CHAPTER 14

Configuring Failover

This chapter describes the security appliance failover feature, which lets you configure two security appliances so that one takes over operation if the other one fails.

Note The ASA 5505 series adaptive security appliance does not support Stateful Failover or Active/Active failover.

This chapter includes the following sections:

For failover configuration examples, see Appendix B, "Sample Configurations."

Understanding Failover

The failover configuration requires two identical security appliances connected to each other through a dedicated failover link and, optionally, a Stateful Failover link. The health of the active interfaces and units is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

The security appliance supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover.

With Active/Active failover, both units can pass network traffic. This lets you configure load balancing on your network. Active/Active failover is only available on units running in multiple context mode.

With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.

Both failover configurations support stateful or stateless (regular) failover.

Note VPN failover is not supported on units running in multiple context mode. VPN failover available for Active/Standby failover configurations only.

This section includes the following topics:

Failover System Requirements

This section describes the hardware, software, and license requirements for security appliances in a failover configuration. This section contains the following topics:

Hardware Requirements

The two units in a failover configuration must have the same hardware configuration. They must be the same model, have the same number and types of interfaces, and the same amount of RAM.

Note The two units do not have to have the same size Flash memory. If using units with different Flash memory sizes in your failover configuration, make sure the unit with the smaller Flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger Flash memory to the unit with the smaller Flash memory will fail.

Software Requirements

The two units in a failover configuration must be in the operating modes (routed or transparent, single or multiple context). They have the same major (first number) and minor (second number) software version. However, you can use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 7.0(1) to Version 7.0(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

See "Performing Zero Downtime Upgrades for Failover Pairs" section on page 41-6 for more information about upgrading the software on a failover pair.

License Requirements

On the PIX 500 series security appliance, at least one of the units must have an unrestricted (UR) license. The other unit can have a Failover Only (FO) license, a Failover Only Active-Active (FO_AA) license, or another UR license. Units with a Restricted license cannot be used for failover, and two units with FO or FO_AA licenses cannot be used together as a failover pair.

Note The FO license does not support Active/Active failover.

The FO and FO_AA licenses are intended to be used solely for units in a failover configuration and not for units in standalone mode. If a failover unit with one of these licenses is used in standalone mode, the unit reboots at least once every 24 hours until the unit is returned to failover duty. A unit with an FO or FO_AA license operates in standalone mode if it is booted without being connected to a failover peer with a UR license. If the unit with a UR license in a failover pair fails and is removed from the configuration, the unit with the FO or FO_AA license does not automatically reboot every 24 hours; it operates uninterrupted unless the it is manually rebooted.

When the unit automatically reboots, the following message displays on 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....
========================================================
```

The ASA 5500 series adaptive security appliance platform does not have this restriction.

The Failover and Stateful Failover Links

This section describes the failover and the Stateful Failover links, which are dedicated connections between the two units in a failover configuration. This section includes the following topics:

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit. The following information is communicated over the failover link:

- The unit state (active or standby).
- Power status (cable-based failover only—available only on the PIX 500 series security appliance).
- Hello messages (keep-alives).
- Network link status.
- MAC address exchange.
- Configuration replication and synchronization.

Caution All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

On the PIX 500 series security appliance, the failover link can be either a LAN-based connection or a dedicated serial Failover cable. On the ASA 5500 series adaptive security appliance, the failover link can only be a LAN-based connection.

This section includes the following topics:

LAN-Based Failover Link

You can use any unused Ethernet interface on the device as the failover link; however, you cannot specify an interface that is currently configured with a name. The LAN failover link interface is not configured as a normal networking interface. It exists for failover communication only. This interface should only be used for the LAN failover link (and optionally for the stateful failover link).

Connect the LAN failover link in one of the following two ways:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the LAN failover interfaces of the ASA.
- Using a crossover Ethernet cable to connect the appliances directly, without the need for an external switch.

Note When you use a crossover cable for the LAN failover link, if the LAN interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which interface failed and caused the link to come down.

Note The ASA supports Auto-MDI/MDIX on its copper Ethernet ports, so you can either use a crossover cable or a straight-through cable. If you use a straight-through cable, the interface automatically detects the cable and swaps one of the transmit/receive pairs to MDIX.

Serial Cable Failover Link (PIX Security Appliance Only)

The serial Failover cable, or “cable-based failover,” is only available on the PIX 500 series security appliance. If the two units are within six feet of each other, then we recommend that you use the serial Failover cable.

The cable that connects the two units is a modified RS-232 serial link cable that transfers data at 117,760 bps (115 Kbps). One end of the cable is labeled “Primary”. The unit attached to this end of the cable automatically becomes the primary unit. The other end of the cable is labeled “Secondary”. The unit attached to this end of the cable automatically becomes the secondary unit. You cannot override these designations in the PIX 500 series security appliance software. If you purchased a PIX 500 series security appliance failover bundle, this cable is included. To order a spare, use part number PIX-FO=.

The benefits of using cable-based failover include:

- The PIX 500 series security appliance can immediately detect a power loss on the peer unit and differentiate between a power loss from an unplugged cable.
- The standby unit can communicate with the active unit and can receive the entire configuration without having to be bootstrapped for failover. In LAN-based failover you need to configure the failover link on the standby unit before it can communicate with the active unit.
- The switch between the two units in LAN-based failover can be another point of hardware failure; cable-based failover eliminates this potential point of failure.
- You do not have to dedicate an Ethernet interface (and switch) to the failover link.
- The cable determines which unit is primary and which is secondary, eliminating the need to manually enter that information in the unit configurations.

The disadvantages include:

- Distance limitation—the units cannot be separated by more than 6 feet.
- Slower configuration replication.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link to pass all state information. You have three options for configuring a Stateful Failover link:

- You can use a dedicated Ethernet interface for the Stateful Failover link.
- If you are using LAN-based failover, you can share the failover link.
- You can share a regular data interface, such as the inside interface. However, this option is not recommended.

If you are using a dedicated Ethernet interface for the Stateful Failover link, you can use either a switch or a crossover cable to directly connect the units. If you use a switch, no other hosts or routers should be on this link.

Note Enable the PortFast option on Cisco switch ports that connect directly to the security appliance.

If you use a data interface as the Stateful Failover link, you receive the following warning when you specify that interface as the Stateful Failover link:

