

Failover for High Availability

This chapter describes how to configure Active/Standby or Active/Active failover to accomplish high availability of the Cisco ASA.

- About Failover, on page 1
- Licensing for Failover, on page 25
- Guidelines for Failover, on page 27
- Defaults for Failover, on page 28
- Configure Active/Standby Failover, on page 29
- Configure Active/Active Failover, on page 30
- Configure Optional Failover Parameters, on page 31
- Manage Failover, on page 37
- Monitoring Failover, on page 41
- History for Failover, on page 43

About Failover

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 whether they meet the specific failover conditions. If those conditions are met, failover occurs.

Failover Modes

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 device functions as the Active unit and passes traffic. The second device, designated as 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. In addition, for container instances, they must use the same resource profile attributes.

  For the Firepower 9300, High Availability is only supported between same-type modules; but the two chassis can include mixed modules. For example, each chassis has an SM-36 and SM-44. You can create High Availability pairs between the SM-36 modules and between the SM-44 modules.

- Have the same number and types of interfaces.

  For the Firepower 2100 and Firepower 4100/9300 chassis, all interfaces must be preconfigured in FXOS identically before you enable Failover. If you change the interfaces after you enable Failover, make the interface changes in FXOS on the Standby unit, and then make the same changes on the Active unit. If you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

- 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 context mode (single or multiple).

- For single mode: Be in the same firewall mode (routed or transparent).

  In multiple context mode, the firewall mode is set at the context-level, and you can use mixed modes.

- 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.

- 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.”

- Be in the same FIPS mode.

- (Firepower 4100/9300) Have the same flow offload mode, either both enabled or both disabled.

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.

Failover and Stateful Failover Links

The failover link and the optional stateful failover link are dedicated connections between the two units. Cisco recommends to use the same interface between two devices in a failover link or a stateful failover link. For example, in a failover link, if you have used eth0 in device 1, use the same interface (eth0) in device 2 as well.



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 an unused data interface (physical, subinterface, 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 also for the state link). For most models, you cannot use a management interface for failover unless explicitly described below.

The ASA does not support sharing interfaces between user data and the failover link. You also cannot use separate subinterfaces on the same parent for the failover link and for data.

See the following guidelines for the failover link:

- 5506-X through 5555-X—You cannot use the Management interface as the failover link; you must use a data interface. The only exception is for the 5506H-X, where you can use the management interface as the failover link.
- 5506H-X—You *can* use the Management 1/1 interface as the failover link. If you configure it for failover, you must reload the device for the change to take effect. In this case, you cannot also use the ASA Firepower module, because it requires the Management interface for management purposes.
- 5585-X—Do not use the Management 0/0 interface, even though it can be used as a data interface. It does not support the necessary performance for this use.
- Firepower 4100/9300—We recommend that you use a 10 GB data interface for the combined failover and state link. You cannot use the management-type interface for the failover link.
- All other models—1 GB interface is large enough for a combined failover and state link.

For a redundant interface used as the failover link, see the following benefits for added redundancy:

- When a failover unit boots up, it alternates between the member interfaces to detect an active unit.
- If a failover unit stops receiving keepalive messages from its peer on one of the member interfaces, it switches to the other member interface.

The alternation frequency is equal to the unit hold time (the **failover polltime unit** command).

Note If you have a large configuration and a low unit hold time, alternating between the member interfaces can prevent the secondary unit from joining/re-joining. In this case, disable one of the member interfaces until after the secondary unit joins.

For an EtherChannel used as the failover 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.

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.

Shared with the Failover Link

Sharing a failover link is the best way to conserve interfaces. However, you must consider a dedicated interface for the state link and failover link, if you have a large configuration and a high traffic network.

Dedicated Interface

You can use a dedicated data interface (physical, redundant, or EtherChannel) for the state link. See Interface for the Failover Link, on page 3 for requirements for a dedicated state link, and Connecting the Failover Link, on page 4 for information about connecting the state link as well.

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

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 the following figures are NOT recommended.

Figure 1: Connecting with a Single Switch—Not Recommended

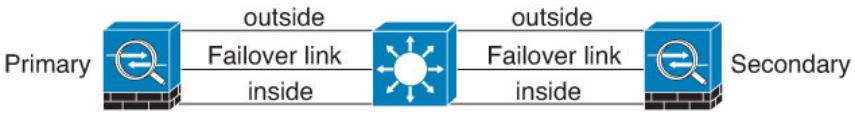
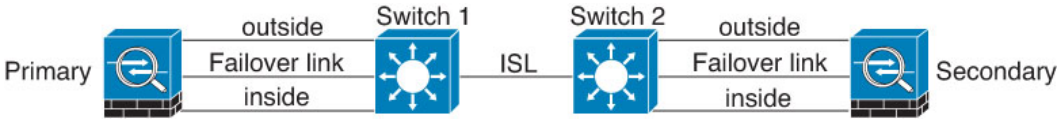


Figure 2: Connecting with a Double-Switch—Not Recommended



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 the following figures.

Figure 3: Connecting with a Different Switch

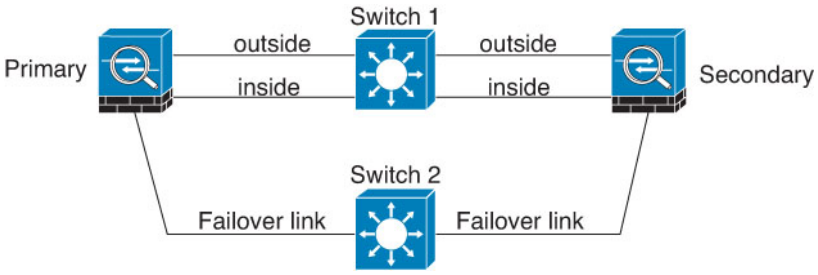
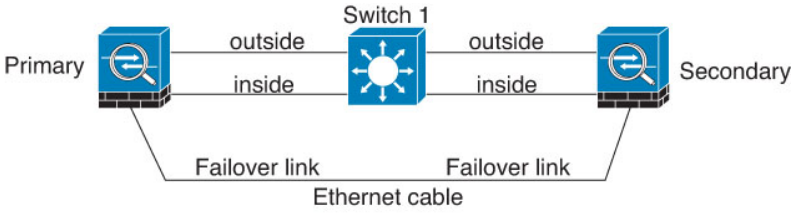


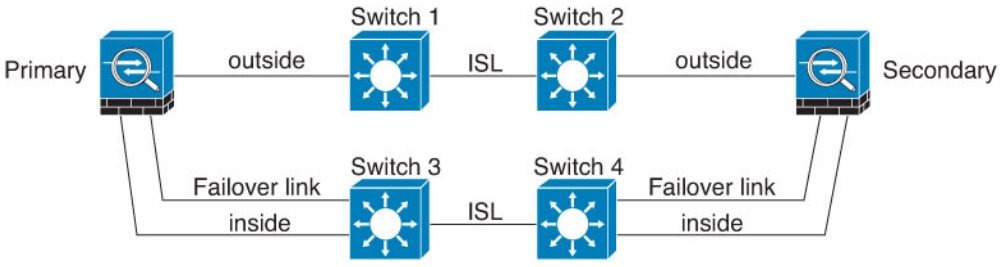
Figure 4: Connecting with a Cable



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 the following figure.

Figure 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 the following figures.

