



Getting Started

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](#)
- [Supporting Specific Protocols and Applications](#)
- [Creating a Virtual Private Network](#)
- [Using PIX Firewall in a Small Office, Home Office Environment](#)
- [Accessing and Monitoring PIX Firewall](#)
- [PIX Firewall Failover](#)
- [Upgrading the PIX Firewall OS and License](#)
- [Using the Command-Line Interface](#)
- [Before You Start Configuring PIX Firewall](#)
- [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](#)
- [Address Translation](#)
- [Cut-Through Proxy](#)
- [Access Control](#)

Chapter 2, “Establishing Connectivity” provides configuration instructions for controlling network connectivity through the PIX Firewall. Chapter 3, “Controlling Network Access and Use” provides configuration instructions for using the PIX Firewall to control network connectivity.

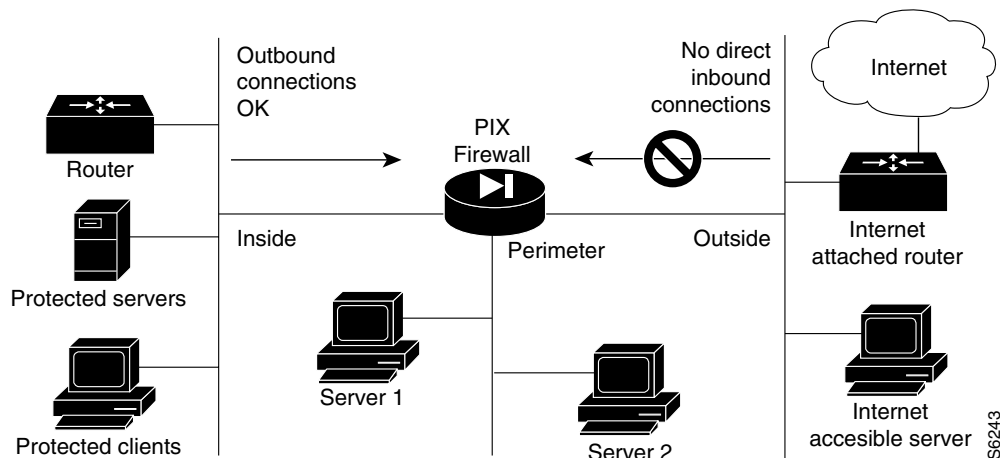
How the PIX Firewall Works

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 inside network, but more restricted than access to the outside 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. Traffic may not exit the PIX Firewall on the same network interface it entered. 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.

For PIX Firewall models with three or more interfaces, 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 501, PIX 506/506E each have two network interfaces, so all systems must 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.



Note

The PIX Firewall checks TCP sequence number and ensures that it fits within an acceptable range.

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.
- Traffic may not exit the PIX Firewall on the same network interface it entered.
- 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 the 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.

For more information about how ASA works and how you can configure application inspection with different types of applications, refer to [Chapter 4, “Configuring Application Inspection \(Fixup\).”](#)

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.



Note

Traffic may not exit the PIX Firewall on the same network interface it entered. This condition results in the following message in the system log:

```
%PIX-7-106011: Deny inbound (No xlate) chars
```

Explanation This is a connection-related message. This message occurs when a packet is sent to the same interface that it arrived on. This usually indicates that a security breach is occurring. When the PIX Firewall receives a packet, it tries to establish a translation slot based on the security policy you set with the global and conduit commands, and your routing policy set with the route command.

Address Translation

The Network Address Translation (NAT) feature works by substituting, or translating, host addresses on one interface with a “global address” associated with another 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.

**Note**

With version 6.2 of the PIX Firewall, NAT is also available for translating outside addresses. This helps to simplify network routing by controlling the addresses that can appear on the inside network.

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 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 lets you 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. A proxy server analyzes every packet at layer seven of the OSI model, which is a time- and processing-intensive function. By contrast, the PIX Firewall uses cut-through proxy to authenticate a connection and then allow traffic to flow quickly and directly.

Cut-through proxy allows a much finer level of administrative control over connections than checking source IP addresses. It allows security policies to be enforced based on individual user accounts. Connections can be authenticated with a user ID and password before are established, and one-time dynamic passwords or security tokens are supported for greater security. Authentication and authorization are supported for HTTP, Telnet, or FTP connections.

