



# CHAPTER 1

## Introduction to the Security Appliance

---

The security appliance combines advanced stateful firewall and VPN concentrator functionality in one device, and for some models, an integrated intrusion prevention module called the AIP SSM or an integrated content security and control module called the CSC SSM. The security appliance includes many advanced features, such as multiple security contexts (similar to virtualized firewalls), transparent (Layer 2) firewall or routed (Layer 3) firewall operation, advanced inspection engines, IPSec and WebVPN support, and many more features. See [Appendix A, “Feature Licenses and Specifications,”](#) for a list of supported platforms and features. For a list of new features, see the *Cisco ASA 5500 Series Release Notes* or the *Cisco PIX Security Appliance Release Notes*.



### Note

---

The Cisco PIX 501 and PIX 506E security appliances are not supported.

---

This chapter includes the following sections:

- [Firewall Functional Overview, page 1-1](#)
- [VPN Functional Overview, page 1-5](#)
- [Intrusion Prevention Services Functional Overview, page 1-5](#)
- [Security Context Overview, page 1-6](#)

## Firewall Functional Overview

Firewalls protect inside networks from unauthorized access by users on an outside network. A firewall can also protect inside networks from each other, for example, by keeping a human resources network separate from a user network. If you have network resources that need to be available to an outside user, such as a web or FTP server, you can place these resources on a separate network behind the firewall, called a *demilitarized zone* (DMZ). The firewall allows limited access to the DMZ, but because the DMZ only includes the public servers, an attack there only affects the servers and does not affect the other inside networks. You can also control when inside users access outside networks (for example, access to the Internet), by allowing only certain addresses out, by requiring authentication or authorization, or by coordinating with an external URL filtering server.

When discussing networks connected to a firewall, the *outside* network is in front of the firewall, the *inside* network is protected and behind the firewall, and a *DMZ*, while behind the firewall, allows limited access to outside users. Because the security appliance lets you configure many interfaces with varied security policies, including many inside interfaces, many DMZs, and even many outside interfaces if desired, these terms are used in a general sense only.

This section includes the following topics:

- [Security Policy Overview, page 1-2](#)
- [Firewall Mode Overview, page 1-3](#)
- [Stateful Inspection Overview, page 1-4](#)

## Security Policy Overview

A security policy determines which traffic is allowed to pass through the firewall to access another network. By default, the security appliance allows traffic to flow freely from an inside network (higher security level) to an outside network (lower security level). You can apply actions to traffic to customize the security policy. This section includes the following topics:

- [Permitting or Denying Traffic with Access Lists, page 1-2](#)
- [Applying NAT, page 1-2](#)
- [Using AAA for Through Traffic, page 1-2](#)
- [Applying HTTP, HTTPS, or FTP Filtering, page 1-3](#)
- [Applying Application Inspection, page 1-3](#)
- [Sending Traffic to the Advanced Inspection and Prevention Security Services Module, page 1-3](#)
- [Sending Traffic to the Content Security and Control Security Services Module, page 1-3](#)
- [Applying QoS Policies, page 1-3](#)
- [Applying Connection Limits and TCP Normalization, page 1-3](#)

### Permitting or Denying Traffic with Access Lists

You can apply an access list to limit traffic from inside to outside, or allow traffic from outside to inside. For transparent firewall mode, you can also apply an EtherType access list to allow non-IP traffic.

### Applying NAT

Some of the benefits of NAT include the following:

- You can use private addresses on your inside networks. Private addresses are not routable on the Internet.
- NAT hides the local addresses from other networks, so attackers cannot learn the real address of a host.
- NAT can resolve IP routing problems by supporting overlapping IP addresses.

### Using AAA for Through Traffic

You can require authentication and/or authorization for certain types of traffic, for example, for HTTP. The security appliance also sends accounting information to a RADIUS or TACACS+ server.

## Applying HTTP, HTTPS, or FTP Filtering

Although you can use access lists to prevent outbound access to specific websites or FTP servers, configuring and managing web usage this way is not practical because of the size and dynamic nature of the Internet. We recommend that you use the security appliance in conjunction with a separate server running one of the following Internet filtering products:

- Websense Enterprise
- Secure Computing SmartFilter

## Applying Application Inspection

Inspection engines are required for services that embed IP addressing information in the user data packet or that open secondary channels on dynamically assigned ports. These protocols require the security appliance to do a deep packet inspection.

## Sending Traffic to the Advanced Inspection and Prevention Security Services Module

If your model supports the AIP SSM for intrusion prevention, then you can send traffic to the AIP SSM for inspection.

## Sending Traffic to the Content Security and Control Security Services Module

If your model supports it, the CSC SSM provides protection against viruses, spyware, spam, and other unwanted traffic. It accomplishes this by scanning the FTP, HTTP, POP3, and SMTP traffic that you configure the adaptive security appliance to send to it.

## Applying QoS Policies

Some network traffic, such as voice and streaming video, cannot tolerate long latency times. QoS is a network feature that lets you give priority to these types of traffic. QoS refers to the capability of a network to provide better service to selected network traffic.