```
******* WARNING ***** WARNING ******* WARNING ****** WARNING  *********
  Sharing Stateful failover interface with regular data interface is not
  a recommended configuration due to performance and security concerns.
******* WARNING ***** WARNING ******* WARNING ****** WARNING  *********
```

Sharing a data interface with the Stateful Failover interface can leave you vulnerable to replay attacks. Additionally, large amounts of Stateful Failover traffic may be sent on the interface, causing performance problems on that network segment.

Note Using a data interface as the Stateful Failover interface is only supported in single context, routed mode.

In multiple context mode, the Stateful Failover link resides in the system context. This interface and the failover interface are the only interfaces in the system context. All other interfaces are allocated to and configured from within security contexts.

Note The IP address and MAC address for the Stateful Failover link does not change at failover unless the Stateful Failover link is configured on a regular data interface.

Caution All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

Failover Interface Speed for Stateful Links

If you use the failover link as the Stateful Failover link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the Stateful Failover interface.

Use the following failover interface speed guidelines for Cisco PIX security appliances and Cisco ASA adaptive security appliances:

- Cisco ASA 5520/5540/5550 and PIX 515E/535
  - The stateful link speed should match the fastest data link
- Cisco ASA 5510 and PIX 525
  - Stateful link speed can be 100 Mbps, even though the data interface can operate at 1 Gigabit due to the CPU speed limitation.

For optimum performance when using long distance LAN failover, the latency for the failover link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is less than 10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

All platforms support sharing of failover heartbeat and stateful link, but we recommend using a separate heartbeat link on systems with high Stateful Failover traffic.

Active/Active and Active/Standby Failover

This section describes each failover configuration in detail. This section includes the following topics:

Active/Standby Failover

This section describes Active/Standby failover and includes the following topics:

Active/Standby Failover Overview

Active/Standby failover lets you use a standby security appliance to take over the functionality of a failed unit. When the active unit fails, it changes to the standby state while the standby unit changes to the active state. The unit that becomes active assumes the IP addresses (or, for transparent firewall, the management IP address) and MAC addresses of the failed unit and begins passing traffic. The unit that is now in standby state takes over the standby IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note For multiple context mode, the security appliance can fail over the entire unit (including all contexts) but cannot fail over individual contexts separately.

Primary/Secondary Status and Active/Standby Status

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit is active, and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both devices in the failover pair boot. Configurations are always synchronized from the active unit to the standby unit. When the standby unit completes its initial startup, it clears its running configuration (except for the failover commands needed to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit and the secondary unit becomes the standby unit.

Note If the secondary unit boots without detecting the primary unit, it becomes the active unit. It uses its own MAC addresses for the active IP addresses. However, when the primary unit becomes available, the secondary unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. To avoid this, configure the failover pair with virtual MAC addresses. See the "Configuring Virtual MAC Addresses" section on page 14-27 for more information.

When the replication starts, the security appliance console on the active unit displays the message "Beginning configuration replication: Sending to mate," and when it is complete, the security appliance displays the message "End Configuration Replication to mate." During replication, commands entered on the active unit may not replicate properly to the standby unit, and commands entered on the standby unit may be overwritten by the configuration being replicated from the active unit. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

On the standby unit, the configuration exists only in running memory. To save the configuration to Flash memory after synchronization:

- For single context mode, enter the **write memory** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **write memory all** command on the active unit from the system execution space. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Using the **all** keyword with this command causes the system and all context configurations to be saved.

Note Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

Command replication always flows from the active unit to the standby unit. As commands are entered on the active unit, they are sent across the failover link to the standby unit. You do not have to save the active configuration to Flash memory to replicate the commands.

The following commands are replicated to the standby unit:

- all configuration commands except for the **mode**, **firewall**, and **failover lan unit** commands
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands are not replicated to the standby unit:

- all forms of the **copy** command except for **copy running-config startup-config**
- all forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **mode**
- **show**

Note Changes made on the standby unit are not replicated to the active unit. If you enter a command on the standby unit, the security appliance displays the message `**** WARNING **** Configuration Replication is NOT performed from Standby unit to Active unit. Configurations are no longer synchronized.` This message displays even when you enter many commands that do not affect the configuration.

If you enter the **write standby** command on the active unit, the standby unit clears its running configuration (except for the failover commands used to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

For multiple context mode, when you enter the **write standby** command in the system execution space, all contexts are replicated. If you enter the **write standby** command within a context, the command replicates only the context configuration.

Replicated commands are stored in the running configuration. To save the replicated commands to the Flash memory on the standby unit:

- For single context mode, enter the **copy running-config startup-config** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **copy running-config startup-config** command on the active unit from the system execution space and within each context on disk. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory. Contexts with startup configurations on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit.

Failover Triggers

The unit can fail if one of the following events occurs:

- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- The **no failover active** command is entered on the active unit or the **failover active** command is entered on the standby unit.

Failover Actions

In Active/Standby failover, failover occurs on a unit basis. Even on systems running in multiple context mode, you cannot fail over individual or groups of contexts.

Table 14-1 shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 14-1 Failover Behavior

Failure Event	Policy	Active Action	Standby Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover interface as failed	Mark failover interface as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Mark failover interface as failed	Become active	If the failover link is down at startup, both units become active.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None.
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, then the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Active/Active Failover

This section describes Active/Active failover. This section includes the following topics:

Active/Active Failover Overview

Active/Active failover is only available to security appliances in multiple context mode. In an Active/Active failover configuration, both security appliances can pass network traffic.

In Active/Active failover, you divide the security contexts on the security appliance into *failover groups*. A failover group is simply a logical group of one or more security contexts. You can create a maximum of two failover groups on the security appliance. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

The failover group forms the base unit for failover in Active/Active failover. Interface failure monitoring, failover, and active/standby status are all attributes of a failover group rather than the unit. When an active failover group fails, it changes to the standby state while the standby failover group becomes active. The interfaces in the failover group that becomes active assume the MAC and IP addresses of the interfaces in the failover group that failed. The interfaces in the failover group that is now in the standby state take over the standby MAC and IP addresses.

Note A failover group failing on a unit does not mean that the unit has failed. The unit may still have another failover group passing traffic on it.

When creating the failover groups, you should create them on the unit that will have failover group 1 in the active state.

Note Active/Active failover generates virtual MAC addresses for the interfaces in each failover group. If you have more than one Active/Active failover pair on the same network, it is possible to have the same default virtual MAC addresses assigned to the interfaces on one pair as are assigned to the interfaces of the other pairs because of the way the default virtual MAC addresses are determined. To avoid having duplicate MAC addresses on your network, make sure you assign each physical interface a virtual active and standby MAC address.

Primary/Secondary Status and Active/Standby Status

As in Active/Standby failover, one unit in an Active/Active failover pair is designated the primary unit, and the other unit the secondary unit. Unlike Active/Standby failover, this designation does not indicate which unit becomes active when both units start simultaneously. Instead, the primary/secondary designation does two things:

- Determines which unit provides the running configuration to the pair when they boot simultaneously.
- Determines on which unit each failover group appears in the active state when the units boot simultaneously. Each failover group in the configuration is configured with a primary or secondary unit preference. You can configure both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, distributing the traffic across the devices.

  Note The security appliance does not provide load balancing services. Load balancing must be handled by a router passing traffic to the security appliance.

Which unit each failover group becomes active on is determined as follows:

- When a unit boots while the peer unit is not available, both failover groups become active on the unit.
- When a unit boots while the peer unit is active (with both failover groups in the active state), the failover groups remain in the active state on the active unit regardless of the primary or secondary preference of the failover group until one of the following:

  - A failover occurs.
  - You manually force the failover group to the other unit with the **no failover active** command.
  - You configured the failover group with the **preempt** command, which causes the failover group to automatically become active on the preferred unit when the unit becomes available.
- When both units boot at the same time, each failover group becomes active on its preferred unit after the configurations have been synchronized.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both units in a failover pair boot. The configurations are synchronized as follows:

- When a unit boots while the peer unit is active (with both failover groups active on it), the booting unit contacts the active unit to obtain the running configuration regardless of the primary or secondary designation of the booting unit.
- When both units boot simultaneously, the secondary unit obtains the running configuration from the primary unit.

When the replication starts, the security appliance console on the unit sending the configuration displays the message "Beginning configuration replication: Sending to mate," and when it is complete, the security appliance displays the message "End Configuration Replication to mate." During replication, commands entered on the unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

On the unit receiving the configuration, the configuration exists only in running memory. To save the configuration to Flash memory after synchronization enter the **write memory all** command in the system execution space on the unit that has failover group 1 in the active state. The command is replicated to the peer unit, which proceeds to write its configuration to Flash memory. Using the **all** keyword with this command causes the system and all context configurations to be saved.

Note Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts configuration files from the disk on the primary unit to an external server, and then copy them to disk on the secondary unit, where they become available when the unit reloads.

Command Replication

After both units are running, commands are replicated from one unit to the other as follows:

- Commands entered within a security context are replicated from the unit on which the security context appears in the active state to the peer unit.

Note A context is considered in the active state on a unit if the failover group to which it belongs is in the active state on that unit.

- Commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.

- Commands entered in the admin context are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.

All configuration and file commands (**copy**, **rename**, **delete**, **mkdir**, **rmdir**, and so on) are replicated, with the following exceptions. The **show**, **debug**, **mode**, **firewall**, and **failover lan unit** commands are not replicated.

Failure to enter the commands on the appropriate unit for command replication to occur causes the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

The following commands are replicated to the standby unit:

- all configuration commands except for the **mode**, **firewall**, and **failover lan unit** commands
- **copy running-config startup-config**
- **delete**
- **mkdir**
- **rename**
- **rmdir**
- **write memory**

The following commands are not replicated to the standby unit:

- all forms of the **copy** command except for **copy running-config startup-config**
- all forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **mode**
- **show**

You can use the **write standby** command to resynchronize configurations that have become out of sync. For Active/Active failover, the **write standby** command behaves as follows:

- If you enter the **write standby** command in the system execution space, the system configuration and the configurations for all of the security contexts on the security appliance is written to the peer unit. This includes configuration information for security contexts that are in the standby state. You must enter the command in the system execution space on the unit that has failover group 1 in the active state.

  Note If there are security contexts in the active state on the peer unit, the **write standby** command causes active connections through those contexts to be terminated. Use the **failover active** command on the unit providing the configuration to make sure all contexts are active on that unit before entering the **write standby** command.

- If you enter the **write standby** command in a security context, only the configuration for the security context is written to the peer unit. You must enter the command in the security context on the unit where the security context appears in the active state.

Replicated commands are not saved to the Flash memory when replicated to the peer unit. They are added to the running configuration. To save replicated commands to Flash memory on both units, use the **write memory** or **copy running-config startup-config** command on the unit that you made the changes on. The command is replicated to the peer unit and cause the configuration to be saved to Flash memory on the peer unit.

Failover Triggers

In Active/Active failover, failover can be triggered at the unit level if one of the following events occurs:

- The unit has a hardware failure.
- The unit has a power failure.
- The unit has a software failure.
- The **no failover active** or the **failover active** command is entered in the system execution space.

Failover is triggered at the failover group level when one of the following events occurs:

- Too many monitored interfaces in the group fail.
- The **no failover active group** *group_id* or **failover active group** *group_id* command is entered.

You configure the failover threshold for each failover group by specifying the number or percentage of interfaces within the failover group that must fail before the group fails. Because a failover group can contain multiple contexts, and each context can contain multiple interfaces, it is possible for all interfaces in a single context to fail without causing the associated failover group to fail.

See the "Failover Health Monitoring" section on page 14-17 for more information about interface and unit monitoring.

Failover Actions

In an Active/Active failover configuration, failover occurs on a failover group basis, not a system basis. For example, if you designate both failover groups as active on the primary unit, and failover group 1 fails, then failover group 2 remains active on the primary unit while failover group 1 becomes active on the secondary unit.

Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.

Table 14-2 shows the failover action for each failure event. For each failure event, the policy (whether or not failover occurs), actions for the active failover group, and actions for the standby failover group are given.

Table 14-2 Failover Behavior for Active/Active Failover

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.

Table 14-2 Failover Behavior for Active/Active Failover (continued)

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, the active failover group does not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless configured with the **preempt** command, the failover groups remain active on their current unit.
Failover link failed at startup	No failover	Become active	Become active	If the failover link is down at startup, both failover groups on both units become active.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the failover interface as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Determining Which Type of Failover to Use

The type of failover you choose depends upon your security appliance configuration and how you plan to use the security appliances.

If you are running the security appliance in single mode, then you can only use Active/Standby failover. Active/Active failover is only available to security appliances running in multiple context mode.

If you are running the security appliance in multiple context mode, then you can configure either Active/Active failover or Active/Standby failover.

- To provide load balancing, use Active/Active failover.
- If you do not want to provide load balancing, use Active/Standby or Active/Active failover.

Table 14-3 provides a comparison of some of the features supported by each type of failover configuration:

Table 14-3 Failover Configuration Feature Support

Feature	Active/Active	Active/Standby
Single Context Mode	No	Yes
Multiple Context Mode	Yes	Yes
Load Balancing Network Configurations	Yes	No
Unit Failover	Yes	Yes

Table 14-3 Failover Configuration Feature Support

Feature	Active/Active	Active/Standby
Failover of Groups of Contexts	Yes	No
Failover of Individual Contexts	No	No

Regular and Stateful Failover

The security appliance supports two types of failover, regular and stateful. This section includes the following topics:

- Regular Failover, page 14-16
- Stateful Failover, page 14-16

Regular Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Stateful Failover

When Stateful Failover is enabled, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

The state information passed to the standby unit includes the following:

- NAT translation table.
- TCP connection states.
- UDP connection states.
- The ARP table.
- The Layer 2 bridge table (when running in transparent firewall mode).
- The HTTP connection states (if HTTP replication is enabled).
- The ISAKMP and IPSec SA table.
- GTP PDP connection database.

The information that is not passed to the standby unit when Stateful Failover is enabled includes the following:

- The HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- The routing tables. After a failover occurs, some packets may be lost our routed out of the wrong interface (the default route) while the dynamic routing protocols rediscover routes.
- State information for Security Service Modules.
- DHCP server address leases.
- L2TP over IPSec sessions.

Note If failover occurs during an active Cisco IP SoftPhone session, the call remains active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client loses connection with the Call Manager. This occurs because there is no session information for the CTIQBE hangup message on the standby unit. When the IP SoftPhone client does not receive a response back from the Call Manager within a certain time period, it considers the Call Manager unreachable and unregisters itself.

Transparent Firewall Mode Requirements

When the active unit fails over to the standby unit, the connected switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss while the port is in a blocking state, you can configure one of the following workarounds depending on the switch port mode:

- Access mode—Enable the STP PortFast feature on the switch:

