



# Using PIX Firewall

---

The Cisco PIX Firewall lets you establish stateful firewall protection and secure VPN access with a single device. PIX Firewall provides a scalable security solution with failover support available for selected models to provide maximum reliability. PIX Firewall uses a specialized operating system that is more secure and easier to maintain than software firewalls that use a general-purpose operating system, which are subject to frequent threats and attacks.

This chapter describes how you can use the PIX Firewall to protect your network assets and to establish secure VPN access. It contains the following sections:

- [Controlling Network Access](#)
- [Protecting Your Network from Attack](#)
- [Enabling Specific Protocols and Applications](#)
- [Creating a Virtual Private Network](#)
- [PIX Firewall System Management](#)
- [PIX Firewall Failover](#)
- [Where to Go from Here](#)

## Controlling Network Access

This section describes the network firewall functionality provided by PIX Firewall. It includes the following topics:

- [How the PIX Firewall Works](#)
- [Adaptive Security Algorithm](#)
- [Multiple Interfaces and Security Levels](#)
- [How Data Moves Through the PIX Firewall](#)
- [Translation of Internal Addresses](#)
- [Cut-Through Proxy](#)
- [Access Control](#)

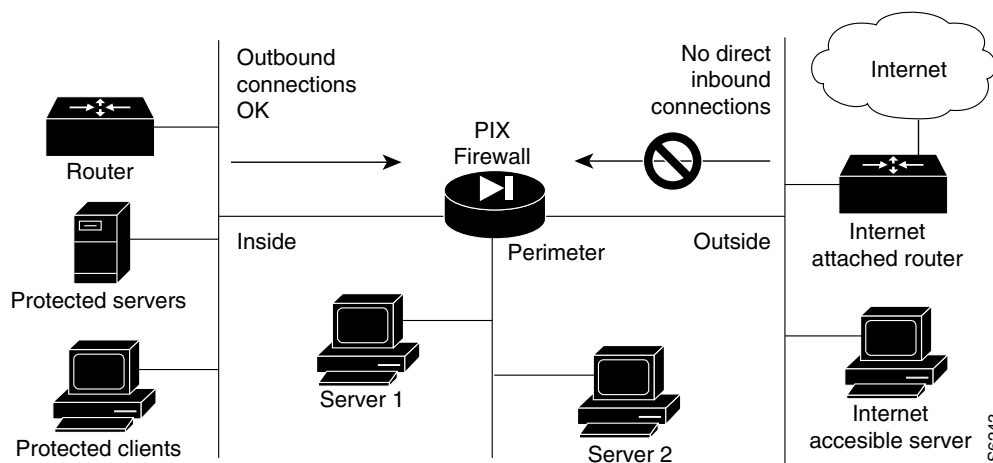
## How the PIX Firewall Works

The PIX Firewall protects an inside network from unauthorized access by users on an outside network, such as the public Internet. Most PIX Firewall models can optionally protect one or more perimeter networks, also known as demilitarized zones (DMZs). Access to the perimeter network is typically less restricted than access to the outside network, but more restricted than access to the inside network. Connections between the inside, outside, and perimeter networks are controlled by the PIX Firewall.

To effectively use a firewall in your organization, you need a security policy to ensure that all traffic from the protected networks passes only through the firewall to the unprotected network. You can then control who may access the networks with which services, and how to implement your security policy using the features that the PIX Firewall provides.

Figure 1-1 shows how a PIX Firewall protects a network while allowing outbound connections and secure access to the Internet.

**Figure 1-1 The PIX Firewall in a Network**



Within this architecture, the PIX Firewall forms the boundary between the protected networks and the unprotected networks. All traffic between the protected and unprotected networks flows through the firewall to maintain security. The unprotected network is typically accessible to the Internet. The PIX Firewall lets you locate servers such as those for Web access, SNMP, electronic mail (SMTP) in the protected network, and control who on the outside can access these servers.

Alternatively, for all PIX Firewall models except the PIX 506 and PIX 501, server systems can be located on a perimeter network as shown in Figure 1-1, and access to the server systems can be controlled and monitored by the PIX Firewall. The PIX 506 and PIX 501 each have two network interfaces, so all systems need to be located either on the inside or the outside interfaces.

The PIX Firewall also lets you implement your security policies for connection to and from the inside network.

Typically, the inside network is an organization's own internal network, or intranet, and the outside network is the Internet, but the PIX Firewall can also be used within an intranet to isolate or protect one group of internal computing systems and users from another.

The perimeter network can be configured to be as secure as the inside network or with varying security levels. Security levels are assigned numeric values from 0, the least secure, to 100, the most secure. The outside interface is always 0 and the inside interface is always 100. The perimeter interfaces can be any security level from 1 to 99.

