

How Failover Works on the Cisco Secure PIX Firewall

Document ID: 5220

Interactive: This document offers customized analysis of your Cisco device.

Contents

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Background Information

Failover Cable

Replicate the PIX Configuration

Failover Monitoring

Fail Back

Interface Testing

Hardware Decision Table

Operating Failover in Switched Environments

Stateful Failover

Failover Commands

Sample Output from the show failover Command

- Example: Normal Failover

- Example: Failover Monitoring Has Not Begun

- Example: Unit Failure

- Example: Stateful Failover

LAN-Based Failover

- LAN-Based Failover Diagram

- Minimum Initial Configuration on the Primary PIX

- Minimum Initial Configuration on the Secondary PIX

- Remaining Configuration – Stateful Failover

Frequently Asked Questions

Information to Collect if You Open a Technical Support Case

Related Information

Introduction

The use of a pair of identical PIX devices (model, memory, network interface cards (NICs), operating system versions), high availability can be provided with no operator intervention.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software versions. All PIX models support failover except 501 and 506E models.

Note: This document does not cover software versions 7.0 and later on Cisco PIX 500 Series Security Appliances. Refer to the Configuring Failover chapter of Cisco Security Appliance CLI Configuration Guide, Version 7.0.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to the Cisco Technical Tips Conventions for more information on document conventions.

Background Information

One PIX is considered the "Active" unit while the other is the "Standby" unit. As the name implies, the active unit performs normal network functions while the standby unit monitors, ready to take control if the active unit fails to perform its functionality. If the **show version** command does not show that failover is enabled and you attempt to do failover, contact your local Cisco account team in order to purchase a license upgrade.

For more information on the license upgrade, refer to License Key Upgrade on a Failover Pair.

The two units each have a presence on the network. The active unit uses the system IP address and the MAC addresses of the Primary unit. The Primary unit is determined by the unit that has the end of the failover cable marked "Primary" plugged into it, or the PIX which is configured with the **failover lan unit primary** command. This command is introduced in PIX OS version 6.2. The standby unit uses the failover IP address and the MAC addresses of the Secondary unit. If a switchover occurs, the units swap the IP address and MAC addresses they use in order to replace each other's presence on the network. This action is invisible to the network. The IP to MAC address relationships remain exactly the same. Therefore, no ARP tables in the network need to time out or be changed. No other piece of network equipment needs to know about the redundancy or that a switchover occurred. Note that the system IP and the failover IP addresses must be on the same subnet, so there is the possibility that there might not be a router between the two units.

Failover Cable

The failover cable is the only additional hardware required to support PIX failover. In PIX 6.2 and later, you can also achieve failover with or without a failover cable. The failover cable is a modified RS-232 serial link cable with a speed setting of 9600 baud.

Note: In PIX Software Release 5.2 (5.1.2.201), the speed is changed to 115.2 K baud. Also, the failover cable cannot be extended.

The basic failover communication is through the failover cable or through the LAN interface which is configured with the **failover lan interface interface_name** command in PIX OS version 6.2 and later. Failover communication through the failover cable is message-based and needs to be reliable. Every message sent is acknowledged (ACKed). If a message is not ACKed by the other PIX in three seconds, the message is retransmitted. After five retransmissions without an ACK (for a total of 15 seconds), a failover condition is triggered by the standby PIX.

Typical failover communication through the failover cable includes:

- MAC addresses exchange
- Hello (a keep-alive)
- State (Active/Standby)
- Network Link Status
- Configuration Replication

Replicate the PIX Configuration

The two units must have the exact same configuration and must run the same software version. This is easily accomplished, since configuration replication occurs over the failover cable, or from the LAN interface configured with **failover lan interface interface_name** command, from the active unit to the standby unit in these ways:

- When the standby unit completes its initial boot-up, the active unit replicates its entire configuration to the standby unit. This occurs if you use a failover cable because you need the initial configuration on both the primary and secondary units in order to identify them as primary and secondary units. This feature has been introduced to overcome the serial cable length and speed.
- As commands are entered on the active unit, they are sent across to the standby unit.
- When you enter the **write standby** command on the active unit, you force the entire configuration to memory on the standby unit.

The configuration replication does a "memory-to-memory" copy. Once this completes, you need to issue a **write memory** command on the active unit in order to write the configuration into the Flash memory of the standby unit. Both "sync started" and "sync completed" console messages are displayed during this operation. Large configurations can take awhile to transfer. If a switchover occurs during replication, the new active PIX has only a partial configuration. The unit then reboots itself to recover the configuration from the Flash or re-sync by the other unit.

The configuration replication only occurs from the active unit to the standby unit. Changes made to the standby unit do not pass to the active unit.

Failover Monitoring

There is a failover poll interval of 15 seconds (after version 5.0 it is configurable) to monitor network activity, failover communications, and the power status. A failure of any of these parameters on the active unit causes the standby unit to take active control. Whenever a unit is determined to have failed, it shuts down its network interfaces.

The two units send special failover "hello" packets to each other over the failover cable and all interfaces every 15 seconds (excludes those that are administratively shutdown). If either unit does not hear the "hello" on an interface for two consecutive poll checks, the PIX puts that LAN interface into testing mode in order to determine where the fault lies. If a standby PIX does not receive a "hello" from the failover cable for two consecutive poll checks, the standby PIX initiates a switchover and declares the other PIX failed. If the active PIX does not hear the "hello" messages, it stays active and sets the other PIX as failed.

