



Advanced Configurations

This chapter includes the following sections:

- [DHCP](#)
- [Failover](#)
- [IDS Syslog Messages](#)
- [PPTP Virtual Private Networks](#)
- [SNMP](#)
- [SSH](#)



Note

For IPSec and L2TP configuration information, refer to the [Cisco PIX Firewall IPSec User Guide, Version 6.0](#).

DHCP

Support for DHCP (Dynamic Host Configuration Protocol) server and DHCP client within the PIX Firewall is now available. DHCP is a protocol that supplies automatic configuration parameters to Internet hosts. This protocol has two components:

- a protocol for delivering host-specific configuration parameters from a DHCP server to a host (DHCP client)
- a 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.

The primary purpose of implementing the DHCP server and DHCP client features into the PIX Firewall is to significantly simplify the configuration of a PIX Firewall unit.

This section includes the following topics:

- [DHCP Client](#)
- [DHCP Server](#)

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. 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 unit's enabled 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.


Note

The PIX Firewall DHCP client does not support **failover** configurations.

To support the DHCP client feature within the PIX Firewall, the following enhancements were made:

- enhanced the **ip address** and the **show ip address** commands
 - **ip address *if_name* dhcp** [setroute] [retry *retry_cnt*]
 - **ip address outside dhcp** [setroute] [retry *retry_cnt*]
 - **show ip address *if_name* dhcp**
- added new **debug** commands:
 - **debug dhcpc packet**
 - **debug dhcpc detail**
 - **debug dhcpc error**

The **ip address dhcp** command enables the DHCP client feature on the specified PIX Firewall interface. The optional **setroute** argument tells the PIX Firewall to set the default route using the default gateway parameter the DHCP server returns.

The **debug dhcpc** commands provide debugging tools for the enabled DHCP client feature.

The PIX Firewall commands used to implement the DHCP client are described in the **ip address** command page and the **debug** command page within [Chapter 5, “Command Reference.”](#) Refer to these command pages for more information.


Note

The DHCP-acquired IP address of the outside interface can also be used as the PAT global address. This makes it unnecessary for the ISP to assign a static IP address to PIX Firewall. Use the **global** command with **interface** keyword to enable PAT to use the DHCP-acquired IP address of outside interface. For more information about the **global** command, see the **global** command page in the [Chapter 5, “Command Reference.”](#)

Enabling the DHCP Client Feature and Setting Default Route

To enable the DHCP client feature on a given PIX Firewall interface and set the default route via the DHCP server, configure the **ip address dhcp setroute** command as part of your entire PIX Firewall configuration, including the **setroute** option. Specify the name of the interface on which the DHCP client will be enabled.

DHCP Server

DHCP server support within the PIX Firewall is designed for use within a remote home or branch office (ROBO) environment using a PIX 506 unit. Connecting to the PIX Firewall are PC clients and other network devices (DHCP clients) that establish network connections that are either nonsecure (not encrypted) or secure (encrypted using IPSec) to access an enterprise or corporate network. As a DHCP server, the PIX Firewall provides network configuration parameters to the DHCP clients through the use of DHCP. These configuration parameters provide a DHCP client the networking parameters used to access the enterprise network, and once in the network, the network services to use, such as the DNS server.


Note

PIX Firewall version 5.3 supports 32 DHCP clients. Prior to the version 5.3 software release, PIX Firewall DHCP servers supported 10 DHCP clients. PIX Firewall version 5.3 and later supports 32 clients on PIX Firewall and 256 on other platforms.


Note

In version 6.0, PIX Firewall DHCP server supports up to 256 DHCP clients. You cannot configure a DHCP server for 256 clients, using a Class C netmask. For example, if a company has a Class C network address of 172.17.1.0 with netmask 255.255.255.0, then 172.17.1.0 (network IP) and 172.17.1.255 (broadcast) cannot be in the DHCP address pool range. Further, one address is used up for the PIX Firewall interface. Thus, if a user uses a Class C netmask, they can only have up to 253 DHCP Clients. To have 256 clients configured, they cannot use a Class C netmask.


Note

The PIX Firewall DHCP server does not support BOOTP requests and failover configurations.

The PIX Firewall commands used to implement the DHCP server feature are described in the [dhcpcd](#) command page and the [debug](#) command page within [Chapter 5, “Command Reference.”](#) Refer to these command pages for more information.

Configuring the DHCP Server Feature

Be sure to configure the IP address and the subnet mask of the **inside** interface using the **ip address** command prior to enabling the DHCP server feature.

Follow these steps to enable the DHCP server feature on a given PIX Firewall interface. (Steps 1 and 6 are required.)

Step 1 Specify a DHCP address pool using the **dhcpcd address** command. The PIX Firewall will assign to a client one of the addresses from this pool to use for a given length of time. The default is the **inside** interface. For example:

```
dhcpcd address 10.0.1.101-10.0.1.110 inside
```

Step 2 (Optional) Specify the IP address(es) of the DNS server(s) the client will use. You can specify up to two DNS servers. For example:

```
dhcpcd dns 209.165.201.2 209.165.202.129
```

Step 3 (Optional) Specify the IP address(es) of the WINS server(s) the client will use. You can specify up to two WINS servers. For example:

```
dhcpcd wins 209.165.201.5
```

- Step 4** Specify the lease length to grant the client. This lease equals the amount of time (in seconds) the client can use its allocated IP address before the lease expires. The default value is 3,600 seconds. For example:

```
dhcpd lease 3000
```

- Step 5** (Optional) Configure the domain name the client will use. For example:

```
dhcpd domain example.com
```

- Step 6** Enable the DHCP daemon within the PIX Firewall to listen for DHCP client requests on the enabled interface. Currently, you can only enable the DHCP server feature on the **inside** interface, which is the default. For example:

```
dhcpd enable inside
```

Here is the sample configuration stemming from the steps configured previously:

```
! set the ip address of the inside interface
ip address inside 10.0.1.2 255.255.255.0
! configure the network parameters the client will use once in the corporate network and
dhcpd address 10.0.1.101-10.0.1.110
dhcpd dns 209.165.201.2 209.165.202.129
dhcpd wins 209.165.201.5
dhcpd lease 3000
dhcpd domain example.com
! enable dhcp server daemon on the inside interface
dhcpd enable inside
```

The following example shows the configuration of a DHCP address pool and a DNS server address with the inside interface being enabled for the DHCP server feature:

```
dhcpd address 10.0.1.100-10.0.1.108
dhcpd dns 209.165.200.227
dhcpd enable
```

The following example shows the configuration of a DHCP address pool and uses the **auto_config** command to configure the dns, wins, and domain parameters:

```
dhcpd address 10.0.1.100-10.0.1.108
dhcpd auto_config
dhcpd enable
```

The following is a partial configuration example of the DHCP server and IPsec features configured on a PIX Firewall that is within a remote office. The PIX 506 unit's VPN peer is another PIX Firewall that has an outside interface IP address of 209.165.200.228 and functions as a gateway for a corporate network.

```
! configure interface ip address
ip address outside 209.165.202.129 255.255.255.0
ip address inside 172.17.1.1 255.255.255.0
! configure ipsec with corporate pix
access-list ipsec-peer permit ip 172.17.1.0 255.255.255.0 192.168.0.0 255.255.255.0
ipsec transform-set myset esp-des esp-sha-hmac
crypto map mymap 10 ipsec-isakmp
crypto map mymap 10 match address ipsec-peer
crypto map mymap 10 set transform-set myset
crypto map mymap 10 set peer 209.165.200.228
crypto map mymap interface outside
sysopt connection permit-ipsec
nat (inside) 0 access-list ipsec-peer
isakmp policy 10 authentication preshare
isakmp policy 10 encryption des
```

```
isakmp policy 10 hash sha
isakmp policy 10 group 1
isakmp policy 10 lifetime 3600
isakmp key 12345678 address 0.0.0.0 netmask 0.0.0.0
isakmp enable outside
!configure dhcp server address
dhcpd address 172.17.1.100-172.17.1.109
dhcpd dns 192.168.0.20
dhcpd wins 192.168.0.10
dhcpd lease 3000
dhcpd domain example.com
! enable dhcp server on inside interface
dhcpd enable
! use outside interface ip as PAT global address
nat (inside) 1 0 0
global (outside) 1 interface
```

Failover

Failover allows you to add a second PIX Firewall unit that takes control if the primary unit fails.

This section includes the following topics:

- [Understanding Failover](#)
- [Configuring Failover](#)
- [Upgrading Failover from a Previous Version](#)
- [Additional Failover Information](#)

A failover configuration example is provided in “[Failover Configuration](#)” in [Chapter 4](#), “[Configuration Examples](#).”

