

Advanced Configurations

This chapter describes how to configure:

- Failover
- WebSENSE URL Filtering
- FTP and URL Logging
- SNMP Traps
- Private Link

Failover

Use the **failover** command without an argument after you connect the optional failover cable between your primary firewall and a secondary firewall. The default is **failover off**. Enter **no failover** in the configuration file for the PIX Firewall if you will not be using the failover feature. Use the **show failover** command to verify the status of the connection and to determine which unit is active.

Note Remove the failover cable before upgrading to a new version of PIX Firewall. Once the new software is installed, configured, and operating properly, then reconnect the cable and reboot the two systems, first the primary and then the secondary. The primary will then automatically update the secondary unit.

Failover is supported only between identical PIX Firewall models running the same software version. For example, failover is not supported between a PIX10000 and a PIX 520.

In PIX Firewall release 4.2, failover IP addresses must be configured on each interface card. The Active unit of the failover pair uses the system IP addresses and the Primary unit's MAC address, while the Standby unit uses the failover IP addresses and the secondary unit's MAC address. The system IP addresses and the failover IP addresses must be on the same subnet with no router between them.

When a failover occurs, each unit changes state. The newly Active unit 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 timeout anywhere on the network.

Both PIX Firewall units in a failover pair must have the same configuration. To accomplish this, always enter configuration changes on the Active unit in a PIX Firewall failover configuration. Use the **write memory** command on the Active unit to save configuration changes to Flash memory (non-volatile memory) on both the active and Standby units. Changes made on the Standby unit are not replicated on the Active unit.

Use the **write standby** command to manually save the configuration of the active failover unit to the standby failover unit from RAM to RAM. The Standby unit must not be configured individually. Commands entered on the Active unit are automatically replicated on the Standby unit. Only use the default configuration initially. You can force an update by using the **write standby** command on the Active unit. If you make changes to the Standby unit, it displays a warning but does not update the Active unit.

To save the configuration of the Active unit to Flash memory (permanent memory) on the Standby unit, use the **write memory** command on the Active unit. The **write memory** command results are replicated on the Standby unit.

Both units in a failover pair communicate through the failover cable. The two units send special failover “hello” packets to each other over all network interfaces and the failover cable every 15 seconds. 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.

The Standby unit does not maintain the state information of each connection. This means that all active connections will be dropped when failover occurs. Client systems must reestablish connections. Additionally, no RIP information is available on the newly Active unit. The newly active PIX Firewall must wait for up to 30 seconds to learn the routing information from the network.

When a failover occurs, syslog messages are generated indicating what happened.

Failover works by passing control to the Standby unit should the Active unit fail. For Ethernet, failover detection should occur within 30 seconds. Token Ring requires additional time for failover.

The markings on the cable let you choose which PIX Firewall unit is primary and which is secondary. You need only connect the failover cable between the PIX Firewall units.

Note The active PIX Firewall does not maintain a copy of the connection state in the Standby unit. If the Active unit fails, network traffic must re-establish previous connections.

SYSLOG messages indicate whether the primary or secondary unit failed. Use the **show failover** command to verify which unit is active.

If you want to force a PIX Firewall to be active or go to standby you can use the **failover active** or **no failover active** command. Use this feature to force a PIX Firewall offline for maintenance or to return a failed unit to service.

Upgrading from PIX Firewall Version 4.1 to Version 4.2

- Step 1** Establish a console connection into the two failover units and perform all configuration changes. Save the PIX Firewall version 4.1 configuration to a blank DOS formatted diskette; write-protect and label it.

Note Keep the configuration diskette in a safe place. The PIX Firewall version 4.1 configuration is required if a downgrade is ever necessary.

- Step 2** Determine which PIX Firewall unit is the Primary unit by inspecting the failover cable. The primary and secondary assignments can be read from labels on each end of the failover cable connected to each PIX Firewall unit.
- Step 3** Power down the Primary unit. The secondary PIX Firewall becomes the Active unit and begins passing traffic using the primary PIX Firewall's IP and MAC addresses.
- Step 4** Remove the failover and network cables from the Primary unit. Do not remove the console cable.
- Step 5** Insert the PIX Firewall version 4.2 diskette into the Primary unit and power up the unit. The PIX Firewall automatically boots from the version 4.2 diskette.
- Step 6** Enter configuration mode and configure failover IP addresses for each interface on the Primary unit.
- Step 7** After configuring the failover IP addresses, enter the **write memory** command to save the configuration changes to Flash memory and to complete the upgrade on the Primary unit.
- Step 8** Power down the primary PIX Firewall and reconnect the failover and network cables.