A network interface is placed in testing mode if a "hello" packet is not received. A network interface test is non-intrusive. This means that while it is in testing mode, it still attempts to pass normal traffic. The testing process consists of four individual tests (NIC status test, network activity test, Address Resolution Protocol (ARP) test and PING test) geared towards the stimulation of network traffic. If an interface that is in testing mode can receive traffic, it is considered operational. If it can hear other network traffic, it is assumed the error must be with the other unit not able to send the "hello" packet. This results in failing the other unit. If it

is determined that the testing unit cannot receive network traffic while the other can, the testing unit fails itself.

In addition to monitoring all network interfaces, failover also monitors the power status of the other unit, as well as the status of the failover cable itself. The failover cable provides the ability to detect if the other unit is plugged in and powered on. If the cable is unplugged from either unit, switching is disabled. If an active unit loses power, the standby unit takes over within 15 seconds. A unit in the "failed" state waits 15 seconds, and then tries to transition to the "standby" state. If the transition triggers a failure, the unit fails again. You can use the **failover reset** command in order to manually reset the PIX from the failed to standby state. If the transition triggers a failure, the unit fails again. A PIX in the failed state cannot switch into active state.

If the failure is due to a "link down" condition on an interface, a "link up" condition clears the failed state (for example, if an interface is unplugged and then later plugged in).

Note: Refer to Configuring Failover for detailed information on PIX version 7.0 failover features.

Fail Back

Whenever a failure or switch occurs, syslog messages are generated that indicate what happened. Fail back to the primary unit is not forced. Fail back is not a forced activity as there is no reason to switch active and standby roles. Therefore, when a failed primary unit is fixed and brought back on line, it does not automatically resume as the active unit. In order to force a unit to be the active unit, use the **failover active** command on the standby unit or the **no failover active** command on the active unit. If Stateful Failover is used, then connection state information passes from the active unit to the standby unit. Otherwise, the state information is not tracked and sessions must be reestablished by applications. This means all active connections drop after a switchover. Because the newly active unit assumes the same IP and MAC address as the previously active unit, no ARP entries need to change or timeout anywhere in the network.

In EFT 5.0 and later, the 15 seconds failover poll interval is changed to be configurable. The interval can be set between 3 to 15 seconds, recognizing the variance in detecting failure of different interface cards.

Interface Testing

In the event the "hello" packets are not received on an interface, or an interface waits for "hello" more than 2.5 minutes after the other interface went into normal state, the interface is placed in "testing" mode (if the interface is not shutdown and link status is up). When this occurs, the other unit is informed through the failover cable that the interface is in testing mode. While an interface is in testing mode, normal traffic can flow, provided the interface functions properly. Testing is started only if an error condition has occurred and is therefore based on the idea that "if I am okay, then you must be failed." Testing consists of four consecutive tests:

1. NIC Status Test

This test is a Link Up/Down check of the NIC itself. If an interface card is not plugged into an operational network, it is considered failed.

2. Network Activity Test

This test is a "received network activity" test. The unit 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 unit performs an ARP test.

3. ARP Test

In the ARP test, the ARP cache of the unit is read for the ten most recently acquired entries. Then, one at a time, the unit sends ARP requests to these machines, in an attempt 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 unit performs the Ping test.

4. Ping Test

In order to perform the Ping test, the unit sends 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.

At the beginning of each test, both units clear the receive count for the interface. At the conclusion of each test, the testing unit first checks to see if it has received any traffic. If so, it considers itself operational and fails the other unit. If not, it asks the other unit if it has received any traffic. If it has, the testing unit considers itself failed. If neither unit has received any traffic, the testing moves to the next test. If at any time the asking unit does not hear the results of the test from the other unit, it considers the other unit to be failed just as if it didn't hear the "hello" message over the failover cable. If an active unit is determined failed, a switchover occurs. If a standby is determined failed, it is not allowed to ever take active control. The results of these tests are sent through syslog by both the active and standby units.

Hardware Decision Table

On each PIX, failover maintains a hardware decision table (HDT) of both PIX devices in order to decide which PIX is the appropriate active device. When there is a NIC interface state change, failover polls the device driver, the local HDT is updated, and the change is sent to the other PIX. Though HDT tables are compared in each poll cycle, it is the second poll that can cause a standby to usurp control from the active by initiating switchover.

Operating Failover in Switched Environments

There are two issues to address in switched environments. First, the switch needs to learn that a particular MAC address has moved from one port to another. Each unit (unless the unit is failed) transmits a series of failover messages on each interface with the use of its new MAC and IP addresses. This allows the switch to update its internal MAC tables. Cisco strongly recommends that customers enable portfast on all switch ports that connect to PIX interfaces. In addition, channeling and trunking need to be disabled on these ports. Thus, if the interface of the PIX goes down during failover, the switch does not have to wait 30 seconds while the port transitions from a listening to learning to forwarding state.

This blocking of network traffic brings us to the second issue. If the "hello" packets that are sent by failover do not get forwarded, each unit thinks something is wrong and begins to test its interfaces. This results in the failure of one unit because the test results are "if I am okay, then you must be failed." In order to get around this problem, any time a switchover takes place, the units enter a "waiting" state. In this state network traffic is free to flow through the active unit, but failover waits for two "hello" messages to be received before it monitors interfaces again. This allows the switch to enter a blocking state without disrupting failover. Once the second "hello" message is heard, failover resumes normal monitoring of its interfaces.

For PIX Software version 5.2 and later, when a device changes state from standby to active, or from active to standby, a gratuitous ARP is set to each network interface to rebroadcast the new IP and MAC addresses.

Stateful Failover

Without retaining PIX stateful information, after a switchover, all existing connections are dropped and the application is required to reinitiate. In the PIX Software 5.0 release, PIX provides stateful failover so that an existing connection can stay up after a switchover.

