Configuring Failover

This chapter describes how to configure Active/Standby or Active/Active failover, and includes the following sections:

Introduction to Failover

Failover Overview

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

The ASA supports two failover modes, Active/Active failover and Active/Standby failover. Each failover mode has its own method for determining and performing failover.

- In Active/Standby failover, one unit is the active unit. It passes traffic. The standby unit does not actively pass traffic. When a failover occurs, the active unit fails over to the standby unit, which then becomes active. You can use Active/Standby failover for ASAs in single or multiple context mode.
- In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into 2 *failover groups*. A failover group is simply a logical group of one or more security contexts. One group is assigned to be active on the primary ASA, and the other group is assigned to be active on the secondary ASA. When a failover occurs, it occurs at the failover group level.

Both failover modes support stateful or stateless failover.

Failover System Requirements

This section describes the hardware, software, and license requirements for ASAs in a failover configuration.

Hardware Requirements

The two units in a failover configuration must:

- Be the same model.
- Have the same number and types of interfaces.
- Have the same modules installed (if any)
- Have the same RAM installed.

If you are 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 same firewall mode (routed or transparent).
- Be in the same context mode (single or multiple).

- Have the same major (first number) and minor (second number) software version. However, you can temporarily use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 8.3(1) to Version 8.3(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

  See the "Upgrading a Failover Pair or ASA Cluster" section on page 42-5 for more information about upgrading the software on a failover pair.
- Have the same AnyConnect images. If the failover pair has mismatched images when a hitless upgrade is performed, then the clientless SSL VPN connection terminates in the final reboot step of the upgrade process, the database shows an orphaned session, and the IP pool shows that the IP address assigned to the client is "in use."

License Requirements

The two units in a failover configuration do not need to have identical licenses; the licenses combine to make a failover cluster license. See the "Failover or ASA Cluster Licenses" section on page 4-30 for more information.

Failover and Stateful Failover Links

The failover link and the optional Stateful Failover link are dedicated connections between the two units.

Caution All information sent over the failover and state links is sent in clear text unless you secure the communication with an IPsec tunnel or a failover key. If the ASA 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 an IPsec tunnel or a failover key if you are using the ASA to terminate VPN tunnels.

Failover Link

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

Failover Link Data

The following information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keep-alives)
- Network link status

- MAC address exchange
- Configuration replication and synchronization

Interface for the Failover Link

You can use any unused interface (physical, redundant, or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface can only be used for the failover link (and optionally also for the state link).

Connecting the Failover Link

Connect the 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 failover interfaces of the ASA.
- Using an Ethernet cable to connect the units directly, without the need for an external switch.

If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface and caused the link to come down.

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.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link (also known as the state link) to pass connection state information.

You have three interface options for the state link:

- Dedicated Interface (Recommended), page 7-4
- Shared with the Failover Link, page 7-5
- Shared with a Regular Data Interface (Not Recommended), page 7-5

Note Do not use a management interface for the state link.

Dedicated Interface (Recommended)

You can use a dedicated interface (physical, redundant, or EtherChannel) for the state link. Connect a dedicated state 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 failover interfaces of the ASA.
- Using an Ethernet cable to connect the appliances directly, without the need for an external switch.

  If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface and caused the link to come down.

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.

For optimum performance when using long distance failover, the latency for the failover link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is more than10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

Shared with the Failover Link

Sharing a failover link might be necessary if you do not have enough interfaces. If you use the failover link as the state link, you should use the fastest Ethernet interface available. If you experience performance problems on that interface, consider dedicating a separate interface for the state link.

Shared with a Regular Data Interface (Not Recommended)

Sharing a data interface with the state link 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.

Using a data interface as the state link is supported in single context, routed mode only.

Avoiding Interrupted Failover and Data Links

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the ASA can use the data interfaces to determine if a failover is required. Subsequently, the failover operation is suspended until the health of the failover link is restored.

See the following connection scenarios to design a resilient failover network.

Scenario 1—Not Recommended

If a single switch or a set of switches are used to connect both failover and data interfaces between two ASAs, then when a switch or inter-switch-link is down, both ASAs become active. Therefore, the following two connection methods shown in Figure 7-1 and Figure 7-2 are NOT recommended.

Figure 7-1 Connecting with a Single Switch—Not Recommended

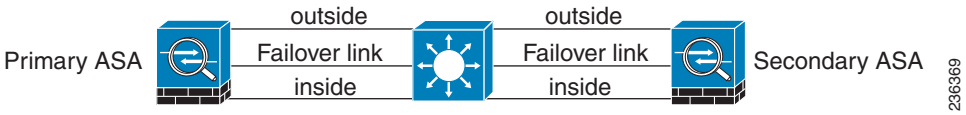

Figure 7-2 Connecting with a Double Switch—Not Recommended

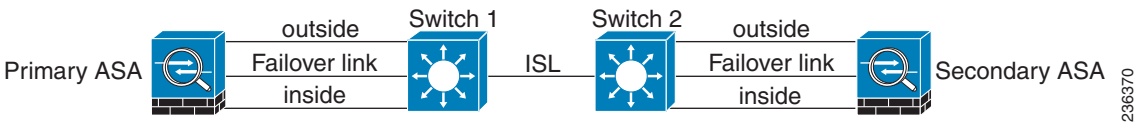

Scenario 2—Recommended

We recommend that failover links NOT use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in Figure 7-3 and Figure 7-4.

Figure 7-3 Connecting with a Different Switch

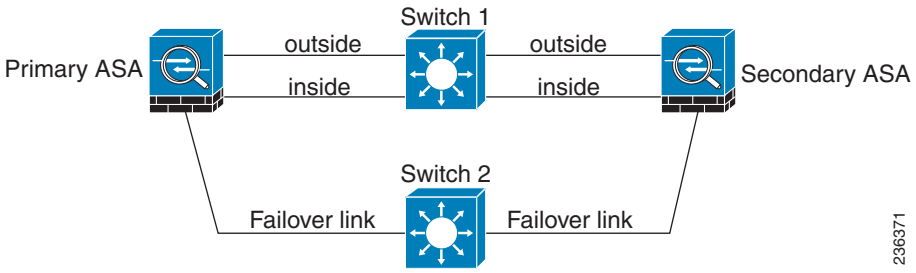

Figure 7-4 Connecting with a Cable

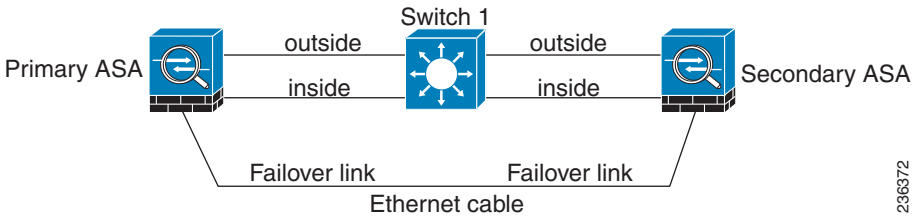

Scenario 3—Recommended

If the ASA data interfaces are connected to more than one set of switches, then a failover link can be connected to one of the switches, preferably the switch on the secure (inside) side of network, as shown in Figure 7-5.

Figure 7-5 Connecting with a Secure Switch

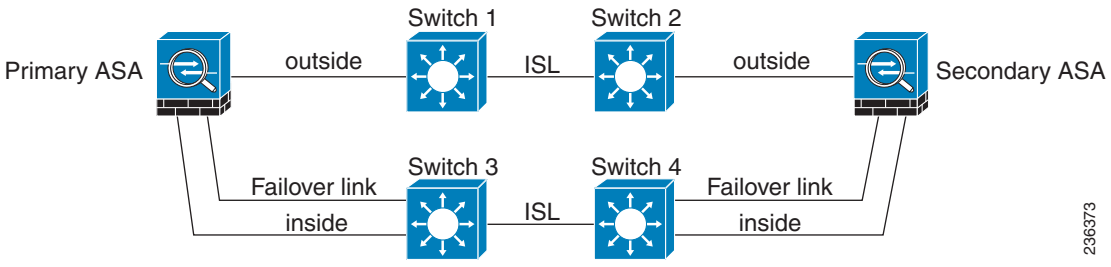

Scenario 4—Recommended

The most reliable failover configurations use a redundant interface on the failover link, as shown in Figure 7-6 and Figure 7-7.