**Note**

The PIX 506 or PIX 506E cannot be used for failover in any configuration.

**Note**

The primary unit in the PIX 515, PIX 525, or PIX 535 failover pair must have an Unrestricted (UR) license. The secondary unit can have a Failover (FO) or UR license. However, the failover pair must be two otherwise identical units with the same PIX Firewall hardware and software.

Understanding Failover

Failover allows you to connect a second PIX Firewall unit to your network to protect your network should the first unit go offline. If you use Stateful Failover, you can maintain operating state for the TCP connection during the failover from the primary unit to the standby unit.

When a failover occurs, each unit changes state. The unit that activates assumes the IP and MAC addresses of the previously Active unit and begins accepting traffic. The new Standby unit assumes the failover IP and MAC addresses of the unit that was previously the Active unit. Because network devices see no change in these addresses, no ARP entries change or time out anywhere on the network.

Failover requires you to purchase a second PIX Firewall unit sold as a failover unit that only works as a failover unit. You need to ensure that both units have the same software version, activation key type, Flash memory, and the same RAM. Once you configure the primary unit and attach the necessary cabling, the primary unit automatically copies the configuration over to the Standby unit.

The ACT indicator light on the front of the PIX 515, PIX 525, and PIX 535 is on when the unit is the Active failover unit. If failover is not enabled, this light is on. If failover is present, the light is on when the unit is the Active unit and off when the unit is the Standby unit.

Failover works over Ethernet. However, Stateful Failover interfaces must have a connection speed of a minimum of 100 Mbps full-duplex.

The failover feature causes the PIX Firewall to ARP for itself every 15 seconds (depending on the time set with the **failover poll** command). This ARPing can only be stopped by disabling failover.

Cabling the two PIX Firewall units together requires connecting a high-speed serial cable for communications. The minimum requirement for Stateful Failover is a dedicated 100 Mbps full-duplex connection.

Configuring the primary PIX Firewall for failover requires you to configure the **failover** command to enable failover, the **failover ip address** command to assign IP addresses to the Standby unit, and the **failover link** command to enable Stateful Failover.

**Note**

Refer to [“Additional Failover Information”](#) for information on Stateful Failover, how failover occurs, and frequently asked questions.

Configuring Failover

For failover, both PIX Firewall units must be the same model number, have at least as much RAM, have the same Flash memory size, and be running the same software version.

**Note**

If you have already powered on the Standby unit, power it off and leave it off until instructed in the steps that follow.

Follow these steps to configure failover:

-
- Step 1** Because the PIX Firewall clock is stored in the CMOS, if you have not done so already, specify the **clock set time** command on the Active PIX Firewall to synchronize the time on both PIX Firewall units. If you are using IPSec with digital certificates, set the time appropriate to the GMT timezone (this is done because the PIX Firewall does not use timezones.)
 - Step 2** Attach a network cable between the primary and secondary units for each network interface to which you have configured an IP address.
 - Step 3** Connect the failover cable to the primary PIX Firewall unit ensuring that the end of the cable marked “primary” attaches to the primary unit and that the end marked “secondary” connects to the secondary unit.
 - Step 4** *Only configure the primary unit.* Changes made to the Standby unit are not copied to the primary unit and are lost during the next reboot. When you are done configuring the PIX Firewall and enter the **write memory** command to save the configuration to Flash memory, the primary unit automatically updates the secondary unit.



Note Do not power on the secondary unit instructed. First configure the primary unit and then power on the secondary unit only when instructed to do so.

- Step 5** Enter configuration mode with the **configure terminal** command.
- Step 6** Ensure that you have not used the **auto** or the **1000auto** option in any **interface** command in your configuration. To view **interface** commands in your configuration, use the **write terminal** command. Reenter an interface with new information to correct a command you wish to change. Always specify the speed for the interface, such as **10baset** for 10 Mbps or **100basetx** for 100 Mbps. Ensure that the same speeds and duplexes are the same for any devices on the subnets including switches and routers.
- Step 7** If you are using Stateful Failover, set the Stateful Failover dedicate interface speed using the **100fullor 1000sxfull** option to the interface command. This is extremely important and must be performed even if you are using a crossover connector to connect the PIX Firewall units directly to each other.
- Step 8** Use the **clear xlate** command after changing the **interface** command.
- Step 9** If you have not done so already, use the **ip address** command statement to assign IP addresses to each interface on the primary unit. If you make a mistake while entering an **ip address** command, reenter the command again correctly.

Use the **show ip address** command to view the addresses you specified:

```
show ip address
System IP Addresses:
  ip address outside 192.168.1.1 255.255.255.0
  ip address inside 10.1.1.1 255.255.255.0
  ip address intf2 192.168.2.1 255.255.255.0
  ip address intf3 192.168.3.1 255.255.255.0
  ip address 4th 172.16.1.1 255.255.255.0
Current IP Addresses:
  ip address outside 192.168.1.1 255.255.255.0
  ip address inside 10.1.1.1 255.255.255.0
  ip address intf2 192.168.2.1 255.255.255.0
  ip address intf3 192.168.3.1 255.255.255.0
  ip address 4th 172.16.1.1 255.255.255.0
```

The Current IP Addresses are the same as the System IP Addresses on the failover Active unit. When the primary unit fails, the Current IP Addresses become those of the Standby unit.

- Step 10** Use the **failover** command statement to enable failover on the primary unit.
- Step 11** Use the **show failover** command to verify that the primary unit is enabled by checking for the following statement:

```
This host: Primary - Active
```

An example of the **show failover** command is as follows:

```
show failover
Failover On
Cable status: Other side powered off
Reconnect timeout 0:00:00
Poll frequency 15 seconds
  This host: Primary - Active
    Active time: 225 (sec)
    Interface 4th (172.16.1.1): Normal (Waiting)
    Interface intf3 (192.168.3.1): Normal (Waiting)
    Interface intf2 (192.168.2.1): Normal (Waiting)
    Interface outside (192.168.1.1): Normal (Waiting)
    Interface inside (10.1.1.1): Normal (Waiting)
  Other host: Secondary - Standby
```

```

Active time: 0 (sec)
Interface 4th (0.0.0.0): Unknown (Waiting)
Interface intf3 (0.0.0.0): Unknown (Waiting)
Interface intf2 (0.0.0.0): Unknown (Waiting)
Interface outside (0.0.0.0): Unknown (Waiting)
Interface inside (0.0.0.0): Unknown (Waiting)

```

The Cable Status that displays with the **show failover** command has these values:

- My side not connected—Indicates that the serial cable is not connected to the unit on which you entered the **show failover** command.
- Normal—Indicates that the Active unit is working and that the Standby unit is ready.
- Other side is not connected—Indicates that the serial cable is not connected to the other unit (the unit *opposite* from where you entered the **show failover** command).
- Other side powered off—Indicates that the unit not shown as Active is powered off.

The failover interface flags appear to the right of each interface's IP address in the **show failover** display. The failover flags indicate the following:

- Failed—The interface has failed.
- Link Down—The interface line protocol is down.
- Normal—The interface is working correctly.
- Shut Down—The interface has been administratively shut down (the **shutdown** option is enabled in the **interface** command statement in the configuration).
- Unknown—The IP address for the interface has not been configured and failover cannot determine the status of the interface.
- Waiting—Monitoring of the other unit's network interface has not yet started.