Both the inside and perimeter networks are protected with the PIX Firewall's Adaptive Security Algorithm(ASA). The inside, perimeter, and outside interfaces can listen to RIP routing updates, and all interfaces can broadcast a RIP default route if required.

## Adaptive Security Algorithm

The Adaptive Security Algorithm (ASA) is a stateful approach to security. Every inbound packet is checked against the Adaptive Security Algorithm and against connection state information in memory. This stateful approach to security is regarded in the industry as being far more secure than a stateless packet screening approach.

ASA allows one way (inside to outside) connections without an explicit configuration for each internal system and application. ASA is always in operation, monitoring return packets to ensure they are valid. It actively randomizes TCP sequence numbers to minimize the risk of TCP sequence number attack.

ASA applies to the dynamic translation slots and static translation slots. You create static translation slots with the **static** command and dynamic translation slots with the **global** command. Collectively, both types of translation slots are referred to as “xlates.” ASA follows these rules:

- No packets can traverse the PIX Firewall without a connection and state.
- Outbound connections or states are allowed, except those specifically denied by access control lists. An outbound connection is one where the originator or client is on a higher security interface than the receiver or server. The highest security interface is always the inside interface and the lowest is the outside interface. Any perimeter interfaces can have security levels between the inside and outside values.
- Inbound connections or states are denied, except those specifically allowed. An inbound connection or state is one where the originator or client is on a lower security interface/network than the receiver or server. You can apply multiple exceptions to a single xlate (translation). This lets you permit access from an arbitrary machine, network, or any host on the Internet to the host defined by the xlate.
- All ICMP packets are denied unless specifically permitted.
- All attempts to circumvent the previous rules are dropped and a message is sent to syslog.

PIX Firewall handles UDP data transfers in a manner similar to TCP. Special handling allows DNS, archie, StreamWorks, H.323, and RealAudio to work securely. The PIX Firewall creates UDP “connection” state information when a UDP packet is sent from the inside network. Response packets resulting from this traffic are accepted if they match the connection state information. The connection state information is deleted after a short period of inactivity.

## Multiple Interfaces and Security Levels

All PIX Firewalls provide at least two interfaces, which by default, are called outside and inside, and are assigned a security level of 0 and 100, respectively. A lower security level indicates that the interface is relatively less protected than the higher security level. Typically, the outside interface is connected to the public Internet, while the inside interface is connected to your private network and is protected from public access.

Many PIX Firewall models provide up to eight interfaces, to allow you to create one or more perimeter networks, also called bastion networks or demilitarized zones (DMZs). A DMZ is a network that is more secure than the outside interface but less secure than the inside interface. You can assign security levels

to your perimeter networks from 0 to 100. Typically, you put mail servers or web servers that need to be accessed by users on the public Internet in a DMZ to provide some protection, but without jeopardizing the resources on your internal network.

## How Data Moves Through the PIX Firewall

When an outbound packet arrives at a PIX Firewall higher security level interface (security levels can be viewed with the **show nameif** command), the PIX Firewall checks to see if the packet is valid based on the Adaptive Security Algorithm, and then whether or not previous packets have come from that host. If not, then the packet is for a new connection, and PIX Firewall creates a translation slot in its state table for the connection. The information that PIX Firewall stores in the translation slot includes the inside IP address and a globally unique IP address assigned by Network Address Translation (NAT), Port Address Translation (PAT), or Identity (which uses the inside address as the outside address). The PIX Firewall then changes the packet's source IP address to the globally unique address, modifies the checksum and other fields as required, and forwards the packet to the lower security level interface.

When an inbound packet arrives at an external interface such as the outside interface, it first passes the PIX Firewall Adaptive Security criteria. If the packet passes the security tests, the PIX Firewall removes the destination IP address, and the internal IP address is inserted in its place. The packet is forwarded to the protected interface.

## Translation of Internal Addresses

The Network Address Translation (NAT) feature works by substituting, or translating, host addresses on an internal interface with a “global address” associated with an outside interface. This protects internal host addresses from being exposed on other network interfaces. To understand whether you want to use NAT, decide if you want to expose internal addresses on other network interfaces connected to the PIX Firewall. If you choose to protect internal host addresses using NAT, you identify the pool of addresses you want to use for translation.

If the addresses that you want to protect access only other networks within your organization, you can use any set of “private” addresses for the pool of translation addresses. For example, if you want to protect the host addresses on the Finance Department’s network (connected to the inside interface on the PIX Firewall) from exposure when connecting to the Sales Department network (connected to the perimeter interface on the PIX Firewall), you can set up translation using any available set of addresses on the Sales network. The effect is that hosts on the Finance network appear as local addresses on the Sales network.