Figure 7-6 Connecting with Redundant Interfaces

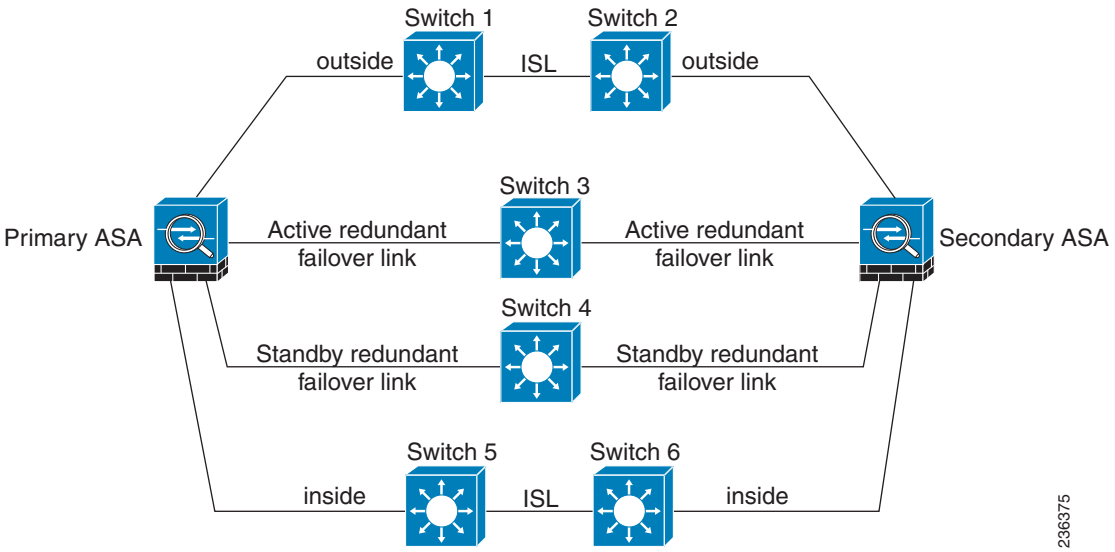

Figure 7-7 Connecting with Inter-switch Links

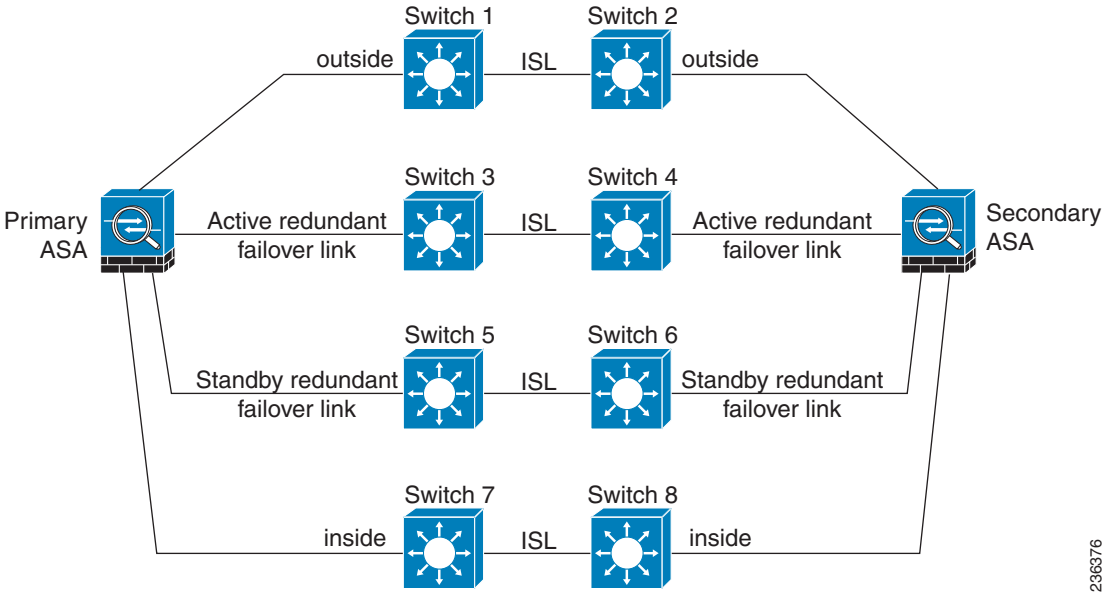

MAC Addresses and IP Addresses

When you configure your interfaces, you must specify an active IP address and a standby IP address on the same network.

1. When the primary unit or failover group fails over, the secondary unit assumes the IP addresses and MAC addresses of the primary unit and begins passing traffic.
2. 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 If the secondary unit boots without detecting the primary unit, the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. However, when the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic.Similarly, if you swap out the primary unit with new hardware, a new MAC address is used.

Virtual MAC addresses guard against this disruption because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. In multiple context mode, the ASA generates virtual active and standby MAC addresses by default. See the "Information About MAC Addresses" section on page 6-11 for more information. In single context mode, you can manually configure virtual MAC addresses; see the "Configuring Active/Active Failover" section on page 7-30 for more information.

If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The ASA does not send gratuitous ARPs for static NAT addresses when the MAC address changes, so connected routers do not learn of the MAC address change for these addresses.

Note The IP address and MAC address for the state link do not change at failover; the only exception is if the state link is configured on a regular data interface.

Intra- and Inter-Chassis Module Placement for the ASA Services Module

You can place the primary and secondary ASASMs within the same switch or in two separate switches. The following sections describe each option:

- Intra-Chassis Failover, page 7-8
- Inter-Chassis Failover, page 7-9

Intra-Chassis Failover

If you install the secondary ASASM in the same switch as the primary ASASM, you protect against module-level failure. To protect against switch-level failure, as well as module-level failure, see the "Inter-Chassis Failover" section on page 7-9.

Even though both ASASMs are assigned the same VLANs, only the active module takes part in networking. The standby module does not pass any traffic.

Figure 7-8 shows a typical intra-switch configuration.

Figure 7-8 Intra-Switch Failover

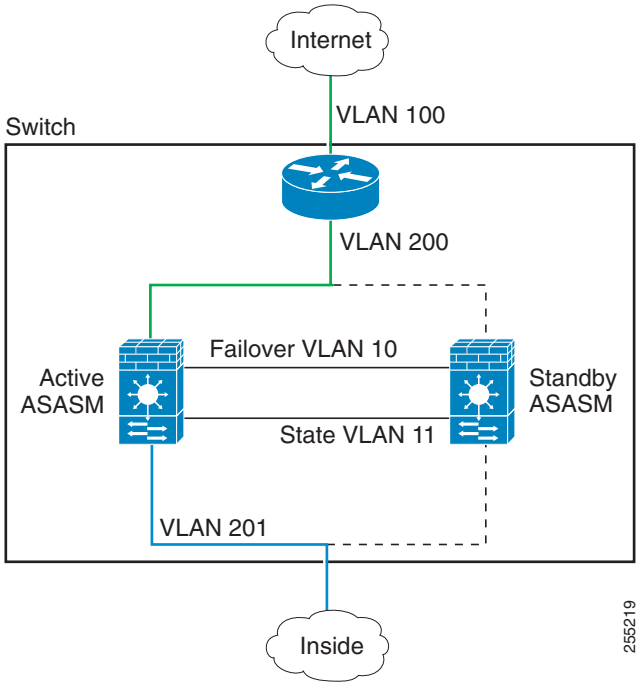

Inter-Chassis Failover

To protect against switch-level failure, you can install the secondary ASASM in a separate switch. The ASASM does not coordinate failover directly with the switch, but it works harmoniously with the switch failover operation. See the switch documentation to configure failover for the switch.

For the best reliability of failover communications between ASASMs, we recommend that you configure an EtherChannel trunk port between the two switches to carry the failover and state VLANs.

For other VLANs, you must ensure that both switches have access to all firewall VLANs, and that monitored VLANs can successfully pass hello packets between both switches.

Figure 7-9 shows a typical switch and ASASM redundancy configuration. The trunk between the two switches carries the failover ASASM VLANs (VLANs 10 and 11).

Note ASASM failover is independent of the switch failover operation; however, ASASM works in any switch failover scenario.