Figure 6: Connecting with Redundant Interfaces

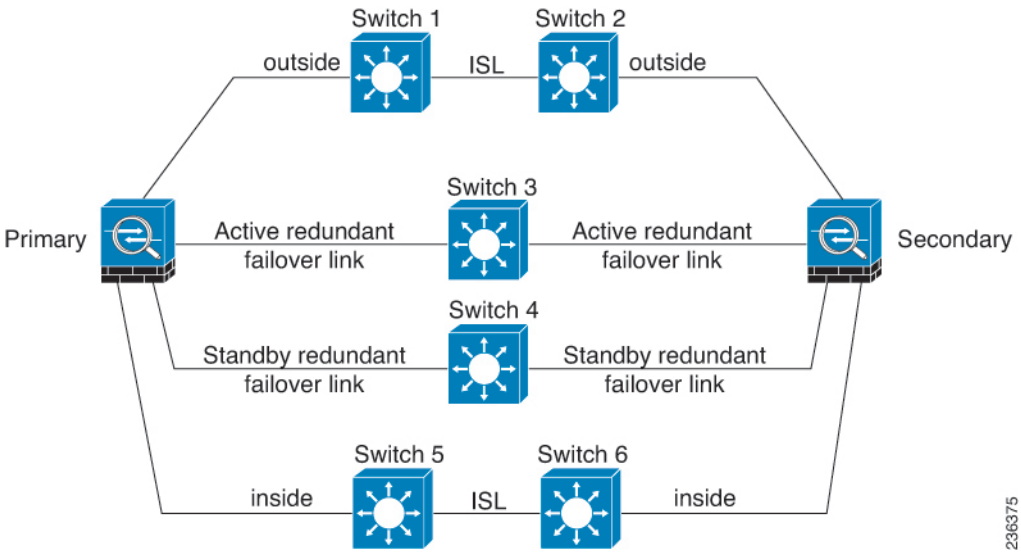
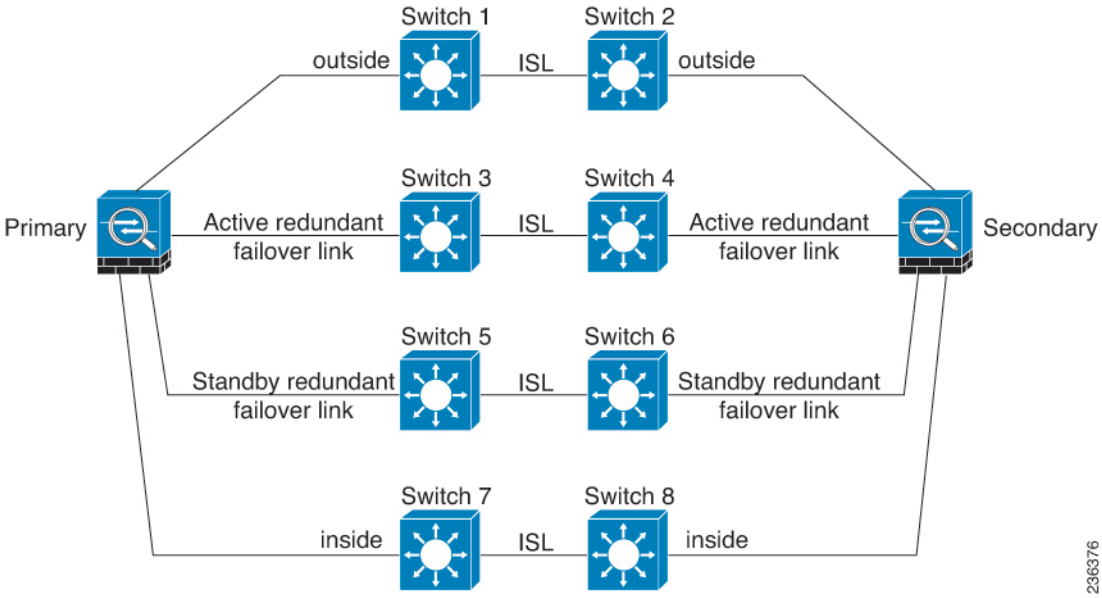


Figure 7: Connecting with Inter-switch Links



MAC Addresses and IP Addresses in Failover

When you configure your interfaces, you can specify an active IP address and a standby IP address on the same network. Generally, when a failover occurs, the new active unit takes over the active 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 Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state. You also cannot connect to the standby unit on that interface for management purposes.

The IP address and MAC address for the state link do not change at failover.

Active/Standby IP Addresses and MAC Addresses

For Active/Standby Failover, see the following for IP address and MAC address usage during a failover event:

1. The active unit always uses the primary unit's IP addresses and MAC addresses.
2. When the active unit fails over, the standby unit assumes the IP addresses and MAC addresses of the failed unit and begins passing traffic.
3. When the failed unit comes back online, it is now in a standby state and takes over the standby IP addresses and MAC addresses.

However, if the secondary unit boots without detecting the primary unit, then the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. 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. 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.

Active/Active IP Addresses and MAC Addresses

For Active/Active failover, see the following for IP address and MAC address usage during a failover event:

1. The primary unit autogenerates active and standby MAC addresses for all interfaces in failover group 1 and 2 contexts. You can also manually configure the MAC addresses if necessary, for example, if there are MAC address conflicts.
2. Each unit uses the active IP addresses and MAC addresses for its active failover group, and the standby addresses for its standby failover group. For example, the primary unit is active for failover group 1, so it uses the active addresses for contexts in failover group 1. It is standby for the contexts in failover group 2, where it uses the standby addresses.
3. When a unit fails over, the other unit assumes the active IP addresses and MAC addresses of the failed failover group and begins passing traffic.
4. When the failed unit comes back online, and you enabled the preempt option, it resumes the failover group.

Virtual MAC Addresses

The ASA has multiple methods to configure virtual MAC addresses. 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. Manual methods include the interface mode **mac-address** command, the **failover mac address** command, and for Active/Active failover, the failover group mode **mac address** command, in addition to autogeneration methods described below.

In multiple context mode, you can configure the ASA to generate virtual active and standby MAC addresses automatically for shared interfaces, and these assignments are synced to the secondary unit (see the **mac-address auto** command). For non-shared interfaces, you can manually set the MAC addresses for Active/Standby mode (Active/Active mode autogenerates MAC addresses for all interfaces).

For Active/Active failover, virtual MAC addresses are always used, either with default values or with values you can set per 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.

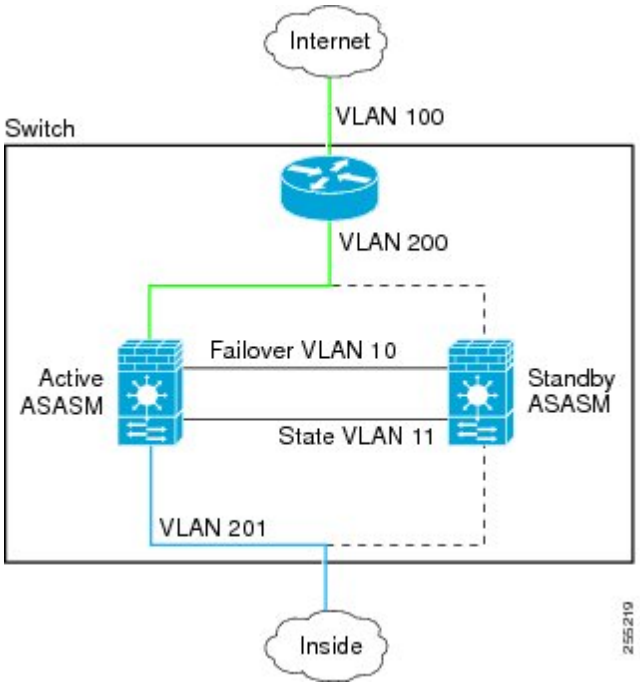
Intra-Chassis Failover

If you install the secondary ASASM in the same switch as the primary ASASM, you protect against module-level failure.

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.

The following figure shows a typical intra-switch configuration.

Figure 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.