If the addresses that you want to protect require Internet access, you use only NIC-registered addresses (official Internet addresses registered with the Network Information Center for your organization) for the pool of translation addresses. For example, if you want to protect host addresses on the Sales network (connected to a perimeter interface of the PIX Firewall) from exposure when making connections to the Internet (accessible through the outside interface of the PIX Firewall), you can set up translation using a pool of registered addresses on the outside interface. The effect is that hosts on the Internet see the only the Internet addresses for the Sales network, not the addresses on the perimeter interface.

If you are installing the PIX Firewall in an established network that has host- or network-registered addresses, you might not want to do translation for those hosts or networks because that would require using another registered address for the translation.

When considering NAT, it is also important to consider whether you have an equal number of addresses for internal hosts. If not, some internal hosts might not get network access when making a connection. In this case you can either apply for additional NIC-registered addresses or use Port Address Translation (PAT). PAT uses a single external address to manage up to 64,000 concurrent connections.

For inside systems, NAT translates the source IP address of outgoing packets (defined in RFC 1631). It supports both dynamic and static translation. NAT allows inside systems to be assigned private addresses (defined in RFC 1918), or to retain existing invalid addresses. NAT also provides additional security by hiding the real network identity of internal systems from the outside network.

PAT uses port remapping, which allows a single valid IP address to support source IP address translation for up to 64,000 active xlate objects. PAT minimizes the number of globally valid IP addresses required to support private or invalid internal addressing schemes. PAT does not work with multimedia applications that have an inbound data stream different from the outgoing control path. PAT provides additional security by hiding the real network identity of internal systems from the outside network.

Another class of address translation on the PIX Firewall is static translation. Static translation allows you to substitute a fixed external IP address for an internal address. This is useful for servers that require fixed IP addresses for access from the public Internet.

The PIX Firewall Identify feature allows address translation to be disabled. If existing internal systems have valid globally unique addresses, the Identity feature allows NAT and PAT to be selectively disabled for these systems. This feature makes internal network addresses visible to the outside network.

## Cut-Through Proxy

Cut-through proxy is a feature unique to PIX Firewall that allows user-based authentication of inbound or outbound connections. Unlike a proxy server that analyzes every packet at layer seven of the OSI model, a time- and processing-intensive function, the PIX Firewall first queries an authentication server, and when the connection is approved, establishes a data flow. All traffic thereafter flows directly and quickly between the two parties.

This feature allows security policies to be enforced on a per-user ID basis. Connections have to be authenticated with a user ID and password before they can be established. Supports authentication and authorization. The user ID and password are entered via an initial HTTP, Telnet, or FTP connection.

Cut-through proxy allows a much finer level of administrative control over connections compared to checking source IP addresses. When providing inbound authentication, appropriate controls need to be applied to the user ID and passwords used by external users (one-time passwords are recommended in this instance).

## Access Control

This section describes the features implemented by the PIX Firewall to support authentication and authorization of network users. It includes the following topics:

- [AAA Integration](#)
- [Access Lists](#)
- [Conduits](#)

## AAA Integration

PIX Firewall provides integration with AAA (Authentication, Accounting, and Authorization) services. AAA services are provided by TACACS+ or RADIUS servers.

PIX Firewall lets you define separate groups of TACACS+ or RADIUS servers for specifying different types of traffic; such as, a TACACS+ server for inbound traffic and another for outbound traffic.

AAA server groups are defined by a tag name that directs different types of traffic to each authentication server. If accounting is in effect, the accounting information goes to the active server.

The PIX Firewall allows a RADIUS server to send user group attributes to the PIX Firewall in the RADIUS authentication response message. The PIX Firewall then matches an access list to the attribute and determines RADIUS authorization from the access list. After the PIX Firewall authenticates a user, it uses the CiscoSecure **acl** attribute returned by the authentication server to identify an access list for a given user group.

## Access Lists

Beginning with version 5.3, the PIX Firewall uses access lists to control connections between inside and outside networks. Access lists are implemented with the **access-list** and **access-group** commands. These commands are used instead of the **conduit** and **outbound** commands, which were used in earlier versions of PIX Firewall, and which are maintained in current versions for backward compatibility.

You can use access lists to control connections based on source address, destination address, or protocol. Configure access lists carefully to allow the minimum access required. When possible, make access lists more restrictive by specifying a remote source address, local destination address, and protocol. The **access-list** and **access-group** command statements take precedence over the **conduit** and **outbound** command statements in your configuration.