Figure 7-9 Normal Operation

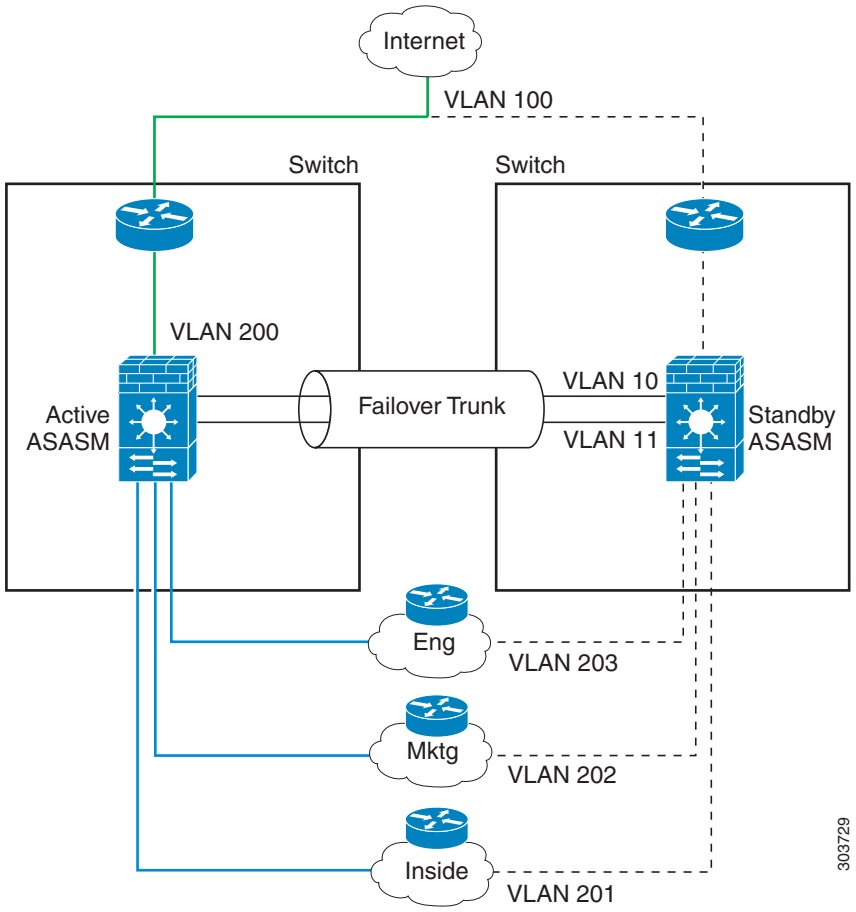

If the primary ASASM fails, then the secondary ASASM becomes active and successfully passes the firewall VLANs (Figure 7-10).

Figure 7-10 ASASM Failure

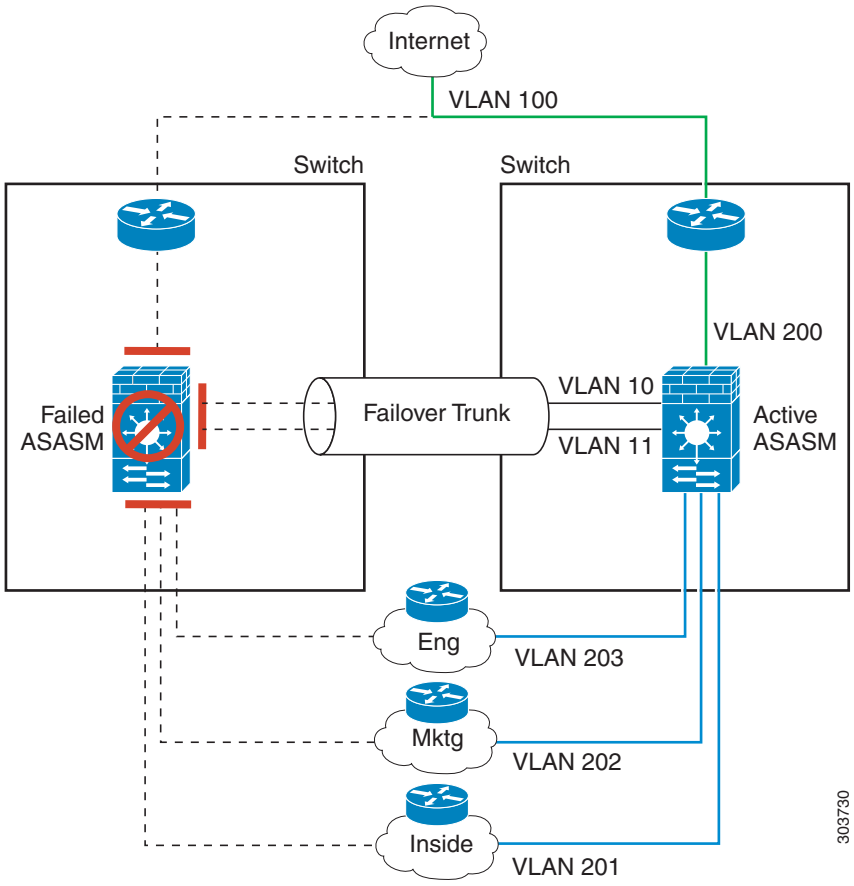

If the entire switch fails, as well as the ASASM (such as in a power failure), then both the switch and the ASASM fail over to their secondary units (Figure 7-11).

Figure 7-11 Switch Failure

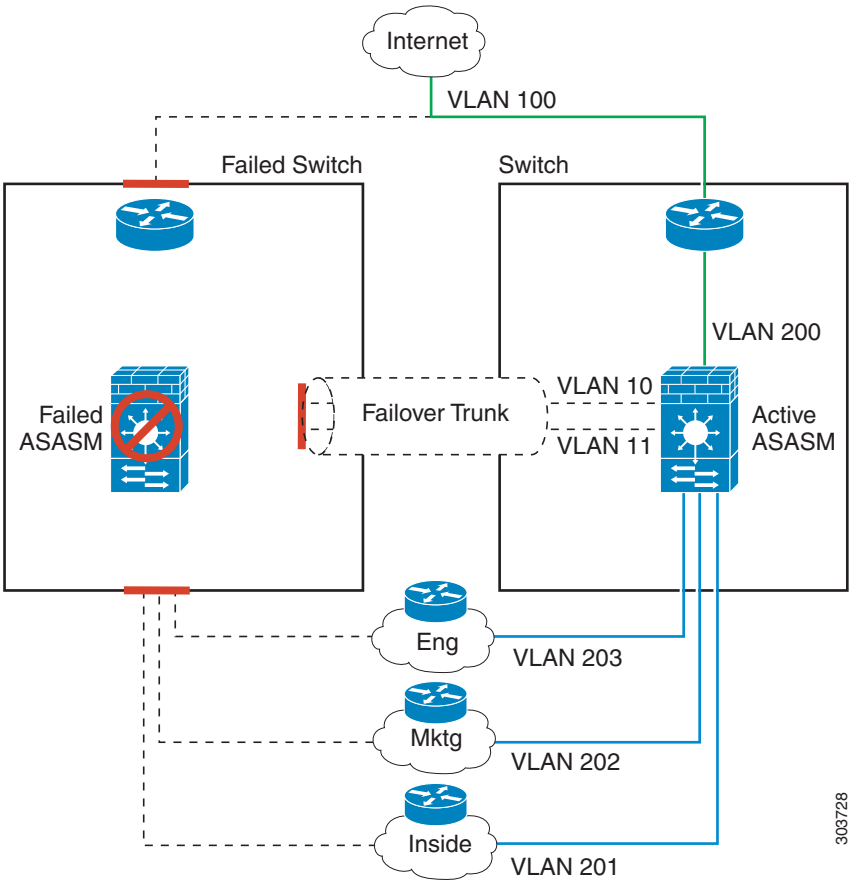

Stateless and Stateful Failover

The ASA supports two types of failover, stateless and stateful for both the Active/Standby and Active/Active modes.

- Stateless Failover, page 7-13
- Stateful Failover, page 7-13

Note Some configuration elements for clientless SSL VPN (such as bookmarks and customization) use the VPN failover subsystem, which is part of Stateful Failover. You must use Stateful Failover to synchronize these elements between the members of the failover pair. Stateless failover is not recommended for clientless SSL VPN.

Stateless Failover

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

Note Some configuration elements for clientless SSL VPN (such as bookmarks and customization) use the VPN failover subsystem, which is part of Stateful Failover. You must use Stateful Failover to synchronize these elements between the members of the failover pair. Stateless (regular) failover is not recommended for clientless SSL VPN.

Stateful Failover

When Stateful Failover is enabled, the active unit continually passes per-connection state information to the standby unit, or in Active/Active failover, between the active and standby failover groups. 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.