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](#)
- [TurboACL](#)
- [Downloadable ACLs](#)
- [Object Grouping](#)
- [Conduits](#)

Chapter 3, “Controlling Network Access and Use” provides configuration instructions for using the features mentioned in this section.

AAA Integration

PIX Firewall provides integration with AAA (authentication, accounting, and authorization) services. AAA services are provided by Terminal Access Controller Access Control System Plus (TACACS+) or Remote Authentication Dial-In User Service (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 will apply an access list for the user that was returned by the AAA server using the Cisco **acl** attribute (**acl=<acl_name>**).

For additional information about configuring AAA servers for use with the PIX Firewall see Authentication and Command Authorization for PIX at the following URL:

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_tech_note09186a00800949d6.shtml

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.

TurboACL

A feature called TurboACL is introduced in PIX Firewall version 6.2 that improves the way that the PIX Firewall processes large access control lists. The method by which the PIX Firewall searches for an access list entry has been improved to reduce the time spent searching large access lists. The implementation of TurboACL in PIX Firewall version 6.2 supports access lists with up to 16,000 access list entries.

Downloadable ACLs

When used with a AAA server, PIX Firewall lets you create access lists that control connections on a per-user basis. Creating per-user access lists requires creating a user profile for the user on a RADIUS server. In previous versions of PIX Firewall, you also had to configure an access list for each user locally on each PIX Firewall. Beginning with PIX Firewall version 6.2, the required per-user access list is downloaded from the AAA server based on the user profile. No additional access list configuration is required on any PIX Firewall. This new feature greatly reduces the complexity and improves the scalability of per-user access lists.

Object Grouping

Object Grouping is another new feature in PIX Firewall version 6.2 that reduces the complexity of configuration and improves scalability for large or complex networks. Object Grouping lets you apply access rules to logical groups of network objects. When you apply a PIX Firewall command to an object group, the command affects all network objects defined within the group. This can reduce a very large number of access rules to a manageable number, which reduces time spent configuring and troubleshooting access rules in large or complex networks.

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, TCP/UDP port, and IP 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](#)
- [Flood Defender](#)
- [FragGuard and Virtual Re-Assembly](#)
- [DNS Control](#)
- [ActiveX Blocking](#)
- [Java Filtering](#)
- [URL Filtering](#)

For more information about the PIX Firewall features used to protect your network against specific attacks, refer to [Chapter 4, “Configuring Application Inspection \(Fixup\).”](#) For information about configuring ActiveX Blocking, Java Filtering, and URL Filtering, refer to the [“Filtering Outbound Connections”](#) section on page 3-26.

For information about features that allow using specific protocols and applications across the firewall, refer to [“Supporting 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 attack.

DNS Control

The PIX Firewall identifies each outbound DNS (Domain Name System) 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

You can use access control lists to prevent outbound access to specific web sites, but configuring and managing web usage this way is not very practical because of the size and dynamic nature of the Internet. The recommended solution is to use the PIX Firewall in conjunction with a separate server running one of the following Internet filtering products:

- Websense Enterprise web filtering application (supported by PIX Firewall version 5.3 and higher)
- Filtering by N2H2 for IFP-Enabled Devices (supported by PIX Firewall version 6.2)

Compared to using access control lists, this reduces the administrative task and improves filtering effectiveness. Also, because URL filtering is handled on a separate platform, the performance of the PIX Firewall is much less affected.

The PIX Firewall checks outgoing URL requests with the policy defined on the URL filtering server. PIX Firewall either permits or denies the connection, based on the response from the filtering server.

For further information, refer to either of the following websites:

<http://www.websense.com>

<http://www.n2h2.com>

**Note**

PIX Firewall version 6.2 also adds support for filtering of long URLs, such as those generated by search engines.

Supporting Specific Protocols and Applications

This section describes how the PIX Firewall enables the secure use of specific protocols and applications. It includes the following sections:

- [How Application Inspection Works](#)
- [Mail Guard](#)
- [Multimedia Applications](#)
- [RAS Version 2](#)
- [RTSP](#)
- [Voice over IP](#)
- [LDAP Version 2 and ILS](#)
- [NetBIOS over IP](#)
- [RIP Version 2](#)
- [Configurable Proxy Pinging](#)
- [Forwarding Multicast Transmissions](#)

For further information about application inspection and how it works with different applications, refer to [Chapter 4, “Configuring Application Inspection \(Fixup\).”](#)

How Application Inspection Works

The behavior of certain Internet applications, such as FTP or multimedia applications, requires PIX Firewall to make some adjustments to how it performs NAT or PAT, and for the ports it opens to receive replies to outbound requests for services. Application inspection provides PIX Firewall with the information it needs to make these adjustments.