## Conduits

Prior to version 5.3, PIX Firewall used the **conduit** and **outbound** commands to control connections between external and internal networks. With PIX Firewall version 6.0 and later, these commands continue to be supported for backward compatibility, but the **access-list** and **access-group** commands are now the preferred method of implementing this functionality.

Each conduit is a potential hole through the PIX Firewall and hence their use should be limited as your security policy and business needs require. When possible, make conduits more restrictive by specifying a remote source address, local destination address, and protocol.

# Protecting Your Network from Attack

This section describes the firewall features provided by PIX Firewall. These firewall features control network activity associated with specific kinds of attacks. This section includes the following topics:

- [Unicast Reverse Path Forwarding](#)
- [Flood Guard](#)
- [FragGuard and Virtual Re-Assembly](#)
- [DNS Control](#)
- [ActiveX Blocking](#)
- [Java Filtering](#)
- [URL Filtering](#)

For information about features that allow using specific protocols and applications across the firewall, refer to “[Enabling Specific Protocols and Applications](#).”

## Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (Unicast RPF), also known as “reverse route lookup,” provides inbound and outbound filtering to help prevent IP spoofing. This feature checks inbound packets for IP source address integrity, and verifies that packets destined for hosts outside the managed domain have IP source addresses verifiable by routes in the enforcing entities local routing table.

Unicast RPF is limited to addresses for networks in the enforcing entities local routing table. If the incoming packet does not have a source address represented by a route, it is impossible to know whether the packet arrived on the best possible path back to its origin.

## Flood Guard

The Flood Guard feature controls the AAA service's tolerance for unanswered login attempts. This helps to prevent a denial of service (DoS) attack on AAA services in particular. This feature optimizes AAA system use. It is enabled by default and can be controlled with the **floodguard 1** command.

## Flood Defender

The Flood Defender feature protects inside systems from a denial of service attack perpetrated by flooding an interface with TCP SYN packets. Enable this feature by setting the maximum embryonic connections option to the **nat** and **static** commands.

The TCP Intercept feature protects systems reachable via a static and TCP conduit. This feature ensures that once the optional embryonic connection limit is reached, and until the embryonic connection count falls below this threshold, every SYN bound for the affected server is intercepted. For each SYN, PIX Firewall responds on behalf of the server with an empty SYN/ACK segment. PIX Firewall retains pertinent state information, drops the packet, and waits for the client's acknowledgment.

## FragGuard and Virtual Re-Assembly

FragGuard and Virtual Re-assembly is a feature that provides IP fragment protection. This feature performs full reassembly of all ICMP error messages and virtual-reassembly of the remaining IP fragments that are routed through the PIX Firewall. Virtual reassembly is currently enabled by default. This feature uses syslog to log any fragment overlapping and small fragment offset anomalies, especially those caused by a teardrop.c attack.

## DNS Control

The PIX Firewall identifies each outbound DNS (Domain Name Service) resolve request, and only allows a single DNS response. A host may query several servers for a response (in the case that the first server is slow in responding), but only the first answer to the request is allowed. All additional responses to the request are dropped by the firewall. This feature is always enabled.

## ActiveX Blocking

ActiveX controls, formerly known as OLE or OCX controls, are components that can be inserted into a web page or other application. The PIX Firewall ActiveX blocking feature blocks HTML <object> commands and comments them out of the HTML web page. As a technology, ActiveX creates many potential problems for the network clients including causing workstations to fail, introducing network security problems, being used to attack servers, or being used to host attacks against servers.

## Java Filtering

The Java Filtering feature lets you prevent Java applets from being downloaded by a system on a protected network. Java applets are executable programs that may be prohibited by some security policies because they can enable certain methods of attacking a protected network.

## URL Filtering

The PIX Firewall URL filtering is provided in partnership with the NetPartners Websense product. The PIX Firewall checks outgoing URL requests with the policy defined on the Websense server, which runs either on Windows NT or UNIX. Websense version 4 is supported in PIX Firewall version 5.3 and later.

PIX Firewall either permits or denies the connection, based on the response from the NetPartners Websense server. This server matches a request against a list of 17 website characteristics deemed inappropriate for business use. Because URL filtering is handled on a separate platform, no additional performance burden is placed on the PIX Firewall. For further information, refer to the following website:

<http://www.websense.com>

# Enabling Specific Protocols and Applications

This section describes the features provided by the PIX Firewall that control and enable the secure use of specific protocols and applications. It contains the following sections:

- [RIP Version 2](#)
- [Configurable Proxy Pinging](#)
- [Mail Guard](#)
- [Multimedia Support](#)
- [Cisco IP Telephony](#)
- [NETBIOS over IP](#)

## RIP Version 2

Routing Information Protocol (RIP) version 2 provides MD5 authentication of encryption keys. The PIX Firewall only listens in passive mode and/or broadcasts a default route. The PIX Firewall supports Cisco IOS software standards, which conform to RFC 1058, RFC 1388, and RFC 2082 of RIPv2 with text and keyed MD5 authentication. The PIX Firewall supports one key and key ID per interface. While the key has an infinite lifetime, for best security, you should change the key every two weeks or sooner.