- Supported Features, page 7-13
- Unsupported Features, page 7-14

Supported Features

The following state information is passed to the standby ASA when Stateful Failover is enabled:

- 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)—By default, the ASA does not replicate HTTP session information when Stateful Failover is enabled. Because HTTP sessions are typically short-lived, and because HTTP clients typically retry failed connection attempts, not replicating HTTP sessions increases system performance without causing serious data or connection loss.
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signalling sessions
- ICMP connection state—ICMP connection replication is enabled only if the respective interface is assigned to an asymmetric routing group.
- Dynamic Routing Protocols—Stateful Failover participates in dynamic routing protocols, like OSPF and EIGRP, so routes that are learned through dynamic routing protocols on the active unit are maintained in a Routing Information Base (RIB) table on the standby unit. Upon a failover event, packets travel normally with minimal disruption to traffic because the active secondary ASA initially has rules that mirror the primary ASA. Immediately after failover, the re-convergence timer starts on the newly Active unit. Then the epoch number for the RIB table increments. During re-convergence, OSPF and EIGRP routes become updated with a new epoch number. Once the timer is expired, stale route entries (determined by the epoch number) are removed from the table. The RIB then contains the newest routing protocol forwarding information on the newly Active unit.

Note Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

- Cisco IP SoftPhone sessions—If a 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 Cisco Call Manager. This connection loss 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.
- VPN—VPN end-users do not have to reauthenticate or reconnect the VPN session after a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.

Unsupported Features

The following state information is *not* passed to the standby ASA when Stateful Failover is enabled:

- The HTTP connection table (unless HTTP replication is enabled)
- The user authentication (uauth) table
- Application inspections that are subject to advanced TCP-state tracking—The TCP state of these connections is not automatically replicated. While these connections are replicated to the standby unit, there is a best-effort attempt to re-establish a TCP state.
- DHCP server address leases
- State information for modules, such as the ASA IPS SSP or ASA CX SSP.
- Phone proxy connections—When the active unit goes down, the call fails, media stops flowing, and the phone should unregister from the failed unit and reregister with the active unit. The call must be re-established.
- Selected clientless SSL VPN features:
  - Smart Tunnels
  - Port Forwarding
  - Plugins
  - Java Applets
  - IPv6 clientless or Anyconnect sessions
  - Citrix authentication (Citrix users must reauthenticate after failover)

Transparent Firewall Mode Requirements

- Transparent Mode Requirements for Appliances, page 7-15
- Transparent Mode Requirements for Modules, page 7-15