Step 12 Enter a **failover ip address** command statement for each interface to specify the Standby unit's interface addresses. It is *not* necessary for the two units to be configured for this command to work correctly. The IP addresses on the Standby unit are different from the Active unit's addresses, but must be in the same subnet for each interface. The following example sets the IP addresses for the interfaces on the Standby unit:

```

failover ip address inside 10.1.1.2
failover ip address outside 192.168.1.2
failover ip address intf2 192.168.2.2
failover ip address intf3 192.168.3.2
failover ip address 4th 172.16.1.2

```

Use the **show failover** command and you can see that the secondary unit now has IP addresses for each interface:

```

show failover
Failover On
Cable status: Other side powered off
Reconnect timeout 0:00:00
Poll frequency 15 seconds
  This host: Primary - Active
    Active time: 510 (sec)
    Interface 4th (172.16.1.1): Normal (Waiting)
    Interface intf3 (192.168.3.1): Normal (Waiting)
    Interface intf2 (192.168.2.1): Normal (Waiting)
    Interface outside (192.168.1.1): Normal (Waiting)
    Interface inside (10.1.1.1): Normal (Waiting)
  Other host: Secondary - Standby
    Active time: 0 (sec)
    Interface 4th (172.16.1.2): Unknown (Waiting)

```

```

Interface intf3 (192.168.3.2): Unknown (Waiting)
Interface intf2 (192.168.2.2): Unknown (Waiting)
Interface outside (192.168.1.2): Unknown (Waiting)
Interface inside (10.1.1.2): Unknown (Waiting)

```

Step 13 If you are configuring Stateful Failover, use the **failover link** command to specify the name of the dedicated interface you are using. For example, assume the “4th” interface will be used for Stateful Failover and enter the following command:

```
failover link 4th
```

Step 14 After enabling Stateful Failover, use the **show failover** command and additional information is provided as follows:

```

show failover
Failover On
Cable status: Other side powered off
Reconnect timeout 0:00:00
Poll frequency 15 seconds
  This host: Primary - Active
    Active time: 510 (sec)
    Interface 4th (172.16.1.1): Normal (Waiting)
    Interface intf3 (192.168.3.1): Normal (Waiting)
    Interface intf2 (192.168.2.1): Normal (Waiting)
    Interface outside (192.168.1.1): Normal (Waiting)
    Interface inside (10.1.1.1): Normal (Waiting)
  Other host: Secondary - Standby
    Active time: 0 (sec)
    Interface 4th (172.16.1.2): Unknown (Waiting)
    Interface intf3 (192.168.3.2): Unknown (Waiting)
    Interface intf2 (192.168.2.2): Unknown (Waiting)
    Interface outside (192.168.1.2): Unknown (Waiting)
    Interface inside (10.1.1.2): Unknown (Waiting)

```

Stateful Failover Logical Update Statistics

```

Link : 4th
Stateful Obj   xmit   xerr   rcv   rerr
General       0       0       0     0
sys cmd       0       0       0     0
up time       0       0       0     0
xlate         0       0       0     0
tcp conn      0       0       0     0
udp conn      0       0       0     0
ARP tbl       0       0       0     0
RIP Tbl       0       0       0     0

```

Logical Update Queue Information

	Cur	Max	Total
Recv Q:	0	0	0
Xmit Q:	0	0	0

The items in the top row of the “Stateful Failover Logical Update Statistics” section of the **show failover** command are as follows:

- Stateful Obj—PIX Firewall stateful object
- xmit—Number of transmitted packets to the other unit
- xerr—Number of errors that occurred while transmitting packets to the other unit
- rcv—Number of received packets
- rerr—Number of errors that occurred while receiving packets from the other unit

The items in the first column provide an object static count for each statistic:

- General—Sum of all stateful objects
- sys cmd—Logical update system commands; for example, LOGIN and Stay Alive
- up time—Up time, which the Active unit passes to the Standby unit
- xlate—Translation information
- tcp conn—CTCP connection information
- udp conn—Dynamic UDP connection information
- ARP tbl—Dynamic ARP table information
- RIF Tbl—Dynamic router table information

The items in the “Logical Update Queue Information” list the current, maximum, and total number of packets in the receive (Recv) and transmit (Xmit) queues.

- Step 15** If you want to set a time shorter than 15 seconds for the units to exchange “hello” packets to ensure each unit is available, use the **failover poll seconds** command. The default is 15 seconds. The minimum value is 3 seconds and the maximum is 15 seconds. Set to a lower value for Stateful Failover. With a faster poll time, PIX Firewall can detect failure and trigger failover faster. However, faster detection may cause unnecessary switchovers when the network is temporarily congested or a network card starts slowly.
- Step 16** Power on the secondary unit. As soon as the secondary unit starts, the primary unit recognizes it and starts synchronizing the configurations. As the configurations synchronize, the messages “Sync Started” and “Sync Completed” appear.
- Step 17** After the Standby unit comes up, use the **show failover** command on the primary unit to verify status:

```
show failover
Failover On
Cable status: Other side powered off
Reconnect timeout 0:00:00
Poll frequency 15 seconds
  This host: Primary - Active
    Active time: 510 (sec)
    Interface 4th (172.16.1.1): Normal
    Interface intf3 (192.168.3.1): Normal
    Interface intf2 (192.168.2.1): Normal
    Interface outside (192.168.1.1): Normal
    Interface inside (10.1.1.1): Normal
  Other host: Secondary - Standby
    Active time: 0 (sec)
    Interface 4th (172.16.1.2): Normal
    Interface intf3 (192.168.3.2): Normal
    Interface intf2 (192.168.2.2): Normal
    Interface outside (192.168.1.2): Normal
    Interface inside (10.1.1.2): Normal
Stateful Failover Logical Update Statistics
Link : 4th
Stateful Obj   xmit      xerr      rcv       rerr
General        0          0          0          0
sys cmd        0          0          0          0
up time        0          0          0          0
xlate          0          0          0          0
tcp conn       0          0          0          0
udp conn       0          0          0          0
ARP tbl        0          0          0          0
RIP Tbl        0          0          0          0

Logical Update Queue Information
          Cur      Max      Total
Recv Q:   0        0        0
Xmit Q:   0        0        0
```

- Step 18** Use the **write memory** to save the configuration to Flash memory and to synchronize the configuration on the Standby unit with the primary unit.

Verifying That the Configuration Was Successful

Follow these steps to verify that the configuration was successful:

- Step 1** If you have access to a syslog server, such as a UNIX system, enable logging so you can view the syslog messages as you proceed with the steps that follow. For information on syslog messages, refer to [System Log Messages for the Cisco Secure PIX Firewall Version 6.0](#), which is available online at the following site:

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_v53/index.htm

Additional information on syslog is available on the **logging** command page in [Chapter 5, “Command Reference.”](#)

To enable logging to a server with the example address 10.1.1.5 on the inside interface, use these commands:

```
logging host inside 10.1.1.5
logging trap debugging
logging on
```

The **logging trap debugging** command statement lets all levels of syslog messages be sent, which can produce a large number of messages on a system in production, but is very helpful for debugging. When you are done testing failover, use the **logging trap error** command to reduce the number of syslog messages to only those messages displaying an alert, critical condition, or error.

- Step 2** Test the secondary unit by turning off the primary unit.
- Step 3** Use the **show failover** command to verify that the secondary unit is now active.
- Step 4** Use FTP to send a file between hosts on different interfaces.
- Step 5** Use the **show interface** command to verify that traffic is being processed.
- Step 6** You can also use the **debug fover option** command. Choose an option from the following [Table 3-1](#):

Table 3-1 *debug failover Options*

Option	Description
cable	Failover cable status
fail	Failover internal exception
fmsg	Failover message
get	IP network packet received
ifc	Network interface status trace
open	Failover device open
put	IP network packet transmitted
rx	Failover cable receive
rxdmp	Cable recv message dump (serial console only)
rxip	IP network failover packet received

Table 3-1 *debug failover Options (continued)*

Option	Description
tx	Failover cable transmit
txdmp	Cable xmit message dump (serial console only)
txip	IP network failover packet transmit
verify	Failover message verify
switch	Failover Switching status

Step 7 When ready, power on the primary unit and it will take over automatically as the active unit.

Upgrading Failover from a Previous Version

Use the steps that follows to upgrade failover for Pix that uses floppy drive to boot image:

- Step 1** Connect a separate console to the primary unit and one to the secondary unit.
- Step 2** Insert the PIX Firewall version's diskette into the primary unit. Enter the **reload** command at the primary unit.
- Step 3** As the primary unit reboots, PIX Firewall prompts you to write the image to Flash memory. Before entering a reply, read the next three substeps and be ready to move quickly to complete them. When ready, enter **y** for yes at the prompt.
- Immediately remove the diskette from the primary unit and insert it into the Standby unit. Locate the reset button on the front of the Standby unit.
 - When the PIX Firewall Cisco banner appears on the primary unit's console, press the reset button on the Standby unit to load the new image.
 - On the primary unit, enter the **show failover** command to make sure the primary unit is active and the secondary unit is in Standby mode after the upgrade of the primary unit.
- Step 4** Wait for the Standby unit to finish booting. Once the Standby unit is up, the two units synchronize during which time the primary unit's console does not accept input. On the Standby unit, use the **show failover** command to monitor progress. When both PIX Firewall units report Normal, the replication is done.
-

Complete the following steps for a PIX Firewall with a bios extension installed, which can TFTP from the monitor mode:

- Step 1** Connect a separate console to the primary unit and one to the secondary unit.
- Step 2** Reload both PIX Firewall units, and bring them to monitor mode.
- Step 3** On the primary unit, use monitor mode TFTP to load the new PIX Firewall image. You will want to save the image to Flash memory and let it boot up. Enter a **show failover** command to ensure everything looks fine.
- Step 4** Repeat Step 3 on the secondary unit.