As described in the “[Address Translation](#)” section, PIX Firewall applies NAT or PAT to the source address of IP packets from hosts for which it is enabled. However, “badly behaved” applications create IP packets with network addresses and other information in the user data portion of the packet. If this information is left unchanged, the application will not work because the address in the source address field will not match the address embedded in the user data field.

To solve this problem, when NAT or PAT is applied to these packets, the application inspection function helps the PIX Firewall find the extra address information so address translation can be applied to it. After changing this addressing information, the PIX Firewall uses application inspection to adjust other fields in the packet that are affected, such as those containing packet length and checksum information.

By default, the PIX Firewall allows replies to outbound requests using many Internet applications, such as HTTP. These services send requests and replies on well-known TCP ports.

However, some applications, such as FTP, use a well-known TCP port to negotiate the use of secondary ports, which are used for the actual exchange of user data. To support the secure use of these applications, PIX Firewall must monitor the negotiation that occurs on the first port to determine on which port replies will be received. Again, it is application inspection that provides the information required to identify and open ports required to receive replies from these applications.

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. This eliminates 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 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
- Intel Internet Phone
- IRC
- Vxtreme
- VDO Live



Note

Traffic using specific protocols can be prevented using access lists.

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.

Voice over IP

This section describes the support provided by the PIX Firewall for the transmission of Voice over IP (VoIP) traffic and includes the following topics:

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

**Note**

Version 6.2 of the PIX Firewall adds PAT support for H.323 and SIP. This helps to expand your address space to accommodate the large number of endpoints involved when implementing VoIP networks.

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 the following:

- 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

SCCP

Skinny (or Simple) Client Control Protocol (SCCP) is a simplified protocol used in VoIP networks. Secure handling of this protocol is required when using Cisco CallManager, Cisco IP Phones, and other Cisco IP Telephony products.

When coupled with an H.323 Proxy, an SCCP client can interoperate with H.323 compliant terminals. Application inspection in the PIX Firewall works with SCCP version 3.1.1. The functionality of PIX Firewall application inspection ensures that all SCCP signalling and media packets can traverse the Firewall by providing NAT of the SCCP signaling packets.

SIP

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.

LDAP Version 2 and ILS

PIX Firewall version 6.2 adds support for using NAT with Lightweight Directory Access Protocol (LDAP) version 2, used by the Internet Locator Service (ILS). Applications that depend on ILS include Microsoft NetMeeting and SiteServer Active Directory. These applications use ILS to provide registration and location of end points in the ILS directory.

Earlier versions of PIX Firewall supported NetMeeting, but did not provide support for using NAT with ILS. With the addition of NAT support for LDAP version 2, PIX Firewall supports NAT for H.323 sessions establishment by NetMeeting.

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.

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 3. Denying ICMP unreachable messages, disables ICMP Path MTU discovery, which can halt IPsec and PPTP traffic.

Forwarding Multicast Transmissions

The Internet Group Management Protocol (IGMP) is used to dynamically register specific hosts in a multicast group on a particular LAN with a multicast (MC) router. MC routers efficiently route multicast data transmissions to the hosts on each LAN in an internetwork that are registered to receive specific multimedia or other broadcasts.

PIX Firewall version 6.2 introduces the Stub Multicast Routing (SMR) feature. SMR allows the PIX Firewall to function as a “stub router,” which is a device that acts as an IGMP proxy agent. A stub router does not operate as a full MC router, but simply forwards IGMP messages between hosts and MC routers.

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](#)

For basic configuration instructions for using IPSec to create a VPN, refer to [Chapter 6, “Configuring IPSec and Certification Authorities.”](#) For configuration instructions and examples to establish site-to-site VPNs and using certification authorities, refer to [Chapter 7, “Site-to-Site VPN Configuration Examples.”](#) For configuration examples and instructions for creating a remote access VPN, refer to [Chapter 8, “Configuring VPN Client Remote Access.”](#)

What is a VPN?