Transparent Mode Requirements for Appliances

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 ASA on both the inside and outside interfaces with an EtherType access rule.

  ```
  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 ASA 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 interface monitoring.
- Increase interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the interface holdtime.

Transparent Mode Requirements for Modules

To avoid loops when you use failover in transparent mode, you should allow BPDUs to pass (the default), and you must use switch software that supports BPDU forwarding.

Loops can occur if both modules are active at the same time, such as when both modules are discovering each other's presence, or due to a bad failover link. Because the ASASMs bridge packets between the same two VLANs, loops can occur when inside packets destined for the outside get endlessly replicated by both ASASMs (see Figure 7-12). The spanning tree protocol can break such loops if there is a timely exchange of BPDUs. To break the loop, BPDUs sent between VLAN 200 and VLAN 201 need to be bridged.

Figure 7-12 Transparent Mode Loop

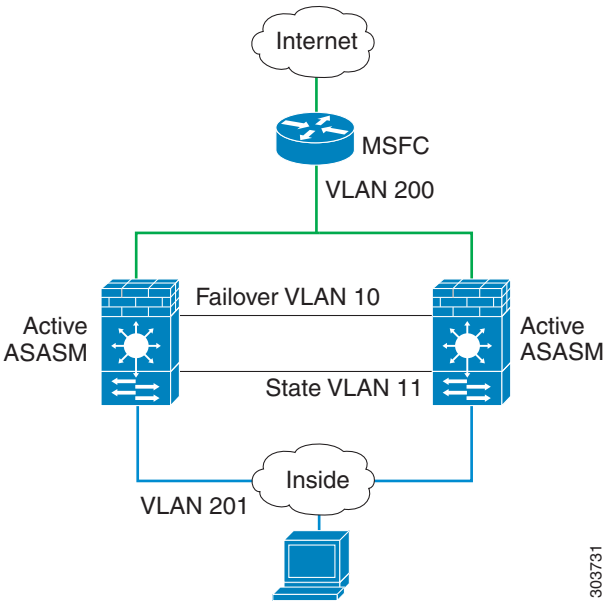

Failover Health Monitoring

The ASA monitors each unit for overall health and for interface health. This section includes information about how the ASA performs tests to determine the state of each unit.

- Unit Health Monitoring, page 7-16
- Interface Monitoring, page 7-17

Unit Health Monitoring

The ASA 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 data interface, including the failover link, to validate whether or not the peer is responsive. The action that the ASA takes depends on the response from the other unit. See the following possible actions:

- If the ASA receives a response on the failover link, then it does not fail over.
- If the ASA does not receive a response on the failover link, but it does receive a response on a data 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 ASA does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

Interface Monitoring

You can monitor up to 250 interfaces (in multiple mode, divided between all contexts). You should monitor important interfaces. For example in multiple mode, 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 ASA 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.

Monitored interfaces can have the following status:

- Unknown—Initial status. This status can also mean the status cannot be determined.
- Normal—The interface is receiving traffic.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface or VLAN is administratively down.
- No Link—The physical link for the interface is down.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

If an interface has IPv4 and IPv6 addresses configured on it, the ASA uses the IPv4 addresses to perform the health monitoring.

If an interface has only IPv6 addresses configured on it, then the ASA uses IPv6 neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the ASA uses the IPv6 all nodes address (FE02::1).

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 ASA 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 Times

Table 7-1 shows the minimum, default, and maximum failover times.

Table 7-1 ASA 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 module interface link down.	2 seconds	5 seconds	15 seconds
Active unit IPS or CSC module fails.	2 seconds	2 seconds	2 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

Configuration Synchronization

Failover includes two types of configuration synchronization:

- Running Configuration Replication, page 7-18
- Command Replication, page 7-19

Running Configuration Replication

Running configuration replication 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.

When the replication starts, the ASA console on the active unit displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the ASA displays the message “End Configuration Replication to mate.” Depending on 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. You should save the configuration to flash memory according to the “Saving Configuration Changes” section on page 3-25.

Note 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 during the configuration replication process.

Note The **crypto ca server** command and related sub commands are not synchronized to the failover peer.

Note Configuration syncing does not replicate the following files and configuration components, so you must copy these files manually so they match:

- AnyConnect images
- CSD images
- AnyConnect profiles
- Local Certificate Authorities (CAs)
- ASA images
- ASDM images

Command Replication

After startup, commands that you enter on the active unit are immediately replicated to the standby unit. You do not have to save the active configuration to flash memory to replicate the commands.

In Active/Active failover, commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state.

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 ASA:

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

The following commands are *not* replicated to the standby ASA:

- 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**
- **show**
- **terminal pager** and **pager**

Information About Active/Standby Failover

Active/Standby failover lets you use a standby ASA 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.

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

Primary/Secondary Roles 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.

Active Unit Determination at Startup

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.

Failover Events

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 7-2 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 7-2 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 link as failed	Mark failover link 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 link as failed	Become active	If the failover link is down at startup, both units become active.
State 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.

Information About Active/Active Failover

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

Active/Active Failover Overview

In an Active/Active failover configuration, both ASAs can pass network traffic. Active/Active failover is only available to ASAs in multiple context mode. In Active/Active failover, you divide the security contexts on the ASA into a maximum of 2 failover groups.

A failover group is simply a logical group of one or more security contexts. You can assign failover group to be active on the primary ASA, and failover group 2 to be active on the secondary ASA. When a failover occurs, it occurs at the failover group level. For example, depending on interface failure patterns, it is possible for failover group 1 to fail over to the secondary ASA, and subsequently failover group 2 to fail over to the primary ASA. This event could occur if the interfaces in failover group 1 are down on the primary ASA but up on the secondary ASA, while the interfaces in failover group 2 are down on the secondary ASA but up on the primary ASA.

The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default. If you want Active/Active failover, but are otherwise uninterested in multiple contexts, the simplest configuration would be to add one additional context and assign it to failover group 2.

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

Note You can assign both failover groups to one ASA if desired, but then you are not taking advantage of having two active ASAs.

Primary/Secondary Roles and Active/Standby Status for a Failover Group

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:

- The primary unit provides the running configuration to the pair when they boot simultaneously.
- Each failover group in the configuration is configured with a primary or secondary unit preference.

Active Unit Determination for Failover Groups at Startup

The unit on which a failover group becomes active 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 occurs:
  - A failover occurs.
  - You manually force a failover.
  - You configured preemption for the failover group, 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.

Failover Events

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.

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.

Table 7-3 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 7-3 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.
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 failover group preemption is configured, 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.
State 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 link 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.

Licensing Requirements Failover

Active/Standby Failover

Model	License Requirement
ASA 5505	Security Plus License. (Stateful failover is not supported).
ASA 5510, ASA 5512-X	Security Plus License.
All other models	Base License.

Failover units do not require the same license on each unit. If you have licenses on both units, they combine into a single running failover cluster license. The exceptions to this rule include:

- Security Plus license for the ASA 5505, 5510, and 5512-X—The Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.
- IPS module license for the ASA 5500-X—You must purchase an IPS module license for each unit, just as you would need to purchase a hardware module for each unit for other models.
- Encryption license—Both units must have the same encryption license.

Active/Active Failover

Model	License Requirement
ASA 5505	No support.
ASA 5510, ASA 5512-X	Security Plus License.
All other models	Base License.

Failover units do not require the same license on each unit. If you have licenses on both units, they combine into a single running failover cluster license. The exceptions to this rule include:

- Security Plus license for the ASA 5510 and 5512-X—The Base license does not support failover, so you cannot enable failover on a standby unit that only has the Base license.
- IPS module license for the ASA 5500-X—You must purchase an IPS module license for each unit, just as you would need to purchase a hardware module for each unit for other models.
- Encryption license—Both units must have the same encryption license.

Prerequisites for Failover

See the "Failover System Requirements" section on page 7-2.

Guidelines and Limitations

For Auto Update guidelines with failover, see the "Auto Update Server Support in Failover Configurations" section on page 42-36.

Context Mode Guidelines

- Active/Standby mode is supported in single and multiple context mode.
- Active/Active mode is supported only in multiple context mode.
- For multiple context mode, perform all steps in the system execution space unless otherwise noted.
- ASA failover replication fails if you try to make a configuration change in two or more contexts at the same time. The workaround is to make configuration changes in each context sequentially.

Firewall Mode Guidelines

Supported in transparent and routed firewall mode.

IPv6 Guidelines

IPv6 is supported.

Model Guidelines

Stateful failover is not supported on the ASA 5505. See the "Licensing Requirements Failover" section on page 7-24 for other guidelines.

Additional Guidelines and Limitations

- Configuring port security on the switch(es) connected to an ASA failover pair can cause communication problems when a failover event occurs. This problem occurs when a secure MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.
- You can monitor up to 250 interfaces on a unit, across all contexts.
- For Active/Active failover, no two interfaces in the same context should be configured in the same ASR group.
- For Active/Active failover, you can define a maximum of two failover groups.
- For Active/Active failover, when removing failover groups, you must remove failover group 1 last. Failover group1 always contains the admin context. Any context not assigned to a failover group defaults to failover group 1. You cannot remove a failover group that has contexts explicitly assigned to it.

Default Settings

By default, the failover policy consists of the following:

- No HTTP replication in Stateful Failover.
- A single interface failure causes failover.
- The interface poll time is 5 seconds.
- The interface hold time is 25 seconds.
- The unit poll time is 1 second.

- The unit hold time is 15 seconds.
- Virtual MAC addresses are enabled in multiple context mode; in single context mode, they are disabled.
- Monitoring on all physical interfaces, or for the ASA 5505 and ASASM, all VLAN interfaces.

Configuring Active/Standby Failover

Configuring the Primary Unit for Active/Standby Failover

Follow the steps in this section to configure the primary in an Active/Standby failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit.

Prerequisites

- Configure standby IP addresses for all interfaces except for the failover and state links according to Chapter 11, "Completing Interface Configuration (Routed Mode)," or Chapter 12, "Completing Interface Configuration (Transparent Mode)."
- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	**failover lan unit primary**	Designates this unit as the primary unit.
Step 2	**failover lan interface** *if_name interface_id* **Example:** `ciscoasa(config)# failover lan interface folink gigabitethernet0/3`	Specifies the interface to be used as the failover link. This interface cannot be used for any other purpose (except, optionally, the state link). The *if_name* argument assigns a name to the interface. The *interface_id* argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the *interface_id* specifies a VLAN ID. **Note** Although you can use an EtherChannel as a failover or state link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

	Command	Purpose

Step 3	**failover interface ip** *failover_if_name* {*ip_address mask* | *ipv6_address*/*prefix*} **standby** *ip_address*
	
	Example:
	`ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2`
	
	Or:
	
	`ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71`
		Assigns the active and standby IP addresses to the failover link. This address should be on an unused subnet.
		
		The standby IP address must be in the same subnet as the active IP address.

Step 4	**interface** *failover_interface_id*
	no shutdown
	
	Example:
	`ciscoasa(config)# interface gigabitethernet 0/3`
	`ciscoasa(config-if)# no shutdown`
		Enables the failover link.

Step 5	**failover link** *if_name interface_id* **Example:** `ciscoasa(config)# failover link statelink gigabitethernet0/4`	(Optional) Specifies the interface you want to use as the state link. We recommend specifying a separate interface from the failover link or data interfaces. The *if_name* argument assigns a name to the interface. The *interface_id* argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the *interface_id* specifies a VLAN ID. **Note** Although you can use an EtherChannel as a failover or state link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Step 6	**failover interface ip** *state_if_name* {*ip_address mask* | *ipv6_address*/*prefix*} **standby** *ip_address*
	
	Example:
	`ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2`
	
	Or:
	
	`ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby 2001:a0a:b00:a::a0a:b71`
		If you specified a separate state link, assigns the active and standby IP addresses to the state link. This address should be on an unused subnet, different from the failover link.
		
		The standby IP address must be in the same subnet as the active IP address.
		
		Skip this step if you are sharing the state link.

	Command	Purpose

Step 7	**interface** *state_interface_id*
	no shutdown
	
	Example:
	`ciscoasa(config)# interface gigabitethernet 0/4`
	`ciscoasa(config-if)# no shutdown`
		If you specified a separate state link, enables the state link.
		Skip this step if you are sharing the state link.

Step 8	(Optional) Do one of the following to encrypt communications on the failover and state links:	
	failover ipsec pre-shared-key [**0** | **8**] *key* **Example:** `ciscoasa(config)# failover ipsec pre-shared-key a3rynsun`	(Preferred) Establishes IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications. The *key* can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels. If you use a master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password. **Note** The **failover ipsec pre-shared-key** shows as ***** in **show running-config** output; this obscured key is not copyable. If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text. You cannot use both IPsec encryption and the legacy **failover key** encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), you must first remove the failover key using the **no failover key** command before you configure IPsec encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.

	Command	Purpose
	failover key [**0** | **8**] {**hex** *key* | *shared_secret*} **Example:** `ciscoasa(config)# failover key johncr1cht0n`	(Optional) Encrypts failover communication on the failover and state links using a *shared_secret*, from 1 to 63 characters, or a 32-character **hex** *key*. For the *shared_secret*, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units. If you use a master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the shared secret or hex key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.
Step 9	**failover** **Example:** `ciscoasa(config)# failover`	Enables failover.
Step 10	**write memory** **Example:** `ciscoasa(config)# write memory`	Saves the system configuration to flash memory.

Examples

The following example configures the failover parameters for the primary unit:

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3
failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
interface gigabitethernet 0/3
  no shutdown
failover link statelink gigabitethernet0/4
failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2
interface gigabitethernet 0/4
  no shutdown
failover ipsec pre-shared-key a3rynsun
failover
```

Configuring the Secondary Unit for Active/Standby Failover

The only configuration required on the secondary unit is for the failover link. The secondary unit requires these commands to communicate initially 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.

Prerequisites

- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

Step 1 Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. See the "Configuring the Primary Unit for Active/Standby Failover" section on page 7-26.

