

Failover for High Availability

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

- About Failover, on page 1
- Licensing for Failover, on page 25
- Guidelines for Failover, on page 26
- Defaults for Failover, on page 29
- Configure Active/Standby Failover, on page 29
- Configure Active/Active Failover, on page 34
- Configure Optional Failover Parameters, on page 40
- Manage Failover, on page 48
- Monitoring Failover, on page 54
- History for Failover, on page 55

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

Failover system requirements are specifications that define the hardware, software, and license prerequisites for ASAs to function in a Failover configuration.

Hardware requirements

Hardware requirements are specifications that ensure two units in a Failover configuration can

- maintain identical hardware configurations including model, interfaces, modules, and RAM
- support proper synchronization and failover functionality, and
- ensure reliable high availability operation.

Hardware requirement specifications

The two units in a Failover configuration must meet these requirements:

- Be the same model.

  For the Firepower 9300, High Availability is only supported between same-type modules. However, the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.

- Have the same number and types of interfaces.

  For the Firepower 2100 in Platform mode 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

Software requirements are configuration specifications that Failover units must meet to function properly in a failover configuration.

Required software specifications

The two units in a Failover configuration must meet these requirements:

- 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 Secure Client 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

A failover link and stateful failover link are dedicated connections that

- establish communication between the two units in a failover configuration
- require the same interface to be used between corresponding devices for consistency, and
- transmit critical failover and state information between the paired units.

Configuration recommendations

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

A failover link is a communication mechanism that enables two units in a failover pair to constantly communicate to determine the operating status of each unit.

Failover link data

This below 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

A failover link interface is a dedicated communication channel that

- exists solely for failover communication between high availability units
- uses an unused data interface (physical, subinterface, or EtherChannel)
- cannot be shared with normal networking interfaces or user data traffic, and
- requires specific interface sizing based on the device model.

Interface restrictions and requirements

You can use an unused data interface (physical, subinterface, 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 these 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.
- Firepower 4100/9300—You cannot use the management-type interface for the failover link.
- See these guidelines for sizing the link.

Table 1: Failover link size

Model	Interface Size for Combined Failover and State Link
Firepower 1010	1 Gbps
Firepower 1100	1 Gbps
Firepower 2100	1 Gbps
Secure Firewall 3100	Secure Firewall 3105—1 Gbps Secure Firewall 3110—1 Gbps Secure Firewall 3120—1 Gbps Secure Firewall 3130—10 Gbps Secure Firewall 3140—10 Gbps
Firepower 4100	10 Gbps
Secure Firewall 4200	10 Gbps
Firepower 9300	10 Gbps

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.

Connect the failover link

The failover link enables communication and coordination between units in a failover configuration.

You need to establish a physical connection between the failover units to enable proper failover functionality.

Procedure

Connect the failover link using one of the following methods:

- 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 and 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.

The failover link is physically connected and ready for failover configuration.

Stateful failover links

A stateful failover link is a network connection that

- passes connection state information between failover units, and
- must be configured to enable stateful failover functionality.

State link configuration requirements

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

Shared failover links

A shared failover link is an interface conservation method that enables multiple network functions to use the same failover connection, though dedicated interfaces for state and failover links should be considered for large configurations and high traffic networks.

Dedicated interface

A dedicated interface is a data interface configuration that uses a dedicated physical or EtherChannel interface specifically for the state link in failover deployments.

Performance considerations

You can use a dedicated data interface—such as a physical or EtherChannel interface—for the state link. For requirements about a dedicated state link, refer to Interface for the failover link, on page 4. See the section on connecting the state link for additional information Connect the failover link, on page 5.