In order to support the stateful failover, a dedicated LAN interface between the two PIX devices is required. The Logical Update (LU) is the software module that provides transport to PIX applications that support stateful failover. The state update occurs from the active to standby through the LAN interface. The state update sent to the standby PIX is triggered by the application. The LU transport is UDP-like, with no retransmission and no blocking applications to delay normal packet processing. The state update packets are transmitted asynchronously in the background. Nevertheless, the LU protocol is real-time, and it provides error notification and reports missing state updates for monitoring purposes.

Initial state synchronization is performed after configuration replication. This is done by walking through the translation and connection table records. After that, a state update can be triggered.

PIX address translation (xlate, static and dynamic) and connection (conn) records are essential state data, and are passed to the standby unit from the active unit along with other state information. Since failover cannot be prescheduled, the state update for the connection is packet-based. This means every packet passes through the PIX and changes the state of a connection, which can trigger a state update.

TCP state tables are transferred. However, by default HTTP (TCP port 80) is not replicated. In version 6.0 and later, you can use the command **failover replicate http** in order to enforce TCP port 80 state replication. Most UDP state tables are not transferred, with the exception of dynamically opened ports that correspond to multi-channel protocols such as H.323 and VoIP. Therefore, DNS resolves are not transferred as it is a single channel port.

There are applications that are latency sensitive, and in some cases the application times out before the failover sequence is completed. In these cases, the application must reestablish the session.

Note: There is no comprehensive list of applications that can be dropped because of the time it takes for the standby to take over. A good rule of thumb is to expect the standby to take 10 seconds to take over using stateful failover. Without stateful failover it can take up to a minute to reestablish connections.

Note: The only caveat about the stateful failover is what causes the failover. If you have **failover hello** set to the maximum of 15 seconds and the inside interface goes bad, then the standby does not declare that the primary has failed until it misses at least two hellos, 30 seconds. Some people set the failover hellos to the minimum of 3 seconds but then the PIX can failover unnecessarily. Cisco recommends that you set the hello to the maximum of 15 seconds.

Failover Commands

- **[no] failover** – Enables or disables failover.
- **[no] failover active** – Causes a unit to become active/standby.
- **failover ip address #.#.#.#** – Sets the failover IP address.
- **failover reset** – Clears the failed state of both units and restarts failover.
- **[no] failover link interface** – Specifies which interface to be used for transmit state update from active to standby PIX in stateful failover.
- **failover poll seconds** – Specifies the failover poll interval (PIX Software version 5.2 and later).
- **failover lan unit primary | secondary** – Used in LAN-based failover to define primary/secondary (PIX version 6.2 and later).
- **failover lan enable** – Specifies LAN-based failover (PIX version 6.2 and later).

- **failover lan interface** *lan_if_name* – The name of the firewall interface dedicated to LAN-based failover (PIX version 6.2 and later).
- **failover lan key** *key_secret* – Enables encryption and authentication of LAN-based failover messages between PIX firewalls using the secret key (PIX version 6.2 and later).
- **failover mac address** *mif_name act_mac stn_mac* – Enables you to configure a virtual MAC address for a PIX firewall failover pair instead of contacting the other peer to get the MAC address (PIX version 6.2 and later).

Sample Output from the show failover Command

These examples assume that the failover cable has been installed and is operational. They also assume that the units have been configured with a System IP address of 192.168.89.1 and a Failover IP address of 192.168.89.2.

Example: Normal Failover

This example is the normal output of the **show failover** command. Note that the IP address of each unit is displayed. If no failover IP address has been entered, it displays 0.0.0.0 and monitoring of the interfaces remains in the "waiting" state. See the Example: Failover Monitoring Has Not Begun section for an explanation of the "waiting" state.

```
Failover On
  Cable status: Normal
  Reconnect timeout 0:00:00
    This host: Primary - Active
      Active time: 6885 (sec)
      Interface 0 (192.168.89.1): Normal
      Interface 1 (192.168.89.1): Normal
    Other host: Secondary - Standby
      Active time: 0 (sec)
      Interface 0 (192.168.89.2): Normal
      Interface 1 (192.168.89.2): Normal
```

Example: Failover Monitoring Has Not Begun

This examples demonstrates what happens when failover has not started to monitor the network interfaces. Failover does not start to monitor the network interfaces until it has heard the second "hello" packet from the other unit on that interface. This takes about 30 seconds. If the unit is attached to a network switch that runs Spanning Tree Protocol (STP), this takes twice the "forward delay" time configured in the switch (typically configured as 15 seconds), plus this 30 second delay. This is because at PIX bootup and immediately following a failover event, the network switch detects a temporary bridge loop. Upon detection of this loop, it stops forwarding packets on these interfaces for the "forward delay" time. It then enters the "listen" mode for an additional "forward delay" time, during which time the switch listens for bridge loops but not forwarding traffic (and thus not forwarding failover "hello" packets). After twice the forward delay time (30 seconds), traffic resumes flowing. Each PIX remains in "waiting" mode until it hears 30 seconds worth of "hello" packets from the other unit. During the time the PIX is passing traffic, it does not fail the other unit based on not hearing the "hello" packets. All other failover monitoring still occurs (that is, Power, Interface Loss of Link, and Failover Cable "hello").

```
Failover On
  Cable status: Normal
  Reconnect timeout 0:00:00
    This host: Primary - Active
      Active time: 6930 (sec)
      Interface 0 (192.168.89.1): Normal (Waiting)
      Interface 1 (192.168.89.1): Normal (Waiting)
    Other host: Secondary - Standby
```

```
Active time: 15 (sec)
Interface 0 (192.168.89.2): Normal (Waiting)
Interface 1 (192.168.89.2): Normal (Waiting)
```

