new chains can be created in the reassembly module.
- Queue—*Display only*. Displays the number of IP packets waiting in the queue for reassembly.
- Assembled—*Display only*. Displays the number of IP packets successfully reassembled.

- Fail—*Display only*. Displays the number of failed reassembly attempts.
- Overflow—*Display only*. Displays the number of IP packets in the overflow queue.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Edit Fragment

The Edit Fragment dialog box lets you configure the IP fragment database of the selected interface.

Fields

- Interface—Displays the interface you selected in the Fragment panel. Changes made in the Edit Fragment dialog box are applied to the interface displayed.
- Size—Sets the maximum number of packets that can be in the IP reassembly database waiting for reassembly.
- Chain Length—Sets the maximum number of packets into which a full IP packet can be fragmented.
- Timeout—Sets the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded.
- Restore Defaults—Restores the factory default settings:
 - Size is 200.
 - Chain is 24 packets.
 - Timeout is 5 seconds.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Configuring Anti-Spoofing

The Anti-Spoofing window lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the security appliance only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the security appliance to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the security appliance, the security appliance routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the security appliance can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the security appliance uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the security appliance drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the security appliance drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fields

- Interface—Lists the interface names.
- Anti-Spoofing Enabled—Shows whether an interface has Unicast RPF enabled, Yes or No.
- Enable—Enables Unicast RPF for the selected interface.
- Disable—Disables Unicast RPF for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	•	—

Configuring TCP Options

The TCP Options pane lets you set parameters for TCP connections.

Fields

- Inbound and Outbound Reset—Sets whether to reset denied TCP connections for inbound and outbound traffic.
 - Interface—Shows the interface name.
 - Inbound Reset—Shows the interface reset setting for inbound TCP traffic, Yes or No. Enabling this setting causes the security appliance to send TCP resets for all inbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.
 - Outbound Reset—Shows the interface reset setting for outbound TCP traffic, Yes or No. Enabling this setting causes the security appliance to send TCP resets for all outbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.
 - Edit—Sets the inbound and outbound reset settings for the interface.
- Other Options—Sets additional TCP options.
 - Send Reset Reply for Denied Outside TCP Packets—Enables resets for TCP packets that terminate at the least secure interface and are denied by the security appliance based on access lists or AAA settings. When this option is not enabled, the security appliance silently discards denied packets. If you enable Inbound Resets for the least secure interface (see [TCP Reset Settings](#)), then you do not also have to enable this setting; Inbound Resets handle to-the-security appliance traffic as well as through the security appliance traffic.
 - Force Maximum Segment Size for TCP—Sets the maximum TCP segment size in bytes, between 48 and any maximum number. The default value is 1380 bytes. You can disable this feature by setting the bytes to 0. Both the host and the server can set the maximum segment size when they first establish a connection. If either maximum exceeds the value you set here, then the security appliance overrides the maximum and inserts the value you set. For example, if you set a maximum size of 1200 bytes, when a host requests a maximum size of 1300 bytes, then the security appliance alters the packet to request 1200 bytes.
 - Force Minimum Segment Size for TCP—Overrides the maximum segment size to be no less than the number of bytes you set, between 48 and any maximum number. This feature is disabled by default (set to 0). Both the host and the server can set the maximum segment size when they first establish a connection. If either maximum is less than the value you set for the Force Minimum Segment Size for TCP Proxy field, then the security appliance overrides the maximum and inserts the “minimum” value you set (the minimum value is actually the smallest maximum allowed). For example, if you set a minimum size of 400 bytes, if a host requests a maximum value of 300 bytes, then the security appliance alters the packet to request 400 bytes.
 - Force TCP Connection to Linger in TIME_WAIT State for at Least 15 Seconds—Forces each TCP connection to linger in a shortened TIME_WAIT state of at least 15 seconds after the final normal TCP close-down sequence. You might want to use this feature if an end host application default TCP terminating sequence is a simultaneous close. The default behavior of the security appliance is to track the shutdown sequence and release the connection after two FINs and the ACK of the last FIN segment. This quick release heuristic enables the security appliance to sustain a high connection rate, based on the most common closing sequence, known as the normal close sequence. However, in a simultaneous close, both ends of the transaction initiate the closing sequence, as opposed to the normal close sequence where one end closes and the other end acknowledges prior to initiating its own closing sequence (see RFC 793). Thus, in a simultaneous close, the quick release forces one side of the connection to linger in the

CLOSING state. Having many sockets in the CLOSING state can degrade the performance of an end host. For example, some WinSock mainframe clients are known to exhibit this behavior and degrade the performance of the mainframe server. Using this feature creates a window for the simultaneous close down sequence to complete.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

TCP Reset Settings

This dialog box sets the inbound and outbound reset settings for an interface.