For optimum performance with long-distance failover, the state link latency should be less than 10 milliseconds but no more than 250 milliseconds. If latency exceeds 10 milliseconds, performance might degrade 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

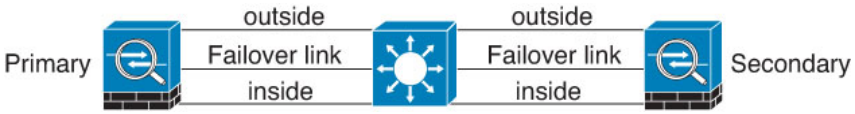
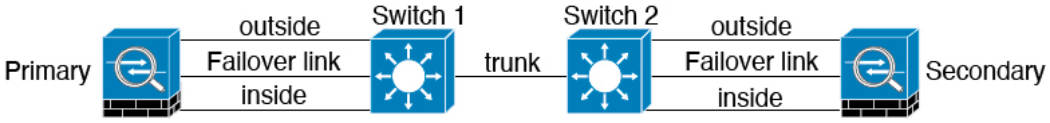


Figure 2: Connecting with a Double-Switch—Not Recommended



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

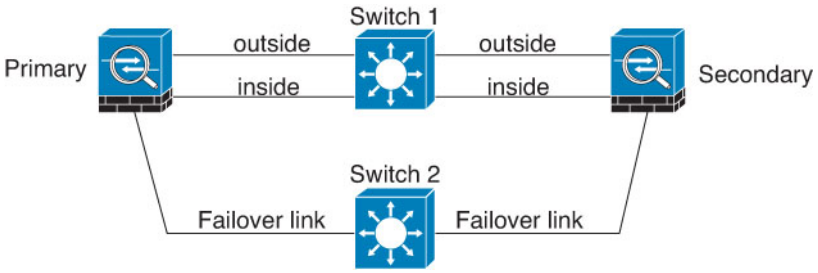
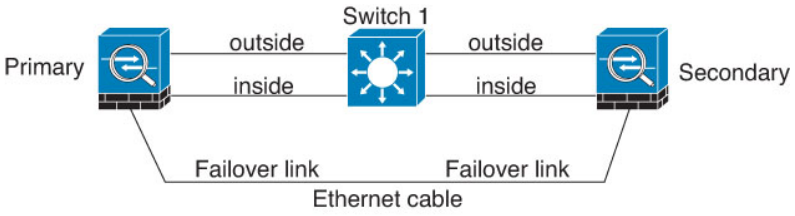


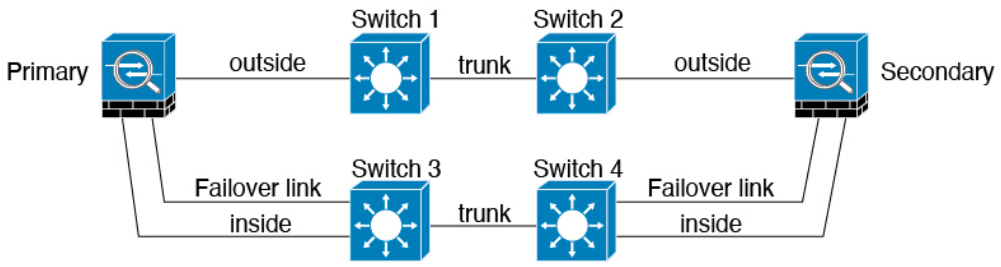
Figure 4: Connecting with a Cable



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



MAC addresses and IP addresses in Failover

MAC addresses and IP addresses in Failover are network addressing mechanisms that

- provide unique identification for network interfaces during failover scenarios
- maintain network connectivity when primary devices become unavailable, and
- ensure seamless traffic flow between active and standby configurations.

Address configuration types

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 configurations use specific address handling methods:

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.

If you disable failover and set the failover configurations to a disabled state, you will need to manually resume failover, or reboot the device. It is recommended to use the command **failover reset** and resume the failover instead of rebooting the device. If you reload the standby unit with the failover configuration disabled, the standby unit boots up as the active unit and uses the primary unit's IP addresses and MAC addresses. This leads to duplicate IP addresses and causes network traffic disruptions. Use the command **failover reset** to enable failover and restore the traffic flow.



Note If you enable failover on a standalone device, the data interfaces go down at negotiation state of failover, interrupting traffic.

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. We recommend that you configure the virtual MAC address on both the primary and secondary units to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not 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.

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.

Active/active configurations implement additional address management:

Virtual MAC addresses provide specialized functionality:

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.

MAC address table updates occur during failover events:

During failover, the device that is designated as the new active device generates multicast packets for each MAC address entry in the MAC table and sends them to all the bridge group interfaces. This action prompts the upstream switches in the bridge group to update their routing tables with the new active device's interface to ensure accurate traffic forwarding.