VPNs let you 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 lets you 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 created in pairs. Data encrypted with one key of this pair can only be unencrypted using the other key. One key is kept secret (called a 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 6, “Configuring IPSec and Certification Authorities.”](#)

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 small office, home office (SOHO) environments. 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 6, “Configuring IPsec and Certification Authorities,”](#) and for example configurations, refer to [Chapter 7, “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 higher
- Cisco VPN Client version 3.x

**Note**

We strongly suggest that you use Cisco VPN Client version 3.x.

PIX Firewall can function as an Easy VPN Server in relation to an Easy VPN Remote device, such as a PIX 501 or PIX 506/506E, or in relation to Cisco VPN software clients. When used as an Easy VPN Remote device, the PIX Firewall can push VPN configuration to the VPN client or Easy VPN Remote device, which greatly simplifies configuration and administration.

Using PIX Firewall in a Small Office, Home Office Environment

This section describes features provided by the PIX Firewall that support its use in a small office, home office (SOHO) environment. It includes the following topics:

- [Using the PIX Firewall as an Easy VPN Remote Device](#)
- [PPPoE](#)
- [DHCP Server](#)
- [DHCP Client](#)

For information about configuring the features in this section, refer to [Chapter 5, “Using PIX Firewall in SOHO Networks.”](#)

Using the PIX Firewall as an Easy VPN Remote Device

PIX Firewall version 6.2 lets you use PIX Firewall as an Easy VPN Remote device when connecting to a Easy VPN Server, such as an Cisco VPN 3000 Concentrator or a PIX Firewall. This functionality, sometimes called a “hardware client,” allows the PIX Firewall to act as a VPN client when initiating a secure connection to an Easy VPN Server. After completing this connection, hosts running on the LAN behind the PIX Firewall can connect through the Easy VPN Server without individually running any VPN client software.

PPPoE

Point-to-Point Protocol over Ethernet (PPPoE) combines two widely accepted standards, Ethernet and the Point-to-Point Protocol (PPP), to provide an authenticated method of assigning IP addresses to client systems. PPPoE clients are typically personal computers connected to an ISP over a remote broadband connection, such as DSL or cable service. ISPs deploy PPPoE because it provides a method of supporting high-speed broadband access using the existing remote access infrastructure and that provides superior ease of use to customers.

PIX Firewall version 6.2 introduces PPPoE client functionality. This allows small office, home office (SOHO) users of the PIX Firewall to connect to ISPs using DSL modems.

**Note**

The PIX Firewall PPPoE client can only be enabled on the outside interface.

By using PPPoE, ISPs can deploy DSL without changing their existing infrastructure, which is typically based on the use of PPP over dial-up connections.

DHCP Server

Dynamic Host Configuration Protocol (DHCP) is a protocol that supplies automatic configuration parameters to Internet hosts. This protocol has two components:

- Protocol for delivering host-specific configuration parameters from a DHCP server to a host (DHCP client)
- Mechanism for allocating network addresses to hosts

A DHCP server is simply a computer that provides configuration parameters to a DHCP client, and a DHCP client is a computer or network device that uses DHCP to obtain network configuration parameters.

When functioning as a DHCP server, the PIX Firewall dynamically assigns IP addresses to DHCP clients from a pool of designated IP addresses.

**Note**

The PIX Firewall DHCP server can only be enabled on the inside interface.

PIX Firewall version 6.2 adds support for DHCP option 66 and DHCP option 150 requests. This allows DHCP clients, such as Cisco IP Phones, to obtain the address of a designated TFTP server. Cisco IP Phones typically obtain the configuration information required to connect to a Cisco CallManager server from a TFTP server. A DHCP Option 66 request causes the DHCP server to provide the address of a single TFTP server; an Option 150 request obtains a list of TFTP servers.

DHCP Client

DHCP client support within the PIX Firewall is designed for use within a small office, home office (SOHO) environment using a PIX Firewall that is directly connected to a DSL or cable modem that supports the DHCP server function.

**Note**

The PIX Firewall DHCP client can only be enabled on the outside interface.

With the DHCP client feature enabled on a PIX Firewall, the PIX Firewall functions as a DHCP client to a DHCP server allowing the server to configure the outside interface with an IP address, subnet mask, and optionally a default route.

**Note**

Use of the DHCP client feature to acquire an IP address from a generic DHCP server is not supported. Also, the PIX Firewall DHCP client does not support failover configurations.