## Applying Connection Limits and TCP Normalization

You can limit TCP and UDP connections and embryonic connections. Limiting the number of connections and embryonic connections protects you from a DoS attack. The security appliance uses the embryonic 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 normalization is a feature consisting of advanced TCP connection settings designed to drop packets that do not appear normal.

## Firewall Mode Overview

The security appliance runs in two different firewall modes:

- Routed
- Transparent

In routed mode, the security appliance is considered to be a router hop in the network.

In transparent mode, the security appliance acts like a “bump in the wire,” or a “stealth firewall,” and is not considered a router hop. The security appliance connects to the same network on its inside and outside interfaces.

You might use a transparent firewall to simplify your network configuration. Transparent mode is also useful if you want the firewall to be invisible to attackers. You can also use a transparent firewall for traffic that would otherwise be blocked in routed mode. For example, a transparent firewall can allow multicast streams using an EtherType access list.

## Stateful Inspection Overview

All traffic that goes through the security appliance is inspected using the Adaptive Security Algorithm and either allowed through or dropped. A simple packet filter can check for the correct source address, destination address, and ports, but it does not check that the packet sequence or flags are correct. A filter also checks every packet against the filter, which can be a slow process.

A stateful firewall like the security appliance, however, takes into consideration the state of a packet:

- Is this a new connection?

If it is a new connection, the security appliance has to check the packet against access lists and perform other tasks to determine if the packet is allowed or denied. To perform this check, the first packet of the session goes through the “session management path,” and depending on the type of traffic, it might also pass through the “control plane path.”

The session management path is responsible for the following tasks:

- Performing the access list checks
- Performing route lookups
- Allocating NAT translations (xlates)
- Establishing sessions in the “fast path”




---

**Note** The session management path and the fast path make up the “accelerated security path.”

---

Some packets that require Layer 7 inspection (the packet payload must be inspected or altered) are passed on to the control plane path. Layer 7 inspection engines are required for protocols that have two or more channels: a data channel, which uses well-known port numbers, and a control channel, which uses different port numbers for each session. These protocols include FTP, H.323, and SNMP.

- Is this an established connection?

If the connection is already established, the security appliance does not need to re-check packets; most matching packets can go through the fast path in both directions. The fast path is responsible for the following tasks:

- IP checksum verification
- Session lookup
- TCP sequence number check
- NAT translations based on existing sessions
- Layer 3 and Layer 4 header adjustments

For UDP or other connectionless protocols, the security appliance creates connection state information so that it can also use the fast path.

Data packets for protocols that require Layer 7 inspection can also go through the fast path.

Some established session packets must continue to go through the session management path or the control plane path. Packets that go through the session management path include HTTP packets that require inspection or content filtering. Packets that go through the control plane path include the control packets for protocols that require Layer 7 inspection.

## VPN Functional Overview

A VPN is a secure connection across a TCP/IP network (such as the Internet) that appears as a private connection. This secure connection is called a tunnel. The security appliance uses tunneling protocols to negotiate security parameters, create and manage tunnels, encapsulate packets, transmit or receive them through the tunnel, and unencapsulate them. The security appliance functions as a bidirectional tunnel endpoint: it can receive plain packets, encapsulate them, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets, unencapsulate them, and send them to their final destination. The security appliance invokes various standard protocols to accomplish these functions.

The security appliance performs the following functions:

- Establishes tunnels
- Negotiates tunnel parameters
- Authenticates users
- Assigns user addresses
- Encrypts and decrypts data
- Manages security keys
- Manages data transfer across the tunnel
- Manages data transfer inbound and outbound as a tunnel endpoint or router

The security appliance invokes various standard protocols to accomplish these functions.

## Intrusion Prevention Services Functional Overview

The Cisco ASA 5500 series adaptive security appliance supports the AIP SSM, an intrusion prevention services module that monitors and performs real-time analysis of network traffic by looking for anomalies and misuse based on an extensive, embedded signature library. When the system detects unauthorized activity, it can terminate the specific connection, permanently block the attacking host, log the incident, and send an alert to the device manager. Other legitimate connections continue to operate independently without interruption. For more information, see *Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface*.

# Security Context Overview

You can partition a single security appliance into multiple virtual devices, known as security contexts. Each context is an independent device, with its own security policy, interfaces, and administrators. Multiple contexts are similar to having multiple standalone devices. Many features are supported in multiple context mode, including routing tables, firewall features, IPS, and management. Some features are not supported, including VPN and dynamic routing protocols.

In multiple context mode, the security appliance includes a configuration for each context that identifies the security policy, interfaces, and almost all the options you can configure on a standalone device. The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single mode configuration, is the startup configuration. The system configuration identifies basic settings for the security appliance. The system configuration does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses one of the contexts that is designated as the admin context.

The admin context is just like any other context, except that when a user logs into the admin context, then that user has system administrator rights and can access the system and all other contexts.

**Note**

---

You can run all your contexts in routed mode or transparent mode; you cannot run some contexts in one mode and others in another.

Multiple context mode supports static routing only.

---