Example: Unit Failure

In this example, failover detects a failure. Note that Interface 1 on the primary unit is the source of the failure. The units are back in "waiting" mode because of the failure. The failed unit has removed itself from the network (interfaces are down) and no longer sends "hello" packets on the network. The active unit remains in a "waiting" state until the failed unit is replaced and failover communications start again.

```
Failover On
Cable status: Normal
Reconnect timeout 0:00:00
  This host: Primary - Standby (Failed)
    Active time: 7140 (sec)
    Interface 0 (192.168.89.2): Normal (Waiting)
    Interface 1 (192.168.89.2): Failed (Waiting)
  Other host: Secondary - Active
    Active time: 30 (sec)
    Interface 0 (192.168.89.1): Normal (Waiting)
    Interface 1 (192.168.89.1): Normal (Waiting)
```

Example: Stateful Failover

This example shows output of the **show failover** command with stateful failover enabled. Note that the "Poll frequency 4 seconds" is displayed. Network interface 4th is administratively shutdown. The network interface FailLink is the stateful failover link.

```
Failover On
Cable status: Normal
Reconnect timeout 0:00:00
Poll frequency 4 seconds
  This host: Secondary - Active
    Active time: 167464 (sec)
    Interface gb (7.7.7.1): Normal
    Interface 4th (172.1.1.3): Link Down (Shutdown)
    Interface FailLink (8.8.8.1): Normal
    Interface pix/intf2 (100.2.1.3): Normal
    Interface outside (100.1.1.3): Normal
    Interface inside (10.1.1.3): Normal
  Other host: Primary - Standby
    Active time: 0 (sec)
    Interface gb (7.7.7.2): Normal
    Interface 4th (172.1.1.4): Link Down (Shutdown)
    Interface FailLink (8.8.8.2): Normal
    Interface pix/intf2 (100.2.1.4): Normal
    Interface outside (100.1.1.4): Normal
    Interface inside (10.1.1.4): Normal
```

Stateful Failover Logical Update Statistics

```
Link : FailLink
Stateful Obj  xmit      xerr      rcv      rerr
General      22501      0         34259      0
sys cmd      16007      0         33961      13
up time       4          0          2          0
xlate        5094      0          6          0
tcp conn     514        0          290        0
udp conn      0          0          0          0
ARP tbl      882        0          0          0
RIP Tbl      0          0          0          0
```

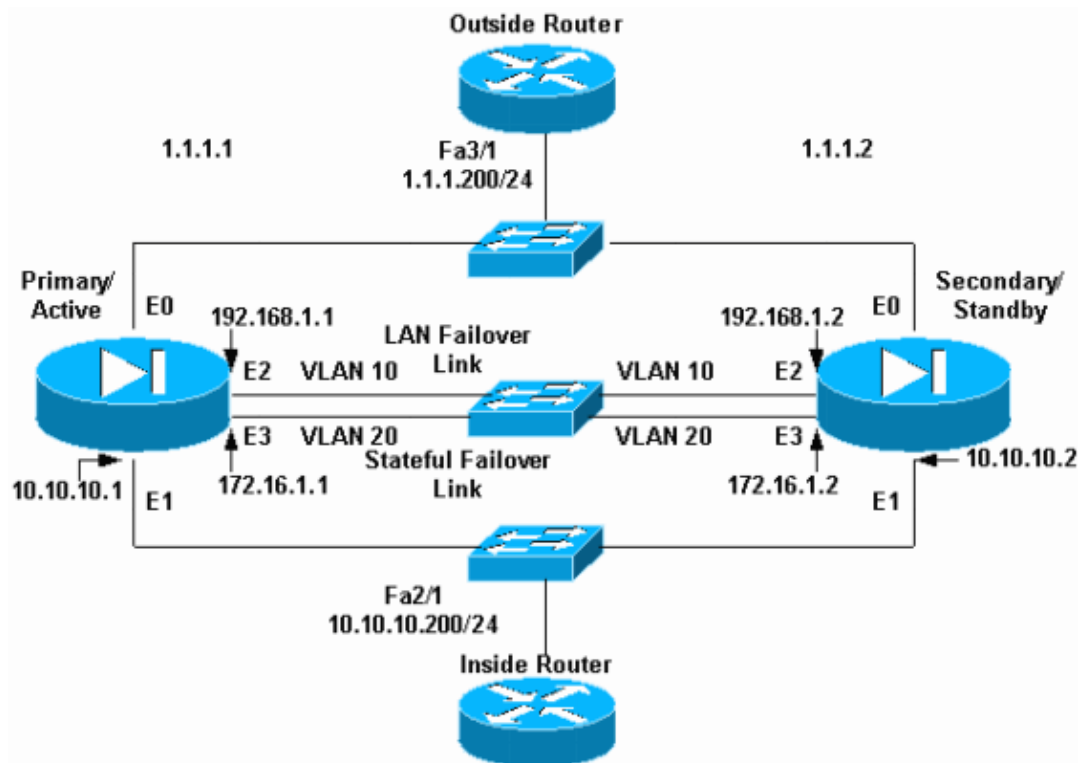
Logical Update Queue Information			
	Cur	Max	Total
Recv Q:	0	3	34259
Xmit Q:	0	7	22504

Other failover information is available in the Configuration Guide for the PIX Firewall.

If you have the output of a **show failover** command from your Cisco device, you can use Output Interpreter (registered customers only) to display potential issues and fixes.