Accessing and Monitoring PIX Firewall

This section describes how you access and monitor the PIX Firewall system. It contains the following topics:

- [Cisco PIX Device Manager \(PDM\)](#)
- [Command Authorization](#)
- [Telnet Interface](#)
- [SSH Version 1](#)
- [NTP](#)
- [Auto Update](#)
- [Capturing Packets](#)
- [Using SNMP](#)
- [XDMCP](#)
- [Using a Syslog Server](#)
- [FTP and URL Logging](#)
- [Integration with Cisco IDS](#)

For information about configuring the features described in this section, refer to [Chapter 9, “Accessing and Monitoring PIX Firewall.”](#)

Cisco PIX Device Manager (PDM)

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 graphical 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 access is password protected, uses Secure Socket Layer (SSL) for encryption, and restricts access to client systems with designated IP addresses.

Command Authorization

PIX Firewall version 6.2 introduces a more flexible method of authenticating and authorizing administrative access to the PIX Firewall. Similar to Cisco IOS software command authorization, PIX Firewall now supports up to 16 privilege levels to be assigned to CLI commands. You can create

user accounts or login contexts tied to these privilege levels either locally or using a TACACS+ server. Additional information is also now provided regarding the usage of CLI commands, such as command tracing by means of syslog messages.

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.

NTP

PIX Firewall version 6.2 allows the PIX Firewall to function as a client for Network Time Protocol (NTP) version 3.0 servers. As an NTP client, the PIX Firewall can synchronize its time to a set of distributed time servers operating in a self-organizing, hierarchical configuration. A precisely coordinated time is required for validating certificate revocation lists (CRLs) when implementing a VPN using Public Key Infrastructure (PKI). A more precise time also improves the accuracy of log entries used for troubleshooting or monitoring security threats.

Auto Update

Auto Update is a protocol specification introduced with PIX Firewall version 6.2. This specification provides the infrastructure necessary for remote management applications to download PIX Firewall configurations, software images, and to perform basic monitoring from a centralized location.

Capturing Packets

PIX Firewall version 6.2 introduces an enhanced and improved packet capture capability that lets you capture packets, including ARP packets, to a linear buffer. You can use access lists to define packets to capture on specific interfaces of the PIX Firewall. You can then display the captured packets on any console or transfer the contents of the packet capture buffer to a TFTP server.

Using SNMP

The PIX Firewall provides support for network monitoring using Simple Network Management Protocol (SNMP). The SNMP interface lets you 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

**Note**

PIX Firewall version 6.2 introduces support for monitoring CPU utilization through SNMP. This feature allows network administrators to monitor the PIX Firewall CPU usage using SNMP management software, such as HP OpenView, for capacity planning. This CPU usage information is the same as that shown by the **show cpu usage** command.

XDMCP

The PIX Firewall supports connections using XDMCP (X Display Manager Control Protocol) using the **established** command. This feature negotiates an XWindows session and creates an embryonic connection at destination port 6000. XDMCP handling is enabled by default, like other UDP application inspection functions.

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 Cisco IDS

The PIX Firewall is interoperable with the Cisco Intrusion Detection System (Cisco IDS). 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 Firewall failover feature lets you connect two identical PIX Firewall units with a special failover cable to achieve a fully redundant firewall solution.

To configure the PIX Firewall failover feature, refer to [Chapter 10, “Using PIX Firewall Failover.”](#) For instructions about upgrading failover from a previous version, refer to “[Upgrading Failover Systems from a Previous Version](#)” in [Chapter 11, “Changing Feature Licenses and System Software.”](#)

[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 501	Not supported
PIX 506/506E	Not supported
PIX 515/515E	Requires additional license
PIX 525	Ships with full support
PIX 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 require the same configuration and run the same software version.

PIX Firewall version 6.2 adds support for failover between two units connected over a dedicated Ethernet interface (LAN-based failover). LAN-based failover eliminates the need for a special failover cable and overcomes the distance limitations imposed by the failover cable required to implement failover on earlier versions of PIX Firewall.

With failover, two PIX Firewall units 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.

Upgrading the PIX Firewall OS and License

The PIX Firewall software is a specialized, hardened operating system that is continuously being improved to provide greater performance, security, and interoperability with Internet devices and applications. For information about obtaining and installing the latest software release, refer to [Chapter 11, “Changing Feature Licenses and System Software.”](#)

With PIX Firewall version 6.2, you can now upgrade your license without reinstalling the operating system software. A new CLI command has been added to let you upgrade your activation key from the command-line interface without reinstalling the software image and without entering monitor mode. For detailed instructions, refer to [Chapter 11, “Changing Feature Licenses and System Software.”](#)

You can use a Trivial File Transfer Protocol (TFTP) configuration server to obtain configuration for multiple PIX Firewall units 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 CCO 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.