The time taken to generate multicast packets and update the routing tables of the upstream switches depends on the number of entries in the MAC address table and the number of bridge group interfaces. Use the **show failover statistics state-switch-delay** command to display statistics related to the delays encountered during failover events.

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

A stateful failover is a high availability mechanism that

- continuously passes per-connection state information from the active unit to the standby unit
- maintains connection information after a failover occurs, and
- allows supported end-user applications to continue communication sessions without reconnecting.

Stateful failover operation

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

Supported features are high availability capabilities that

- enable state information to be passed to the standby ASA device during Stateful Failover
- maintain connection states and critical information for seamless traffic flow, and
- ensure minimal disruption during failover events.

State information types

For Stateful Failover, these state information types are passed to the standby device:

- 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 are included. SCTP inspection stateful failover is best effort. If any SACK packets are lost during failover, the new active unit drops 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.
- From all the connections, only established ones will be replicated on the Standby device.

Unsupported features

Unsupported features are device capabilities that

- are not passed to the standby ASA in Stateful Failover
- do not maintain state information during failover events, and
- may require re-establishment after a failover occurs.

State information limitations

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

- The user authentication (uauth) table
- Multicast routing.

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.

Failover health monitoring

Failover health monitoring is a system capability that monitors each unit for overall health and interface health.

Health monitoring tests

The ASA performs tests to determine the state of each unit. This section includes information about how the ASA performs tests to determine the state of each unit.

Unit health monitoring

Unit health monitoring is a failover mechanism that

- determines the health of peer units by monitoring the failover link with hello messages
- sends LANTEST messages on each data interface when three consecutive hello messages are missed, and
- initiates appropriate failover actions based on peer unit responsiveness.

Unit health monitoring behavior

The ASA determines the health of the other unit by monitoring the failover link with hello messages. If a unit does not receive three consecutive hello messages on the failover link, the sends LANTEST messages on each data interface, including the failover link, to validate whether 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. These are the 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.

Heartbeat module redundancy

Heartbeat module redundancy is a high availability feature that

- sends heartbeat messages using the data plane transport infrastructure to supplement control plane heartbeat packets
- increments a counter when the peer receives heartbeat packets in the data plane
- prevents false failover or split-brain scenarios when the control plane is congested with traffic.

Heartbeat module redundancy behavior

Each unit in the HA periodically sends a broadcast keepalive heartbeat packet over the failover link. If the control plane is too busy handling traffic, sometimes the heartbeat packets do not reach the peers, or the peers do not process the heartbeat packets due to CPU overloading. When peers cannot communicate the keepalive status within the configurable timeout period, a false failover or split-brain scenario occurs.

The heartbeat module in the data plane helps to avoid the occurrence of false failover or split-brain due to traffic congestion in the control plane.

- The additional heartbeat module works similarly to the control plane module but sends and receives heartbeat messages using the data plane transport infrastructure.
- When the peer receives heartbeat packets in the data plane, a counter gets incremented.
- If the heartbeat transfer in the control plane fails, the node checks the heartbeat counter in the data plane. If the counter is incrementing, then the peer is alive, and the cluster does not perform a failover in this situation.



Note

- The additional heartbeat module is enabled by default whenever HA is enabled. You do not have to set a poll interval for the additional heartbeat module in the data plane. This module uses the same heartbeat interval that you set for the control plane.
- This feature is not available in Version 7.3.

Interface monitoring

Interface monitoring is a failover capability that

- tracks the health of up to 1025 interfaces to detect interface failures
- performs interface tests when hello messages are not received within 15 seconds, and

- triggers failover when the number of failed interfaces meets the defined threshold.

Interface monitoring operation

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 the **failover polltime interface** command, or for Active/Active failover, the **polltime interface** command) 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 the **failover interface-policy** command, or for Active/Active failover, the **interface-policy** command), 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 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

Interface tests are monitoring mechanisms that:

- evaluate interface operational status during failover scenarios
- run sequentially with approximately 1.5 second duration per test by default, or 1/16 of the failover interface holdtime(see the **failover polltime interface** command, or for Active/Active failover, the **interface-policy** command), and
- determine interface failure status on ASA units.