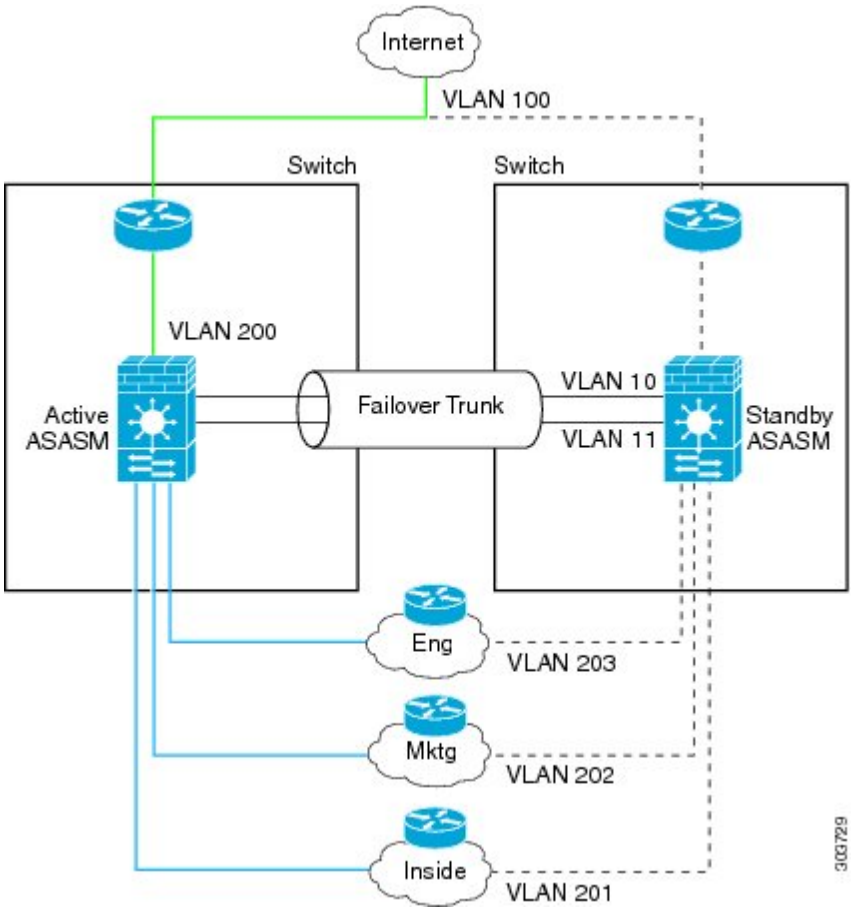
For the best reliability of failover communications between ASASMs, we recommend that you configure an EtherChannel trunk port between the two switches to carry the failover and state VLANs.

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.

The following figure shows a typical switch and ASASM redundancy configuration. The trunk between the two switches carries the failover ASASM VLANs (VLANs 10 and 11).

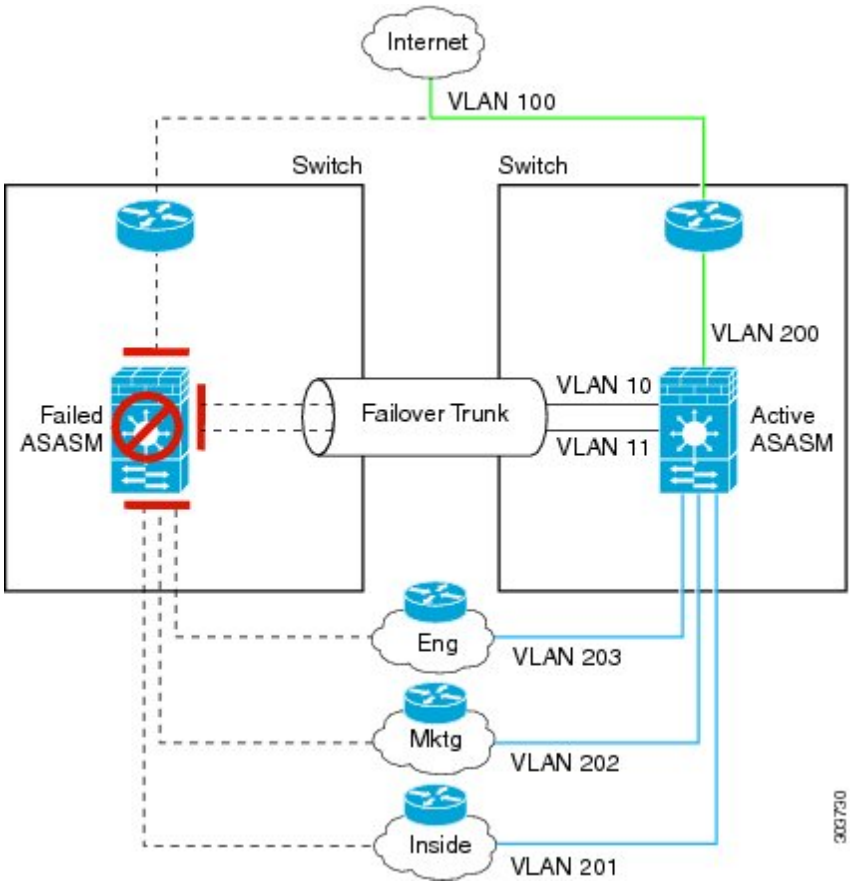
Note ASASM failover is independent of the switch failover operation; however, ASASM works in any switch failover scenario.

Figure 9: Normal Operation



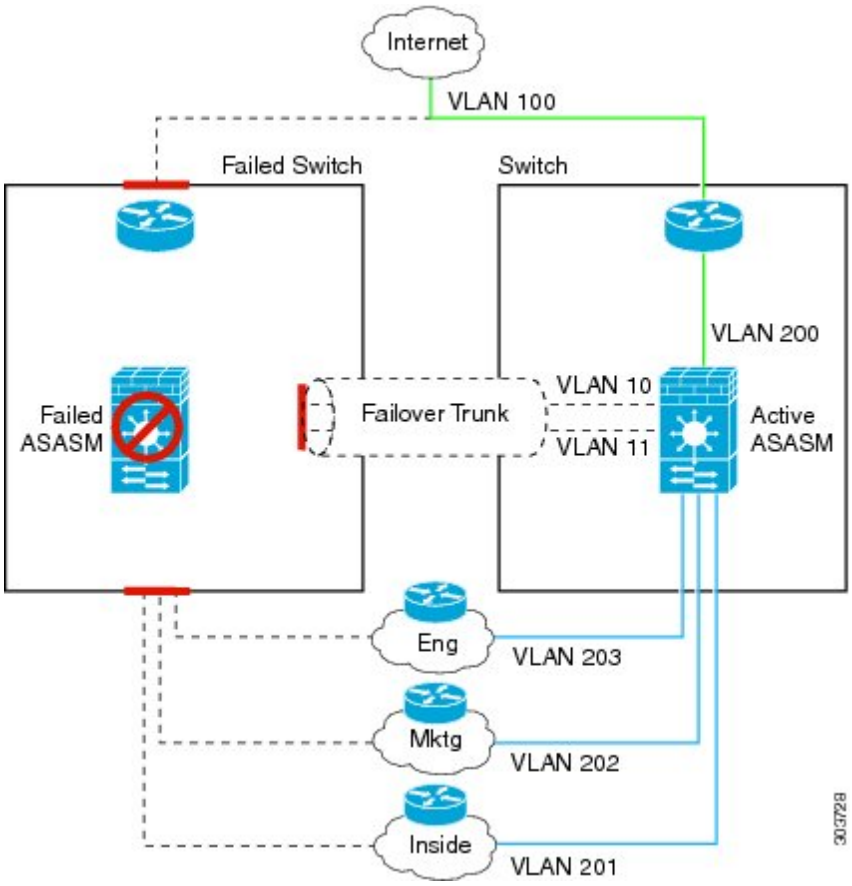
If the primary ASASM fails, then the secondary ASASM becomes active and successfully passes the firewall VLANs.

Figure 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 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.