  ```
  interface interface_id
  spanning-tree portfast
  ```

  The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.

- Trunk mode—Block BPDUs on the security appliance on both the inside and outside interfaces:

  ```
  access-list id ethertype deny bpdu
  access-group id in interface inside_name
  access-group id in interface outside_name
  ```

  Blocking BPDUs disables STP on the switch. Be sure not to have any loops involving the security appliance in your network layout.

If neither of the above options are possible, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:

- Disable failover interface monitoring.
- Increase failover interface holdtime to a high value that will allow STP to converge before the security appliances fail over.
- Decrease STP timers to allow STP to converge faster than the failover interface holdtime.

Failover Health Monitoring

The security appliance monitors each unit for overall health and for interface health. See the following sections for more information about how the security appliance performs tests to determine the state of each unit:

- Unit Health Monitoring, page 14-18
- Interface Monitoring, page 14-18

Unit Health Monitoring

The security appliance determines the health of the other unit by monitoring the failover link. When a unit does not receive three consecutive hello messages on the failover link, the unit sends interface hello messages on each interface, including the failover interface, to validate whether or not the peer interface is responsive. The action that the security appliance takes depends upon the response from the other unit. See the following possible actions:

- If the security appliance receives a response on the failover interface, then it does not fail over.
- If the security appliance does not receive a response on the failover link, but receives a response on another interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the security appliance does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

Note If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

You can configure the frequency of the hello messages and the hold time before failover occurs. A faster poll time and shorter hold time speed the detection of unit failures and make failover occur more quickly, but it can also cause "false" failures due to network congestion delaying the keepalive packets. See Configuring Unit Health Monitoring, page 14-40 for more information about configuring unit health monitoring.

Interface Monitoring

You can monitor up to 250 interfaces divided between all contexts. You should monitor important interfaces, for example, you might configure one context to monitor a shared interface (because the interface is shared, all contexts benefit from the monitoring).

When a unit does not receive hello messages on a monitored interface for half of the configured hold time, it runs the following tests:

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is operational, then the security appliance performs network tests. The purpose of these tests is to generate network traffic 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 the next test is used.
2. Network Activity test—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 ARP test begins.
3. ARP test—A reading of the unit ARP cache for the 2 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.

4. Broadcast Ping test—A ping test that 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 all network tests fail for an interface, but this interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed. If the threshold for failed interfaces is met, then a failover occurs. If the other unit interface also fails all the network tests, then both interfaces go into the "Unknown" state and do not count towards the failover limit.

An interface becomes operational again if it receives any traffic. A failed security appliance returns to standby mode if the interface failure threshold is no longer met.

Note If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Failover Feature/Platform Matrix

Table 14-4 shows the failover features supported by each hardware platform.

Table 14-4 Failover Feature Support by Platform

Platform	Cable-Base Failover	LAN-Based Failover	Stateful Failover
ASA 5505 series adaptive security appliance	No	Yes	No
ASA 5500 series adaptive security appliance (other than the ASA 5505)	No	Yes	Yes
PIX 500 series security appliance	Yes	Yes	Yes

Failover Times by Platform

Table 14-5 shows the minimum, default, and maximum failover times for the PIX 500 series security appliance.

Table 14-5 PIX 500 series security appliance failover times.

Failover Condition	Minimum	Default	Maximum
Active unit loses power or stops normal operation.	800 milliseconds	45 seconds	45 seconds
Active unit interface link down.	500 milliseconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Table 14-6 shows the minimum, default, and maximum failover times for the ASA 5500 series adaptive security appliance.

Table 14-6 ASA 5500 series adaptive security appliance failover times.

Failover Condition	Minimum	Default	Maximum
Active unit loses power or stops normal operation.	800 milliseconds	15 seconds	45 seconds
Active unit main board interface link down.	500 milliseconds	5 seconds	15 seconds
Active unit 4GE card interface link down.	2 seconds	5 seconds	15 seconds
Active unit IPS or CSC card fails.	2 seconds	2 seconds	2 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Configuring Failover

This section describes how to configure failover and includes the following topics:

Failover Configuration Limitations

You cannot configure failover with the following type of IP addresses:

- IP addresses obtained through DHCP
- IP addresses obtained through PPPoE
- IPv6 addresses

Additionally, the following restrictions apply:

- Stateful Failover is not supported on the ASA 5505 adaptive security appliance.
- Active/Active failover is not supported on the ASA 5505 adaptive security appliance.
- You cannot configure failover when Easy VPN Remote is enabled on the ASA 5505 adaptive security appliance.
- VPN failover is not supported in multiple context mode.

Configuring Active/Standby Failover

This section provides step-by-step procedures for configuring Active/Standby failover. This section includes the following topics:

- Configuring LAN-Based Active/Standby Failover, page 14-22
- Configuring Optional Active/Standby Failover Settings, page 14-26

Prerequisites

Before you begin, verify the following:

- Both units have the same hardware, software configuration, and proper license.
- Both units are in the same mode (single or multiple, transparent or routed).

Configuring Cable-Based Active/Standby Failover (PIX Security Appliance Only)

Follow these steps to configure Active/Standby failover using a serial cable as the failover link. The commands in this task are entered on the *primary* unit in the failover pair. The primary unit is the unit that has the end of the cable labeled "Primary" plugged into it. For devices in multiple context mode, the commands are entered in the system execution space unless otherwise noted.

You do not need to bootstrap the secondary unit in the failover pair when you use cable-based failover. Leave the secondary unit powered off until instructed to power it on.

Cable-based failover is only available on the PIX 500 series security appliance.

To configure cable-based Active/Standby failover, perform the following steps:

Step 1 Connect the Failover cable to the PIX 500 series security appliances. Make sure that you attach the end of the cable marked "Primary" to the unit you use as the primary unit, and that you attach the end of the cable marked "Secondary" to the other unit.

Step 2 Power on the primary unit.

Step 3 If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.

Step 4 (Optional) To enable Stateful Failover, configure the Stateful Failover link.

Note Stateful Failover is not available on the ASA 5505 series adaptive security appliance.

a. Specify the interface to be used as the Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose.

b. Assign an active and standby IP address to the Stateful Failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

Note If the Stateful Failover link uses a data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover unless it uses a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 5 Enable failover:

```
hostname(config)# failover
```

Step 6 Power on the secondary unit and enable failover on the unit if it is not already enabled:

```
hostname(config)# failover
```

The active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: sending to mate.” and “End Configuration Replication to mate” appear on the primary console.

Step 7 Save the configuration to Flash memory on the primary unit. Because the commands entered on the primary unit are replicated to the secondary unit, the secondary unit also saves its configuration to Flash memory.

```
hostname(config)# copy running-config startup-config
```

Configuring LAN-Based Active/Standby Failover

This section describes how to configure Active/Standby failover using an Ethernet failover link. When configuring LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

Note If you are changing from cable-based failover to LAN-based failover, you can skip any steps, such as assigning the active and standby IP addresses for each interface, that you completed for the cable-based failover configuration.

This section includes the following topics:

Configuring the Primary Unit

Follow these steps to configure the primary unit in a LAN-based, Active/Standby failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit. For multiple context mode, all steps are performed in the system execution space unless otherwise noted.

To configure the primary unit in an Active/Standby failover pair, perform the following steps:

Step 1 If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.

Step 2 (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

Step 3 Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

Step 4 Define the failover interface:

a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN.

b. Assign the active and standby IP address to the failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 5 (Optional) To enable Stateful Failover, configure the Stateful Failover link.

Note Stateful Failover is not available on the ASA 5505 series adaptive security appliance.

a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

Note If the Stateful Failover link uses the failover link or a data interface, then you only need to supply the *if_name* argument.

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

b. Assign an active and standby IP address to the Stateful Failover link.

Note If the Stateful Failover link uses the failover link or data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover unless it uses a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

c. Enable the interface.

Note If the Stateful Failover link uses the failover link or data interface, skip this step. You have already enabled the interface.

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 6 Enable failover:

```
hostname(config)# failover
```

Step 7 Save the system configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Configuring the Secondary Unit

The only configuration required on the secondary unit is for the failover interface. The secondary unit requires these commands to initially communicate with the primary unit. After the primary unit sends its configuration to the secondary unit, the only permanent difference between the two configurations is the **failover lan unit** command, which identifies each unit as primary or secondary.

For multiple context mode, all steps are performed in the system execution space unless noted otherwise.

To configure the secondary unit, perform the following steps:

Step 1 (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

Step 2 Define the failover interface. Use the same settings as you used for the primary unit.

a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a name to the interface specified by the *phy_if* argument.

b. Assign the active and standby IP address to the failover link. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface on the primary unit.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 3 (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```

Note This step is optional because by default units are designated as secondary unless previously configured.

Step 4 Enable failover:

```
hostname(config)# failover
```

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the active unit console.

Step 5 After the running configuration has completed replication, save the configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Configuring Optional Active/Standby Failover Settings

You can configure the following optional Active/Standby failover setting when you are initially configuring failover or after failover has already been configured. Unless otherwise noted, the commands should be entered on the active unit.

This section includes the following topics:

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information replication, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information.

Enter the following command in global configuration mode to enable HTTP state replication when Stateful Failover is enabled:

```
hostname(config)# failover replication http
```

Disabling and Enabling Interface Monitoring

By default, monitoring physical interfaces is enabled and monitoring subinterfaces is disabled. You can monitor up to 250 interfaces on a unit. You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This lets you exclude interfaces attached to less critical networks from affecting your failover policy.

For units in multiple configuration mode, use the following commands to enable or disable health monitoring for specific interfaces:

- To disable health monitoring for an interface, enter the following command within a context:

  ```
  hostname/context(config)# no monitor-interface if_name
  ```

- To enable health monitoring for an interface, enter the following command within a context:

  ```
  hostname/context(config)# monitor-interface if_name
  ```

For units in single configuration mode, use the following commands to enable or disable health monitoring for specific interfaces:

- To disable health monitoring for an interface, enter the following command in global configuration mode:

  ```
  hostname(config)# no monitor-interface if_name
  ```

- To enable health monitoring for an interface, enter the following command in global configuration mode:

```
hostname(config)# monitor-interface if_name
```

Configuring Interface Health Monitoring

The security appliance sends hello packets out of each data interface to monitor interface health. If the security appliance does not receive a hello packet from the corresponding interface on the peer unit for over half of the hold time, then the additional interface testing begins. If a hello packet or a successful test result is not received within the specified hold time, the interface is marked as failed. Failover occurs if the number of failed interfaces meets the failover criteria.

Decreasing the poll and hold times enables the security appliance to detect and respond to interface failures more quickly, but may consume more system resources.

To change the interface poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime interface [msec] time [holdtime time]
```

Valid values for the poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.

Note If the interface link is down, interface testing is not conducted and the standby unit could become active in just one interface polling period if the number of failed interface meets or exceeds the configured failover criteria.

Configuring Failover Criteria

By default, a single interface failure causes failover. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs.

To change the default failover criteria, enter the following command in global configuration mode:

```
hostname(config)# failover interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Configuring Virtual MAC Addresses

In Active/Standby failover, the MAC addresses for the primary unit are always associated with the active IP addresses. If the secondary unit boots first and becomes active, it uses the burned-in MAC address for its interfaces. When the primary unit comes online, the secondary unit obtains the MAC addresses from the primary unit. The change can disrupt network traffic.

You can configure virtual MAC addresses for each interface to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not specify virtual MAC addresses the failover pair uses the burned-in NIC addresses as the MAC addresses.

Note You cannot configure a virtual MAC address for the failover or Stateful Failover links. The MAC and IP addresses for those links do not change during failover.

Enter the following command on the active unit to configure the virtual MAC addresses for an interface:

```
hostname(config)# failover mac address phy_if active_mac standby_mac
```

The *phy_if* argument is the physical name of the interface, such as Ethernet1. The *active_mac* and *standby_mac* arguments are MAC addresses in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.

The *active_mac* address is associated with the active IP address for the interface, and the *standby_mac* is associated with the standby IP address for the interface.

There are multiple ways to configure virtual MAC addresses on the security appliance. When more than one method has been used to configure virtual MAC addresses, the security appliance uses the following order of preference to determine which virtual MAC address is assigned to an interface:

1. The **mac-address** command (in interface configuration mode) address.
2. The **failover mac address** command address.
3. The **mac-address auto** command generated address.
4. The burned-in MAC address.

Use the **show interface** command to display the MAC address used by an interface.

Configuring Active/Active Failover

This section describes how to configure Active/Active failover.

Note Active/Active failover is not available on the ASA 5505 series adaptive security appliance.

This section includes the following topics:

Prerequisites

Before you begin, verify the following:

- Both units have the same hardware, software configuration, and proper license.
- Both units are in multiple context mode.

Configuring Cable-Based Active/Active Failover (PIX security appliance)

Follow these steps to configure Active/Active failover using a serial cable as the failover link. The commands in this task are entered on the *primary* unit in the failover pair. The primary unit is the unit that has the end of the cable labeled “Primary” plugged into it. For devices in multiple context mode, the commands are entered in the system execution space unless otherwise noted.

You do not need to bootstrap the secondary unit in the failover pair when you use cable-based failover. Leave the secondary unit powered off until instructed to power it on.

Cable-based failover is only available on the PIX 500 series security appliance.

To configure cable-based, Active/Active failover, perform the following steps:

Step 1 Connect the failover cable to the PIX 500 series security appliances. Make sure that you attach the end of the cable marked "Primary" to the unit you use as the primary unit, and that you attach the end of the cable marked "Secondary" to the unit you use as the secondary unit.

Step 2 Power on the primary unit.

Step 3 If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

You must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. You must enter a management IP address for each context in transparent firewall multiple context mode.

Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname/context(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

Step 4 (Optional) To enable Stateful Failover, configure the Stateful Failover link.

a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

b. Assign an active and standby IP address to the Stateful Failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask.

The Stateful Failover link IP address and MAC address do not change at failover except for when Stateful Failover uses a regular data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 5 Configure the failover groups. You can have at most two failover groups. The **failover group** command creates the specified failover group if it does not exist and enters the failover group configuration mode.

For each failover group, you need to specify whether the failover group has primary or secondary preference using the **primary** or **secondary** command. You can assign the same preference to both failover groups. For load balancing configurations, you should assign each failover group a different unit preference.

The following example assigns failover group 1 a primary preference and failover group 2 a secondary preference:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# exit
```

Step 6 Assign each user context to a failover group using the **join-failover-group** command in context configuration mode.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.

Enter the following commands to assign each context to a failover group:

```
hostname(config)# context context_name
hostname(config-context)# join-failover-group {1 | 2}
hostname(config-context)# exit
```

Step 7 Enable failover:

```
hostname(config)# failover
```

Step 8 Power on the secondary unit and enable failover on the unit if it is not already enabled:

```
hostname(config)# failover
```

The active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the primary console.

Step 9 Save the configuration to Flash memory on the Primary unit. Because the commands entered on the primary unit are replicated to the secondary unit, the secondary unit also saves its configuration to Flash memory.

```
hostname(config)# copy running-config startup-config
```

Step 10 If necessary, force any failover group that is active on the primary to the active state on the secondary. To force a failover group to become active on the secondary unit, issue the following command in the system execution space on the primary unit:

```
hostname# no failover active group group_id
```

The *group_id* argument specifies the group you want to become active on the secondary unit.

Configuring LAN-Based Active/Active Failover

This section describes how to configure Active/Active failover using an Ethernet failover link. When configuring LAN-based failover, you must bootstrap the secondary device to recognize the failover link before the secondary device can obtain the running configuration from the primary device.

This section includes the following topics:

Configure the Primary Unit

To configure the primary unit in an Active/Active failover configuration, perform the following steps:

Step 1 If you have not done so already, configure the active and standby IP addresses for each data interface (routed mode), for the management IP address (transparent mode), or for the management-only interface.To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.

You must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context. In transparent firewall mode, you must enter a management IP address for each context.

Note Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

```
hostname/context(config-if)# ip address active_addr netmask standby standby_addr
```

In routed firewall mode and for the management-only interface, this command is entered in interface configuration mode for each interface. In transparent firewall mode, the command is entered in global configuration mode.

Step 2 Configure the basic failover parameters in the system execution space.

a. (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# hostname(config)# failover lan enable
```

b. Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

c. Specify the failover link:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN. This interface should not be used for any other purpose (except, optionally, the Stateful Failover link).

d. Specify the failover link active and standby IP addresses:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask. The failover link IP address and MAC address do not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

Step 3 (Optional) To enable Stateful Failover, configure the Stateful Failover link:

a. Specify the interface to be used as Stateful Failover link:

```
hostname(config)# failover link if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

Note If the Stateful Failover link uses the failover link or a regular data interface, then you only need to supply the *if_name* argument.

b. Assign an active and standby IP address to the Stateful Failover link.

Note If the Stateful Failover link uses the failover link or a regular data interface, skip this step. You have already defined the active and standby IP addresses for the interface.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The state link IP address and MAC address do not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

c. Enable the interface.

Note If the Stateful Failover link uses the failover link or regular data interface, skip this step. You have already enabled the interface.

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 4 Configure the failover groups. You can have at most two failover groups. The **failover group** command creates the specified failover group if it does not exist and enters the failover group configuration mode.

For each failover group, specify whether the failover group has primary or secondary preference using the **primary** or **secondary** command. You can assign the same preference to both failover groups. For load balancing configurations, you should assign each failover group a different unit preference.

The following example assigns failover group 1 a primary preference and failover group 2 a secondary preference:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# exit
```

Step 5 Assign each user context to a failover group using the **join-failover-group** command in context configuration mode.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.

Enter the following commands to assign each context to a failover group:

```
hostname(config)# context context_name
hostname(config-context)# join-failover-group {1 | 2}
hostname(config-context)# exit
```

Step 6 Enable failover:

```
hostname(config)# failover
```

Configure the Secondary Unit

When configuring LAN-based Active/Active failover, you need to bootstrap the secondary unit to recognize the failover link. This allows the secondary unit to communicate with and receive the running configuration from the primary unit.

To bootstrap the secondary unit in an Active/Active failover configuration, perform the following steps:

Step 1 (PIX security appliance only) Enable LAN-based failover:

```
hostname(config)# failover lan enable
```

Step 2 Define the failover interface. Use the same settings as you used for the primary unit:

a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name phy_if
```

The *if_name* argument assigns a logical name to the interface specified by the *phy_if* argument. The *phy_if* argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. On the ASA 5505 adaptive security appliance, the *phy_if* specifies a VLAN.

b. Assign the active and standby IP address to the failover link. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface.

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

c. Enable the interface:

```
hostname(config)# interface phy_if
hostname(config-if)# no shutdown
```

Step 3 (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```

Note This step is optional because by default units are designated as secondary unless previously configured otherwise.

Step 4 Enable failover:

```
hostname(config)# failover
```

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages `Beginning configuration replication: Sending to mate` and `End Configuration Replication to mate` appear on the active unit console.

Step 5 After the running configuration has completed replication, enter the following command to save the configuration to Flash memory:

```
hostname(config)# copy running-config startup-config
```

Step 6 If necessary, force any failover group that is active on the primary to the active state on the secondary unit. To force a failover group to become active on the secondary unit, enter the following command in the system execution space on the primary unit:

```
hostname# no failover active group group_id
```

The *group_id* argument specifies the group you want to become active on the secondary unit.

Configuring Optional Active/Active Failover Settings

The following optional Active/Active failover settings can be configured when you are initially configuring failover or after you have already established failover. Unless otherwise noted, the commands should be entered on the unit that has failover group 1 in the active state.

This section includes the following topics:

Configuring Failover Group Preemption

Assigning a primary or secondary priority to a failover group specifies which unit the failover group becomes active on when both units boot simultaneously. However, if one unit boots before the other, then both failover groups become active on that unit. When the other unit comes online, any failover groups that have the unit as a priority do not become active on that unit unless manually forced over, a failover occurs, or the failover group is configured with the **preempt** command. The **preempt** command causes a failover group to become active on the designated unit automatically when that unit becomes available.

Enter the following commands to configure preemption for the specified failover group:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# preempt [delay]
```

You can enter an optional *delay* value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit.

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information. You can use the **replication http** command to cause a failover group to replicate HTTP state information when Stateful Failover is enabled.

To enable HTTP state replication for a failover group, enter the following command. This command only affects the failover group in which it was configured. To enable HTTP state replication for both failover groups, you must enter this command in each group. This command should be entered in the system execution space.

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# replication http
```

Disabling and Enabling Interface Monitoring

You can monitor up to 250 interfaces on a unit. By default, monitoring of physical interfaces is enabled and the monitoring of subinterfaces is disabled. You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This lets you exclude interfaces attached to less critical networks from affecting your failover policy.

To disable health monitoring on an interface, enter the following command within a context:

```
hostname/context(config)# no monitor-interface if_name
```

To enable health monitoring on an interface, enter the following command within a context:

```
hostname/context(config)# monitor-interface if_name
```

Configuring Interface Health Monitoring

The security appliance sends hello packets out of each data interface to monitor interface health. If the security appliance does not receive a hello packet from the corresponding interface on the peer unit for over half of the hold time, then the additional interface testing begins. If a hello packet or a successful test result is not received within the specified hold time, the interface is marked as failed. Failover occurs if the number of failed interfaces meets the failover criteria.