Interface test sequence

The ASA uses these interface tests in sequence:

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

Interface status is a monitoring indicator that

- shows the current operational state of monitored interfaces in a failover configuration
- tracks traffic reception and hello packet communication between peer units, and
- determines whether an interface is actively monitored by the failover process.

Interface status types

Monitored interfaces can have these status types:

- **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

Failover times are high availability mechanisms that:

- initiate automatic switchover when specific failure conditions occur on the active unit,
- monitor system health through configurable detection thresholds and timing parameters, and
- ensure service continuity by transferring control to the standby unit when failures exceed defined limits.

Failover triggering events and timing parameters

These 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 is exceeded on the active device, failover occurs. If the threshold is exceeded on the standby device, the unit moves to **Fail** state.

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

Table 2: Interface policy command

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 3: 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 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 according to Save Configuration Changes. 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.

File Replication

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

- Secure Client images
- CSD images
- Secure Client profiles

  The ASA uses a cached file for the Secure Client profile stored in cache:/stc/profiles, and not the file stored in the flash file system. To replicate the Secure 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, commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state.

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

The following commands are replicated to the standby ASA:

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

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

- All forms of the **copy** command except for **copy running-config startup-config**
- All forms of the **write** command except for **write memory**
- **debug**
- **failover lan unit**
- **firewall**
- **show**
- **terminal pager** and **pager**

Config-sync optimization

Config-sync optimization is a configuration synchronization enhancement that

- compares configuration hash values between active and joining devices to determine if full synchronization is necessary
- enables faster device rejoining by skipping full configuration synchronization when hash values match, and

- reduces maintenance window and upgrade time during failover operations.

Guidelines and limitations of configuration sync optimization

When a device reboots or rejoins after a suspend or resume failover, the joining device clears its running configuration. The active device then sends its entire configuration to the joining device for a full configuration synchronization. If the active device has a large configuration, this process can take several minutes.

The configuration sync optimization functionality enables comparing the configuration of the joining device and the active device by exchanging configuration hash values. If the hash computed on both active and joining devices match, the joining device skips full configuration synchronization and rejoin the failover configuration. This functionality ensures faster peering and reduces maintenance window and upgrade time.

Guidelines and limitations:

- The configuration sync optimization functionality is enabled by default.
- ASA multiple context mode supports configuration sync optimization by sharing the context order during full configuration synchronization, allowing comparison of context order during subsequent node-rejoin.
- If you configure passphrase and failover IPsec key, then configuration sync optimization is not effective as the hash value computed in the active and standby devices differs.
- If you configure the device with dynamic ACL or SNMPv3, configuration sync optimization is not effective.
- Active device synchronizes full configuration with flapping LAN links as default behavior. During failover flaps between active and standby devices, configuration sync optimization is not triggered and devices perform a full configuration synchronization.
- Configuration sync optimization gets triggered when the failover configuration recovers from an interruption or loss of network communication between the active and standby devices.

Monitoring config-sync optimization:

- Syslog messages are generated displaying whether the hash values computed on the active and joining unit match, does not match, or if the operation timeout expires.
- The syslog message also displays the time elapsed, from the time of sending the hash request to the time of getting and comparing the hash response.

Monitoring commands for configuration sync optimization.

- **show failover config-sync checksum**: Displays information about the device status and checksum.
- **show failover config-sync configuration**: Displays information about the device configuration and checksum.
- **show failover config-sync status**: Displays status of configuration sync optimization functionality.



Note The Configuration State in **show failover state** command displays the config sync state status during HA joining. This state does not reflect the later config deployments or changes, and replication on device until a resync has been initiated.

Active/standby failover

Active/standby failover is a high availability configuration that

- lets you use a standby ASA device to take over the functionality of a failed unit,
- enables the standby unit to become the active unit when the active unit fails, and
- allows the failed unit to come back online as the standby unit if the problem is temporary.

The failed unit must be replaced if the problem is not temporary. Set the standby unit to primary before replacing the failed unit to retain the configuration of the secondary unit.