- Step 5** Once the standby (secondary) unit completes booting and is up, the active (primary) unit will start to synchronize the configuration from the primary unit to the secondary. Wait until the configuration replication is finished, then use the **show failover** command on both PIX Firewall units to ensure the failover is running correctly.

**Note**

Steps above can be done using boot-helper disk as well. Just insert the boot-helper disks on both PIX Firewalls, reload, then complete steps 3-5.

Additional Failover Information

This section includes the following topics:

- [Failover Communication](#)
- [What Causes Failover?](#)
- [Configuration Replication](#)
- [Stateful Failover](#)
- [Disabling Failover](#)
- [Failover Usage Notes](#)
- [Frequently Asked Failover Questions](#)
- [Stateful Failover Questions](#)

Failover Communication

Both units in a failover pair communicate through the failover cable, which is a modified RS-232 serial link cable that transfers data at 117,760 baud (115K). The data provides the unit identification of primary or secondary, the power status of the other unit, and serves as a communication link for various failover communications between the two units.

The two units send special failover “hello” packets to each other over all network interfaces and the failover cable every 15 seconds. The **failover poll seconds** command allows you to determine how long failover waits before sending special failover “hello” packets between the primary and Standby units over all network interfaces and the failover cable. The default is 15 seconds. The minimum value is 3 seconds and the maximum is 15 seconds. Set to a lower value for Stateful Failover. With a faster poll time, PIX Firewall can detect failure and trigger failover faster. However, faster detection may cause unnecessary switchovers when the network is temporarily congested or a network card starts slowly.

The failover feature in PIX Firewall monitors failover communication, the power status of the other unit, and hello packets received at each interface. If two consecutive hello packets are not received within a time determined by the failover feature, failover starts testing the interfaces to determine which unit has failed, and transfers active control to the Standby unit.

You can choose the Stateful Failover option if you have 100 Mbps LAN interfaces so that connection states are automatically relayed between the two units. If you are using Stateful Failover, connection states are relayed from the primary unit to the secondary unit. Without Stateful Failover, the Standby unit does not maintain the state information of each connection. This means that all active connections will be dropped when failover occurs and that client systems must reestablish connections.

What Causes Failover?

If a failure is due to a condition other than a loss of power on the other unit, failover will begin a series of tests to determine which unit failed. This series of tests will begin when “hello” messages are not heard for two consecutive 15-second intervals (the interval depends on how you set the **failover poll** command). Hello messages are sent over both network interfaces and the failover cable.

The purpose of these tests is to generate network traffic in order to determine which (if either) unit has failed. At the start of each test, each unit clears its received packet count for its interfaces. At the conclusion of each test, each unit looks to see if it has received any traffic. If it has, the interface is considered operational. If one unit receives traffic for a test and the other unit does not, the unit that received no traffic is considered failed. If neither unit has received traffic, then go to the next test.



Note

If the failover IP address has not been set, failover does not work, and the Network Activity, ARP, and Broadcast ping tests are not performed.

Failover uses the following tests to determine if the other unit is available:

- **Link Up/Down test**—This is a test of the NIC card itself. If an interface card is not plugged into an operational network, it is considered failed (for example, a switch failed, has a failed port, or a cable is unplugged).
- **Network Activity test**—This is a received network activity test. The unit will count all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the ARP test begins.
- **ARP test**—The ARP test consists of reading the unit’s ARP cache for the 10 most recently acquired entries. One at a time the unit sends ARP requests to these machines attempting to stimulate network traffic. After each request the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next machine. If at the end of the list no traffic has been received, the ping test begins.
- **Broadcast Ping test**—The ping test consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the testing starts over again with the ARP test.

Configuration Replication

The two PIX Firewall units must be configured exactly the same and running the same software release. Configuration replication occurs over the failover cable from the Active unit to the Standby unit in three ways:

- When the Standby unit completes its initial bootup, the Active unit replicates its entire configuration to the Standby unit.
- As commands are entered on the Active unit they are sent across the Failover Cable to the Standby unit.
- Entering the **write standby** command on the Active unit forces the entire configuration in memory to be sent to the Standby unit.

The configuration replication only occurs from memory to memory. After replication, use the write memory command to write the configuration into Flash memory. Because the failover cable is used, the replication can take a long time to complete with a large configuration. If a switchover occurs during the replication, the new Active unit will have a partial configuration. The unit will reboot itself to recover

the configuration from the Flash or re-synchronize with the other unit. When the replication starts, the PIX Firewall console displays the message “Sync Started,” and when complete, displays the message “Sync Completed.” During the replication, information cannot be entered on the PIX Firewall console.

Stateful Failover

The Stateful Failover feature passes per-connection stateful information to the Standby unit. After a failover occurs, the same connection information is available at the new Active unit. End user applications are not required to do a reconnect to keep the same communication session.

The state information passed to the Standby unit includes the global pool addresses and status, connection and translation information and status, the negotiated H.323 UDP ports, the port allocation bit map for PAT, and other details necessary to let the Standby unit take over processing if the primary unit fails.

Depending on the failure, the PIX Firewall takes from 15 to 45 seconds to cause a switchover. Applications not handled by Stateful Failover will then require time to reconnect before the Active unit becomes fully functional.

Stateful Failover requires a dedicated interface on each PIX Firewall, with a minimum connection speed of 100 Mbps full-duplex, to be used exclusively for passing state information between the two PIX Firewall units.


The Stateful Failover interface can be connected to any of the following:

- Cat 5 crossover cable directly connecting the primary unit to the secondary unit.
- 100BaseTX half duplex switch using straight Cat 5 cables.
- 100BaseTX full duplex on a dedicated switch or dedicated VLAN of a switch.
- 1000BaseTx full duplex on a dedicated switch or dedicated VLAN of a switch.

Data is passed over the dedicated interface using IP protocol 105. No hosts or routers should be on this interface.

Figure 3-1 shows two PIX Firewall units connected for use with Stateful Failover.

Figure 3-1 Stateful Failover Minimum Setup



Note

All enabled interfaces must be connected between the Active and Standby units. If an interface is not in use, use the **shutdown** option to the **interface** command to disable the interface.

Disabling Failover

You can disable failover with the **no failover** command. If failover is disabled, the following messages display when you enter the **show failover** command:

```
show failover
Failover Off
Cable Status: My side not connected
Reconnect timeout: 0:00:00
```

Failover Usage Notes

The following notes apply to the use of failover on the PIX Firewall:

1. When a failover cable connects two PIX Firewall units, the **no failover** command disables failover until you enter the **failover** command to explicitly enable failover. Previously, when the failover cable connected two PIX Firewall units and you entered the **no failover** command, failover would automatically re-enable after 15 seconds.
If you reboot the PIX Firewall without entering the **write memory** command and the failover cable is connected, failover mode automatically enables.
2. Perform the following on any switch that connects to the PIX Firewall:
 - a. Enable portfast on all ports on the switch that connect directly to the PIX Firewall.

- b. Turn off trunking on all ports on the switch that connect directly to the PIX Firewall.
 - c. Turn off channeling on all ports on the switch that connect directly to the PIX Firewall.
 - d. Ensure the MSFC is not running a deferred Cisco IOS software version.
3. The PIX Firewall failover unit is intended to be used solely for failover and not in standalone mode. If a failover unit is used in standalone mode, the unit will reboot at least once every 24 hours until the unit is returned to failover duty. When the unit reboots, the following message displays at the console:

```

=====NOTICE =====
      This machine is running in secondary mode without
      a connection to an active primary PIX. Please
      check your connection to the primary system.

                        REBOOTING...
=====

```

- 4. If a failover-only PIX Firewall is not attached to a failover cable or is attached to the primary end of a failover cable, then it will hang at boot time. It must be a secondary unit.
- 5. Changes made on the Standby unit are not replicated on the Active unit.
- 6. Failover messages always have a syslog priority level of 2, which indicates a critical condition. Refer to the **logging** command description in [Chapter 5, “Command Reference,”](#) for more information on syslog messages. Refer to *System Log Messages for the Cisco Secure PIX Firewall Version 6.0*, available online. PIX Firewall documentation is available online at the following site:

<http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/index.htm>

To receive SNMP syslog traps (SNMP failover traps), you must configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and also have compiled the Cisco syslog MIB into your SNMP management station. See the **snmp-server** and **logging** command pages in [Chapter 5, “Command Reference,”](#) for more information.

Frequently Asked Failover Questions

This section contains some frequently asked questions about the failover features.

- What happens when failover is triggered?

A switch can be initiated by either unit. When a switch takes place, each unit changes state. The newly Active unit assumes the IP address and MAC address of the previously Active unit and begins accepting traffic for it. The new Standby unit assumes the IP address and MAC address of the unit that was previously the Standby unit.