Decreasing the poll and hold times enables the security appliance to detect and respond to interface failures more quickly, but may consume more system resources.

To change the default interface poll time, enter the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# polltime interface seconds
```

Valid values for the poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.

Configuring Failover Criteria

By default, if a single interface fails failover occurs. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs. The failover criteria is specified on a failover group basis.

To change the default failover criteria for the specified failover group, enter the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Configuring Virtual MAC Addresses

Active/Active failover uses virtual MAC addresses on all interfaces. If you do not specify the virtual MAC addresses, then they are computed as follows:

- Active unit default MAC address: 00a0.c9*physical_port_number.failover_group_id*01.
- Standby unit default MAC address: 00a0.c9*physical_port_number.failover_group_id*02.

Note If you have more than one Active/Active failover pair on the same network, it is possible to have the same default virtual MAC addresses assigned to the interfaces on one pair as are assigned to the interfaces of the other pairs because of the way the default virtual MAC addresses are determined. To avoid having duplicate MAC addresses on your network, make sure you assign each physical interface a virtual active and standby MAC address for all failover groups.

You can configure specific active and standby MAC addresses for an interface by entering the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# mac address phy_if active_mac standby_mac
```

The *phy_if* argument is the physical name of the interface, such as Ethernet1. The *active_mac* and *standby_mac* arguments are MAC addresses in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.

The *active_mac* address is associated with the active IP address for the interface, and the *standby_mac* is associated with the standby IP address for the interface.

There are multiple ways to configure virtual MAC addresses on the security appliance. When more than one method has been used to configure virtual MAC addresses, the security appliance uses the following order of preference to determine which virtual MAC address is assigned to an interface:

1. The **mac-address** command (in interface configuration mode) address.
2. The **failover mac address** command address.
3. The **mac-address auto** command generate address.
4. The automatically generated failover MAC address.

Use the **show interface** command to display the MAC address used by an interface.

Configuring Asymmetric Routing Support

When running in Active/Active failover, a unit may receive a return packet for a connection that originated through its peer unit. Because the security appliance that receives the packet does not have any connection information for the packet, the packet is dropped. This most commonly occurs when the two security appliances in an Active/Active failover pair are connected to different service providers and the outbound connection does not use a NAT address.

You can prevent the return packets from being dropped using the **asr-group** command on interfaces where this is likely to occur. When an interface configured with the **asr-group** command receives a packet for which it has no session information, it checks the session information for the other interfaces that are in the same group. If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all of the layer 2 header is rewritten and the packet is reinjected into the stream.

Note Using the **asr-group** command to configure asymmetric routing support is more secure than using the **static** command with the **nailed** option.

The **asr-group** command does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

Prerequisites

You must have to following configured for asymmetric routing support to function properly:

- Active/Active Failover
- Stateful Failover—passes state information for sessions on interfaces in the active failover group to the standby failover group.
- **replication http**—HTTP session state information is not passed to the standby failover group, and therefore is not present on the standby interface. For the security appliance to be able re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.

You can configure the **asr-group** command on an interface without having failover configured, but it does not have any effect until Stateful Failover is enabled.

Configuring Support for Asymmetrically Routed Packets

To configure support for asymmetrically routed packets, perform the following steps:

Step 1 Configure Active/Active Stateful Failover for the failover pair. See Configuring Active/Active Failover, page 14-28.

Step 2 For each interface that you want to participate in asymmetric routing support enter the following command. You must enter the command on the unit where the context is in the active state so that the command is replicated to the standby failover group. For more information about command replication, see Command Replication, page 14-12.

```
hostname/ctx(config)# interface phy_if
hostname/ctx(config-if)# asr-group num
```

Valid values for *num* range from 1 to 32. You need to enter the command for each interface that participates in the asymmetric routing group. You can view the number of ASR packets transmitted, received, or dropped by an interface using the **show interface detail** command. You can have more than one ASR group configured on the security appliance, but only one per interface. Only members of the same ASR group are checked for session information.

Example

Figure 14-1 shows an example of using the **asr-group** command for asymmetric routing support.

Figure 14-1 ASR Example

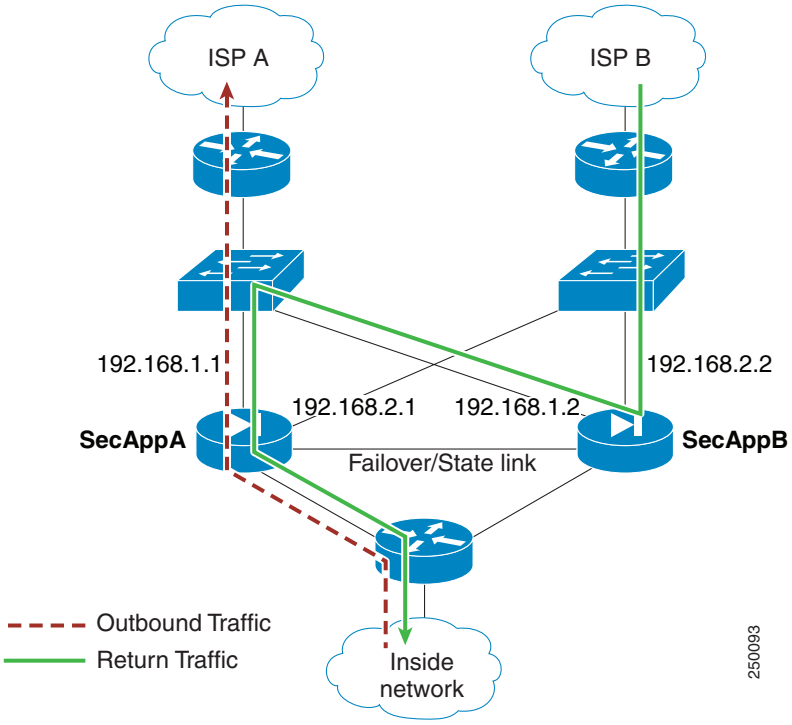

The two units have the following configuration (configurations show only the relevant commands). The device labeled SecAppA in the diagram is the primary unit in the failover pair.

Example 14-1 Primary Unit System Configuration

```
hostname primary
interface GigabitEthernet0/1
description LAN/STATE Failover Interface
interface GigabitEthernet0/2
no shutdown
interface GigabitEthernet0/3
no shutdown
interface GigabitEthernet0/4
no shutdown
interface GigabitEthernet0/5
no shutdown
failover
failover lan unit primary
failover lan interface folink GigabitEthernet0/1
failover link folink
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
primary
failover group 2
secondary
admin-context admin
context admin
description admin
```

```
allocate-interface GigabitEthernet0/2
allocate-interface GigabitEthernet0/3
config-url flash:/admin.cfg
join-failover-group 1
context ctx1
description context 1
allocate-interface GigabitEthernet0/4
allocate-interface GigabitEthernet0/5
config-url flash:/ctx1.cfg
join-failover-group 2
```

Example 14-2 admin Context Configuration

```
hostname SecAppA
interface GigabitEthernet0/2
nameif outsideISP-A
security-level 0
ip address 192.168.1.1 255.255.255.0 standby 192.168.1.2
asr-group 1
interface GigabitEthernet0/3
nameif inside
security-level 100
ip address 10.1.0.1 255.255.255.0 standby 10.1.0.11
monitor-interface outside
```

Example 14-3 ctx1 Context Configuration

```
hostname SecAppB
interface GigabitEthernet0/4
nameif outsideISP-B
security-level 0
ip address 192.168.2.2 255.255.255.0 standby 192.168.2.1
asr-group 1
interface GigabitEthernet0/5
nameif inside
security-level 100
ip address 10.2.20.1 255.255.255.0 standby 10.2.20.11
```

Figure 14-1 on page 14-38 shows the ASR support working as follows:

1. An outbound session passes through security appliance SecAppA. It exits interface outsideISP-A (192.168.1.1).
2. Because of asymmetric routing configured somewhere upstream, the return traffic comes back through the interface outsideISP-B (192.168.2.2) on security appliance SecAppB.
3. Normally the return traffic would be dropped because there is no session information for the traffic on interface 192.168.2.2. However, the interface is configure with the command **asr-group 1**. The unit looks for the session on any other interface configured with the same ASR group ID.
4. The session information is found on interface outsideISP-A (192.168.1.2), which is in the standby state on the unit SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is re-written with information for interface 192.168.1.1 and the traffic is redirected out of the interface 192.168.1.2, where it can then return through the interface on the unit from which it originated (192.168.1.1 on SecAppA). This forwarding continues as needed until the session ends.