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

For Stateful Failover, the following state information is passed to the standby ASA:

- NAT translation table.
- TCP and UDP connections and states. Other types of IP protocols, and ICMP, are not parsed by the active unit, because they get established on the new active unit when a new packet arrives.
- The HTTP connection table (unless you enable HTTP replication).
- The HTTP connection states (if HTTP replication is enabled)—By default, the ASA does not replicate HTTP session information when Stateful Failover is enabled. We suggest that you enable HTTP replication.
- SCTP connection states. However, SCTP inspection stateful failover is best effort. During failover, if any SACK packets are lost, the new active unit will drop all other out of order packets in the queue until the missing packet is received.
- The ARP table
- The Layer 2 bridge table (for bridge groups)
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signaling sessions and pin holes.
- ICMP connection state—ICMP connection replication is enabled only if the respective interface is assigned to an asymmetric routing group.
- Static and dynamic routing tables—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 unit initially has rules that mirror the primary unit. 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.

- DHCP Server—DHCP address leases are not replicated. However, a DHCP server configured on an interface will send a ping to make sure an address is not being used before granting the address to a DHCP client, so there is no impact to the service. State information is not relevant for DHCP relay or DDNS.
- 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.
- RA VPN—Remote access 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

For Stateful Failover, the following state information is not passed to the standby ASA:

- The user authentication (uauth) table
- TCP state bypass connections
- Multicast routing.
- State information for modules, such as the ASA FirePOWER module.
- 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)

Bridge Group Requirements for Failover

There are special considerations for failover when using bridge groups.

Bridge Group Requirements for Appliances, ASAv

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 a bridge group's member interfaces with an EtherType access rule.

  ```
  access-list id ethertype deny bpdu
  access-group id in interface name1
  access-group id in interface name2
  ```

  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:

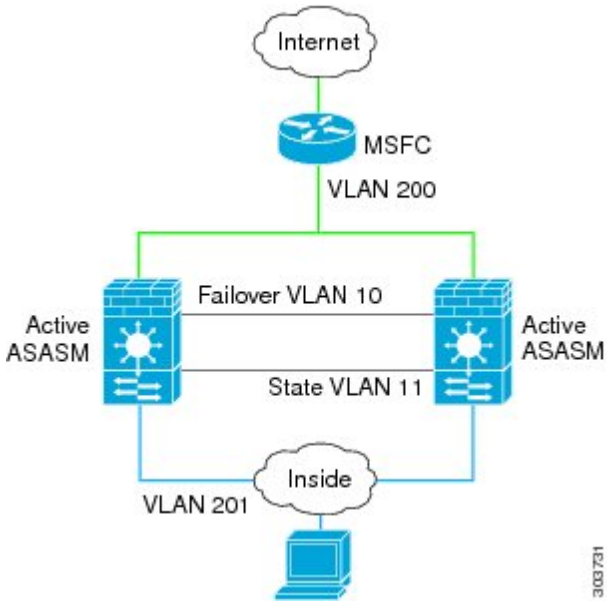
- Disable interface monitoring.
- Increase interface holdtime to a high value that will allow STP to converge before the ASAs fail over.
- Decrease STP timers to allow STP to converge faster than the interface holdtime.

Bridge Group Requirements for the ASA Services Module

To avoid loops when you use failover with bridge groups, 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 packets between bridge group member interfaces get endlessly replicated by both ASASMs. 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 12: Bridge Group 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

The ASA determines the health of the other unit by monitoring the failover link with hello messages. When a unit does not receive three consecutive hello messages on the failover link, the unit sends LANTEST messages on each data interface, including the failover link, to validate whether or not the peer is responsive. For the Firepower 9300 and 4100 series, you can enable Bidirectional Forwarding Detection (BFD) monitoring, which is more reliable than hello messages. 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 1025 interfaces (in multiple context mode, divided between all contexts). You should monitor important interfaces. For example in multiple context 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 15 seconds (the default), it runs interface tests. (To change the period, see **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria** > **Failover Poll Times**.) If one of the interface tests fails for an interface, but this same interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed, and the ASA stops running tests.

If the threshold you define for the number of failed interfaces is met (see **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria** > **Interface Policy**), and the active unit has more failed interfaces than the standby unit, then a failover occurs. If an interface fails on both units, then both interfaces go into the "Unknown" state and do not count towards the failover limit defined by failover interface policy.

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.

If you have an ASA FirePOWER module, then the ASA also monitors the health of the module over the backplane interface. Failure of the module is considered a unit failure and will trigger failover. This setting is configurable.

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).



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.

Interface Tests

The ASA uses the following interface tests. The duration of each test is approximately 1.5 seconds by default, or 1/16 of the failover interface holdtime(see **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria** > **Failover Poll Times**).

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is down, then the ASA considers it failed, and testing stops. If the status is Up, then the ASA performs the Network Activity test.
2. Network Activity test—A received network activity test. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any eligible packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the ASA starts the ARP test.
3. ARP test—A test for successful ARP replies. Each unit sends a single ARP request for the IP address in the most recent entry in its ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If the unit does not receive an ARP reply, then the ASA sends a single ARP request for the IP address in the *next* entry in the ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the ASA starts the Broadcast Ping test.

4. Broadcast Ping test—A test for successful ping replies. Each unit sends a broadcast ping, and then counts all received packets. If the unit receives any packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then testing starts over again with the ARP test. If both units continue to receive no traffic from the ARP and Broadcast Ping tests, then these tests will continue running in perpetuity.

Interface Status

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.

Failover Times

The following events trigger failover in a Firepower high availability pair:

- More than 50% of the Snort instances on the active unit are down.
- Disk space on the active unit is more than 90% full.
- The **no failover active** command is run on the active unit or the **failover active** command is run on the standby unit.
- The active unit has more failed interfaces than the standby unit.
- Interface failure on the active device exceeds the threshold configured.

  By default, failure of a single interface causes failover. You can change the default value by configuring a threshold for the number of interfaces or a percentage of monitored interfaces that must fail for the failover to occur. If the threshold breaches on the active device, failover occurs. If the threshold breaches on the standby device, the unit moves to **Fail** state.

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

  Table 1:

Command	Purpose
failover interface-policy num [%] `hostname (config)# failover interface-policy 20%`	Changes the default failover criteria. 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.

Note If you manually fail over using the CLI or ASDM, or you reload the ASA, the failover starts immediately and is not subject to the timers listed below.

Table 2: ASA

Failover Condition	Minimum	Default	Maximum
Active unit loses power, hardware goes down, or the software reloads or crashes. When any of these occur, the monitored interfaces or failover link do not receives any hello message.	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 FirePOWER 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 various types of configuration synchronization.

Running Configuration Replication

Running configuration replication occurs when any one or both of the devices in the failover pair boot.

In Active/Standby failover, configurations are always synchronized from the active unit to the standby unit.

In Active/Active failover, whichever unit boots second obtains the running configuration from the unit that boots first, regardless of the primary or secondary designation of the booting unit. After both units are up, commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state.

When the standby/second 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/second 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 unit receiving the configuration, the configuration exists only in running memory. You should save the configuration to flash memory. For example, in Active/Active failover, enter the **write memory all** command in the system execution space on the unit that has failover group 1 in the active state. The command is replicated to the peer unit, which proceeds to write its configuration to flash memory.



Note During replication, commands entered on the unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process.



Note The **crypto ca server** command and related subcommands are not supported with failover; you must remove them using the **no crypto ca server** command.

File Replication

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

  The ASA uses a cached file for the AnyConnect client profile stored in cache:/stc/profiles, and not the file stored in the flash file system. To replicate the AnyConnect client profile to the standby unit, perform one of the following:

  - Enter the **write standby** command on the active unit.
  - Reapply the profile on the active unit.
  - Reload the standby unit.

- Local Certificate Authorities (CAs)
- ASA images
- ASDM images

Command Replication

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

In Active/Active failover,changes 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 changes 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**

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, the standby unit becomes the active unit.



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 becomes 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.