- How is startup initialization accomplished between two units?

When a unit boots up, it defaults to Failover Off and secondary, unless the failover cable is present or failover has been saved in the configuration. The configuration from the Active unit is also copied to the Standby unit. If the cable is not present, the unit automatically becomes the Active unit. If the cable is present, the unit that has the primary end of the failover cable plugged into it becomes the primary unit by default.

- How can both units be configured the same without manually entering the configuration twice?

Commands entered on the Active unit are automatically replicated on the Standby unit.

- What happens if a primary unit has a power failure?

When the primary PIX Firewall unit experiences a power failure, the Standby PIX Firewall comes up in active mode. If the primary unit is powered on again it will become the Standby unit.

- What constitutes a failure?

Fault detection is based on the following:

- Failover hello packets are received on each interface. If hello packets are not heard for two consecutive 15 second intervals, the interface will be tested to determine which unit is at fault. (You can change this duration with the **failover poll** command.)
- Cable errors. The cable is wired so that each unit can distinguish between a power failure in the other unit, and an unplugged cable. If the Standby unit detects that the Active unit is turned off (or resets), it will take active control.

If the cable is unplugged, a syslog is generated but no switching occurs. An exception to this is at bootup, at which point an unplugged cable will force the unit active. If both units are powered on without the failover cable installed they will both become active creating a duplicate IP address conflict on your network. The failover cable must be installed for failover to work correctly.

- Failover communication. The two units share information every 15 seconds, but you can change this duration with the **failover poll** command. If the Standby unit does not hear from the Active unit in two communication attempts (and the cable status is OK) the Standby unit will take over as active.
- How long does it take to detect a failure?
 - Network errors are detected within 30 seconds (two consecutive 15-second intervals).
 - Power failure (and cable failure) is detected within 15 seconds.
 - Failover communications errors are detected within 30 seconds (two consecutive 15-second intervals).
- What maintenance is required?

Syslog messages will be generated when any errors or switches occur. Evaluate the failed unit and fix or replace it.

Stateful Failover Questions

- What causes Stateful Failover to occur?
 - A power off or a power down condition on the Active PIX Firewall.
 - Reboot of the Active PIX Firewall.
 - A link goes down on the Active PIX Firewall for more than twice the configured poll time or a maximum of 30 seconds.
 - “Failover active” on the Standby PIX Firewall.
 - Block memory exhaustion for 15 consecutive seconds or more on the Active unit.
- What information is replicated to the Standby PIX Firewall on Stateful Failover?
 - The configuration.
 - TCP connection table including timeout information of each connection.
 - Translation (xlate) table.
 - System up time; that is, the system clock is synchronized on both PIX Firewall units.
- What information is not replicated to the Standby PIX Firewall on Stateful Failover?
 - The user authentication (uauth) table.
 - The ISAKMP and IPsec SA table.

- The ARP table.
- Routing information.
- What are Stateful Failover hardware requirements?
 - Two identical PIX Firewall units are required.
 - You must connect the LAN ports for Stateful Failover on both PIX Firewall units with a crossover cable or through a switch. A minimum interface speed of 100 Mbps full-duplex is required for the Stateful Failover ports.
 - For better performance, a PIX 520 or later model of PIX Firewall is recommended.
 - You need a failover cable to connect the two failover ports on both PIX Firewall units.
- What are Stateful Failover hardware restrictions?
 - The failover cable must be installed and work correctly.
 - The dedicated failover ports on both PIX Firewall units must be connected and fully functional.
- What are Stateful Failover software requirements?
 - PIX Firewall version 5.1 or later is required for Stateful Failover.
 - Both PIX Firewall units must run the same version of PIX Firewall software.
- What are Stateful Failover license requirements?
 - Stateful Failover requires a feature-based license key with failover feature support or connection-based license key.

IDS Syslog Messages

PIX Firewall lists single-packet (called “atomic”) Cisco Secure Intrusion Detection System (formerly known as NetRanger) signature messages via syslog. Refer to *PIX Firewall System Log Messages, Version 6.0* for a list of the supported messages. You can view this document online at the following site:

http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_60/syslog/index.htm

All signature messages are not supported by PIX Firewall in this release. IDS syslog messages all start with `%PIX-4-4000nn` and have the following format:

```
%PIX-4-4000nn IDS:sig_num sig_msg from ip_addr to ip_addr on interface int_name
```

For example:

```
%PIX-4-400013 IDS:2003 ICMP redirect from 10.4.1.2 to 10.2.1.1 on interface dmz
%PIX-4-400032 IDS:4051 UDP Snork attack from 10.1.1.1 to 192.168.1.1 on interface outside
```

Options:

sig_num The signature number. Refer to the *Cisco Secure Intrusion Detection System Version 2.2.1 User Guide* for more information. You can view the “NSDB and Signatures” chapter from this guide at the following site:

<http://www.cisco.com/univercd/cc/td/doc/product/iaabu/csids/csids1/csidsug/index.htm>

sig_msg The signature message—approximately the same as the NetRanger signature message.

ip_addr The local to remote address to which the signature applies.

int_name The name of the interface on which the signature originated.

You can determine which messages display with the following commands:

ip audit signature *signature_number* **disable**

Attaches a global policy to a signature. Used to disable or exclude a signature from auditing.

no ip audit signature *signature_number*

Removes the policy from a signature. Used to reenable a signature.

show ip audit signature [*signature_number*]

Displays disabled signatures.

ip audit info [**action** [**alarm**] [**drop**] [**reset**]]

Specifies the default action to be taken for signatures classified as informational signatures.

The **alarm** option indicates that when a signature match is detected in a packet, PIX Firewall reports the event to all configured syslog servers. The **drop** option drops the offending packet. The **reset** option drops the offending packet and closes the connection if it is part of an active connection. The default is **alarm**. To cancel event reactions, specify the **ip audit info** command without an **action** option.

no ip audit info

Sets the action to be taken for signatures classified as informational and reconnaissance to the default action.

show ip audit info

Displays the default informational actions.

ip audit attack [**action** [**alarm**] [**drop**] [**reset**]]

Specifies the default actions to be taken for attack signatures. The **action** options are as previously described.

no ip audit attack

Sets the action to be taken for attack signatures to the default action.

show ip audit attack

Displays the default attack actions. An audit policy (audit rule) defines the attributes for all signatures that can be applied to an interface along with a set of actions. Using an audit policy the user may limit the traffic that is audited or specify actions to be taken when the signature matches. Each audit policy is identified by a name and can be defined for informational or attack signatures. Each interface can have two policies; one for informational signatures and one for attack signatures. If a policy is defined without actions, then the configured default actions will take effect. Each policy requires a different name.

ip audit name *audit_name* **info** [**action** [**alarm**] [**drop**] [**reset**]]

All informational signatures except those disabled or excluded by the **ip audit signature** command are considered part of the policy. The actions are the same as described previously.

no ip audit name *audit_name* [**info**]

Remove the audit policy *audit_name*.

ip audit name *audit_name* **attack** [**action** [**alarm**] [**drop**] [**reset**]]

All attack signatures except those disabled or excluded by the **ip audit signature** command are considered part of the policy. The actions are the same as described previously.

no ip audit name *audit_name* [**attack**]

Removes the audit specification *audit_name*.

show ip audit name [**name** [**info** | **attack**]]

Displays all audit policies or specific policies referenced by name and possibly type.

ip audit interface *if_name* *audit_name*

Applies an audit specification or policy (via the **ip audit name** command) to an interface.

no ip audit interface [*if_name*]

Removes a policy from an interface.

show ip audit interface

Displays the interface configuration.

PPTP Virtual Private Networks

PIX Firewall provides support for Microsoft PPTP, which is an alternative to IPSec handling for VPN Clients. While PPTP is less secure than IPSec, PPTP is easier to implement and maintain.

This section contains the following topics:

- [Introduction to PPTP Configuration](#)
- [vpdn Command with PPTP](#)
- [vpdn Command Example](#)

Introduction to PPTP Configuration

The **vpdn** command implements the PPTP feature for inbound connections between the PIX Firewall and a Windows client. Point-to-Point Tunneling Protocol (PPTP) is a layer 2 tunneling protocol which lets a remote client use a public IP network to communicate securely with servers at a private corporate network. PPTP tunnels the IP protocol. RFC 2637 describes the PPTP protocol.

Support is provided for only inbound PPTP and only one PIX Firewall interface can have the **vpdn** command enabled.