Configuring Unit Health Monitoring

The security appliance sends hello packets over the failover interface to monitor unit health. If the standby unit does not receive a hello packet from the active unit for two consecutive polling periods, it sends additional testing packets through the remaining device interfaces. If a hello packet or a response to the interface test packets is not received within the specified hold time, the standby unit becomes active.

You can configure the frequency of hello messages when monitoring unit health. Decreasing the poll time allows a unit failure to be detected more quickly, but consumes more system resources.

To change the unit poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime [msec] time [holdtime [msec] time]
```

You can configure the polling frequency from 1 to 15 seconds or, if the optional **msec** keyword is used, from 200 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when failover occurs. The hold time must be at least 3 times the poll time. You can configure the hold time from 1 to 45 seconds or, if the optional **msec** keyword is used, from 800 to 990 milliseconds.

Setting the security appliance to use the minimum poll and hold times allows it to detect and respond to unit failures in under a second, but it also increases system resource usage and can cause false failure detection in cases where the networks are congested or where the security appliance is running near full capacity.

Configuring Failover Communication Authentication/Encryption

You can encrypt and authenticate the communication between failover peers by specifying a shared secret or hexadecimal key.

Note On the PIX 500 series security appliance, if you are using the dedicated serial failover cable to connect the units, then communication over the failover link is not encrypted even if a failover key is configured. The failover key only encrypts LAN-based failover communication.

Caution All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

Enter the following command on the active unit of an Active/Standby failover pair or on the unit that has failover group 1 in the active state of an Active/Active failover pair:

```
hostname(config)# failover key {secret | hex key}
```

The *secret* argument specifies a shared secret that is used to generate the encryption key. It can be from 1 to 63 characters. The characters can be any combination of numbers, letters, or punctuation. The **hex** *key* argument specifies a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

Note To prevent the failover key from being replicated to the peer unit in clear text for an existing failover configuration, disable failover on the active unit (or in the system execution space on the unit that has failover group 1 in the active state), enter the failover key on both units, and then re-enable failover. When failover is re-enabled, the failover communication is encrypted with the key.

For new LAN-based failover configurations, the **failover key** command should be part of the failover pair bootstrap configuration.

Verifying the Failover Configuration

This section describes how to verify your failover configuration. This section includes the following topics:

Using the show failover Command

This section describes the **show failover** command output. On each unit you can verify the failover status by entering the **show failover** command. The information displayed depends upon whether you are using Active/Standby or Active/Active failover.

This section includes the following topics:

show failover—Active/Standby

The following is sample output from the **show failover** command for Active/Standby Failover. Table 14-7 provides descriptions for the information shown.

```
hostname# show failover

Failover On
Cable status: N/A - LAN-based failover enabled
Failover unit Primary
Failover LAN Interface: fover Ethernet2 (up)
Unit Poll frequency 1 seconds, holdtime 3 seconds
Interface Poll frequency 15 seconds
Interface Policy 1
Monitored Interfaces 2 of 250 maximum
failover replication http
Last Failover at: 22:44:03 UTC Dec 8 2004
        This host: Primary - Active
                Active time: 13434 (sec)
                Interface inside (10.130.9.3): Normal
                Interface outside (10.132.9.3): Normal
        Other host: Secondary - Standby Ready
                Active time: 0 (sec)
                Interface inside (10.130.9.4): Normal
                Interface outside (10.132.9.4): Normal
```

```
Stateful Failover Logical Update Statistics
        Link : fover Ethernet2 (up)
        Stateful Obj    xmit       xerr       rcv        rerr
        General         1950       0          1733       0
        sys cmd         1733       0          1733       0
        up time         0          0          0          0
        RPC services    0          0          0          0
        TCP conn        6          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         106        0          0          0
        Xlate_Timeout   0          0          0          0
        VPN IKE upd     15         0          0          0
        VPN IPSEC upd   90         0          0          0
        VPN CTCP upd    0          0          0          0
        VPN SDI upd     0          0          0          0
        VPN DHCP upd    0          0          0          0

        Logical Update Queue Information
                        Cur     Max     Total
        Recv Q:         0       2       1733
        Xmit Q:         0       2       15225
```

In multiple context mode, using the **show failover** command in a security context displays the failover information for that context. The information is similar to the information shown when using the command in single context mode. Instead of showing the active/standby status of the unit, it displays the active/standby status of the context. Table 14-7 provides descriptions for the information shown.

```
Failover On
Last Failover at: 04:03:11 UTC Jan 4 2003
        This context: Negotiation
                Active time: 1222 (sec)
                Interface outside (192.168.5.121): Normal
                Interface inside (192.168.0.1): Normal
        Peer context: Not Detected
                Active time: 0 (sec)
                Interface outside (192.168.5.131): Normal
                Interface inside (192.168.0.11): Normal

Stateful Failover Logical Update Statistics
        Status: Configured.
        Stateful Obj    xmit       xerr       rcv        rerr
        RPC services    0          0          0          0
        TCP conn        99          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         22          0          0          0
        Xlate_Timeout   0          0          0          0
        GTP PDP         0          0          0          0
        GTP PDPMCB      0          0          0          0
```

Table 14-7 Show Failover Display Description

Field	Options
Failover	• On • Off
Cable status:	• Normal—The cable is connected to both units, and they both have power. • My side not connected—The serial cable is not connected to this unit. It is unknown if the cable is connected to the other unit. • Other side is not connected—The serial cable is connected to this unit, but not to the other unit. • Other side powered off—The other unit is turned off. • N/A—LAN-based failover is enabled.
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the logical and physical name of the failover link.
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	*n* seconds The number of seconds you set with the **failover polltime interface** command. The default is 15 seconds.
Interface Policy	Displays the number or percentage of interfaces that must fail to trigger failover.
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
failover replication http	Displays if HTTP state replication is enabled for Stateful Failover.
Last Failover at:	The date and time of the last failover in the following form: *hh*:*mm*:*ss* **UTC** *DayName Month Day yyyy* UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host: Other host:	For each host, the display shows the following information.
Primary or Secondary	• Active • Standby
Active time:	*n* (sec) The amount of time the unit has been active. This time is cumulative, so the standby unit, if it was active in the past, also shows a value.
slot *x*	Information about the module in the slot or empty.

Table 14-7 Show Failover Display Description (continued)

Field	Options
Interface *name* (*n.n.n.n*):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: • Failed—The interface has failed. • No Link—The interface line protocol is down. • Normal—The interface is working correctly. • Link Down—The interface has been administratively shut down. • Unknown—The security appliance cannot determine the status of the interface. • Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	• *interface_name*—The interface used for the Stateful Failover link. • Unconfigured—You are not using Stateful Failover. • up—The interface is up and functioning. • down—The interface is either administratively shutdown or is physically down. • failed—The interface has failed and is not passing stateful data.
Stateful Obj	For each field type, the following statistics are shown. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit. • 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.
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.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.

Table 14-7 Show Failover Display Description (continued)

Field	Options
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPMCB	GTP PDPMCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	For each field type, the following statistics are used: • Cur—Current number of packets • Max—Maximum number of packets • Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Show Failover—Active/Active

The following is sample output from the **show failover** command for Active/Active Failover. Table 14-8 provides descriptions for the information shown.

```
hostname# show failover

Failover On
Failover unit Primary
Failover LAN Interface: third GigabitEthernet0/2 (up)
Unit Poll frequency 1 seconds, holdtime 15 seconds
Interface Poll frequency 4 seconds
Interface Policy 1
Monitored Interfaces 8 of 250 maximum
failover replication http
Group 1 last failover at: 13:40:18 UTC Dec 9 2004
Group 2 last failover at: 13:40:06 UTC Dec 9 2004

  This host:    Primary
  Group 1       State:          Active
                Active time:    2896 (sec)
  Group 2       State:          Standby Ready
                Active time:    0 (sec)

                slot 0: ASA-5530 hw/sw rev (1.0/7.0(0)79) status (Up Sys)
                slot 1: SSM-IDS-20 hw/sw rev (1.0/5.0(0.11)S91(0.11)) status (Up)
                admin Interface outside (10.132.8.5): Normal
                admin Interface third (10.132.9.5): Normal
                admin Interface inside (10.130.8.5): Normal
                admin Interface fourth (10.130.9.5): Normal
                ctx1 Interface outside (10.1.1.1): Normal
                ctx1 Interface inside (10.2.2.1): Normal
                ctx2 Interface outside (10.3.3.2): Normal
                ctx2 Interface inside (10.4.4.2): Normal