The following table 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 3: Failover Events

Failure Event	Policy	Active Unit Action	Standby Unit 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	Become active Mark failover link as failed	Become active Mark failover link as failed	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.

Failure Event	Policy	Active Unit Action	Standby Unit Action	Notes
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.

About Active/Active Failover

This section describes Active/Active failover.

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. When used with preemption, this preference ensures that the failover group runs on the correct unit after it starts up. Without preemption, both groups run on the first unit to boot up.

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.
  - A failover is manually forced.
  - A preemption for the failover group is configured, which causes the failover group to automatically become active on the preferred unit when the unit becomes available.

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.

The following table 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 4: Failover Events

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.

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, the active failover group does not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless 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 for Failover

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. There are some exceptions to this rule. See the following table for precise licensing requirements for failover.

Model	License Requirement
ASA 5506-X and ASA 5506W-X	• Active/Standby—Security Plus License. • Active/Active—No Support. **Note** Each unit must have the same encryption license.

Model	License Requirement

ASA 5512-X through ASA 5555-X	• ASA 5512-X—Security Plus License.
	• Other models—Base License.
	Note
	• Each unit must have the same encryption license.
	• In multiple context mode, each unit must have the the same AnyConnect Apex license.
	• Each unit must have the same IPS module license. You also need the IPS signature subscription on the IPS side for both units. See the following guidelines:
	  • To buy the IPS signature subscription you need to have the ASA with IPS pre-installed (the part number must include "IPS", for example ASA5525-IPS-K9); you cannot buy the IPS signature subscription for a non-IPS part number ASA.
	  • You need the IPS signature subscription on both units; this subscription is not shared in failover, because it is not an ASA license.
	  • The IPS signature subscription requires a unique IPS module license per unit. Like other ASA licenses, the IPS module license is technically shared in the failover cluster license. However, because of the IPS signature subscription requirements, you must buy a separate IPS module license for each unit in.

ASAv	See Failover Licenses for the ASAv.
Firepower 2100	See Failover Licenses for the Firepower 2100.
Firepower 4100/9300	See Failover Licenses for the ASA on the Firepower 4100/9300 Chassis.
All other models	Base License or Standard License. **Note** • Each unit must have the same encryption license. • In multiple context mode, each unit must have the the same AnyConnect Apex license.

Note A valid permanent key is required; in rare instances, your PAK authentication key can be removed. If your key consists of all 0's, then you need to reinstall a valid authentication key before failover can be enabled.

Guidelines for Failover

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.

Model Support

- ASA 5506W-X—You must disable interface monitoring for the internal GigabitEthernet 1/9 interface. These interfaces will not be able to communicate to perform the default interface monitoring checks, resulting in a switch from active to standby and back again because of expected interface communication failures.
- Firepower 9300—We recommend that you use inter-chassis Failover for the best redundancy.
- The ASAv on public cloud networks such as Microsoft Azure and Amazon Web Services are not supported with regular Failover because Layer 2 connectivity is required. Instead, see Failover for High Availability in the Public Cloud.
- The ASA FirePOWER module does not support failover directly; when the ASA fails over, any existing ASA FirePOWER flows are transferred to the new ASA. The ASA FirePOWER module in the new ASA begins inspecting the traffic from that point forward; old inspection states are not transferred.

  You are responsible for maintaining consistent policies on the ASA FirePOWER modules in the high-availability ASA pair to ensure consistent failover behavior.

  Note Create the failover pair before you configure the ASA FirePOWER modules. If the modules are already configured on both devices, clear the interface configuration on the standby device before creating the failover pair. From the CLI on the standby device, enter the **clear configure interface** command.

ASAv Failover for High Availability

When creating a failover pair with the ASAv, it is necessary to add the data interfaces to each ASAv in the same order. If the exact same interfaces are added to each ASAv, but in different order, errors may be presented at the ASAv Console. Failover functionality may also be affected

Additional Guidelines

- 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 enable the STP PortFast feature on the switch:

interface *interface_id* **spanning-tree portfast**

This workaround applies to switches connected to both routed mode and bridge group interfaces. 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.

- You cannot enable failover if a local CA server is configured. Remove the CA configuration using the **no crypto ca server** command.
- Configuring port security on the switches connected to the 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 1025 interfaces on a unit, across all contexts.
- For Active/Standby Failover and a VPN IPsec tunnel, you cannot monitor both the active and standby units using SNMP over the VPN tunnel. The standby unit does not have an active VPN tunnel, and will drop traffic destined for the NMS. You can instead use SNMPv3 with encryption so the IPsec tunnel is not required.
- 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.
- Immediately after failover, the source address of syslog messages will be the failover interface address for a few seconds.
- When using SNMPv3 with failover, if you replace a failover unit, then SNMPv3 users are not replicated to the new unit. You must re-add the SNMPv3 users to the active unit to force the users to replicate to the new unit; or you can add the users directly on the new unit. Reconfigure each user by entering the **snmp-server user** *username group-name* **v3** command on the active unit or directly to the standby unit with the *priv-password* option and *auth-password* option in their unencrypted forms.

Defaults for Failover

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 disabled in multiple context mode, except for the ASASM, where they are enabled by default.
- Monitoring on all physical interfaces, or for the ASASM, all VLAN interfaces.

Configure Active/Standby Failover

To configure Active/Standby failover, configure basic failover settings on both the primary and secondary units. All other configuration occurs only on the primary unit, and is then synched to the secondary unit.

The **High Availability and Scalability Wizard** guides you through a step-by-step process of creating an Active/Standby failover configuration.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability**. See select wizard guidelines in the following steps.

Step 2 On the **Failover Peer Connectivity and Compatibility** screen, enter the IP address of the peer unit. This address must be an interface that has ASDM access enabled on it.

By default, the peer address is assigned to be the standby address for the ASDM management interface.

Step 3 On the **LAN Link Configuration** screen:

- **Interface**—The interface can be a data physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASASM, the interface is a VLAN ID. For the ASA 5506H-X only, you can specify the Management 1/1 interface as the failover link. If you do so, you must save the configuration, and then reload the device. You then cannot use this interface for failover and also use the ASA Firepower module; the module requires the interface for management, and you can only use it for one function. For the Firepower 4100/9300, you can use any data-type interface.
- **Active IP Address**—This IP address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.
- (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see Modify the Failover Setup, on page 37).

Step 4 On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:

- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.

Step 5 After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.

After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting** *xx* **more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.

Configure Active/Active Failover

This section tells how to configure Active/Active failover.

The **High Availability and Scalability Wizard** guides you through a step-by-step process of creating an Active/Active failover configuration.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability**. See select wizard guidelines in the following steps.

Step 2 In the **Failover Peer Connectivity and Compatibility Check** screen, the peer IP address must be an interface that has ASDM access enabled on it.

By default, the peer address is assigned to be the standby address for the interface to which ASDM is connected.

Step 3 In the **Security Context Configuration** screen, if you converted to multiple context mode as part of the wizard, you will only see the admin context. You can add other contexts after you exit the wizard.

Step 4 On the **LAN Link Configuration** screen:

- **Interface**—The interface can be a data physical interface, subinterface, redundant interface, or EtherChannel interface ID. On the ASASM, the interface ID is a VLAN ID. For the ASA 5506H-X only, you can specify the Management 1/1 interface as the failover link. If you do so, you must save the configuration, and then reload the device. You then cannot use this interface for failover and also use the ASA Firepower module; the module requires the interface for management, and you can only use it for one function. For the Firepower 4100/9300, you can use any data-type interface.
- **Active IP Address**—This IP address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.
- (Optional) **Communications Encryption**—Encrypt communications on the failover link. **Note**: Instead of a Secret Key, we recommend using an IPsec preshared key, which you can configure after you exit the wizard (see Modify the Failover Setup, on page 37).