Fields

- Send Reset Reply for Denied Inbound TCP Packets—Sends TCP resets for all inbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets.

You might want to explicitly send resets for inbound traffic if you need to reset identity request (IDENT) connections. When you send a TCP RST (reset flag in the TCP header) to the denied host, the RST stops the incoming IDENT process so that you do not have to wait for IDENT to time out. Waiting for IDENT to time out can cause traffic to slow because outside hosts keep retransmitting the SYN until the IDENT times out, so the **service resetinbound** command might improve performance.

- Send Reset Reply for Denied Outbound TCP Packets—Sends TCP resets for all outbound TCP sessions that attempt to transit the security appliance and are denied by the security appliance based on access lists or AAA settings. Traffic between same security level interfaces is also affected. When this option is not enabled, the security appliance silently discards denied packets. This option is enabled by default. You might want to disable outbound resets to reduce the CPU load during traffic storms, for example.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Configuring Global Timeouts

The Timeouts pane lets you set the timeout durations for use with the security appliance. All durations are displayed in the format hh:mm:ss. It sets the idle time for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool. TCP_connection slots are freed approximately 60 seconds after a normal connection close sequence.



Note

It is recommended that you do not change these values unless advised to do so by Customer Support.

Fields

In all cases, except for Authentication absolute and Authentication inactivity, unchecking the check boxes means there is no timeout value. For those two cases, clearing the check box means to reauthenticate on every new connection.

- **Connection**—Modifies the idle time until a connection slot is freed. Enter 0:0:0 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
- **Half-closed**—Modifies the idle time until a TCP half-closed connection closes. The minimum is 5 minutes. The default is 10 minutes. Enter 0:0:0 to disable timeout for a half-closed connection.
- **UDP**—Modifies the idle time until a UDP protocol connection closes. This duration must be at least 1 minute. The default is 2 minutes. Enter 0:0:0 to disable timeout.
- **ICMP**—Modifies the idle time after which general ICMP states are closed.
- **H.323**—Modifies the idle time until an H.323 media connection closes. The default is 5 minutes. Enter 0:0:0 to disable timeout.
- **H.225**—Modifies the idle time until an H.225 signaling connection closes. The H.225 default timeout is 1 hour (01:00:00). Setting the value of 00:00:00 means never close this connection. To close this connection immediately after all calls are cleared, a value of 1 second (00:00:01) is recommended.
- **MGCP**—Modifies the timeout value for MGCP which represents the idle time after which MGCP media ports are closed. The MGCP default timeout is 5 minutes (00:05:00). Enter 0:0:0 to disable timeout.
- **MGCP PAT**—Modifies the idle time after which an MGCP PAT translation is removed. The default is 5 minutes (00:05:00). The minimum time is 30 seconds. Uncheck the check box to return to the default value.
- **SUNRPC**—Modifies the idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes. Enter 0:0:0 to disable timeout.
- **SIP**—Modifies the idle time until an SIP signalling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- **SIP Media**—Modifies the idle time until an SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- **SIP Provisional Media**—Modifies the timeout value for SIP provisional media connections, between 0:1:0 and 1193:0:0. The default is 2 minutes.
- **SIP Invite**—Modifies the idle time after which pinholes for PROVISIONAL responses and media xlates will be closed. The minimum value is 0:1:0, the maximum value is 0:30:0. The default value is 0:03:00.

- SIP Disconnect—Modifies the idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message. The minimum value is 0:0:1, the maximum value is 0:10:0. The default value is 0:02:00.
- Authentication absolute—Modifies the duration until the authentication cache times out and you have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value. The system waits until you start a new connection to prompt you again. Enter 0:0:0 to disable caching and reauthenticate on every new connection.



Note Do not set this value to 0:0:0 if passive FTP is used on the connections.



Note When Authentication Absolute = 0, HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. This workaround opens a 1-second window of opportunity that might allow non-authenticated users to go through the firewall if they are coming from the same source IP address.

- Authentication inactivity—Modifies the idle time until the authentication cache times out and users have to reauthenticate a new connection. This duration must be shorter than the Translation Slot value.
- Translation Slot—Modifies the idle time until a translation slot is freed. This duration must be at least 1 minute. The default is 3 hours. Enter 0:0:0 to disable timeout.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