For example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
INFO: Non-failover interface config is cleared on GigabitEthernet0/3 and its
sub-interfaces
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby
172.27.48.2
ciscoasa(config)# interface gigabitethernet 0/3
  no shutdown
ciscoasa(config)# failover link statelink gigabitethernet0/4
INFO: Non-failover interface config is cleared on GigabitEthernet0/4 and its
sub-interfaces
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
ciscoasa(config)# interface gigabitethernet 0/4
  no shutdown
ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
ciscoasa(config)# failover
```

Step 2 After the failover configuration syncs, save the configuration to flash memory:

```
ciscoasa(config)# write memory
```

Configuring Active/Active Failover

Configuring the Primary Unit for Active/Active Failover

Follow the steps in this section to configure the primary unit in an Active/Active failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit.

Prerequisites

- Enable multiple context mode according to the "Enabling or Disabling Multiple Context Mode" section on page 6-16.
- Configure standby IP addresses for all interfaces except for the failover and state links according to Chapter 11, "Completing Interface Configuration (Routed Mode)," or Chapter 12, "Completing Interface Configuration (Transparent Mode)."
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	**failover lan unit primary**	Designates this unit as the primary unit.
Step 2	**failover lan interface** *if_name interface_id* **Example:** `ciscoasa(config)# failover lan interface folink gigabitethernet0/3`	Specifies the interface to be used as the failover link. This interface cannot be used for any other purpose (except, optionally, the state link). The *if_name* argument assigns a name to the interface. The *interface_id* argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the *interface_id* specifies a VLAN ID. **Note** Although you can use an EtherChannel as a failover or state link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Step 3	**failover interface ip** *if_name* {*ip_address mask* | *ipv6_address/prefix*} **standby** *ip_address*
	
	Example:
	`ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2`
	
	Or:
	
	`ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71`
		Assigns the active and standby IP addresses to the failover link. This address should be on an unused subnet.
		
		The standby IP address must be in the same subnet as the active IP address.

	Command	Purpose

Step 4	**interface** *failover_interface_id*
	no shutdown
	
	Example:
	`ciscoasa(config)# interface gigabitethernet 0/3`
	`ciscoasa(config-if)# no shutdown`
		Enables the failover link.

Step 5	**failover link** *if_name interface_id* **Example:** `ciscoasa(config)# failover link statelink gigabitethernet0/4`	(Optional) Specifies the interface you want to use as the state link. We recommend specifying a separate interface from the failover link or data interfaces. The *if_name* argument assigns a name to the interface. The *interface_id* argument can be a physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASA 5505 or ASASM, the *interface_id* specifies a VLAN ID. **Note** Although you can use an EtherChannel as a failover or state link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Step 6	**failover interface ip** *state_if_name* {*ip_address mask* | *ipv6_address*/*prefix*} **standby** *ip_address*
	
	Example:
	`ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2`
	Or:
	`ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby 2001:a0a:b00:a::a0a:b71`
		If you specified a separate state link, assigns the active and standby IP addresses to the state link. This address should be on an unused subnet, different from the failover link.
		The standby IP address must be in the same subnet as the active IP address.
		Skip this step if you are sharing the state link.

Step 7	**interface** *state_interface_id*
	no shutdown
	
	Example:
	`ciscoasa(config)# interface gigabitethernet 0/4`
	`ciscoasa(config-if)# no shutdown`
		If you specified a separate state link, enables the state link.
		Skip this step if you are sharing the state link.

Step 8	(Optional) Do one of the following to encrypt communications on the failover and state links:	

Command	Purpose
failover ipsec pre-shared-key [**0** | **8**] *key* **Example:** `ciscoasa(config)# failover ipsec pre-shared-key a3rynsun`	(Preferred) Establishes IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications. The *key* can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels. If you use a master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password. **Note** The **failover ipsec pre-shared-key** shows as ***** in **show running-config** output; this obscured key is not copyable. If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text. You cannot use both IPsec encryption and the legacy **failover key** encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), you must first remove the failover key using the **no failover key** command before you configure IPsec encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.
failover key [**0** | **8**] {**hex** *key* | *shared_secret*} **Example:** `ciscoasa(config)# failover key johncr1cht0n`	(Optional) Encrypts failover communication on the failover and state links using a *shared_secret*, from 1 to 63 characters, or a 32-character **hex** *key*. For the *shared_secret*, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units. If you use a master passphrase (see the "Configuring the Master Passphrase" section on page 13-8), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the shared secret or hex key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

	Command	Purpose
Step 9	**failover group 1** **Example:** `ciscoasa(config)# failover group 1`	Creates failover group 1. By default, this group is assigned to the primary unit. Typically, you assign group 1 to the primary unit, and group 2 to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired using the **primary** or **secondary** subcommands.

Step 10	**failover group 2**
	secondary
	
	Example:
	`ciscoasa(config)# failover group 2`
	`ciscoasa(config-fover-group)# secondary`
		Creates failover group 2 and assigns it to the secondary unit.

Step 11	**context** *name*
	join-failover-group {**1** | **2**}
	
	Example:
	`ciscoasa(config)# context Eng`
	`ciscoasa(config-ctx)# join-failover-group 2`
		Enters the context configuration mode for a given context, and assigns the context to a failover group. Repeat this command for each context.
		
		Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1; you cannot assign it to group 2.

Step 12	**failover** **Example:** `ciscoasa(config)# failover`	Enables failover.
Step 13	**write memory** **Example:** `ciscoasa(config)# write memory`	Saves the system configuration to flash memory.

Examples

The following example configures the failover parameters for the primary unit:

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3
failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
interface gigabitethernet 0/3
  no shutdown
failover link statelink gigabitethernet0/4
failover interface ip statelink 172.27.49.1 255.255.255.0 standby 172.27.49.2
interface gigabitethernet 0/4
  no shutdown
failover group 1
failover group 2
  secondary
context admin
  join-failover-group 1
failover ipsec pre-shared-key a3rynsun
failover
```

Configuring the Secondary Unit for Active/Active Failover

The only configuration required on the secondary unit is for the failover link. The secondary unit requires these commands to communicate initially 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.

Prerequisites

- Enable multiple context mode according to the "Enabling or Disabling Multiple Context Mode" section on page 6-16.
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

Step 1 Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. You also do not need to enter the **failover group** and **join-failover-group** commands, as they are replicated from the primary unit. See the "Configuring the Primary Unit for Active/Active Failover" section on page 7-31.

For example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
INFO: Non-failover interface config is cleared on GigabitEthernet0/3 and its
sub-interfaces
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby
172.27.48.2
ciscoasa(config)# interface gigabitethernet 0/3
  no shutdown
ciscoasa(config)# failover link statelink gigabitethernet0/4
INFO: Non-failover interface config is cleared on GigabitEthernet0/4 and its
sub-interfaces
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
ciscoasa(config)# interface gigabitethernet 0/4
  no shutdown
ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
ciscoasa(config)# failover
```

Step 2 After the failover configuration syncs from the primary unit, save the configuration to flash memory:

```
ciscoasa(config)# write memory
```

Step 3 If necessary, force failover group 2 to be active on the secondary unit:

```
failover active group 2
```

Configuring Optional Failover Parameters

You can customize failover settings as desired.

Configuring Failover Criteria, HTTP Replication, Group Preemption, and MAC Addresses

See the "Default Settings" section on page 7-25 for the default settings for many parameters that you can change in this section. For Active/Active mode, you set most criteria per failover group.

Prerequisites

Configure these settings in the system execution space in multiple context mode.

Detailed Steps

	Command	Purpose
Step 1	**failover polltime** [**unit**] [**msec**] *poll_time* [**holdtime** [**msec**] *time*] **Example:** `ciscoasa(config)# failover polltime unit msec 200 holdtime msec 800`	Changes the unit poll and hold times. In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group. You cannot enter a holdtime value that is less than 3 times the unit poll time. With a faster poll time, the ASA can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested. If a unit does not hear hello packet on the failover communication interface for one polling period, additional testing occurs through the remaining interfaces. If there is still no response from the peer unit during the hold time, the unit is considered failed and, if the failed unit is the active unit, the standby unit takes over as the active unit.
Step 2	**failover replication rate** *conns* **Example:** `ciscoasa(config)# failover replication rate 20000`	Sets the HTTP replication rate in connections per second, between 8341 and 50000. The default is 50000. In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.

Step 3	(Active/Active mode only)
	failover group {**1** | **2**}
	
	Example:
	`ciscoasa(config)# failover group 1`
	`ciscoasa(config-fover-group)#`
		Specifies the failover group you want to customize.

	Command	Purpose
Step 4	(Active/Active Mode Only) **preempt** [*delay*] **Example:** `ciscoasa(config-fover-group)# preempt 1200`	Configures failover group preemption for failover group 1. If one unit boots before the other, then both failover groups become active on that unit, despite the primary or secondary setting. This command causes the failover group to become active on the designated unit automatically when that unit becomes available. 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. Valid values are from 1 to 1200. **Note** If Stateful Failover is enabled, the preemption is delayed until the connections are replicated from the unit on which the failover group is currently active.

Step 5	For Active/Standby mode:
	failover replication http
	For Active/Active mode:
	replication http
	
	Example:
	`ciscoasa(config)# failover replication http`
	Or
	`ciscoasa(config-fover-group)# replication http`
		Enables HTTP state replication. 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.

Step 6	For Active/Standby mode:
	failover interface-policy *num*[%]
	For Active/Active mode:
	interface-policy *num*[%]
	
	Example:
	`ciscoasa (config)# failover interface-policy 20%`
	Or
	`ciscoasa(config-fover-group)# interface-policy 20%`
		Sets the threshold for failover when interfaces fail. By default, one interface failure causes failover.
		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.

	Command	Purpose

Step 7	For Active/Standby mode:
	failover polltime interface [**msec**] *time* [**holdtime** *time*]
	For Active/Active mode:
	polltime interface [**msec**] *time* [**holdtime** *time*]
	Example:
	`ciscoasa(config)# failover polltime interface msec 500 holdtime 5`
	Or
	`ciscoasa(config-fover-group)# polltime interface msec 500 holdtime 5`
		Changes the interface poll and hold times.
		Valid values for 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.
		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 interfaces meets or exceeds the configured failover criteria.

Step 8	For Active/Standby mode:
	failover mac address *phy_if active_mac standby_mac*
	For Active/Active mode:
	mac address *phy_if active_mac standby_mac*
	Example:
	`ciscoasa(config)# failover mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8`
	Or
	`ciscoasa(config-fover-group)# mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8`
		Configures the virtual MAC address for an interface.
		The *phy_if* argument is the physical name of the interface, such as gigabitethernet0/1.
		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.
		You can also set the MAC address using other commands or methods, but we recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.
		Use the **show interface** command to display the MAC address used by an interface.

Step 9	(Active/Active mode only) Repeat this procedure for the other failover group, if desired.	

Configuring Interface Monitoring

By default, monitoring is enabled on all physical interfaces, or for the ASA 5505 and ASASM, all VLAN interfaces. You might want to exclude interfaces attached to less critical networks from affecting your failover policy.

Guidelines

- You can monitor up to 250 interfaces on a unit (across all contexts in multiple context mode).
- In multiple context mode, configure interfaces within each context.

Detailed Steps

[**no**] **monitor-interface** *if_name*
 	
Example:
`ciscoasa(config)# monitor-interface inside`
`ciscoasa(config)# no monitor-interface eng1`
	Enables or disables health monitoring for an interface.

Configuring Support for Asymmetrically Routed Packets (Active/Active Mode)

When running in Active/Active failover, a unit may receive a return packet for a connection that originated through its peer unit. Because the ASA that receives the packet does not have any connection information for the packet, the packet is dropped. This drop most commonly occurs when the two ASAs 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 by allowing asymmetrically routed packets. To do so, you assign the similar interfaces on each ASA to the same ASR group. For example, both ASAs connect to the inside network on the inside interface, but connect to separate ISPs on the outside interface. On the primary unit, assign the active context outside interface to ASR group 1; on the secondary unit, assign the active context outside interface to the same ASR group 1. When the primary unit outside interface receives a packet for which it has no session information, it checks the session information for the other interfaces in standby contexts that are in the same group; in this case, ASR group 1. 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 This feature does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

Figure 7-13 shows an example of an asymmetrically routed packet.

Figure 7-13 ASR Example

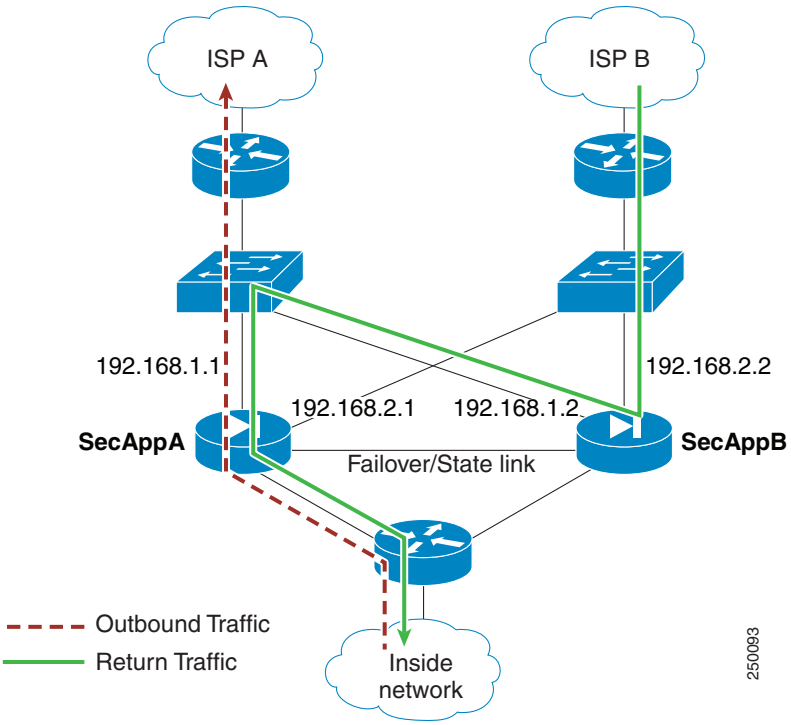

1. An outbound session passes through the ASA with the active SecAppA context. It exits interface outsideISP-A (192.168.1.1).
2. Because of asymmetric routing configured somewhere upstream, the return traffic comes back through the interface outsideISP-B (192.168.2.2) on the ASAwith the active SecAppB context.
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 configured as part of 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 with SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is rewritten 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.

Prerequisites

- 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 ASA to be able to re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.
- Perform this procedure within each active context on the primary and secondary units.

Detailed Steps

	Command	Purpose
Step 1	On the primary unit: **interface** *phy_if* **Example:** `primary/admin(config)# interface gigabitethernet 0/0`	Specifies the interface on the primary unit for which you want to allow asymmetrically routed packets.
Step 2	**asr-group** *num* **Example:** `primary/admin(config-ifc)# asr-group 1`	Sets the ASR group number for the interface. Valid values for *num* range from 1 to 32.
Step 3	On the secondary unit: **interface** *phy_if* **Example:** `secondary/ctx1(config)# interface gigabitethernet 0/1`	Specifies the similar interface on the secondary unit for which you want to allow asymmetrically routed packets.
Step 4	**asr-group** *num* **Example:** `secondary/ctx1(config-ifc)# asr-group 1`	Sets the ASR group number for the interface to match the primary unit interface.

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 7-1 Primary Unit System Configuration

```
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 SecAppA
context admin
```

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

Example 7-2 SecAppA Context Configuration

```
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 7-3 SecAppB Context Configuration

```
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
```

Managing Failover

Forcing Failover

To force the standby unit to become active, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Command	Purpose

For Active/Standby mode on the standby unit:
failover active
For Active/Active mode on the standby unit:
failover active [**group** *group_id*]
 	
Example:
`standby# failover active`
Or:
`standby# failover active group 1`
	Forces a failover when entered on the *standby* unit. The standby unit becomes the active unit.
	If you specify the **group** *group_id*, then this command forces a failover when entered on the *standby* unit for the specified Active/Active failover group. The standby unit becomes the active unit for the failover group.

For Active/Standby mode on the active unit:
no failover active
For Active/Active mode on the active unit:
no failover active [**group** *group_id*]
 	
Example:
`active# no failover active`
Or:
`active# no failover active group 1`
	Forces a failover when entered on the *active* unit. The active unit becomes the standby unit.
	If you specify the **group** *group_id*, then this command forces a failover when entered on the *active* unit for the specified failover group. The active unit becomes the standby unit for the failover group.

Disabling Failover

To disable failover, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Command	Purpose
no failover **Example:** `ciscoasa(config)# no failover`	Disables failover. Disabling failover on an Active/Standby pair causes the active and standby state of each unit to be maintained until you reload. 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 7-42. Disabling failover on an Active/Active failover pair causes the failover groups to remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer.

Restoring a Failed Unit

To restore a failed unit to an unfailed state, perform the following procedure.

Prerequisites

In multiple context mode, perform this procedure in the System execution space.