**Note**

---

The use of Telnet to change the configuration may expose the key and key ID on the network.

---

## Configurable Proxy Pinging

The Configurable Proxy Pinging feature lets you control ICMP access to PIX Firewall interfaces. This feature shields PIX Firewall interfaces from detection by users on an external network.



**Note**

---

We recommend that you grant permission for ICMP unreachable message type (type 3). Denying ICMP unreachable messages, disables ICMP Path MTU discovery, which can halt IPsec and PPTP traffic.

---

## Mail Guard

The Mail Guard feature provides safe access for Simple Mail Transfer Protocol (SMTP) connections from the outside to an inside messaging server. This feature allows a single mail server to be deployed within the internal network without it being exposed to known security problems with some SMTP server implementations. Avoids the need for an external mail relay (or bastion host) system. Mail Guard enforces a safe minimal set of SMTP commands to avoid an SMTP server system from being compromised. This feature also logs all SMTP connections.

## Multimedia Support

The following paragraphs describe the features provided by the PIX Firewall that support multimedia applications:

- [Supported Multimedia Applications](#)
- [RAS Version 2](#)
- [RTSP](#)

### Supported Multimedia Applications

Users increasingly make use of a wide range of multimedia applications, many of which require special handling in a firewall environment. The PIX Firewall handles these without requiring client reconfiguration and without becoming a performance bottleneck. The specific multimedia applications supported by the PIX Firewall include the following:

- RealAudio
- Streamworks
- CU-SeeMe
- Internet Phone
- IRC
- Vxtreme
- VDO Live

**Note**

---

Support for specific protocols can be disabled using access-lists if required.

---

### RAS Version 2

The Registration, Admission, and Status (RAS) protocol is required by multimedia applications such as video conferencing and Voice over IP that require video and audio encoding. A RAS channel carries bandwidth change, registration, admission, and status messages (following the recommendations in H.225) between endpoints and gatekeepers. Multimedia applications use a large number of dynamically negotiated data and control channels to handle the various visual and auditory streams.

### RTSP

The PIX Firewall allows the secure forwarding of Real Time Streaming Protocol (RTSP) packets. RTSP is used by RealAudio, RealNetworks, Apple QuickTime 4, RealPlayer, and Cisco IP/TV connections. This feature lets the firewall handle multimedia applications including Cisco IP/TV connections.

**Note**

---

PIX Firewall does not yet have the ability to recognize HTTP cloaking where an RTSP message is hidden within an HTTP message. Also, RTSP is not supported with NAT.

---

## Cisco IP Telephony

Cisco IP Telephony allows the integration of VoIP (Voice over IP) networks and Public Switched Telephone Networks (PSTN). The transmission of voice traffic between internal and external voice networks requires support for the following protocols:

- [H.323](#)
- [SIP](#)

### H.323

The PIX Firewall supports the secure use of H.323 Version 2. H.323 is a suite of protocols defined by the International Telecommunication Union (ITU) for multimedia conferences over LANs. Some of the features provided include:

- Fast Connect or Fast Start Procedure for faster call setup
- H.245 tunneling for resource conservation, call synchronization, and reduced set up time
- Call redirection
- Conferencing—The conference is not established until both endpoints agree to participate

### SIP

The Session Initiation Protocol (SIP) enables call handling sessions—particularly two-party audio conferences, or “calls.” The PIX Firewall supports SIP VoIP gateways and VoIP proxy servers. It also supports definition using SDP for dynamically allocated UDP ports.

## NETBIOS over IP

The PIX Firewall supports NETBIOS over IP connections from the internal network to the external network. This allows Microsoft client systems on the internal network, possibly using NAT, to access servers, such as Windows NT, located on the external network. This allows security policies to encompass Microsoft environments across the Internet and inside an intranet. It allows the use of access controls native to the Microsoft environment.