Supported authentication protocols include: PAP, CHAP, and MS-CHAP using external AAA (RADIUS or TACACS+) servers or the PIX Firewall local username and password database. Through the PPP IPCP protocol negotiation, PIX Firewall assigns a dynamic internal IP address to the PPTP client allocated from a locally defined IP address pool.

PIX Firewall PPTP VPN supports standard PPP CCP negotiations with Microsoft Point-To-Point Encryption (MPPE) extensions using RSA/RC4 algorithm. MPPE currently supports 40-bit and 128-bit session keys. MPPE generates an initial key during user authentication and refreshes the key regularly. In this release, compression is not supported.

When you specify MPPE, you must use the MS-CHAP PPP authentication protocol. If you are using an external AAA server, the protocol must be RADIUS and the external RADIUS server must be able to return the Microsoft MSCHAP_MPPE_KEY attribute to the PIX Firewall in the RADIUS Authentication Accept packet. See RFC 2548, “Microsoft Vendor Specific RADIUS Attributes,” for more information on the MSCHAP_MPPE_KEY attribute.

Cisco Secure ACS 2.5 and later releases support the MSCHAP/MPPE encryption.

PIX Firewall PPTP VPN has been tested with the following Microsoft Windows products: Windows 95 with DUN1.3, Windows 98, Windows NT 4.0 with SP6, and Windows 2000.



Note

If you configure PIX Firewall for 128-bit encryption and if a Windows 95 or Windows 98 client does not support 128-bit or greater encryption, then the connection to the PIX Firewall is refused. When this occurs, the Windows client moves the dial-up connection menu down to the screen corner while the PPP negotiation is in progress. This gives the appearance that the connection is accepted when it is not. When the PPP negotiation completes, the tunnel terminates and PIX Firewall ends the connection. The Windows client eventually times out and disconnects.

vpdn Command with PPTP

Use the **vpdn** command with the **sysopt connection permit-pptp** command to allow PPTP traffic to bypass checking of **access-list** command statements.

The **show vpdn** command lists tunnel and session information.

The **clear vpdn** command removes all **vpdn** commands from the configurations and stops all the active PPTP tunnels. The **clear vpdn all** command allows you to remove all tunnels, and the **clear vpdn id tunnel_id** command allows you to remove tunnels associated with *tunnel_id*. (You can view the *tunnel_id* with the **show vpdn** command.)

The **clear vpdn group** command removes all the **vpdn group** commands from the configuration. The **clear vpdn username** command removes all the **vpdn username** commands from the configuration. The **clear vpdn** command removes all **vpdn** commands from the configuration.

You can troubleshoot PPTP traffic with the **debug ppp** and **debug vpdn** commands.

vpdn Command Example

The following example shows a simple configuration, which lets a Windows PPTP client dial in without any authentication (not recommended). Refer to the **vpdn** command page in [Chapter 5, “Command Reference,”](#) for more examples and descriptions of the **vpdn** commands and the command syntax.

The **ip local pool** command specifies the IP addresses assigned to each VPN Client as they log in to the network. The Windows client can Telnet to host 192.168.0.2 through the global IP address 209.165.201.2 in the **static** command statement. The **access-list** command statement permits Telnet access to the host.

```
ip local pool my-addr-pool 10.1.1.1-10.1.1.254
vpdn group 1 accept dialin pptp
vpdn group 1 client configuration address local my-addr-pool
vpdn enable outside
static (inside, outside) 209.165.201.2 192.168.0.2 netmask 255.255.255.255
access-list acl_out permit tcp any host 209.165.201.2 eq telnet
access-group acl_out in interface outside
```

SNMP

The **snmp-server** command causes the PIX Firewall to send SNMP traps so that the PIX Firewall can be monitored remotely. Use **snmp-server host** command to specify which systems receive the SNMP traps.

This section includes the following topics:

- [Introduction](#)
- [MIB Support](#)
- [SNMP Usage Notes](#)
- [SNMP Traps](#)
- [Compiling Cisco Syslog MIB Files](#)
- [Using the Firewall and Memory Pool MIBs](#)

Introduction

The PIX Firewall SNMP MIB-II groups available are System and Interfaces. The Cisco Firewall MIB and Cisco Memory Pool MIB are also available.

All SNMP values are read only (RO).

Using SNMP, you can monitor system events on the PIX Firewall. SNMP events can be read, but information on the PIX Firewall cannot be changed with SNMP.

The PIX Firewall SNMP traps available to an SNMP management station are as follows:

- Generic traps:
 - Link up and link down (cable connected to the interface or not; cable connected to an interface working or not working)
 - Cold start
 - Authentication failure (mismatched community string)
- Security-related events sent via the Cisco Syslog MIB:
 - Global access denied
 - Failover syslog messages
 - syslog messages

Use CiscoWorks for Windows or any other SNMP V1, MIB-II compliant browser to receive SNMP traps and browse an MIB. SNMP traps occur at UDP port 162.

MIB Support

**Note**

The PIX Firewall does not support browsing of the Cisco syslog MIB.

You can browse the System and Interface groups of MIB-II. Browsing an MIB is different from sending traps. Browsing means doing an **snmpget** or **snmpwalk** of the MIB tree from the management station to determine values.

MIB Support

The Cisco Firewall MIB and Cisco Memory Pool MIB are available.

PIX Firewall does not support the following in the Cisco Firewall MIB:

- cfwSecurityNotification NOTIFICATION-TYPE
- cfwContentInspectNotification NOTIFICATION-TYPE
- cfwConnNotification NOTIFICATION-TYPE
- cfwAccessNotification NOTIFICATION-TYPE
- cfwAuthNotification NOTIFICATION-TYPE
- cfwGenericNotification NOTIFICATION-TYPE

For more information on MIB Support, go to:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

SNMP Usage Notes

- The MIB-II ifEntry.ifAdminStatus object returns 1 if the interface is accessible and 2 if you administratively shut down the interface using the **shutdown** option of the **interface** command.
- The SNMP “ifOutUcastPkts” object now correctly returns the outbound packet count.
- Syslog messages generated by the SNMP module now specify the interface name instead of an interface number.

SNMP Traps

Traps are different than browsing; they are unsolicited “comments” from the managed device to the management station for certain events, such as link up, link down, and syslog event generated.

An SNMP object ID (OID) for PIX Firewall displays in SNMP event traps sent from the PIX Firewall. PIX Firewall provides system OID in SNMP event traps & SNMP mib-2.system.sysObjectID variable based on the hardware platform:

[Table 3-2](#) lists the system OID in PIX Firewall platforms.

Table 3-2 System OID in PIX Firewall Platform

PIX Platform	System OID
PIX 506	1.3.6.1.4.1.9.1.389
PIX 515	.1.3.6.1.4.1.9.1.390
PIX 520	.1.3.6.1.4.1.9.1.391
PIX 525	1.3.6.1.4.1.9.1.392
PIX 535	1.3.6.1.4.1.9.1.393
others	1.3.6.1.4.1.9.1.227 (original PIX Firewall OID)

Two mechanisms work with SNMP, PIX Firewall responds to an SNMP request from a management station and the PIX Firewall sends a trap, which is an event notification. PIX Firewall supports two types of traps, generic and syslog traps.

Receiving Requests and Sending Syslog Traps

Follow these steps to receive requests and send traps from the PIX Firewall to an SNMP management station:

-
- Step 1** Identify the IP address of the SNMP management station with the **snmp-server host** command.
- Step 2** Set the **snmp-server** options for **location**, **contact**, and the **community** password as required.
- If you only want to send the cold start, link up, and link down generic traps, no further configuration is required.
- If you only want to receive SNMP requests, no further configuration is required.
- Step 3** Add an **snmp-server enable traps** command statement.
- Step 4** Set the logging level with the **logging history** command; for example:
- ```
logging history debugging
```
- We recommend that you use the **debugging** level during initial set up and during testing. Thereafter, set the level from **debugging** to a lower value for production use.
- (The **logging history** command sets the severity level for SNMP syslog messages.)
- Step 5** Start sending syslog traps to the management station with the **logging on** command.
- Step 6** To disable sending syslog traps, use the **no logging on** command or the **no snmp-server enable traps** command.
- 

## Compiling Cisco Syslog MIB Files

To receive security and failover SNMP traps from the PIX Firewall, compile the Cisco SMI MIB and the Cisco syslog MIB into your SNMP management application. If you do not compile the Cisco syslog MIB into your application, you only receive traps for link up or down, firewall cold start and authentication failure.

You can get the Cisco MIB files on the Web from the following sites:

- <http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

Follow these steps to compile Cisco syslog MIB files into your browser using CiscoWorks for Windows (SNMPc):

- 
- Step 1** Get the Cisco syslog MIB files.
- Step 2** Start SNMPc.
- Step 3** Click **Config>Compile MIB**.
- Step 4** Scroll to the bottom of the list, and click the last entry.
- Step 5** Click **Add**.
- Step 6** Find the Cisco syslog MIB files.