Using the Command-Line Interface

This section includes the following topics, which describe how to use the PIX Firewall command-line interface (CLI):

- [Access Modes](#)
- [Accessing Configuration Mode](#)
- [Abbreviating Commands](#)
- [Backing Up Your PIX Firewall Configuration](#)
- [Command Line Editing](#)
- [Command Output Paging](#)
- [Comments](#)
- [Configuration Size](#)
- [Help Information](#)
- [Viewing the Default Configuration](#)
- [Resetting the Default Configuration](#)
- [Clearing and Removing Configuration Settings](#)



Note

The PIX Firewall CLI uses similar syntax and other conventions to the Cisco IOS CLI, but the PIX Firewall operating system is not a version of Cisco IOS software. Do not assume that a Cisco IOS CLI command works or has the same function with the PIX Firewall.

Access Modes

PIX Firewall version 6.2 introduces support for up to 16 levels of command authorization. This is similar to what is available with Cisco IOS software. With this feature, you can assign specific PIX Firewall commands to one of 16 levels. You can either assign separate passwords for each privilege level or perform authentication using a local or remote AAA database of user accounts.

For information about configuring this feature, refer to the “[Command Authorization and LOCAL User Authentication](#)” section in [Chapter 9, “Accessing and Monitoring PIX Firewall.”](#)

The PIX Firewall provides five administrative access modes:

- Unprivileged mode—Available without entering a password, when you first access the PIX Firewall. In this mode, the PIX Firewall displays the “>” prompt and lets you enter a small number of commands. In PIX Firewall version 6.2, by default, commands in this mode are mapped to privilege Level 0.

- Privileged mode—Displays the “#” prompt and lets you change configuration information. Any unprivileged command also works in privileged mode. Use the **enable** command to start privileged mode and the **disable**, **exit**, or **quit** commands to exit.
In PIX Firewall version 6.2, by default, all privileged mode commands are mapped to privilege Level 15. You can assign enable passwords to other privilege levels and reassign specific commands to each level.
- Configuration mode—Displays the prompt `<pix_name>(config)#`, where *pixname* is the host name assigned to the PIX Firewall. You use configuration mode to change system configuration. All privileged, unprivileged, and configuration commands work in this mode. Use the **configure terminal** command to start configuration mode and the **exit** or **quit** commands to exit.
- Subcommand mode—Displays the prompt `<pix_name>(config-<main_cmd_name>)#`, where *pixname* is the host name assigned to the PIX Firewall and *main_cmd_name* is the Object Grouping command used to enter subcommand mode. Object Grouping is a way to simplify access control by letting you apply access control statements to groups of network objects, such as protocols or hosts. For further information about enabling and using this mode, refer to the “[Simplifying Access Control with Object Grouping](#)” section in Chapter 3, “Controlling Network Access and Use.”
- Monitor mode—This is a special mode that enables you to update the image over the network. While in the monitor mode, you can enter commands specifying the location of the TFTP server and the binary image to download. For information about using monitor mode to upgrade your PIX Firewall software, refer to [Chapter 11, “Changing Feature Licenses and System Software.”](#)

Accessing Configuration Mode

Perform the following steps to access the PIX Firewall configuration mode:

-
- Step 1** Start your terminal emulation program.
- Step 2** Power on the PIX Firewall. On newer models, the switch is at the back, on older models, at the front.
- Step 3** If you are configuring a PIX 506/506E, PIX 515/515E, PIX 525, or PIX 535 and your site downloads configuration images from a central source with TFTP, look for the following prompt in the startup messages:
- Use BREAK or ESC to interrupt flash boot.
- PIX Firewall displays this prompt for 10 seconds. To download an image, press the **Escape** key to start boot mode. If you are not downloading an image, ignore the prompt or press the Space bar to start immediately and PIX Firewall starts normally.
- Step 4** After the startup messages appear, you are prompted with the following unprivileged mode prompt:
- ```
pixfirewall>
```
- Enter the following command:
- ```
enable privilegelevel
```
- Replace *privilegelevel* with a number from 0 to 15, indicating the privilege level to which you require access. If you omit this parameter, the system assumes you are seeking access to privilege Level 15.
- With PIX Firewall version 6.2, you can configure up to fifteen different enable passwords for different privilege levels. By default, all commands are assigned to Level 0 or Level 15, and only Level 15 is preconfigured with a password.
- Step 5** The following prompt appears:

Password:

Press the **Enter** key.