LAN-Based Failover

LAN-Based Failover Diagram



It is recommended that you connect the Primary and Secondary PIXes with a dedicated switch. Do not use crossover cables. In this diagram, a Cisco Catalyst 3500 switch connects the Primary and Secondary PIXes. The LAN failover and stateful failover links are in different VLANs, VLAN 10 and VLAN 20, respectively. The inside-router and outside-router are used only for the sake of testing connectivity.

Minimum Initial Configuration on the Primary PIX

These are the minimum commands that need to be configured on the Primary PIX:

Basic Commands

```

pixfirewall(config)#hostname PIX

!--- Naming the PIX is optional.

PIX(config)#nameif ethernet2 fo security20

!--- Naming the interface is optional. It is recommended that you

```

```
!--- hardcode the speed/duplex.
PIX(config)#interface ethernet2 100full
!--- Bring up the interface.
PIX(config)#ip address fo 192.168.1.1 255.255.255.0
!--- Assign an IP address.
```

Failover Commands

```
PIX(config)#failover ip address fo 192.168.1.2
!--- IP address for the failover link.
PIX(config)#failover lan unit primary
!--- This unit is primary
.
PIX(config)#failover lan interface fo
!--- The 'fo' interface is used for LAN failover.
PIX(config)#failover lan key cisco
!--- The Pre-shared key.
PIX(config)#failover lan enable
!--- Enables failover.
PIX(config)#failover
!--- Start the failover process.
```

This message appears on the console:

```
LAN-based Failover: trying to contact peer
LAN-based Failover: Send hello msg and start failover monitoring
```

Minimum Initial Configuration on the Secondary PIX

These are the minimum commands that need to be configured on the Secondary PIX:

Basic Commands

```
pixfirewall(config)#hostname PIX
PIX(config)#nameif ethernet2 fo security20
!--- It is recommended that you hardcode the speed/duplex.
PIX(config)#interface ethernet2 100full
PIX(config)#ip address fo 192.168.1.1 255.255.255.0
```

Failover Commands

```
PIX(config)#failover ip address fo 192.168.1.2
PIX(config)#failover lan unit secondary
```

```
!--- This unit is secondary.
```

```
PIX(config)#failover lan interface fo  
PIX(config)#failover lan key cisco  
PIX(config)#failover lan enable  
PIX(config)#failover
```

```
!--- This unit is secondary because the "active" keyword is not used.
```

After issuing these commands on the Secondary PIX, these messages appear on the console:

```
LAN-based Failover: trying to contact peer??  
LAN-based Failover: Send hello msg and start failover monitoring
```

Then, on the Primary PIX, these messages appear on the console:

```
LAN-based Failover: Peer is UP  
Sync Started  
Sync Completed
```

Note: If you do not see these messages, then there is something wrong. In order to quickly troubleshoot the issue, turn on ICMP trace debugs on the Secondary unit using the **debug icmp trace** command and ping from the Primary unit to the failover IP address.

Note: On the Primary PIX:

```
PIX(config)#ping 192.168.1.2  
  192.168.1.2 response received -- 0ms  
  192.168.1.2 response received -- 0ms  
  192.168.1.2 response received -- 0ms  
PIX(config)#
```

```
On Secondary Unit  
PIX(config)#debug icmp trace
```

```
!--- Configure this command before you initiate ping.
```

```
ICMP trace on  
Warning: this may cause problems on busy networks  
PIX(config)# 1: ICMP echo request (len 32 id 9233 seq 0)  
  192.168.1.1 > 192.168.1.2  
2: ICMP echo reply (len 32 id 9233 seq 0) 192.168.1.2 > 192.168.1.1  
3: ICMP echo request (len 32 id 9233 seq 1) 192.168.1.1 > 192.168.1.2  
4: ICMP echo reply (len 32 id 9233 seq 1) 192.168.1.2 > 192.168.1.1  
5: ICMP echo request (len 32 id 9233 seq 2) 192.168.1.1 > 192.168.1.2  
6: ICMP echo reply (len 32 id 9233 seq 2) 192.168.1.2 > 192.168.1.1
```

Note: Once you are done, turn off these debugs with the **no debug icmp trace** command.

If you are not able to ping successfully, check the VLAN and port configurations on the intermediate switch. Also make sure that you use reliable Cat 5 cables.

Primary PIX output:

```
PIX(config)#show failover lan  
  
LAN-based Failover is Active  
  interface fo (192.168.1.1): Normal, peer (192.168.1.2): Normal  
  
PIX(config)#show failover  
Failover On
```

Cable status: My side not connected

!--- The failover serial cable is not used.

Reconnect timeout 0:00:00

Poll frequency 15 seconds

This host: Primary - Active

Active time: 4335 (sec)

Interface intf5 (0.0.0.0): Link Down (Shutdown)

Interface intf4 (0.0.0.0): Link Down (Shutdown)

Interface intf3 (0.0.0.0): Link Down (Shutdown)

Interface outside (0.0.0.0): Link Down (Shutdown)

Interface inside (0.0.0.0): Link Down (Shutdown)

Other host: Secondary - Standby

Active time: 30 (sec)

Interface intf5 (0.0.0.0): Link Down (Shutdown)

Interface intf4 (0.0.0.0): Link Down (Shutdown)

Interface intf3 (0.0.0.0): Link Down (Shutdown)

Interface outside (0.0.0.0): Link Down (Shutdown)

Interface inside (0.0.0.0): Link Down (Shutdown)

Stateful Failover Logical Update Statistics

Link : Unconfigured.

!--- Stateful failover is not configured yet.