**Note** With certain applications, only files with a .mib extension may show in the file selection window of the SNMPc. The Cisco syslog MIB files with the .my extension will not be shown. In this case, you should manually change the .my extension to a .mib extension.

- Step 7** Select CISCO-FIREWALL-MIB.my (CISCO-FIREWALL-MIB.mib) and click **OK**.
- Step 8** Scroll to the bottom of the list, and click the last entry.
- Step 9** Click **Add**.
- Step 10** Find the file CISCO-MEMORY-POOL-MIB.my (CISCO-MEMORY-POOL-MIB.mib) and click **OK**.
- Step 11** Scroll to the bottom of the list, and click the last entry.
- Step 12** Click **Add**.
- Step 13** Find the file CISCO-SMI.my (CISCO-SMI.mib) and click **OK**.
- Step 14** Scroll to the bottom of the list, and click the last entry.
- Step 15** Click **Add**.
- Step 16** Find the file CISCO-SYSLOG-MIB.my (CISCO-SYSLOG-MIB.mib) and click **OK**.
- Step 17** Click **Load All**.
- Step 18** If there are no errors, restart SNMPc.

These instructions are only for SNMPc (CiscoWorks for Windows).

## Using the Firewall and Memory Pool MIBs

The Cisco Firewall and Memory Pool MIBs let you poll failover and system status.

This section contains the following topics:

- [ipAddrTable Notes](#)
- [Viewing Failover Status](#)
- [Verifying Memory Usage](#)
- [Viewing The Connection Count](#)
- [Viewing System Buffer Usage](#)

In the tables that follow in each section, the meaning of each returned value is shown in parentheses.

### ipAddrTable Notes

- Use of the SNMP ip.ipAddrTable entry requires that all interfaces have unique addresses. If interfaces have not been assigned IP addresses, by default, their IP addresses are all set to 127.0.0.1. Having duplicate IP addresses causes the SNMP management station to loop indefinitely. The workaround is to assign each interface a different address. For example, you can set one address to 127.0.0.1, another to 127.0.0.2, and so on.

SNMP uses a sequence of GetNext operations to traverse the MIB tree. Each GetNext request is based on the result of the previous request. Therefore, if two consecutive interfaces have the same IP 127.0.0.1 (table index), the GetNext function returns 127.0.0.1, which is correct; however, when SNMP generates the next GetNext request using the same result (127.0.0.1), the request is identical to the previous one, which causes the management station to loop infinitely. For example:

```
GetNext(ip.ipAddrTable.ipAddrEntry.ipAdEntAddr.127.0.0.1)
```

In SNMP protocol, the MIB table index must be unique for the agent to identify a row from the MIB table. The table index for ip.ipAddrTable is the PIX Firewall interface IP address, so the IP address must be unique; otherwise, the SNMP agent will get confused and may return information of another interface (row), which has the same IP (index).

## Viewing Failover Status

The Cisco Firewall MIB's cfwHardwareStatusTable allows you to determine whether failover is enabled and which unit is active. The Cisco Firewall MIB indicates failover status by two rows in the cfwHardwareStatusTable object. From the PIX Firewall command line, you can view failover status with the **show failover** command. You can access the object table from the following path:

```
.iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoFirewallMIB.
ciscoFirewallMIBObjects.cfwSystem.cfwStatus.cfwHardwareStatusTable
```

Table 3-3 lists which objects provide failover information.

**Table 3-3 Failover Status Objects**

| Object                           | Object Type     | Row 1: Returned if Failover is Disabled | Row 1: Returned if Failover is Enabled                             | Row 2: Returned if Failover is Enabled                          |
|----------------------------------|-----------------|-----------------------------------------|--------------------------------------------------------------------|-----------------------------------------------------------------|
| cfwHardwareType<br>(table index) | Hardware        | 6 (If primary unit)                     | 6 (If primary unit)                                                | 7 (If secondary unit)                                           |
| cfwHardwareInformation           | SnmpAdminString | blank                                   | blank                                                              | blank                                                           |
| cfwHardwareStatusValue           | HardwareStatus  | 0 (Not used)                            | active or 9 (If Active unit)<br>or standby or 10 (If Standby unit) | active or 9 (If Active unit) or standby or 10 (If Standby unit) |
| cfwHardwareStatusDetail          | SnmpAdminString | Failover Off                            | blank                                                              | blank                                                           |

In the HP OpenView Browse MIB application's "MIB values" window, if failover is disabled, a sample MIB query yields the following information:

```
cfwHardwareInformation.6 :
cfwHardwareInformation.7 :
cfwHardwareStatusValue.6 :0
cfwHardwareStatusValue.7 :0
cfwHardwareStatusDetail.6 :Failover Off
cfwHardwareStatusDetail.7 :Failover Off
```

From this listing, the table index, cfwHardwareType, appears as either .6 or .7 appended to the end of each of the subsequent objects. The cfwHardwareInformation field is blank, the cfwHardwareStatusValue is 0, and the cfwHardwareStatusDetail contains **Failover Off**, which indicates the failover status.

When failover is enabled, a sample MIB query yields the following information:

```
cfwHardwareInformation.6 :
```

```

cfwHardwareInformation.7 :
cfwHardwareStatusValue.6 : active
cfwHardwareStatusValue.7 : standby
cfwHardwareStatusDetail.6 :
cfwHardwareStatusDetail.7 :

```

In this listing, only the `cfwHardwareStatusValue` contains values, either **active** or **standby** to indicate the status of each unit.

## Verifying Memory Usage

You can determine how much free memory is available with the Cisco Memory Pool MIB. From the PIX Firewall command line, memory usage is viewed with the **show memory** command. The following is sample output from the **show memory** command:

```

show memory
16777216 bytes total, 5595136 bytes free

```

You can access the MIB objects from the following path:

```

.iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoMemoryPoolMIB.
ciscoMemoryPoolObjects.ciscoMemoryPoolTable

```

Table 3-4 lists which objects provide memory usage information.

**Table 3-4** Memory Usage Objects

| Object                               | Object Type          | Returned Value                                                                         |
|--------------------------------------|----------------------|----------------------------------------------------------------------------------------|
| ciscoMemoryPoolType<br>(table index) | CiscoMemoryPoolTypes | 1 (Processor memory)                                                                   |
| ciscoMemoryPoolName                  | DisplayString        | <b>PIX system memory</b>                                                               |
| ciscoMemoryPoolAlternate             | Integer32            | 0 (No alternate memory pool)                                                           |
| ciscoMemoryPoolValid                 | TruthValue           | <b>true</b> (Means that the values of the remaining objects are valid)                 |
| ciscoMemoryPoolUsed                  | Gauge32              | <i>integer</i> (Number of bytes currently in use—the total bytes minus the free bytes) |
| ciscoMemoryPoolFree                  | Gauge32              | <i>integer</i> (Number of bytes currently free)                                        |
| ciscoMemoryPoolLargestFree           | Gauge32              | 0 (Information not available)                                                          |

In the HP OpenView Browse MIB application's "MIB values" window a sample MIB query yields the following information:

```

ciscoMemoryPoolName.1 :PIX system memory
ciscoMemoryPoolAlternate.1 :0
ciscoMemoryPoolValid.1 :true
ciscoMemoryPoolUsed.1 :12312576
ciscoMemoryPoolFree.1 :54796288
ciscoMemoryPoolLargestFree.1 :0

```

From this listing, the table index, `ciscoMemoryPoolName`, appears as the `.1` value at the end of each subsequent object value. The `ciscoMemoryPoolUsed` object lists the number of bytes currently in use, **12312576**, and the `ciscoMemoryPoolFree` object lists the number of bytes currently free **54796288**. The other objects always list the values described in Table 3-4.

## Viewing The Connection Count

You can view the number of connections in use from the `cfwConnectionStatTable` in the Cisco Firewall MIB. From the PIX Firewall command line, you can view the connection count with the **show conn** command. The following is sample output from the **show conn** command to demonstrate where the information in `cfwConnectionStatTable` originates:

```
show conn
15 in use, 88 most used
```

The `cfwConnectionStatTable` object table can be accessed from the following path:

```
.iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoFirewallMIB.
ciscoFirewallMIBObjects.cfwSystem.cfwStatistics.cfwConnectionStatTable
```

Table 3-5 lists which objects provide connection count information.