- Step 6** You are now in privilege Level 15, which lets you use all the commands assigned to this privilege level. The following prompt appears:

```
pixfirewall#
```

Type **configure terminal** and press **Enter**. You are now in configuration mode.



Note

If the Command Authorization feature (new in PIX Firewall version 6.2) is enabled, the commands you are permitted to enter are determined by the administrative privilege level to which your user account has been assigned. For information about configuring this feature, refer to the “[Command Authorization and LOCAL User Authentication](#)” section in [Chapter 9, “Accessing and Monitoring PIX Firewall.”](#)

Abbreviating Commands

You can abbreviate most commands down to the fewest unique characters for a command; for example, you can enter **wr t** to view the configuration instead of entering the full command **write terminal**, or you can enter **en** to start privileged mode and **con te** to start configuration mode. In addition, you can enter **0** to represent **0.0.0.0**.

Backing Up Your PIX Firewall Configuration

You should back up your configuration in at least one of the following ways:

- Store the configuration in Flash memory with the **write memory** command. Should the need arise, you can restore a configuration from Flash memory using the **configure memory** command.
- Use the **write terminal** command to list the configuration. Then cut and paste the configuration into a text file. Then archive the text file. You can restore a configuration from a text file using the **configure terminal** command and pasting the configuration either line by line or as a whole.
- Store the configuration on another system using the **ftp-server** command to initially specify a host and the **write net** command to store the configuration.
- If you have a PIX 520 or older model, store the configuration on a diskette using the **write floppy** command. If you are using Windows, make sure the diskette is IBM formatted. If you are formatting a disk, access the MS-DOS command prompt and use the **format** command. Do not back up your configuration to the PIX Firewall boot disk.

Each image you store overwrites the last stored image.

Should the need arise, you can restore your configuration from Flash memory with the **configure memory** command, or from diskette with the **configure floppy** command.

Command Line Editing

PIX Firewall uses the same command-line editing conventions as Cisco IOS software. You can view all previously entered commands with the **show history** command or individually with the up arrow or **^p** command. Once you have examined a previously entered command, you can move forward in the list with the down arrow or **^n** command. When you reach a command you wish to reuse, you can edit it or press the **Enter** key to start it. You can also delete the word to the left of the cursor with **^w**, or erase the line with **^u**.

PIX Firewall permits up to 512 characters in a command; additional characters are ignored.

Command Output Paging

On commands such as **help** or **show**, **show xlate**, or other commands that provide long listings, you can determine if the information displays a screenful and pauses, or lets the command run to completion. The **pager** command lets you choose the number of lines to display before the More prompt appears.

When paging is enabled, the following prompt appears:

```
<--- More --->
```

The More prompt uses syntax similar to the UNIX **more** command:

- To view another screenful, press the Space bar.
- To view the next line, press the **Enter** key.
- To return to the command line, press the **q** key.

Comments

You can precede a line with a colon (:) or exclamation mark (!) to create a comment. You can view the comment with the **show history** command or by pressing an arrow key to retrieve a previous command, but because the comment is not in the configuration, the **write terminal** command does not display it.

You can also store configurations with comments preceded by a colon or exclamation mark on a server and then use the **configure net** *[[location]:[filename]]* command to load the configuration from a TFTP server to the PIX Firewall. Replace *location* with the TFTP server name and *filename* with the configuration file name. The PIX Firewall will prune the comments and they will not be visible in the PIX Firewall configuration listing.

You can also add comments to object groups using the **description text** parameter after the **object-group** command. For more information about object groups, refer to Chapter 2, “Controlling Network Access and Use.”

Configuration Size

For PIX Firewall version 5.3(2) and higher, the PIX 525 and PIX 535 support configurations up to 2 MB. The maximum size for the PIX 501 is 256 KB. The maximum configuration size for all other PIX Firewall platforms is 1 MB. For PIX Firewall models using software before version 5.3(2), the maximum configuration size is 350 KB.

**Note**

Regardless of the platform, smaller configuration sizes are recommended to ensure optimum performance.