LAN-based Failover is Active

interface fo (192.168.1.1): Normal, peer (192.168.1.2): Normal

Secondary PIX output:

PIX(config)#show failover lan

LAN-based Failover is Active

interface fo (192.168.1.2): Normal, peer (192.168.1.1): Normal

PIX(config)#show failover

Failover On

Cable status: My side not connected

!--- A failover serial cable is not used.

Reconnect timeout 0:00:00

Poll frequency 15 seconds

This host: Secondary - Standby

Active time: 30 (sec)

Interface intf5 (0.0.0.0): Link Down (Shutdown)

Interface intf4 (0.0.0.0): Link Down (Shutdown)

Interface intf3 (0.0.0.0): Link Down (Shutdown)

Interface outside (0.0.0.0): Link Down (Shutdown)

Interface inside (0.0.0.0): Link Down (Shutdown)

Other host: Primary - Active

Active time: 4485 (sec)

Interface intf5 (127.0.0.1): Link Down (Shutdown)

Interface intf4 (127.0.0.1): Link Down (Shutdown)

Interface intf3 (127.0.0.1): Link Down (Shutdown)

Interface outside (127.0.0.1): Link Down (Shutdown)

Interface inside (127.0.0.1): Link Down (Shutdown)

Stateful Failover Logical Update Statistics

Link : Unconfigured.

!--- Stateful failover is not configured yet.

```
LAN-based Failover is Active
  interface fo (192.168.1.2): Normal, peer (192.168.1.1): Normal
```

Remaining Configuration – Stateful Failover

The basic work on both the Primary and Secondary PIXes is complete.

In this example, interface E3 is used to carry state information between the two units. You can use the interface E2 (which is used for health check and configuration replication) for this purpose as well, if your PIX is not heavily loaded. It is recommended to use a separate interface for this purpose.

On the primary PIX, configure these commands:

```
ip address stateful-fo 172.16.1.1 255.255.255.0
interface ethernet3 100full
failover ip address stateful-fo 172.16.1.2
failover link stateful-fo
```

On the secondary PIX, configure this:

```
ip address stateful-fo 172.16.1.2 255.255.255.0
nameif ethernet3 stateful-fo security30
interface ethernet3 100full
```

Check the status.

```
PIX(config)#show failover
Failover On
Cable status: My side not connected
Reconnect timeout 0:00:00
Poll frequency 15 seconds
  This host: Primary - Active
    Active time: 3945 (sec)
    Interface intf5 (0.0.0.0): Link Down (Shutdown)
    Interface intf4 (0.0.0.0): Link Down (Shutdown)
    Interface stateful-fo (172.16.1.1): Normal
    Interface outside (0.0.0.0): Link Down (Shutdown)
    Interface inside (0.0.0.0): Link Down (Shutdown)
  Other host: Secondary - Standby
    Active time: 30 (sec)
    Interface intf5 (0.0.0.0): Link Down (Shutdown)
    Interface intf4 (0.0.0.0): Link Down (Shutdown)
    Interface stateful-fo (172.16.1.2): Normal
    Interface outside (0.0.0.0): Link Down (Shutdown)
    Interface inside (0.0.0.0): Link Down (Shutdown)
```

Stateful Failover Logical Update Statistics

```
Link : stateful-fo
```

```
!--- Interface stateful-fo is used for stateful failover
```

```
.
Stateful Obj   xmit      xerr      rcv       rerr
General        40         0         40        0
sys cmd        40         0         40        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        1       41
Xmit Q:         0        1       41
```

LAN-based Failover is Active

interface fo (192.168.1.1): Normal, peer (192.168.1.2): Normal

PIX(config)#show failover lan detail

LAN-based Failover is Active

This PIX is Primary

Command Interface is fo

My Command Interface IP is 192.168.1.1

Peer Command Interface IP is 192.168.1.2

My interface status is Normal

Peer interface status is Normal

Peer interface down time is 0x0

!--- This is good.

Total cmd msgs sent: 2579, rcvd: 2241, dropped: 2, retrans: 19, send_err: 0

Total secure msgs sent: 2760, rcvd: 2383

bad_signature: 0, bad_authen: 0, bad_hdr: 0, bad_osversion: 0, bad_length: 0

Total failed retx lck cnt: 0

Total/Cur/Max of 1245:0:1 msgs on retransQ, 1239 ack msgs

Cur/Max of 0:21 msgs on txq

Number of blk allocation failure: 0, cmd failure: 0, Flapping: 0

Current cmd window: 1, Slow cmd Ifc cnt: 0

Cmd Link down: 0, down and up: 0, Window Limit: 4301

Number of fmsg allocation failure: 0

Cmd Response Time History stat:

< 100ms: 1237

100 - 250ms: 0

250 - 500ms: 0

500 - 750ms: 0

750 - 1000ms: 0

1000 - 2000ms: 7

2000 - 4000ms: 5

> 4000ms: 9

Cmd Response Retry History stat:

Retry 0 = 1242, 1 = 5, 2 = 5, 3 = 3, 4 = 3

Failover enable state is 0x1

Failover state is 0x7d

Failover peer state is 0x58

Failover switching state is 0x0

Failover config syncing is not in progress

Failover poll cnt is 0

Failover Fmsg cnt is 0

Failover OS version is 6.2(0)243

failover interface 0, tst_mystat = 0x3, tst_peerstat = 0x3

zcnt = 0, hcnt = 0, my_rcnt = 0, peer_rcnt = 0

myflag = 0x0, peer_flag=0x0, dhcp = 0x807696d8

act_ip: 0.0.0.0, stn_ip:0.0.0.0

act_mac: 00d0.b71d.2b4d, stb_mac: 00d0.b780.574f

failover interface 1, tst_mystat = 0x3, tst_peerstat = 0x3

zcnt = 0, hcnt = 0, my_rcnt = 0, peer_rcnt = 0

myflag = 0x0, peer_flag=0x0, dhcp = 0x80769738

act_ip: 0.0.0.0, stn_ip:0.0.0.0

act_mac: 00d0.b71a.e6fb, stb_mac: 00e0.b600.8673

```

failover interface 2, tst_mystat = 0x0, tst_peerstat = 0x2
  zcnt = 0, hcnt = 0, my_rcnt = 2271, peer_rcnt = 0
  myflag = 0x0, peer_flag=0x0, dchp = 0x80769618
  act_ip: 192.168.1.1, stn_ip:192.168.1.2
  act_mac: 00e0.b600.a931, stb_mac: 00e0.b600.a931
LAN-based Failover command link
failover interface 3, tst_mystat = 0x0, tst_peerstat = 0x0
  zcnt = 0, hcnt = 0, my_rcnt = 88, peer_rcnt = 54
  myflag = 0x1, peer_flag=0x1, dchp = 0x80769558
  act_ip: 172.16.1.1, stn_ip:172.16.1.2
  act_mac: 00e0.b600.a930, stb_mac: 00e0.b600.8671
failover interface 4, tst_mystat = 0x3, tst_peerstat = 0x3
  zcnt = 0, hcnt = 0, my_rcnt = 0, peer_rcnt = 0
  myflag = 0x0, peer_flag=0x0, dchp = 0x80769498
act_ip: 0.0.0.0, stn_ip:0.0.0.0
  act_mac: 00e0.b600.a92f, stb_mac: 00e0.b600.8670
failover interface 5, tst_mystat = 0x3, tst_peerstat = 0x3
  zcnt = 0, hcnt = 0, my_rcnt = 0, peer_rcnt = 0
  myflag = 0x0, peer_flag=0x0, dchp = 0x807693d8
  act_ip: 0.0.0.0, stn_ip:0.0.0.0
  act_mac: 00e0.b600.a92e, stb_mac: 00d0.b780.564f

```

Other Configurations on the Primary PIX

This is another configuration for the primary PIX.

1. Hardcode the speed/duplex for other interfaces. You can use "auto," but it recommended that you hardcode the speed/duplex.

```

interface ethernet0 100full
interface ethernet1 100full

```

2. Assign IP addresses to other interfaces.

```

ip address outside 1.1.1.1 255.255.255.0
ip address inside 10.10.10.1 255.255.255.0

```

3. Add the **failover ip address** command for all interfaces excluding ones that are shut down:

```

failover ip address outside 1.1.1.2
failover ip address inside 10.10.10.2

```

Other Configurations on the Secondary PIX

This is how to configure the secondary PIX.

1. Hardcode the speed/duplex for other interfaces. You can use "auto," but it recommended that you hardcode the speed/duplex.

```

interface ethernet0 100full
interface ethernet1 100full

```

2. Assign IP addresses to other interfaces.

```

ip address outside 1.1.1.2 255.255.255.0
ip address inside 10.10.10.1 255.255.255.0

```

This is output from the Secondary PIX after failover occurs.

```

PIX(config)#show failover
Failover On
Cable status: My side not connected
Reconnect timeout 0:00:00
Poll frequency 15 seconds

```

```

This host: Secondary - Active
  Active time: 315 (sec)
  Interface intf5 (127.0.0.1): Link Down (Shutdown)
  Interface intf4 (127.0.0.1): Link Down (Shutdown)
  Interface stateful-fo (172.16.1.2): Normal (Waiting)
  Interface outside (1.1.1.2): Normal (Waiting)
  Interface inside (10.10.10.2): Normal (Waiting)
Other host: Primary - Standby
  Active time: 8025 (sec)
  Interface intf5 (0.0.0.0): Link Down (Shutdown)
  Interface intf4 (0.0.0.0): Link Down (Shutdown)
  Interface stateful-fo (172.16.1.2): Normal (Waiting)
  Interface outside (1.1.1.2): Normal (Waiting)
  Interface inside (10.1.1.2): Link Down (Waiting)

```

Stateful Failover Logical Update Statistics

```

Link : stateful-fo
Stateful Obj   xmit      xerr      rcv        rerr
General        146        0          0          0
sys cmd         146        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	1	146

LAN-based Failover is Active

```
interface fo (192.168.1.2): Normal, peer (192.168.1.1): Normal
```

Other failover commands that can be configured on the PIX:

```

failover mac address <ifc_name> <act_mac> <stn_mac>
failover poll <seconds>
failover replication http

```

```
outside-router#write terminal
```

```

interface FastEthernet3/1
  ip address 1.1.1.200 255.255.255.0
  duplex auto
  speed auto

```