Step 5 On the **State Link Configuration** screen, if you choose another interface for Stateful Failover:

- **Active IP Address**—This IP address should be on an unused subnet, different from the failover link. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.0.0/16 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.
- **Standby IP Address**—This IP address must be on the same network as the active IP address.

Step 6 After you click **Finish**, the wizard shows the **Waiting for Config Sync** screen.

After the specified time period is over, the wizard sends the failover configuration to the secondary unit, and you see an information screen showing that failover configuration is complete.

- If you do not know if failover is already enabled on the secondary unit, then wait for the specified period.
- If you know failover is already enabled, click **Skip configuring peer**.
- If you know the secondary unit is not yet failover-enabled, click **Stop waiting** *xx* **more seconds**, and the failover bootstrap configuration is sent to the secondary unit immediately.

Configure Optional Failover Parameters

You can customize failover settings as desired.

Configure Failover Criteria and Other Settings

See Defaults for Failover, on page 28 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. This section includes enabling HTTP replication per failover group for Active/Active mode; to configure HTTP replication for Active/Standby mode, see Modify the Failover Setup, on page 37.

Before you begin

- Configure these settings in the system execution space in multiple context mode.
- For Bidirectional Forwarding Detection (BFD) for unit health monitoring, see the following limitations:
  - Firepower 9300 and 4100 only.
  - Active/Standby only.
  - Routed mode only

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover**.

Step 2 Disable the ability to make any configuration changes directly on the standby unit or context: on the **Setup** tab check the **Disable configuration changes on the standby unit** check box.

By default, configurations on the standby unit/context are allowed with a warning message.

Step 3 Under **BFD Health Check**, click **Manage** to define a BFD template to be used for failover health detection. The regular unit monitoring can cause false alarms when CPU usage is high. The BFD method is distributed, so high CPU does not affect its operation.

The **Configuration** > **Device Setup** > **Routing** > **BFD** > **Template** page opens. Click **Add** to create a single hop template; multi-hop is not supported. For the interval settings, you can specify milliseconds; microseconds are not supported. For template details, see Create the BFD Template.

Step 4 Click the **Criteria** tab.

Step 5 Configure the unit poll times:

In the **Failover Poll Times** area:

- **Unit Failover**—The amount of time between hello messages among units. The range is between 1 and 15 seconds or between 200 and 999 milliseconds.
- **Unit Hold Time**—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. The range is between 1and 45 seconds or between 800 and 999 milliseconds. You cannot enter a value that is less than 3 times the polltime.

Note Other settings on this pane apply only to Active/Standby mode. In Active/Active mode, you must configure the rest of the parameters per failover group.

Step 6 (Active/Active mode only) Click the **Active/Active** tab, then choose a failover group and click **Edit**.

Step 7 (Active/Active mode only) Change the preferred role of the failover group when used with preemption: click either **Primary** or **Secondary**.

If you used the wizard, failover group 1 is assigned to the primary unit, and failover group 2 is assigned to the secondary unit. If you want a non-standard configuration, you can specify different unit preferences if desired. These settings are only used in conjunction with the preempt setting. Both failover groups become active on the unit that boots first (even if it seems like they boot simultaneously, one unit becomes active first), despite the primary or secondary setting for the group.

Step 8 (Active/Active mode only) Configure failover group preemption: check the **Preempt after booting with optional delay of** check box.

Both failover groups become active on the unit that boots first (even if it seems like they boot simultaneously, one unit becomes active first), despite the primary or secondary setting for the group.

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.

If you manually fail over, the Peempt option is ignored.

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 9 Configure the **Interface Policy**:

- **Number of failed interfaces that triggers failover**—Define a specific number of interfaces that must fail to trigger failover, from 1 to 250. When the number of failed monitored interfaces exceeds the value you specify, the ASA fails over.

- **Percentage of failed interfaces that triggers failover**—Define a percentage of configured interfaces that must fail to trigger failover. When the number of failed monitored interfaces exceeds the percentage you set, the ASA fails over.

Note Do not use the **Use system failover interface policy** option. You can only set the policy per group at this time.

Step 10 (Active/Standby mode) Configure interface poll times:

In the **Failover Poll Time** area:

- **Monitored Interfaces**—Specifies the interface polltime: how long to wait between sending hello packets to the peer. The range is between 1and 15 seconds or 500 to 999 milliseconds. The default is 5 seconds.
- **Link State**—By default, each ASA in a failover pair checks the link state of its interfaces every 500 msec. You can customize the polltime; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. The range is between 300 and 799 milliseconds.
- **Interface Hold Time**—Sets the time (as a calculation) between the last-received hello message from the peer unit and the commencement of interface tests to determine the health of the interface. It also sets the duration of each interface test as *holdtime*/16. Valid values are from 5 to 75 seconds. The default is 5 times the polltime. You cannot enter a holdtime value that is less than five times the polltime.

  To calculate the time before starting interface tests (y):

  a. x = (*holdtime*/*polltime*)/2, rounded to the nearest integer. (.4 and down rounds down; .5 and up rounds up.)

  b. y = x**polltime*

  For example, if you use the default holdtime of 25 and polltime of 5, then y = 15 seconds.

For Active/Active mode, configure interface poll times on the **Add/Edit Failover Group** dialog box.

Step 11 (Active/Active mode only) Enable HTTP replication: check the **Enable HTTP replication** check box.

See Modify the Failover Setup, on page 37 section for the session replication rate.

Note Because of a delay when deleting HTTP flows from the standby unit when using failover, the **show conn count** output might show different numbers on the active unit vs. the standby unit; if you wait several seconds and re-issue the command, you will see the same count on both units.

Step 12 Configure virtual MAC addresses:

- Active/Standby mode—click the **MAC Addresses** tab, and click **Add**.

  The **Add/Edit Interface MAC Address** dialog box appears.

- Active/Active mode—Go to the bottom of the **Active/Active** tab.

You can also set the MAC address using other 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.

a) Choose an interface from the **Physical Interface** drop-down list.
b) In the **Active MAC Address** field, type the new MAC address for the active interface.
c) In the **Standby MAC Address** field, type the new MAC address for the standby interface.

d) Click **OK**. (Active/Active mode only) Click **OK** again.

Step 13 Click **Apply**.

Configure Interface Monitoring and Standby Addresses

By default, monitoring is enabled on all physical interfaces, or for the ASASM, all VLAN interfaces, and on any hardware or software modules installed on the ASA, such as the ASA FirePOWER module.

You might want to exclude interfaces attached to less critical networks from affecting your failover policy.

You can monitor up to 1025 interfaces on a unit (across all contexts in multiple context mode).

If you did not configure the standby IP addresses in the wizard, you can configure them manually.

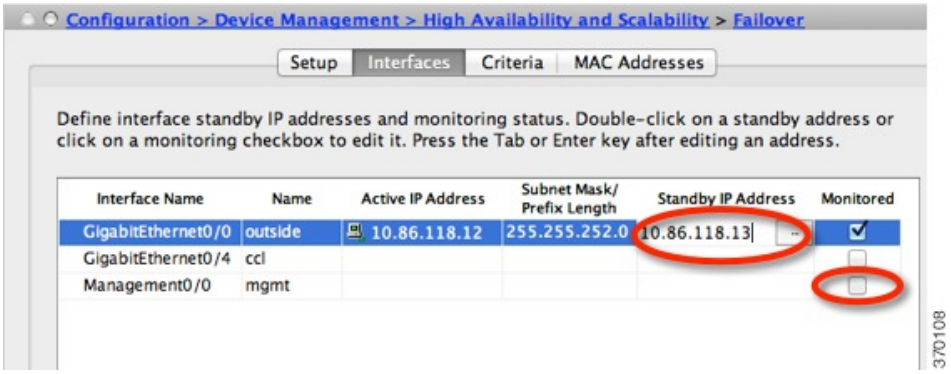
Before you begin

In multiple context mode, configure interfaces within each context.

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability** > **Failover** > **Interfaces**.