Use the UNIX **wc** command or a Windows word processing program, such as Microsoft Word, to view the number of characters in the configuration.

Help Information

Help information is available from the PIX Firewall command line by entering **help** or a question mark to list all commands, or after a command to list command syntax; for example, **arp ?**.

The number of commands listed when you use the question mark or **help** command differs by access mode so that unprivileged mode offers the least commands and configuration mode offers the greatest number of commands.

In addition, you can enter any command by itself on the command line and then press **Enter** to view the command syntax.

Viewing the Default Configuration

When you power on your PIX Firewall for the first time, the configuration comes with many of the basic commands required to get started. The configuration you first receive is known as the default configuration. You can use the **write terminal** command to view your configuration at any time. Also use the **write memory** command frequently to save your configuration to Flash memory.

Resetting the Default Configuration

If you make a mistake configuring a PIX 501 or PIX 506/506E, or need to restore the default configuration for any reason, enter the following command:

```
config factory-default [inside-ip-address [address-mask]]
```

This command writes the factory default configuration to memory. If you specify the optional *inside-ip-address* and *address-mask* parameters, the command adjusts the default configuration based on the specified IP address and subnetwork mask.

If you enter this command on other PIX Firewall platforms that do not support it, you will receive the following message:

The config factory default command is only supported on the PIX 501 or PIX 506/506E.

**Note**

The factory default setting for the DHCP address pool size is determined by your PIX Firewall platform and your feature license. For information about the possible options, refer to “[Using the PIX Firewall DHCP Client](#)” in [Chapter 5](#), “[Using PIX Firewall in SOHO Networks](#).”

Clearing and Removing Configuration Settings

To clear all the configuration for a specified command and all its subcommands, enter the following command:

```
clear configurationcommand [subconfigurationcommand]
```

This command clears all the current configuration for the specified configuration command. If you only want to clear the configuration for a specific subcommand, you can enter a value for *subconfigurationcommand*.

To disable the specific parameters or options of a command or subcommand, enter the **no** form of the command, as follows:

```
no configurationcommand [subconfigurationcommand] qualifier [...]
```

In this case, you use the **no** command to remove the specific configuration identified by *qualifier*.

Before You Start Configuring PIX Firewall

The key to successful implementation of your PIX Firewall is having a clear security policy that describes how to control access and use of your organization's network resources. You need to understand your security policy to ensure that you implement and configure the PIX Firewall in a way that supports this policy. Your security policy should have the support of the various departments and administrators responsible for its implementation and should be well understood by network users.

Before you configure the PIX Firewall, sketch out a network diagram with IP addresses that you will assign to the PIX Firewall and those of routers on each interface. If you have more than two interfaces in the PIX Firewall, note the security level for each interface.

You can use the configuration forms in [Appendix A, "Firewall Configuration Forms,"](#) to help you plan your implementation and to collect the information required. The following are the first four configuration forms in the appendix, which will help you collect the information required to complete the configuration described in this chapter:

- PIX Firewall Network Interface Information
- Routing Information
- Network Address Translation
- Static Address Translation

Where to Go from Here

- To record the information required to implement your PIX Firewall, refer to [Appendix A, "Firewall Configuration Forms."](#)
- To complete the configuration required connect your PIX Firewall to your existing network, refer to [Chapter 2, "Establishing Connectivity."](#)
- To allow or restrict specific types of network activity and access, refer to [Chapter 3, "Controlling Network Access and Use."](#)
- To use the application inspection and the **fixup** command to control the secure use of specific applications and services, refer to [Chapter 4, "Configuring Application Inspection \(Fixup\)."](#)

- To use a PIX Firewall as an Easy VPN Remote device in relation to an Easy VPN Server or to use it with DHCP or PPPoE, refer to [Chapter 5, “Using PIX Firewall in SOHO Networks.”](#)
- To perform basic VPN configuration, refer to [Chapter 6, “Configuring IPsec and Certification Authorities.”](#)
- To configure or use PIX Firewall system management tools, refer to [Chapter 9, “Accessing and Monitoring PIX Firewall.”](#)
- To configure the PIX Firewall failover feature, refer to [Chapter 10, “Using PIX Firewall Failover.”](#)
- To upgrade the software image on your PIX Firewall, refer to [Chapter 11, “Changing Feature Licenses and System Software.”](#)

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/warp/public/cc/pd/fw/sqfw500/index.shtml>