```
inside-router#write terminal
```

```

interface FastEthernet2/1
  ip address 10.10.10.200 255.255.255.0
  duplex auto
  speed auto
!
ip route 0.0.0.0 0.0.0.0 10.10.10.1

```

Continue configuring the primary unit, and the configuration will be replicated to the secondary unit automatically.

Frequently Asked Questions

1. How is startup initialization accomplished between two units?

The default for failover is off (no failover). However, if the failover cable is plugged into a unit at boot time, failover automatically detects the cable and turns failover on and sets the unit's status to Primary or Secondary. This is true at boot time even if the failover IP addresses are not configured properly.

Note: If the cable is installed to a running PIX, you must issue the **failover** command start failover. For PIX software later than the 4.4.3 release, this can be changed. This is because configuration replication can accidentally disable the failover with the **clear config** command. If the failover cable is not present at boot time, the unit immediately becomes the active unit. However, the unit shows as "Secondary."

These discussions assume that failover is enabled and that both units have the failover cable plugged in. When a unit first boots, it enables failover and defaults to standby if power is detected from the other unit. The unit sends runtime status (standby) and requests a MAC address from the other unit. If no unit has taken active control within the failover polltime, the unit switches to active.

Note: For PIX software versions earlier than 5.2.1, the failover polltime was hard coded to 15 seconds.

Normally, the other unit responds to the request or sends out failover HELLO messages for each failover poll. Once the failover cable communication is started, both units check the active/standby status. The primary unit switches to active if the secondary unit is in standby state. This means that if the primary unit and the secondary both complete their bootups within the first failover poll check of each other, the primary unit becomes active. If the secondary is already active, the primary unit remains standby (assume the secondary learned the primary MAC address before. The primary unit does *not* automatically take active control. With failover enabled, do not boot the secondary unit without first booting the primary one, since the MAC address used is from the primary unit. If a unit is booted up without the failover cable, or there is no failover communication through the failover cable, both units can become active and network traffic is interrupted.

With failover enabled, full configuration replication occurs from the active unit to the standby unit when the standby unit first boots up. From that point on, commands are passed from the active to the Standby as they are entered. In order to force a complete configuration replication, use the **write standby** command. Configuration replication only happens from active to standby unit. Commands entered on the standby unit are not replicated to the active unit. A warning message is displayed when you enter commands on the standby unit, telling you that the configurations are no longer synchronized.

2. What constitutes a failure?

Fault detection is based on:

- a. *Network Interface Card (NIC) status.* If the Link Status of a NIC is down, the unit fails. "Down" means that the NIC is not plugged into an operation port. If a NIC has been configured as "down," it does not fail this test.
- b. *Failover Network communications.* The two units send "hello" packets to each other over all network interfaces. If no "hello" packets are heard in 30 seconds, the offending interface is put in testing mode in order to determine who is at fault.
- c. *Failover cable communication.* The two units send "hello" messages to each other over the failover cable. If the standby does not hear from the active within 30 seconds, and the cable status is OK, the standby takes over as active.

Also, if failover commands sent over the failover cable are not acknowledged in 15 seconds, the standby takes over as active.

- d. *Cable errors.* The failover cable is wired so that each unit can distinguish between:

- ◇ A power failure other unit.
- ◇ A cable unplugged this unit.
- ◇ A cable unplugged other unit.

If the standby detects that the active is powered off (or reload/reset), it takes active control. If the failover cable is unplugged, a syslog is generated but *no switching* will occur. An exception to this is at boot up, at which point an unplugged cable forces the unit to become active. If both units are powered up without the failover cable installed, they both become active, creating a duplicate IP address with different MAC addresses, causing conflict on your network. The failover cable must be installed for failover to work correctly.

3. How long does it take to detect a failure with default poll interval values?

- ◆ Network communications errors are detected within 30 seconds.
- ◆ Failover communications errors are detected with 30 seconds.
- ◆ Power failure (and cable failure) is detected within 15 seconds.

4. What happens when failover is triggered?

Either unit can initiate a switchover. When a switch takes place, the units each change their states as well as the IP address and MAC addresses they use. From the network's point of view, the standby transparently replaces the previously Active unit. Because configuration is already complete on the standby, no updates need to be made. Since the two units do not share dynamic connection states. Any active connections will be dropped when a failover occurs. The clients must reestablish the connections through the newly active unit (unless stateful failover is in use). For each switchover, the new active unit sends a syslog for the reason.

For example:

```
Switching to ACTIVE (cause: no power detected from other side).
```

Other reasons:

- ◆ "normal master"
- ◆ "no failover cable"
- ◆ "no power detected from other side"
- ◆ "unable to talk to the other side"
- ◆ "line interface failed at other side"
- ◆ "do not see traffic count change"
- ◆ "the other side wants me to take over"
- ◆ "fail reported by the other side"
- ◆ "state check"
- ◆ "set by the ioctl cmd"

5. What maintenance is required?

Use the **show fail** command to monitor the two units' status. Syslogs are generated when errors and switches occur.

6. How do I disable failover?

Remove the failover cable from the unit and then configure it with the **no failover** command. The failover software detects the absence of the cable and automatically turns off failover.

7. What is the failover bundle?

PIX-5XX-FO-BUN, consisting of a chassis, software, and two 10/100 ports. Customers *do not* need to purchase matching software for PIX 515. This bundle includes unrestricted PIX software. This unit is used strictly for failover. Make sure to use the **write memory** command on the failover unit to save the information. If the configuration is not saved and the standby unit is reloaded, it loses the

configuration copied from the primary unit.

Information to Collect if You Open a Technical Support Case

If you still need assistance after following the troubleshooting steps above and want to open a case with Technical Support, be sure to include the following information for troubleshooting your PIX Firewall.

- Problem description and relevant topology details
- Troubleshooting performed before opening the case
- Output from the **show tech-support** command (from both the Primary and Secondary firewalls)
- Output from the **show log** command after running with the **logging buffered debugging** command, or console captures that demonstrate the problem (if available)

Please attach the collected data to your case in non-zipped, plain text format (.txt). You can attach information to your case by uploading it using the Case Query Tool (registered customers only) . If you cannot access the Case Query Tool, you can send the information in an email attachment to attach@cisco.com with your case number in the subject line of your message.

Related Information

- [PIX Failover Documentation](#)
- [Documentation for PIX Firewall](#)
- [PIX Command Reference](#)
- [PIX Support Page](#)
- [Requests for Comments \(RFCs\)](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2010 – 2011 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Sep 30, 2008

Document ID: 5220