  Other host:   Secondary
```

```
  Group 1       State:          Standby Ready
                Active time:    190 (sec)
  Group 2       State:          Active
                Active time:    3322 (sec)

                slot 0: ASA-5530 hw/sw rev (1.0/7.0(0)79) status (Up Sys)
                slot 1: SSM-IDS-20 hw/sw rev (1.0/5.0(0.1)S91(0.1)) status (Up)
                admin Interface outside (10.132.8.6): Normal
                admin Interface third (10.132.9.6): Normal
                admin Interface inside (10.130.8.6): Normal
                admin Interface fourth (10.130.9.6): Normal
                ctx1 Interface outside (10.1.1.2): Normal
                ctx1 Interface inside (10.2.2.2): Normal
                ctx2 Interface outside (10.3.3.1): Normal
                ctx2 Interface inside (10.4.4.1): Normal

Stateful Failover Logical Update Statistics
        Link : third GigabitEthernet0/2 (up)
        Stateful Obj    xmit       xerr       rcv        rerr
        General         1973       0          1895       0
        sys cmd         380        0          380        0
        up time         0          0          0          0
        RPC services    0          0          0          0
        TCP conn        1435       0          1450       0
        UDP conn        0          0          0          0
        ARP tbl         124        0          65         0
        Xlate_Timeout   0          0          0          0
        VPN IKE upd     15         0          0          0
        VPN IPSEC upd   90         0          0          0
        VPN CTCP upd    0          0          0          0
        VPN SDI upd     0          0          0          0
        VPN DHCP upd    0          0          0          0

        Logical Update Queue Information
                        Cur     Max     Total
        Recv Q:         0       1       1895
        Xmit Q:         0       0       1940
```

The following is sample output from the **show failover group** command for Active/Active Failover. The information displayed is similar to that of the **show failover** command, but limited to the specified group. Table 14-8 provides descriptions for the information shown.

```
hostname# show failover group 1

  Last Failover at: 04:09:59 UTC Jan 4 2005

  This host:    Secondary
                State:          Active
                Active time:    186 (sec)

                admin Interface outside (192.168.5.121): Normal
                admin Interface inside (192.168.0.1): Normal


  Other host:   Primary
                State:          Standby
                Active time:    0 (sec)

                admin Interface outside (192.168.5.131): Normal
                admin Interface inside (192.168.0.11): Normal

  Stateful Failover Logical Update Statistics
        Status: Configured.
```

```
RPC services    0           0           0           0
TCP conn        33          0           0           0
UDP conn        0           0           0           0
ARP tbl         12          0           0           0
Xlate_Timeout   0           0           0           0
GTP PDP         0           0           0           0
GTP PDPMCB      0           0           0           0
```

Table 14-8 Show Failover Display Description

Field	Options
Failover	• On • Off
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the logical and physical name of the failover link.
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	*n* seconds The number of seconds you set with the **failover polltime interface** command. The default is 15 seconds.
Interface Policy	Displays the number or percentage of interfaces that must fail before triggering failover.
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
Group 1 Last Failover at: Group 2 Last Failover at:	The date and time of the last failover for each group in the following form: *hh*:*mm*:*ss* **UTC** *DayName Month Day yyyy* UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host: Other host:	For each host, the display shows the following information.
Role	Primary or Secondary
System State	• Active or Standby Ready • Active Time in seconds
Group 1 State Group 2 State	• Active or Standby Ready • Active Time in seconds
slot *x*	Information about the module in the slot or empty.

Table 14-8 Show Failover Display Description (continued)

Field	Options
context Interface *name* (*n.n.n.n*):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: • Failed—The interface has failed. • No link—The interface line protocol is down. • Normal—The interface is working correctly. • Link Down—The interface has been administratively shut down. • Unknown—The security appliance cannot determine the status of the interface. • Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	• *interface_name*—The interface used for the Stateful Failover link. • Unconfigured—You are not using Stateful Failover. • up—The interface is up and functioning. • down—The interface is either administratively shutdown or is physically down. • failed—The interface has failed and is not passing stateful data.
Stateful Obj	For each field type, the following statistics are used. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit. • 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
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.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.

Table 14-8 Show Failover Display Description (continued)

Field	Options
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPMCB	GTP PDPMCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	For each field type, the following statistics are used: • Cur—Current number of packets • Max—Maximum number of packets • Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Viewing Monitored Interfaces

To view the status of monitored interfaces, enter the following command. In single context mode, enter this command in global configuration mode. In multiple context mode, enter this command within a context.

```
primary/context(config)# show monitor-interface
```

For example:

```
hostname/context(config)# show monitor-interface
        This host: Primary - Active
                Interface outside (192.168.1.2): Normal
                Interface inside (10.1.1.91): Normal
        Other host: Secondary - Standby
                Interface outside (192.168.1.3): Normal
                Interface inside (10.1.1.100): Normal
```

Displaying the Failover Commands in the Running Configuration

To view the failover commands in the running configuration, enter the following command:

```
hostname(config)# show running-config failover
```

All of the failover commands are displayed. On units running multiple context mode, enter this command in the system execution space. Entering **show running-config all failover** displays the failover commands in the running configuration and includes commands for which you have not changed the default value.

Testing the Failover Functionality

To test failover functionality, perform the following steps:

Step 1 Test that your active unit or failover group is passing traffic as expected by using FTP (for example) to send a file between hosts on different interfaces.

Step 2 Force a failover to the standby unit by entering the following command:

- For Active/Standby failover, enter the following command on the active unit:

  ```
  hostname(config)# no failover active
  ```

- For Active/Active failover, enter the following command on the unit where the failover group containing the interface connecting your hosts is active:

  ```
  hostname(config)# no failover active group group_id
  ```

Step 3 Use FTP to send another file between the same two hosts.

Step 4 If the test was not successful, enter the **show failover** command to check the failover status.

Step 5 When you are finished, you can restore the unit or failover group to active status by enter the following command:

- For Active/Standby failover, enter the following command on the active unit:

  ```
  hostname(config)# failover active
  ```

- For Active/Active failover, enter the following command on the unit where the failover group containing the interface connecting your hosts is active:

  ```
  hostname(config)# failover active group group_id
  ```

Controlling and Monitoring Failover

This sections describes how to control and monitor failover. This section includes the following topics:

Forcing Failover

To force the standby unit or failover group to become active, enter one of the following commands:

- For Active/Standby failover:

  Enter the following command on the standby unit:

  ```
  hostname# failover active
  ```

  Or, enter the following command on the active unit:

```
hostname# no failover active
```

- For Active/Active failover:

  Enter the following command in the system execution space of the unit where the failover group is in the standby state:

  ```
  hostname# failover active group group_id
  ```

  Or, enter the following command in the system execution space of the unit where the failover group is in the active state:

  ```
  hostname# no failover active group group_id
  ```

  Entering the following command in the system execution space causes all failover groups to become active:

  ```
  hostname# failover active
  ```

Disabling Failover

To disable failover, enter the following command:

```
hostname(config)# no failover
```

Disabling failover on an Active/Standby pair causes the active and standby state of each unit to be maintained until you restart. For example, the standby unit remains in standby mode so that both units do not start passing traffic. To make the standby unit active (even with failover disabled), see the "Forcing Failover" section on page 14-50.

Disabling failover on an Active/Active pair causes the failover groups to remain in the active state on whichever unit they are currently active on, no matter which unit they are configured to prefer. The no failover command should be entered in the system execution space.

Restoring a Failed Unit or Failover Group

To restore a failed unit to an unfailed state, enter the following command:

```
hostname(config)# failover reset
```

To restore a failed Active/Active failover group to an unfailed state, enter the following command:

```
hostname(config)# failover reset group group_id
```

Restoring a failed unit or group to an unfailed state does not automatically make it active; restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with the **preempt** command. If previously active, a failover group becomes active if it is configured with the **preempt** command and if the unit on which it failed is the preferred unit.

Monitoring Failover

When a failover occurs, both security appliances send out system messages. This section includes the following topics:

- Failover System Messages, page 14-52

Failover System Messages

The security appliance issues a number of system messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the *Cisco Security Appliance Logging Configuration and System Log Messages* to enable logging and to see descriptions of the system messages.

Note During switchover, failover logically shuts down and then bring up interfaces, generating syslog 411001 and 411002 messages. This is normal activity.

Debug Messages

To see debug messages, enter the **debug fover** command. See the *Cisco Security Appliance Command Reference* for more information.

Note Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station. See the **snmp-server** and **logging** commands in the *Cisco Security Appliance Command Reference* for more information.