**Table 3-5 Connection Count Objects**

| Object                                                 | Object Type     | Row 1: Returned Value                                                | Row 2: Returned Value                                                            |
|--------------------------------------------------------|-----------------|----------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <code>cfwConnectionStatService</code><br>(Table index) | Services        | <b>40</b> (IP protocol)                                              | <b>40</b> (IP protocol)                                                          |
| <code>cfwConnectionStatType</code><br>(Table index)    | ConnectionStat  | <b>6</b> (Current connections in use)                                | <b>7</b> (High)                                                                  |
| <code>cfwConnectionStatDescription</code>              | SnmpAdminString | <b>number of connections currently in use by the entire firewall</b> | <b>highest number of connections in use at any one time since system startup</b> |
| <code>cfwConnectionStatCount</code>                    | Counter32       | <b>0</b> (Not used)                                                  | <b>0</b> (Not used)                                                              |
| <code>cfwConnectionStatValue</code>                    | Gauge32         | <i>integer</i> (In use number)                                       | <i>integer</i> (Most used number)                                                |

In the HP OpenView Browse MIB application's "MIB values" window a sample MIB query yields the following information:

```
cfwConnectionStatDescription.40.6 :number of connections currently in use by the entire firewall
cfwConnectionStatDescription.40.7 :highest number of connections in use at any one time since system startup
cfwConnectionStatCount.40.6 :0
cfwConnectionStatCount.40.7 :0
cfwConnectionStatValue.40.6 :15
cfwConnectionStatValue.40.7 :88
```

From this listing, the table index, `cfwConnectionStatService`, appears as the **.40** appended to each subsequent object and the table index, `cfwConnectionStatType`, appears as either **.6** to indicate the number of connections in use or **.7** to indicate the most used number of connections. The `cfwConnectionStatValue` object then lists the connection count. The `cfwConnectionStatCount` object always returns **0** (zero).

## Viewing System Buffer Usage

You can view the system buffer usage from the Cisco Firewall MIB in multiple rows of the `cfwBufferStatsTable`. The system buffer usage provides an early warning of the PIX Firewall reaching the limit of its capacity. On the command line, you can view this information with the **show blocks** command. The following is sample output from the **show blocks** command to demonstrate how `cfwBufferStatsTable` is populated:

```
show blocks
```

| SIZE  | MAX  | LOW  | CNT  |
|-------|------|------|------|
| 4     | 1600 | 1600 | 1600 |
| 80    | 100  | 97   | 97   |
| 256   | 80   | 79   | 79   |
| 1550  | 780  | 402  | 404  |
| 65536 | 8    | 8    | 8    |

You can view `cfwBufferStatsTable` at the following path:

```
.iso.org.dod.internet.private.enterprises.cisco.ciscoMgmt.ciscoFirewallMIB.
ciscoFirewallMIBObjects.cfwSystem.cfwStatistics.cfwBufferStatsTable
```

Table 3-6 lists the objects required to view the system block usage.

**Table 3-6 System Block Usage Objects**

| Object                                          | Object Type        | First Row: Returned Value                                                                                         | Next Row: Returned Value                                                                                                    | Next Row: Returned Value                                                                                          |
|-------------------------------------------------|--------------------|-------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <code>cfwBufferStatSize</code><br>(Table index) | Unsigned32         | <i>integer</i> (SIZE value; for example, <b>4</b> for a 4-byte block)                                             | <i>integer</i> (SIZE value; for example, <b>4</b> for a 4-byte block)                                                       | <i>integer</i> (SIZE value; for example, <b>4</b> for a 4-byte block)                                             |
| <code>cfwBufferStatType</code><br>(Table index) | ResourceStatistics | <b>3</b> (MAX)                                                                                                    | <b>5</b> (LOW)                                                                                                              | <b>8</b> (CNT)                                                                                                    |
| <code>cfwBufferStatInformation</code>           | SnmpAdminString    | <b>maximum number of allocated <i>integer</i> byte blocks</b> ( <i>integer</i> is the number of bytes in a block) | <b>fewest <i>integer</i> byte blocks available since system startup</b> ( <i>integer</i> is the number of bytes in a block) | <b>current number of available <i>integer</i> byte blocks</b> ( <i>integer</i> is the number of bytes in a block) |
| <code>cfwBufferStatValue</code>                 | Gauge32            | <i>integer</i> (MAX number)                                                                                       | <i>integer</i> (LOW number)                                                                                                 | <i>integer</i> (CNT number)                                                                                       |



**Note**

The three rows repeat for every block size listed in the output of the **show blocks** command.

In the HP OpenView Browse MIB application's "MIB values" window a sample MIB query yields the following information:

```
cfwBufferStatInformation.4.3 :maximum number of allocated 4 byte blocks
cfwBufferStatInformation.4.5 :fewest 4 byte blocks available since system startup
cfwBufferStatInformation.4.8 :current number of available 4 byte blocks
cfwBufferStatInformation.80.3 :maximum number of allocated 80 byte blocks
cfwBufferStatInformation.80.5 :fewest 80 byte blocks available since system startup
cfwBufferStatInformation.80.8 :current number of available 80 byte blocks
cfwBufferStatInformation.256.3 :maximum number of allocated 256 byte blocks
cfwBufferStatInformation.256.5 :fewest 256 byte blocks available since system startup
cfwBufferStatInformation.256.8 :current number of available 256 byte blocks
cfwBufferStatInformation.1550.3 :maximum number of allocated 1550 byte blocks
cfwBufferStatInformation.1550.5 :fewest 1550 byte blocks available since system startup
cfwBufferStatInformation.1550.8 :current number of available 1550 byte blocks
cfwBufferStatValue.4.3: 1600
cfwBufferStatValue.4.5: 1600
cfwBufferStatValue.4.8: 1600
cfwBufferStatValue.80.3: 400
cfwBufferStatValue.80.5: 396
cfwBufferStatValue.80.8: 400
cfwBufferStatValue.256.3: 1000
```

```
cfwBufferStatValue.256.5: 997
cfwBufferStatValue.256.8: 999
cfwBufferStatValue.1550.3: 1444
cfwBufferStatValue.1550.5: 928
cfwBufferStatValue.1550.8: 932
```

From this listing, the first table index, `cfwBufferStatSize`, appears as first number appended to the end of each object, such as `.4` or `.256`. The other table index, `cfwBufferStatType`, appears as `.3`, `.5`, or `.8` after the first index. For each block size, the `cfwBufferStatInformation` object identifies the type of value and the `cfwBufferStatValue` object identifies the number of bytes for each value.

## SSH

SSH (Secure Shell) is an application running on top of a reliable transport layer, such as TCP/IP that provides strong authentication and encryption capabilities. PIX Firewall supports the SSH remote shell functionality as provided in SSH version 1. SSH version 1 also works with Cisco IOS software devices. Up to five SSH clients are allowed simultaneous access to the PIX Firewall console.



### Note

You must generate an RSA key-pair for the PIX Firewall before clients can connect to the PIX Firewall console. To use SSH, your PIX Firewall must have a DES or 3DES activation key.

The current method of remotely configuring a PIX Firewall unit involves initiating a Telnet connection to the PIX Firewall to start a shell session and then entering configuration mode. This connection method can only provide as much security as Telnet provides, which is only provided as lower-layer encryption (for example, IPSec) and application security (username/password authentication at the remote host). The PIX Firewall SSH implementation provides a secure remote shell session without IPSec, and only functions as a server, which means the PIX Firewall cannot initiate SSH connections.

For more information, refer to the [aaa](#) and [ssh](#) command pages in [Chapter 5, “Command Reference.”](#)

## Obtaining an SSH Client

The following sites let you download an SSH v1.x client. Because SSH version 1.x and v2 are entirely different protocols and are not compatible, be sure you download a client that supports SSH v1.x.

- Windows 3.1, Windows CE, Windows 95, and Windows NT 4.0—download the free Tera Term Pro SSH v1.x client from the following site:

<http://hp.vector.co.jp/authors/VA002416/teraterm.html>

The TTSSH security enhancement for Tera Term Pro is available at the following site:

<http://www.zip.com.au/~roca/ttssh.html>



### Note

You must download TTSSH to use Tera Term Pro with SSH. TTSSH provides a Zip file you copy to your system. Extract the zipped files into the same folder that you installed Tera Term Pro. For a Windows 95 system, by default, this would be the C:\Program Files\Ttempro folder.

- Linux, Solaris, OpenBSD, AIX, IRIX, HP/UX, FreeBSD, and NetBSD—download the SSH v1.x client from the following site:  
<http://www.openssh.com>
- Macintosh (international users only)—download the Nifty Telnet 1.1 SSH client from the following site:  
<http://www.lysator.liu.se/~jonasw/freeware/niftyssh/>