Note In the following step, you must power up the Primary unit at the same time you power down the secondary unit or a failover negotiation error can occur where both units are competing to be the Active unit. If this occurs, be sure the secondary unit is powered down, and reboot the primary PIX Firewall. Do not power up the secondary unit.

- Step 9** Power up the primary PIX Firewall at the same time you power down the secondary PIX Firewall.
The primary PIX Firewall becomes the Active unit and begins passing network traffic.
- Step 10** Insert the 4.2 disk into the secondary PIX Firewall and power up the unit. Once the system boots, enter configuration mode.
- Step 11** On the secondary unit, enter the **write memory** command to save the configuration changes to Flash memory. This completes the upgrade of the secondary unit.
- Step 12** At the primary (active) PIX Firewall console, enter the **write memory** command to write the Failover IP addresses from the active to the standby PIX Firewall.
- Step 13** On the secondary (standby) unit, enter the **show failover** command to verify the failover IP addresses were written to the standby PIX Firewall unit.

This completes the upgrade procedure.

Upgrading from PIX Firewall Version 4.2(1)

Step 1 Establish a console connection into the two failover units and perform all configuration changes. Save the existing PIX Firewall configuration to a blank DOS formatted diskette; write-protect and label it.

Note Keep the configuration diskette in a safe place. Retaining a copy of the existing PIX Firewall configuration is helpful if a downgrade is ever necessary.


- Step 2** Check the status of the failover pair with the **show failover** command and determine which PIX Firewall unit is the Standby unit.
- Step 3** Insert the latest PIX Firewall software diskette into the Standby unit and enter the **reload** command. This reboots the Standby unit.
- Step 4** After the Standby unit finishes booting, enter the **write memory** command to save the latest software changes to Flash memory. This completes the upgrade on the Standby unit.
- Step 5** Enter configuration mode on the Standby unit, and enter the **failover active** command to force the Standby unit to become the Active unit.
- Step 6** Check the status of the failover pair with the **show failover** command to verify that the previously Standby unit is now the Active unit.
- Step 7** Repeat Steps 3 and 4 for the current Standby unit.
- Step 8** On the Active unit, enter the **write memory** command to save the latest software changes to Flash memory on the both units.

This completes the upgrade procedure.

Failover Cable Pinouts

The pin outs for the cable are shown in Figure 3-1.

Figure 3-1 Failover Cable Pin Outs



12069

Configuring Firewall Units for Failover

Note Always enter configuration changes on the Active unit. Configuration changes entered on the Standby unit are not saved to the Active unit.

The following guidelines apply to configuring failover on the Active unit:

- The unit that has the cable end labeled “primary” becomes the default Primary unit.
- Configure a failover IP address for each interface on the Active unit using the **ip address** command. From the Active unit, configure a failover IP address for each interface on the Standby unit using the **failover ip address** command.

Note When a failover occurs, each unit changes state. The newly Active unit assumes the IP addresses and MAC addresses of the previously Active unit and begins accepting traffic. The new Standby unit assumes the IP addresses and MAC addresses of the unit that was previously the Standby unit. Because network devices see no change in these addresses, no ARP entries change or time out anywhere on the network. The failover feature uses the different IP addresses for the Standby unit as a way of testing the interfaces of both units, ensuring that failover can occur over the physical network.

- Use the **write standby** command to save configuration change to the Standby unit. Saving the configuration changes to memory also saves the changes to Flash memory on the Standby unit.

Frequently Asked Failover Questions

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

- Can the failover feature work without using the failover cable?

No, failover will not work without the cable. If you run without the failover cable you are essentially running two separate PIX Firewall units. This will result in a bridge loop and flood the network. The failover cable is an essential part of failover.

- Can I extend the length of the failover cable with modems or line extenders?

No, the cable cannot be extended using modems or other RS-232 line extenders. Part of what the failover cable does is indicate the presence and power status of the other unit. When you place line extenders in this path you are relaying the status of the line extender rather than of the other PIX Firewall unit.

- 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. The two units do not share connection states. Any active connections will be dropped when a failover switch occurs. The clients must re-establish the connections through the newly Active 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 active PIX Firewall experiences a power failure, the standby PIX Firewall comes up in active mode. If the Primary unit is powered up again it will become the Standby unit.

- What happens if an interface card is disconnected?