In multiple context mode, within a context choose **Configuration > Device Management > Failover > Interfaces**



A list of configured interfaces appears as well as any installed hardware or software modules, such as the ASA FirePOWER module. The **Monitored** column displays whether or not an interface is monitored as part of your failover criteria. If it is monitored, a check appears in the **Monitored** check box.

If you do not want a hardware or software module failure to trigger failover, you can disable module monitoring. Note that for the ASA 5585-X, if you disable monitoring of the service module, you may also want to disable monitoring of the interfaces on the module, which are monitored separately.

The IP address for each interface appears in the **Active IP Address** column. If configured, the standby IP address for the interface appears in the **Standby IP Address** column. The failover link and state link do not display IP address; you cannot change those addresses from this tab.

Step 2 To disable monitoring of a listed interface, uncheck the **Monitored** check box for the interface.

Step 3 To enable monitoring of a listed interface, check the **Monitored** check box for the interface.

Step 4 For each interface that does not have a standby IP address, double-click the **Standby IP Address** field and enter an IP address into the field.

If you use a 31-bit subnet mask for point-to-point connections, do not configure a standby IP address.

Step 5 Click **Apply**.

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

When running in Active/Active failover, a unit might 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:

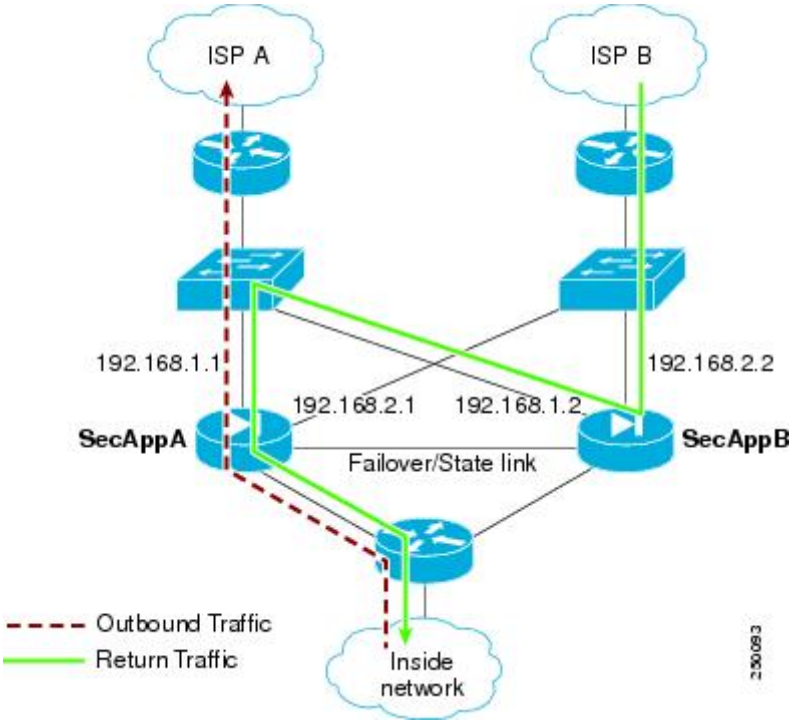
- If the incoming traffic originated on a peer unit, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all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

The following figure shows an example of an asymmetrically routed packet.

Figure 13: ASR Example



1. An outbound session passes through the ASA with the active SecAppA context. It exits interface outside ISP-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 ASA with 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.

Before you begin

- 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.

- You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group.

Procedure

Step 1 On the primary unit active context, choose **Configuration** > **Device Setup** > **Routing** > **ASR Groups**.

Step 2 For the interface that receives asymmetrically routed packets, choose an **ASR Group ID** from the drop-down list.

Step 3 Click **Apply** to save your changes to the running configuration.

Step 4 Connect ASDM to the secondary unit, and choose the active context similar to the primary unit context.

Step 5 Choose **Configuration** > **Device Setup** > **Routing** > **ASR Groups**.

Step 6 For the similar interface on this unit, choose the same **ASR Group ID**.

Step 7 Click **Apply** to save your changes to the running configuration.

Manage Failover

This section describes how to manage Failover units after you enable Failover, including how to change the Failover setup and how to force failover from one unit to another.

Modify the Failover Setup

If you do not use the wizard, or want to change a setting, you can configure the failover setup manually. This section also includes the following options that are not included in the wizard, so you must configure them manually:

- IPsec preshared key for encrypting failover traffic
- HTTP replication rate
- HTTP replication (Active/Standby mode)

Before you begin

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

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.

Step 2 Check the **Enable Failover** check box.

Note Failover is not actually enabled until you apply your changes to the device.

Step 3 To encrypt communications on the failover and state links, use one of the following options:

- **IPsec Preshared Key** (preferred)—The preshared key is used by IKEv2 to establish IPsec LAN-to-LAN tunnels on the failover links between the failover units. Note: failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.
- **Secret Key**—Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

  Use 32 hexadecimal character key—To use a 32-hexadecimal key for the secret key, check this check box.

Step 4 In the **LAN Failover** area, set the following parameters for the failover link:

- **Interface**—Choose the interface to use for the failover link. Failover requires a dedicated interface, however you can share the interface with Stateful Failover.

  Only unconfigured interfaces or subinterfaces are displayed in this list and can be selected as the failover link. Once you specify an interface as the failover link, you cannot edit that interface in the Configuration > Interfaces pane.
- **Logical Name**—Specify the logical name of the interface used for failover communication, such as "failover". This name is informational.
- **Active IP**—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet.
- **Standby IP**—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- **Subnet Mask**—Specify the subnet mask.
- **Preferred Role**—Select **Primary** or **Secondary** to specify whether the preferred role for this ASA is as the primary or secondary unit.

Step 5 (Optional) Configure the state link by doing the following:

- **Interface**—Choose the interface to use for the state link. You can choose an unconfigured interface or subinterface, the failover link, or the **--Use Named--** option.

  Note We recommend that you use two separate, dedicated interfaces for the failover link and the state link.

  If you choose an unconfigured interface or subinterface, you must supply the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** for the interface.

  If you choose the failover link, you do not need to specify the **Active IP**, **Subnet Mask**, **Logical Name**, and **Standby IP** values; the values specified for the failover link are used.

  If you choose the **--Use Named--** option, the Logical Name field becomes a drop-down list of named interfaces. Choose the interface from this list. The **Active IP**, **Subnet Mask/Prefix Length**, and **Standby IP** values do not need to be specified. The values specified for the interface are used.
- **Logical Name**—Specify the logical name of the interface used for state communication, such as "state". This name is informational.

- **Active IP**—Specify the active IP address for the interface. The IP address can be either an IPv4 or an IPv6 address. This IP address should be on an unused subnet, different from the failover link.
- **Standby IP**—Specify the standby IP address for the interface, on the same subnet as the active IP address.
- **Subnet Mask**—Specify the subnet mask.
- (Optional, Active/Standby only) **Enable HTTP Replication**—This option enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover. In Active/Active mode, set the HTTP replication per failover group.

  Note Because of a delay when deleting HTTP flows from the standby unit when using failover, the **show conn count** output might show different numbers on the active unit vs. the standby unit; if you wait several seconds and re-issue the command, you will see the same count on both units.

Step 6 In the **Replication** area, set the session replication rate in connections per second. The minimum and maximum rates are determined by your model. The default is the maximum rate. To use the default, check the **Use Default check** box.

Step 7 Click **Apply**.

The configuration is saved to the device.

Step 8 If you are enabling failover, you see a dialog box to configure the failover peer.

- Click **No** if you want to connect to the failover peer later and configure the matching settings manually.
- Click **Yes** to let ASDM automatically configure the relevant failover settings on the failover peer. Provide the peer IP address in the **Peer IP Address** field.