## Creating a Virtual Private Network

This section introduces Virtual Private Network (VPN) technology and describes how this technology is implemented by the PIX Firewall. It contains the following topics:

- [What is a VPN?](#)
- [IPSec](#)
- [Internet Key Exchange \(IKE\)](#)
- [Certification Authorities](#)
- [Using a Site-to-Site VPN](#)
- [Using a Remote Access VPN](#)

## What is a VPN?

VPNs allow you to securely interconnect geographically distributed users and sites over the public Internet. VPNs can provide lower cost, improved reliability, and easier administration than traditional wide-area networks based on private Frame Relay or dial-up connections. VPNs maintain the same security and management policies as a private network. With a VPN, customers, business partners, and remote users, such as telecommuters, can access enterprise computing resources securely.

IPSec is a standard that defines vendor-independent methods of establishing a VPN. As part of its security functions, the PIX Firewall provides IPSec standards-based VPN capability. With IPSec, data can be transmitted across a public network without fear of observation, modification, or spoofing.

Site-to-site and remote-access VPNs are the two main types of VPN, both of which are supported by the PIX Firewall.

## IPSec

IPSec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices (peers), such as PIX Firewall units.

IPSec provides the following network security services:

- Data Confidentiality—The IPSec sender can encrypt packets before transmitting them across a network.
- Data Integrity—The IPSec receiver can authenticate packets sent by the IPSec sender to ensure that the data has not been altered during transmission.
- Data Origin Authentication—The IPSec receiver can authenticate the source of the IPSec packets sent. This service is dependent upon the data integrity service.
- Anti-Replay—The IPSec receiver can detect and reject replayed packets.

**Note**

---

The term data authentication is generally used to mean data integrity and data origin authentication. Within this chapter, it also includes anti-replay services, unless otherwise specified.

---

IPSec provides secure tunnels between two peers, such as two PIX Firewall units. You define which packets are considered sensitive and should be sent through these secure tunnels, and you define the parameters that should be used to protect these sensitive packets, by specifying the characteristics of these tunnels. Then, when the IPSec peer sees such a sensitive packet, it sets up the appropriate secure tunnel and sends the packet through the tunnel to the remote peer. The secure tunnel used to transmit information is based on encryption keys and other security parameters, described by security associations (SAs).

## Internet Key Exchange (IKE)

The process by which IPsec can automatically establish a secure tunnel is divided into two phases:

- Phase 1—This phase, implemented through the Internet Key Exchange (IKE) protocol, establishes a pair of IKE SAs. IKE SAs are used for negotiating one or more IPsec SAs, which are used for the actual transmission of application data.
- Phase 2—This phase uses the secure channel provided by the IKE SAs to negotiate the IPsec SAs. At the end of this phase both peers have established a pair of IPsec SAs, which provide the secure tunnel used for transmission of application data. One of the SA parameters is its lifetime, which enhances IPsec security by causing the SA to automatically expire after a configurable length of time.

The IKE protocol establishes a secure tunnel for negotiating IPsec SAs. It allows you to implement IPsec without manual configuration of every IPsec peer. Manual configuration of IPsec peers becomes prohibitively complicated as the number of peers increase, because each peer requires a pair of SAs for every other peer with which it communicates using IPsec.

Like IPsec, IKE uses a pair of SAs to establish a secure tunnel for communication between two peers. However, IKE uses its SAs to securely negotiate SAs for IPsec tunnels, rather than for the transmission of user information.

You can manually configure SAs to establish an IPsec tunnel between two peers. However, this method is not as secure, because manually configured SAs do not automatically expire. In addition, a severe problem of scalability occurs as the number of peers increases. A new pair of SAs is required on each existing peer whenever you add a peer that uses IPsec to your network. For this reason, manual configuration is only used when the remote peer does not support IKE.

IKE SAs can be established by using pre-shared keys, in a way similar to manual configuration of IPsec SAs. This method, however, suffers from the same problems of scalability that affects manual configuration of IPsec SAs. A certification authority (CA) provides a scalable method to share keys for establishing IKE SAs.

## Certification Authorities

Understanding how CAs help to configure IKE requires understanding something about public/private key encryption. Public/private keys, also called asymmetric keys, are a pair of keys with the property that data encrypted with one key can only be unencrypted using the other key. This property has been used to solve the scalability problem encountered when sharing secrets over a non-secure network.

After generating a public/private key pair, one key is kept secret (the private key) and the other key is made easily available (the public key). When any peer needs to share a secret with the owner of the private key, it simply encrypts the information using the public key. The only way to unencrypt the original information is by using the private key. Using this method, encrypted information can be shared over a non-secure network without transmitting the secret key required to decipher the encrypted information.

This unique property of public/private key pairs also provides an excellent method of authentication. A public key only unencrypts a message encrypted with the corresponding private key. If a message can be read using a given public key, you know for certain that the sender of the message owns the corresponding private key.