When the active PIX Firewall is failed by disconnecting the interface (cable pull), the standby PIX Firewall becomes the Active unit. When the interface is plugged back in, the unit automatically recovers, and its status is changed from failed to standby.

- Does failover work in a switched environment?

Yes, if you are running PIX Firewall version 4.2.x on *both* units.

- 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.
- 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 powered down (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 boot-up, at which point an unplugged cable will force the unit active. If both units are powered up 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. 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.

Failover Interface Tests

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 is failed. This series of tests will begin when hello messages are not heard for two consecutive 15-second intervals. 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 is 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, they 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.

- 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, the hub or switch is 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.

Failover SYSLOG Messages

In the messages that follow, *PI*S can be either Primary or Secondary depending on which PIX Firewall is sending the message. Failover messages always have a SYSLOG priority level of 2, which indicates a critical condition. Refer to the **logging** command description for more information on SYSLOG messages.

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.

The SYSLOG messages sent to record failover events are:

- System okay messages:
 - “*PI*S: Cable OK.”
 - “*PI*S: Disabling failover.” The **no failover** command was entered.
 - “*PI*S: Enabling Failover.” Either a PIX Firewall is booting that has the **failover** command in its configuration file or the **failover** command was just entered in the current configuration.
 - “*PI*S: Mate ifc *number* OK.” The interface (ifc) is now working correctly after being brought back online after a failure. The *number* is either **0** for the inside network interface or **1** for the outside interface.
- Cabling problem messages:
 - “*PI*S: Bad cable.” The cable is connected on both units, but the failover cable has developed a wiring problem.
 - “*PI*S: Cable not connected my side.” The cable on the current PIX Firewall is not connected.
 - “*PI*S: Cable not connected other side.” The cable on the current unit is connected, but the connector on the other unit is disconnected.
 - “*PI*S: Error reading cable status.” The cable state cannot be determined. Ensure that all connectors are securely attached.
- Failure in process messages:
 - “*PI*S: No response from mate.” The other PIX Firewall has not responded in the last 30 seconds.
 - “*PI*S: Power failure other side.” The other unit has lost power.
 - “*PI*S: Mate ifc *number* failed.” The interface (ifc) for the other unit failed.
- Status messages:
 - “*PI*S: Switching to ACTIVE.” The other unit has brought the network back online and is receiving connections. This message also occurs if you force a unit to active with the **failover active** command, or force the other unit inactive with the **no failover active** command.
 - “*PI*S: Switching to STANDBY.” The unit is inactive as a result of entering **no failover active** on the unit or by entering **failover active** on the other unit.

WebSENSE URL Filtering

If your network has a WebSENSE server on any network interface, you can provide URL filtering through the PIX Firewall.

To configure the PIX Firewall to use WebSENSE:

- Step 1** Specify the interface and IP address of the WebSENSE server with the **url-server** command as shown in this example:

```
url-server (dmz) host 192.168.1.42 timeout 10
```

In this example, the WebSENSE host is on the dmz interface at IP address 192.168.1.42. A timeout value of 10 seconds is specified as maximum allowed idle time before the PIX Firewall switches to the next WebSENSE server.

- Step 2** Use the **filter url http** command to tell the PIX Firewall how to filter requests. For example, to filter requests for all hosts, use:

```
filter url http 0 0 0 0 allow
```

Note The **allow** option in the **filter** command is crucial to the use of PIX Firewall URL filtering feature. If you use the **allow** option and the WebSENSE server goes offline, the PIX Firewall lets all Web requests continue without filtering. If the **allow** option is not specified, all port 80 Web requests are stopped until the server is back online.

- Step 3** If you want to disable URL filtering, use the **no filter url** command.

FTP and URL Logging

You can log FTP commands and WWW URLs when syslog is enabled. FTP and URL messages are logged at syslog level 7.

Refer to the section “Step 15 - Enable Syslog” in Chapter 2, “Configuring the PIX Firewall,” for more information on how to view syslog messages on a server, console session, or via Telnet to the console.

Use the **show fixup** command to ensure that the **fixup protocol** commands for FTP and HTTP are present in the configuration:

```
fixup protocol http 80
fixup protocol ftp 21
```

These commands are in the default configuration.

The sections that follow provide sample output displays for each logging type.

Sample URL Log

The following is an example of a URL logging syslog message:

```
192.168.69.71 accessed URL 10.0.0.1/secrets.gif
```

Sample FTP Log

The following are examples of FTP logging syslog messages:

```
192.168.69.42 Retrieved 10.0.0.42:feathers.tar
192.168.42.54 Stored 10.0.42.69:privacy.zip
```

You can view these messages at the PIX Firewall console with the **show logging** command.