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

Primary/secondary roles and active/standby status are failover configuration concepts that

- define unit hierarchy, with primary units taking precedence during simultaneous startup
- determine IP address and MAC address usage patterns during failover operations, and
- establish traffic flow responsibilities between paired units in Active/Standby failover configurations.

Primary and secondary unit differences

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 depending 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

Active unit determination at startup is a high availability process that

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

Failover events

A failover event is a system condition that

- triggers an active unit to transfer control to a standby unit in Active/Standby failover configurations

- occurs on a unit basis and cannot fail over individual or groups of contexts even on systems running in multiple context mode, and
- follows specific failover policies that determine whether failover occurs and what actions each unit takes.

Failover event policies and actions

This 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 4: 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.
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 5: 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.
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.

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
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

For most models, 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 Virtual	See Failover Licenses for the ASAv.
Firepower 1010	Security Plus license on both units. See Failover Licenses for the Firepower 1010.
Firepower 1100	See Failover Licenses for the Firepower 1100.
Firepower 2100	See Failover Licenses for the Firepower 2100.
Secure Firewall 3100/4200	See Failover Licenses for the Secure Firewall 3100.
Firepower 4100/9300	See Failover Licenses for the Firepower 4100/9300.
ISA 3000	Security Plus license on both units. **Note** Each unit must have the same encryption license.

Note A valid permanent key is required; in rare instances on the ISA 3000, 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

- / 1010/ :
  - You should not use the switch port functionality when using Failover. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. Failover is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal Failover network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Theoretically, you can put a single switch port on a VLAN and successfully use Failover, but a simpler setup is to use physical firewall interfaces instead.
  - You can only use a firewall interface as the failover link.
- Firepower 9300—We recommend that you use inter-chassis Failover for the best redundancy.
- The ASA Virtual 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.

ASA Virtual Failover for High Availability

When creating a failover pair with the ASA Virtual, it is necessary to add the data interfaces to each ASA Virtual in the same order. If the exact same interfaces are added to each ASA Virtual, but in different order, errors may be presented at the ASA Virtual 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.
- 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.
- When no failover is triggered during a healthy active/standby state, none of the interfaces on the standby unit are shut down. However, if this occurs while HA is in an active/failed state, all interfaces—including the management interface—are shut down on the failed unit, except for the failover LAN and state interfaces. This is done to prevent potential traffic issues if the failed interfaces become active after failover has been disabled. The failover LAN and state interfaces remain up to ensure ongoing failover communication.
- 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.
- For better convergence (during a failover), you must shut down the interfaces on a HA pair that are not associated with any configuration or instance.
- In high availability deployments with a shared failover and state interfaces, if the connectivity over failover interface fails while the standby unit is in sync config state, and eventually when the standby unit switches to active state, it results in a split brain. This is observed even when there is a monitored interface.

  The connectivity over failover interface can fail in the following situations:

  - After the standby unit reboots, the connectivity over failover interface fails while the standby unit is in sync config state.
  - After enabling previously disabled failover on the standby unit, the connectivity over failover interface fails while the standby unit is in sync config state.

  To avoid split brain, use failover and state interfaces. If the connectivity failures over the failover interfaces persist, isolate the new active (previously standby) unit from the network.
- If you configure failover encryption in evaluation mode, the systems use DES for the encryption. If you then register the devices using an export-compliant account, the devices will use AES after a reboot. Thus, if a system reboots for any reason, including after installing an upgrade, the peers will be unable to communicate and both units will become the active unit. We recommend that you do not configure encryption until after you register the devices. If you do configure this in evaluation mode, we recommend you remove the encryption before registering the devices.
- 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.

- The device does not share SNMP client engine data with its peer.
- If you have a very large number of access control and NAT rules, the size of the configuration can prevent efficient configuration replication, resulting in the standby unit taking an excessively long time to reach standby ready state. This can also impact your ability to connect to the standby unit during replication through the console or SSH session. To enhance configuration replication performance, enable transactional commit for both access rules and NAT, using the **asp rule-engine transactional-commit access-group** and **asp rule-engine transactional-commit nat** commands.
- A unit in a Failover pair transitioning to the standby role synchronizes its clock with the active unit.