This is where the CA comes in. A public key certificate, or digital certificate, is used to associate a public/private key pair with a given IP address or host name. A certification authority (CA) issues public key certificates for a specific period of time. A CA can be a private (in-house) CA, run by your own organization, or a public CA. A public CA, like VeriSign, is operated by a third-party that you trust to validate the identity of each client or server to which it issues a certificate.

Digital certificates are used by the IKE protocol to create the first pair of SAs, which provide a secure channel for negotiating the IPSec SAs. To use certificates for negotiating IKE SAs, both IPSec peers have to generate public/private key pairs, request and receive public key certificates, and be configured to trust the CA that issues the certificates.

Most browsers, by default, trust certificates from well-known CAs, such as VeriSign, and provide options for adding CAs, and for generating and requesting a digital certificate. You can also preconfigure browser software before it is distributed to users with your CA and the necessary certificates.

The procedure for configuring PIX Firewall to use IKE with digital certificates is described in [“Using Certification Authorities”](#) in [Chapter 4, “Basic VPN Configuration.”](#)

## Using a Site-to-Site VPN

Site-to-site VPNs are an alternative WAN infrastructure that replace and augment existing private networks using leased lines, Frame Relay, or ATM to connect remote and branch offices and central site(s). For site-to-site VPNs, the PIX Firewall can interoperate with any Cisco VPN-enabled network device, such as a Cisco VPN router.

Site-to-site VPNs are established between the PIX Firewall and a remote IPSec security gateway. The remote IPSec security gateway can be a PIX Firewall, a Cisco VPN concentrator or VPN-enabled router, or any IPSec-compliant third-party device. For configuration instructions, refer to [Chapter 4, “Basic VPN Configuration,”](#) and for example configurations, refer to [Chapter 5, “Site-to-Site VPN Configuration Examples.”](#)

## Using a Remote Access VPN

The PIX Firewall supports mixed VPN deployments, including both site-to-site and remote-access traffic. A remote access VPN uses analog, dial, ISDN, DSL, mobile IP, and cable technologies to securely connect mobile users, telecommuters, and other individual systems to a network protected by the PIX Firewall. Use one of the following Cisco remote access VPN applications to gain access into a PIX Firewall-protected network:

- Cisco Secure VPN Client, version 1.1 or later
- Cisco VPN 3000 Client, version 2.5 or later
- Cisco VPN Client, version 3.0

**Note**

---

We strongly suggest that you use the Cisco VPN Client version 3.0.

---

For general configuration instructions, refer to [Chapter 4, “Basic VPN Configuration.”](#) For more specific procedures and example configurations, refer to [Chapter 6, “Configuring VPN Client Remote Access.”](#)

# PIX Firewall System Management

This section describes the features and tools available for managing the PIX Firewall. It contains the following sections:

- [PIX Device Manager](#)
- [Telnet Interface](#)
- [SSH Version 1](#)
- [Using SNMP](#)
- [TFTP Configuration Server](#)
- [XDMCP](#)
- [Using a Syslog Server](#)
- [FTP and URL Logging](#)
- [Integration with IDS](#)

## PIX Device Manager

The Cisco PIX Device Manager (PDM) is a browser-based configuration tool that lets you set up, configure, and monitor your PIX Firewall from a graphic user interface (GUI), without any extensive knowledge of the PIX Firewall command-line interface (CLI). PDM provides a management interface from Windows NT, Windows 95, Windows 2000, or Solaris web browsers. PDM limits access to the HTML interface to specified client systems within the inside network (based on source address) and is password protected.

## Telnet Interface

The PIX Firewall Telnet interface provides a command-line interface similar to Cisco IOS software. The Telnet interface lets you remotely manage the PIX Firewall via the console interface. The Telnet interface limits access of the Telnet interface to specified client systems within the inside network (based on source address) and is password protected. If the inside network is not secure and sessions on the LAN can be snooped, you should limit use of the Telnet interface. If IPSec is configured, you can also access the PIX Firewall console from the outside interface.

## SSH Version 1

PIX Firewall supports the SSH remote shell functionality as provided in SSH version 1. SSH allows secure remote configuration of a PIX Firewall, providing encryption and authentication capabilities.

## Using SNMP

The PIX Firewall provides support for network monitoring using Simple Network Management Protocol (SNMP). The SNMP interface allows you to monitor the PIX Firewall through traditional network management systems. The PIX Firewall only supports the SNMP GET command, which allows read-only access.