SNMP Traps

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

Note The PIX Firewall does not support browsing of the Cisco syslog MIB. The only MIBs you can browse are System and Interfaces.

Browsing a MIB is different from sending traps. Browsing means doing an **snmpget** or **snmpwalk** of the MIB tree from the management station to determine values. Traps are different; they are unsolicited “comments” from the managed device to the management station for certain events, such as link up, link down, syslog event generated, and so on.

To send traps to an SNMP server:

- Step 1** Identify the IP address of the SNMP server with the **snmp-server host** command.
- Step 2** Set the **snmp-server** options for **location**, **contact**, and the **community** password as required.
- Step 3** Set the logging level with the **logging trap** command; for example:

```
logging trap debugging
```

Cisco recommends 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.
- Step 4** Start sending syslog traps to the server with the **logging on** command.

The PIX Firewall SNMP MIB-II groups available are System and Interfaces.

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 server are:

- Link up and link down (cable on outside interface working or not working)
- Warm and cold start
- Failover syslog messages
- Security-related events sent via the Cisco Enterprise MIB:
 - Global access denied
 - syslog messages

Use CiscoWorks for Windows (Product Number CWPC-2.0-WIN) or any other SNMP V1, MIB-II compliant browser to receive SNMP traps and browse a MIB. SNMP traps occur at UDP port 162.

Compiling Cisco Syslog Enterprise MIB Files

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

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

<http://www.cisco.com/public/mibs/v2/CISCO-SYSLOG-MIB.my>

To compile Cisco syslog Enterprise MIB files into your browser using CiscoWorks for Windows (SNMPc), complete the following steps:

- Step 1** Get the Cisco syslog Enterprise MIB files.
- Step 2** Start SNMPc.
- Step 3** Select **Config>Compile MIB**.
- Step 4** Scroll to the bottom of the list, and select the last entry.
- Step 5** Click the **Add** button.
- Step 6** Find the file CISCO-SMI.my and click **OK**.
- Step 7** Scroll to the bottom of the list, and select the last entry.
- Step 8** Click the **Add** button again.
- Step 9** Find the file CISCO-syslog-MIB.my and click **OK**.
- Step 10** Click **Load All**.
- Step 11** If there are no errors, restart SNMPc.

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

Private Link

The **link** command creates an encrypted path between Private Link-equipped PIX Firewall units. You can specify up to seven encryption keys for data access between the local unit and the remote unit. The key-ID and key values must be the same on each side of the Private Link. Once you specify the same keys on both sides of the connection, the systems alert each other when a new key takes effect. You can use the **age** command to specify the number of minutes that a key is in effect.

Note After using the **link** command to add or delete link entries, use the **write memory** command to store the configuration, and then reboot the PIX Firewall.

Specify the **link** command once for each key you want to specify; for example, if you want seven keys, enter the **link** command in the configuration seven times.

The PIX Firewall Private Link consists of an encryption card and software that permits the PIX Firewall units to provide encrypted communications across an unsecure network such as the Internet.

The PIX Firewall allows up to 256 Private Links and up to 512 link paths.

At least two PIX Firewall units are required to use Private Link and each system must have the same hardware and software versions.

Private Link works by checking packets that arrive at the PIX Firewall inside interface. If a route link previously created by the **linkpath** command exists that matches the destination network address, the packet is encrypted and encapsulated in an AH/ESP frame. The frame has a destination address of the remote PIX Firewall and a source address of the local PIX Firewall. When the packet arrives at the remote PIX Firewall unit, the data in the packet is decrypted and then sent through the designated interface to the original IP address specified. No translation takes place on packets that traverse the PIX Firewall Private Link. The addressing and data remains completely unchanged.

If you use the **link** command to change the interface on which a Private Link tunnel terminates, you must reboot the PIX Firewall on which you made the change. For example, if the Private Link tunnel terminates on the perimeter interface of the foreign PIX Firewall and you change it to terminate on the inside interface of the foreign PIX Firewall, you must reboot the local PIX Firewall on which you changed the configuration.