Force Failover

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

Before you begin

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

Procedure

Step 1 To force failover at the unit level:

a) Choose the screen depending on your context mode:

- In single context mode choose **Monitoring > Properties > Failover > Status**.
- In multiple context mode, in the System choose **Monitoring > Failover > System**.

b) Click one of the following buttons:

- Click **Make Active** to make the unit this unit.
- Click **Make Standby** to make the other unit the active unit.

Step 2 (Active/Active mode only) To force failover at the failover group level:

a) In the System choose **Monitoring > Failover > Failover Group** #, where # is the number of the failover group you want to control.

b) Click one of the following buttons

- Click **Make Active** to make the failover group active on this unit.
- Click **Make Standby** to make the failover group active on the other unit.

Disable Failover

Disabling failover on one or both units causes the active and standby state of each unit to be maintained until you reload. For an Active/Active failover pair, the failover groups remain in the active state on whichever unit they are active, no matter which unit they are configured to prefer.

See the following characteristics when you disable failover:

- The standby unit/context remains in standby mode so that both units do not start passing traffic (this is called a pseudo-standby state).
- The standby unit/context continues to use its standby IP addresses even though it is no longer connected to an active unit/context.
- The standby unit/context continues to listen for a connection on the failover link. If failover is re-enabled on the active unit/context, then the standby unit/context resumes ordinary standby status after re-synchronizing the rest of its configuration.
- Do not enable failover manually on the standby unit to make it active; instead see Force Failover, on page 39. If you enable failover on the standby unit, you will see a MAC address conflict that can disrupt IPv6 traffic.
- To truly disable failover, save the no failover configuration to the startup configuration, and then reload.

Before you begin

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

Procedure

Step 1 In single mode, choose **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**.

In multiple context mode, choose **Configuration > Device Management > Failover > Setup** in the System execution space.

Step 2 Uncheck the **Enable Failover** check box.

Step 3 Click **Apply**.

Step 4 To completely disable failover, save the configuration and reload:

a) Click the **Save** button.

b) Choose **Tools > System Reload** and reload the ASA.

Restore a Failed Unit

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

Before you begin

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

Procedure

Step 1 To restore failover at the unit level:

a) Choose the screen depending on your context mode:

- In single context mode choose **Monitoring > Properties > Failover > Status**.
- In multiple context mode, in the System choose **Monitoring > Failover > System**.

b) Click **Reset Failover**.

Step 2 (Active/Active mode only) To reset failover at the failover group level:

a) In the System choose **Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

b) Click **Reset Failover**.

Re-Sync the Configuration

Replicated commands are stored in the running configuration. To save replicated commands to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**.

Monitoring Failover

This section lets you monitor the Failover status.

Failover Messages

When a failover occurs, both ASAs send out system messages.

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. The ranges of message IDs associated with failover are: 101xxx, 102xxx, 103xxx, 104xxx, 105xxx, 210xxx, 311xxx, 709xxx, 727xxx. For example, 105032 and 105043 indicate a problem with the failover link.



Note During failover, the ASA logically shuts down and then brings 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.

Monitoring Failover Status



Note After a failover event, you should either re-launch ASDM or switch to another device in the Devices pane and then come back to the original ASA to continue monitoring the device. This action is necessary because the monitoring connection is not re-established when ASDM is disconnected from and then reconnected to the device.

Choose **Monitoring > Properties > Failover** to monitor Active/Standby failover.

Use the following screens in the **Monitoring > Properties > Failover** area to monitor Active/Active failover.

System

The System pane displays the failover state of the system. You can also control the failover state of the system by:

- Toggling the active/standby state of the device.
- Resetting a failed device.
- Reloading the standby unit.

Fields

Failover state of the system—*Display only*. Displays the failover state of the ASA. The information shown is the same output you would receive from the **show failover** command. Refer to the command reference for more information about the displayed output.

The following actions are available on the System pane:

- Make Active—Click this button to make the ASA the active unit in an active/standby configuration. In an active/active configuration, clicking this button causes both failover groups to become active on the ASA.
- Make Standby—Click this button to make the ASA the standby unit in an active/standby pair. In an active/active configuration, clicking this button causes both failover groups to go to the standby state on the ASA.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Reload Standby—Click this button to force the standby unit to reload.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

Failover Group 1 and Failover Group 2

The Failover Group 1 and Failover Group 2 panes display the failover state of the selected group.You can also control the failover state of the group by toggling the active/standby state of the group or by resetting a failed group.

Fields

Failover state of Group[*x*]—*Display only*. Displays the failover state of the selected failover group. The information shown is the same as the output you would receive from the **show failover group** command.

You can perform the following actions from this pane:

- Make Active—Click this button to make the failover group active unit on the ASA.
- Make Standby—Click this button to force the failover group into the standby state on the ASA.
- Reset Failover—Click this button to reset a system from the failed state to the standby state. You cannot reset a system to the active state. Clicking this button on the active unit resets the standby unit.
- Refresh—Click this button to refresh the status information in the Failover state of the system field.

History for Failover

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.

Feature Name	Releases	Feature Information
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 screen: Configuration > Device Management > High Availability > Failover > Setup.
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. There were no ASDM changes.
IPv6 support for failover added.	8.2(2)	We modified the following screens: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces
Change to failover group unit preference during "simultaneous" bootup.	9.0(1)	Earlier software versions allowed “simultaneous” boot up so that the failover groups did not require the **preempt** command to become active on the preferred unit. However, this functionality has now changed so that both failover groups become active on the first unit to boot up.
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, 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 modified the following screen: **Configuration > Device Management > High Availability > Failover > Setup**.

Feature Name	Releases	Feature Information
Disable health monitoring of a hardware module	9.3(1)	By default, the ASA monitors the health of an installed hardware module such as the ASA FirePOWER module. If you do not want a hardware module failure to trigger failover, you can disable module monitoring. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Interfaces**.
Lock configuration changes on the standby unit or standby context in a failover pair	9.3(2)	You can now lock configuration changes on the standby unit (Active/Standby failover) or the standby context (Active/Active failover) so you cannot make changes on the standby unit outside normal configuration syncing. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Setup**.
Enable use of the Management 1/1 interface as the failover link on the ASA 5506H	9.5(1)	On the ASA 5506H only, you can now configure the Management 1/1 interface as the failover link. This feature lets you use all other interfaces on the device as data interfaces. Note that if you use this feature, you cannot use the ASA Firepower module, which requires the Management 1/1 interface to remain as a regular management interface. We modified the following screen: **Configuration > Device Management > High Availability and Scalability > Failover > Setup**
Carrier Grade NAT enhancements now supported in failover and ASA clustering	9.5(2)	For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). This feature is now supported in failover and ASA cluster deployments. We did not modify any screens.

Feature Name	Releases	Feature Information
Improved sync time for dynamic ACLs from AnyConnect when using Active/Standby failover	9.6(2)	When you use AnyConnect on a failover pair, then the sync time for the associated dynamic ACLs (dACLs) to the standby unit is now improved. Previously, with large dACLs, the sync time could take hours during which time the standby unit is busy syncing instead of providing high availability backup. We did not modify any screens.
Stateful failover for AnyConnect connections in multiple context mode	9.6(2)	Stateful failover is now supported for AnyConnect connections in multiple context mode. We did not modify any screens.
Interface link state monitoring polling for failover now configurable for faster detection	9.7(1)	By default, each ASA in a failover pair checks the link state of its interfaces every 500 msec. You can now configure the polling interval, between 300 msec and 799 msec; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Criteria**
Bidirectional Forwarding Detection (BFD) support for Active/Standby failover health monitoring on the Firepower 9300 and 4100	9.7(1)	You can enable Bidirectional Forwarding Detection (BFD) for the failover health check between two units of an Active/Standby pair on the Firepower 9300 and 4100. Using BFD for the health check is more reliable than the default health check method and uses less CPU. We modified the following screen: **Configuration** > **Device Management** > **High Availability and Scalability** > **Failover** > **Setup**