The SNMP Firewall and Memory Pool MIBs extend the number of traps you can use to discover additional information about the state of the PIX Firewall, including the following events:

- Buffer usage from the **show block** command
- Connection count from the **show conn** command
- Failover status
- Memory usage from the **show memory** command

## TFTP Configuration Server

You can use a Trivial File Transfer Protocol (TFTP) configuration server to obtain configuration for multiple PIX Firewalls from a central source. However, TFTP is inherently insecure so you should not use it over networks where sharing privileged information in clear text is a violation of your network security policy.

You can also use TFTP to download a .bin image from cisco.com to a PIX Firewall to upgrade or replace the software image on the PIX Firewall. TFTP does not perform any authentication when transferring files, so a username and password on the remote host are not required.

## XDMCP

The PIX Firewall supports connections using XDMCP (X Display Manager Control Protocol) using the **established** command. This feature uses an XWindows TCP back connection fixup, which negotiates an XWindows session and creates an embryonic connection at destination port 6000. XDMCP handling is enabled by default, which is the same as other UDP fixups.

## Using a Syslog Server

The PIX Firewall sends messages in TCP and UDP Syslog messages to any existing syslog server and provides a syslog server for use on a Windows NT system. The Windows NT Syslog server can provide time-stamped syslog messages, accept messages on alternate ports, and be configured to stop PIX Firewall traffic if messages cannot be received. You can also configure the Windows NT Syslog server to stop PIX Firewall connections if the Windows NT log disk fills or if the server goes down.

## FTP and URL Logging

The FTP and URL logging feature lets you view inbound and outbound FTP commands entered by your users as well as the URLs they use to access other sites. You can use this feature to monitor user access of internal and external sites. It provides data you can use to block access to problem sites. You enable this feature with the **logging trap debugging** command statement. Note that this feature can generate a huge amount of syslog data on a high-traffic PIX Firewall.

## Integration with IDS

The PIX Firewall is interoperable with the Cisco Intrusion Detection System. The PIX Firewall traps IDS signatures and sends these as syslog messages the Syslog server. This feature supports only single-packet IDS signatures.

# PIX Firewall Failover

The PIX failover feature lets you connect two identical PIX Firewall units with a special failover cable to achieve a fully redundant firewall solution. [Table 1-1](#) summarizes the support for the failover feature provided by different PIX Firewall models.

**Table 1-1 Support for Failover**

PIX Firewall Model	Support for Failover
PIX Firewall 501	Not supported
PIX Firewall 506	Not supported
PIX Firewall 515	Requires additional license
PIX Firewall 525	Ships with full support
PIX Firewall 535	Ships with full support

When implementing failover, one unit functions as the active unit, while the other assumes the role of the standby unit. Both units need the same configuration and run the same software version. The failover cable that connects two PIX Firewall units allows the two units to synchronize configuration and session state information so that if the active unit fails, the standby unit can assume its role without any interruption in network connectivity or security. To configure the PIX Firewall failover feature, refer to [Chapter 8, “Using PIX Firewall Failover.”](#)

## Where to Go from Here

- To complete the basic configuration required regardless of how you implement your PIX Firewall, refer to [Chapter 2, “Basic Firewall Configuration.”](#)
- To allow or restrict specific types of network activity and access, refer to [Chapter 3, “Managing Network Access and Use.”](#)
- To perform basic VPN configuration, refer to [Chapter 4, “Basic VPN Configuration.”](#)
- To configure or use PIX Firewall system management tools, refer to [Chapter 7, “PIX Firewall System Management.”](#)
- To configure the PIX Firewall failover feature, refer to [Chapter 8, “Using PIX Firewall Failover.”](#)
- To upgrade the software image on your PIX Firewall, refer to [Chapter 9, “Upgrading PIX Firewall Software.”](#)
- To record the information required to implement your PIX Firewall, refer to [Appendix A, “Firewall Configuration Forms.”](#)

For more information on firewalls, refer to:

- Bernstein, T., Bhimani, A.B., Schultz, E. and Siegel, C. A. *Internet Security for Business*. Wiley. Information about this book is available at: <http://www.wiley.com>
- Chapman, D. B. & Zwicky, E. D. *Building Internet Firewalls*. O’Reilly. Information on this book is available at: <http://www.ora.com/>
- Cheswick, W. and Bellovin, S. *Firewalls & Internet Security*. Addison-Wesley. Information about this book is available at: <http://www.aw.com>

- Garfinkel, S. and Spafford, G. *Practical UNIX Security*. O'Reilly. Information about this book is available at: <http://www.ora.com/>
- Stevens, W. R. *TCP/IP Illustrated, Volume 1 The Protocols*. Addison-Wesley. Information about this book is available at: <http://www.aw.com>
- Cisco's Products and Technologies information on PIX Firewall is available at: <http://www.cisco.com/en/US/products/sw/secursw/ps2120/index.html>