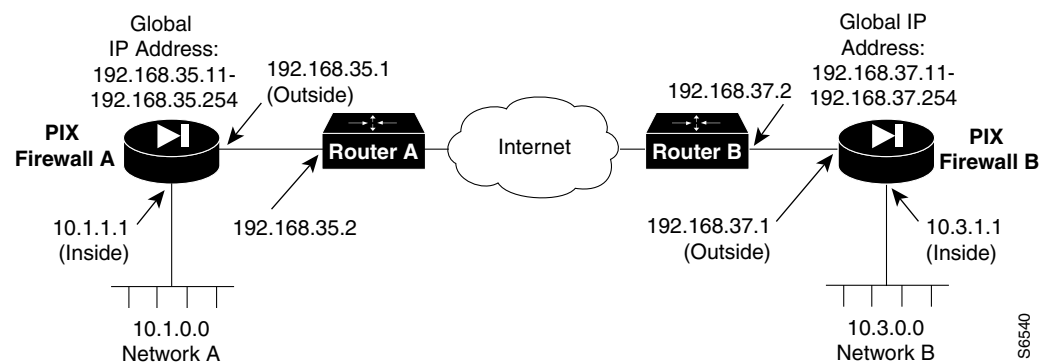
You can manage remote PIX Firewall units through the Private Link interface.

You can use the **linkpath 0.0.0.0 0.0.0.0 foreign_external_ip** command to route all outbound traffic on a foreign PIX Firewall to a central PIX Firewall. However, this use has two caveats: there can be only one central PIX Firewall and the other PIX Firewall units must be satellites to it. This implies that the satellites only relay connections to the central and do not communicate among themselves. The second caveat is that the **linkpath 0 0** command overrides the default route on the outside interface of the satellite PIX Firewall causing all outbound traffic to flow over Private Link to the central PIX Firewall unit. One use of this feature is when access to the Internet is controlled through one PIX Firewall and the other PIX Firewall units feed their Internet traffic to this one site. This could occur when a central processing facility wants to manage all the Internet IP addresses, let the internal networks use any IP numbering scheme, and have local PIX Firewall units protecting individual departments or sites.

Configuring Private Link

To configure a Private Link, refer to the example shown in Figure 3-2.

Figure 3-2 Example Private Link Network Diagram



Before configuring Private Link, you would initially configure the systems using the standard commands.

When you configure a Private Link, follow these steps:

Step 1 Agree on up to seven hexadecimal encryption keys for use between the PIX Firewall Private Link local and remote units; for example, one key could be like the hexadecimal value **fadebacbeebee**. Be sure to select unique keys that are difficult to guess. The key can be up to 56 bits in length (14 hexadecimal digits).

Step 2 Use the **link** command to create an encrypted link for each key you want to specify.

Step 3 Use the **linkpath** command to specify the IP address of the network on the inside of the remote firewall.

Step 4 On PIX Firewall A, in Figure 3-2, enter these commands to configure the Private Link:

```
link 192.168.37.1 1 fadebacfadebac
link 192.168.37.1 2 bacfadefadebac
link 192.168.37.1 3 baabaaafadebac
link 192.168.37.1 4 beebееefadebac
linkpath 10.3.0.0 255.255.255.0 192.168.37.1
```

Step 5 On PIX Firewall B, enter these commands:

```
link 192.168.35.1 1 fadebacfadebac
link 192.168.35.1 2 bacfadefadebac
link 192.168.35.1 3 baabaaafadebac
link 192.168.35.1 4 beebееefadebac
linkpath 10.1.0.0 255.255.255.0 192.168.35.1
```

Note Use random keys, not the ones shown in this guide.

Step 6 Test the connection to each foreign PIX Firewall with the **ping** command.

Note You can ping the PIX Firewall interfaces on the foreign PIX Firewall unit. Use hosts on the network to ensure that PIX Firewall interfaces are reachable.

Step 7 After configuring the **link** and **linkpath** commands, if a **ping inside** command to the inside address of the remote PIX Firewall does not work, enter the **show link** command and look at packets in and out. If both are at 0 that means the link is up, but traffic is not being routed to the inside interface of the local PIX Firewall.

Step 8 Proceed to the router closest to the PIX Firewall on the inside, and look at the routing table. If there is not a route to the remote PIX Firewall network, add a static route, or turn RIP on at the PIX Firewall.

When you Telnet to the PIX Firewall, and perform a **ping inside**, the packet is not simply generated from the inside address of the PIX Firewall and forwarded across the bus to the outside address and out the encrypted tunnel. Instead the ICMP packet is placed on the inside network, picked up by the closest router, and retransmitted to the PIX Firewall, where it is then picked up, encrypted and sent across the link to the remote box.