Detailed Steps

Command	Purpose

For Active/Standby mode:
failover reset
For Active/Active mode:
failover reset [**group** *group_id*]
Example:
`ciscoasa(config)# failover reset`
Or:
`ciscoasa(config)# failover reset group 1`
	Restores a failed unit to an unfailed state. Restoring a failed unit to an unfailed state does not automatically make it active; restored units remain in the standby state until made active by failover (forced or natural). An exception is a failover group (Active/Active mode only) configured with failover preemption. If previously active, a failover group becomes active if it is configured with preemption and if the unit on which it failed is the preferred unit.
	If you specify the **group** *group_id*, this command restores a failed Active/Active failover group to an unfailed state.

Re-Syncing 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.

Testing the Failover Functionality

To test failover functionality, perform the following procedure.

Detailed Steps

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

Step 2 Force a failover by entering the following command on the active unit:

Active/Standby mode:

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

Active/Active mode:

```
ciscoasa(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 to active status by enter the following command on the newly active unit:

Active/Standby mode:

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

Active/Active mode:

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

Note When an ASA interface goes down, for failover it is still considered to be a unit issue. If the ASA detects that an interface is down, failover occurs immediately, without waiting for the interface holdtime. The interface holdtime is only useful when the ASA considers its status to be OK, although it is not receiving hello packets from the peer. To simulate interface holdtime, shut down the VLAN on the switch to prevent peers from receiving hello packets from each other.

Remote Command Execution

Remote command execution lets you send commands entered at the command line to a specific failover peer.

Sending a Command

Because configuration commands are replicated from the active unit or context to the standby unit or context, you can use the **failover exec** command to enter configuration commands on the correct unit, no matter which unit you are logged in to. For example, if you are logged in to the standby unit, you can use the **failover exec active** command to send configuration changes to the active unit. Those changes are then replicated to the standby unit. Do not use the **failover exec** command to send configuration commands to the standby unit or context; those configuration changes are not replicated to the active unit and the two configurations will no longer be synchronized.

Output from configuration, exec, and **show** commands is displayed in the current terminal session, so you can use the **failover exec** command to issue **show** commands on a peer unit and view the results in the current terminal.

You must have sufficient privileges to execute a command on the local unit to execute the command on the peer unit.

Detailed Steps

Step 1 If you are in multiple context mode, use the **changeto context** *name* command to change to the context you want to configure. You cannot change contexts on the failover peer with the **failover exec** command.

Step 2 Use the following command to send commands to he specified failover unit:

```
ciscoasa(config)# failover exec {active | mate | standby}
```

Use the **active** or **standby** keyword to cause the command to be executed on the specified unit, even if that unit is the current unit. Use the **mate** keyword to cause the command to be executed on the failover peer.

Commands that cause a command mode change do not change the prompt for the current session. You must use the **show failover exec** command to display the command mode the command is executed in. See "Changing Command Modes" for more information.

Changing Command Modes

The **failover exec** command maintains a command mode state that is separate from the command mode of your terminal session. By default, the **failover exec** command mode starts in global configuration mode for the specified device. You can change that command mode by sending the appropriate command (such as the **interface** command) using the **failover exec** command. The session prompt does not change when you change modes using **failover exec**.

For example, if you are logged in to global configuration mode of the active unit of a failover pair, and you use the **failover exec active** command to change to interface configuration mode, the terminal prompt remains in global configuration mode, but commands entered using **failover exec** are entered in interface configuration mode.

The following examples show the difference between the terminal session mode and the **failover exec** command mode. In the example, the administrator changes the **failover exec** mode on the active unit to interface configuration mode for the interface GigabitEthernet0/1. After that, all commands entered using **failover exec active** are sent to interface configuration mode for interface GigabitEthernet0/1. The administrator then uses failover exec active to assign an IP address to that interface. Although the prompt indicates global configuration mode, the **failover exec active** mode is in interface configuration mode.

```
ciscoasa(config)# failover exec active interface GigabitEthernet0/1
ciscoasa(config)# failover exec active ip address 192.168.1.1 255.255.255.0 standby
192.168.1.2
ciscoasa(config)# router rip
ciscoasa(config-router)#
```

Changing commands modes for your current session to the device does not affect the command mode used by the **failover exec** command. For example, if you are in interface configuration mode on the active unit, and you have not changed the **failover exec** command mode, the following command would be executed in global configuration mode. The result would be that your session to the device remains in interface configuration mode, while commands entered using **failover exec active** are sent to router configuration mode for the specified routing process.

```
ciscoasa(config-if)# failover exec active router ospf 100
ciscoasa(config-if)#
```

Use the **show failover exec** command to display the command mode on the specified device in which commands sent with the **failover exec** command are executed. The **show failover exec** command takes the same keywords as the **failover exec** command: **active**, **mate**, or **standby**. The **failover exec** mode for each device is tracked separately.

For example, the following is sample output from the **show failover exec** command entered on the standby unit:

```
ciscoasa(config)# failover exec active interface GigabitEthernet0/1
ciscoasa(config)# sh failover exec active
Active unit Failover EXEC is at interface sub-command mode

ciscoasa(config)# sh failover exec standby
Standby unit Failover EXEC is at config mode

ciscoasa(config)# sh failover exec mate
Active unit Failover EXEC is at interface sub-command mode
```

Security Considerations

The **failover exec** command uses the failover link to send commands to and receive the output of the command execution from the peer unit. You should enable encryption on the failover link to prevent eavesdropping or man-in-the-middle attacks.

Limitations of Remote Command Execution

When you use remote commands you face the following limitations:

- If you upgrade one unit using the zero-downtime upgrade procedure and not the other, both units must be running software that supports the **failover exec** command for the command to work.
- Command completion and context help is not available for the commands in the *cmd_string* argument.
- In multiple context mode, you can only send commands to the peer context on the peer unit. To send commands to a different context, you must first change to that context on the unit to which you are logged in.
- You cannot use the following commands with the **failover exec** command:
  - **changeto**
  - **debug** (**undebug**)
- If the standby unit is in the failed state, it can still receive commands from the **failover exe**c command if the failure is due to a service card failure; otherwise, the remote command execution will fail.
- You cannot use the **failover exec** command to switch from privileged EXEC mode to global configuration mode on the failover peer. For example, if the current unit is in privileged EXEC mode, and you enter **failover exec mate configure terminal**, the **show failover exec mate** output will show that the failover exec session is in global configuration mode. However, entering configuration commands for the peer unit using **failover exec** will fail until you enter global configuration mode on the current unit.
- You cannot enter recursive failover exec commands, such as **failover exec mate failover exec mate** command.

- Commands that require user input or confirmation must use the **/nonconfirm** option.

Monitoring Failover

Failover Messages

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

Failover Syslog Messages

The ASA issues a number of syslog messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the syslog messages guide. To enable logging, see Chapter 44, "Configuring Logging."

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

Failover Debug Messages

To see debug messages, enter the **debug fover** command. See the 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 Failover Traps

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 Chapter 45, "Configuring SNMP" for more information.

Monitoring Failover

To monitor failover, enter one of the following commands:

Command	Purpose
`show failover`	Displays information about the failover state of the unit.
`show failover group`	Displays information about the failover state of the failover group. The information displayed is similar to that of the **show failover** command but limited to the specified group.
`show monitor-interface`	Displays information about the monitored interface.
`show running-config failover`	Displays the failover commands in the running configuration.

For more information about the output of the monitoring commands, refer to the command reference.

Feature History for Failover

Table 7-4 lists the release history for this feature.

Table 7-4 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
Active/Standby failover	7.0(1)	This feature was introduced.
Active/Active failover	7.0(1)	This feature was introduced.
Support for a hex value for the failover key	7.0(4)	You can now specify a hex value for failover link encryption. We modified the following command: **failover key hex**.
Support for the master passphrase for the failover key	8.3(1)	The failover key now supports the master passphrase, which encrypts the shared key in the running and startup configuration. If you are copying the shared secret from one ASA to another, for example from the **more system:running-config** command, you can successfully copy and paste the encrypted shared key. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. We modified the following command: **failover key** [**0** | **8**].

Table 7-4 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
IPv6 support for failover added.	8.2(2)	We modified the following commands: **failover interface ip**, **show failover**, **ipv6 address**, **show monitor-interface**.
Support for IPsec LAN-to-LAN tunnels to encrypt failover and state link communications	9.1(2)	Instead of using the proprietary encryption for the failover key (the **failover key** command), you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We introduced or modified the following commands: **failover ipsec pre-shared-key**, **show vpn-sessiondb**.