  Example:

```
firepower#show clock
01:00:52 UTC Mar 1 2022

...
01:01:18 UTC Mar 1 2022 <======= Incorrect (previous) clock
Cold Standby               Sync Config               Detected an Active mate

19:38:21 UTC Apr 9 2022 <======= Updated clock
Sync Config                Sync File System          Detected an Active mate
...
firepower/sec/stby#show clock
19:38:40 UTC Apr 9 2022
```

- The units in Failover do not dynamically synchronize the clock. Here are some examples of events when synchronization takes place:
  - A new Failover pair is created.
  - Failover is broken and re-created.
  - Communication over the failover link was disrupted and reestablished.
  - Failover status was manually changed at the CLI .
- Enabling Failover forces all routes to be deleted and are re-added after the Failover progression changes to the Active state. You could experience connection loss during this phase.
- If you enable failover on a standalone device, the data interfaces go down at negotiation state of failover, interrupting traffic.
- In the Failover configuration, short-lived connections, usually using port 53, are closed quickly and never transferred or synchronized from Active to Standby, so there might be a difference in the number of connections on both Failover devices. This is expected behavior for short-lived connections. You can try to compare the connections that are long-lived (for example, more than 30-60 seconds).
- In the Failover configuration, embryonic connections—connection requests that have not yet completed the three-way handshake process—are closed quickly and not synchronized between the active and standby devices. This design ensures HA system efficiency and security. For this reason, there might be a difference in the number of connections on both Failover devices, which is to be expected.
- If the failover LAN link is not connected back-to-back and instead connected through one or more switches, a failure within the intermediate path can cause the active unit to lose connectivity with the

standby unit, resulting in inconsistent active/standby states. Although this does not impact Failover functionality, it is recommended to check and recover the failover-link path between the active and standby units.

When the failover LAN link is down, it is not recommended to deploy any configuration, as it may not be replicated to the peer unit.

- In OSPF in ASA, if a nearby switch is down and if the ASA interface is connected to the same switch, the interfaces in the firewall also go down with the switch failure. This is an expected behavior. This will trigger an High Availablity Failover, as designed.
- In OSPF, after a failover, the OSPF connection with the peer device becomes invalid. Terminate the invalid connection, and then establish a new OSPF connection.
- If the standby ASA is connected to a different switch, in this case when the interface comes up, the routing tables will be different than that of the active ASA. This will lead to an outage for a short duration (approximately 15-17 seconds) until the routes and adjacency is updated.
- In transparent mode, if you have problems with the active unit losing the MAC address of the hot-standby router (HSRP), create a static mapping for the MAC address.

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.
- Monitoring on all physical 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.

Configure the Primary Unit for Active/Standby Failover

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

Before you begin

- We recommend that you configure standby IP addresses for all interfaces except for the failover and state links. If you use a 31-bit subnet mask for point-to-point connections, do not configure a standby IP address. You will not be able to enable failover if any interfaces are configured for DHCP.
- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Procedure

Step 1 Designate this unit as the primary unit:

failover lan unit primary

Step 2 Specify the interface to be used as the failover link:

failover lan interface *if_name interface_id*

Example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
```

This interface cannot be used for any other purpose (except, optionally, the state link).

The *if_name* argument assigns a name to the interface.

The *interface_id* argument can be a data physical interface, subinterface, or EtherChannel interface ID. On the Firepower 1010, the interface is a firewall interface ID; you cannot specify a switch port ID or VLAN ID. For the Firepower 4100/9300, you can use any data-type interface.

Step 3 Assign the active and standby IP addresses to the failover link:

failover interface ip *failover_if_name* {*ip_address mask* | *ipv6_address*/*prefix*} **standby** *ip_address*

Example:

```
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
```

Or:

```
ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby
2001:a0a:b00::a0a:b71
```

This address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.

The standby IP address must be in the same subnet as the active IP address.

Step 4 Enable the failover link:

interface *failover_interface_id*

no shutdown

Example:

```
ciscoasa(config)# interface gigabitethernet 0/3
ciscoasa(config-if)# no shutdown
```

Step 5 (Optional) If you want to use a separate interface for the state link, specify the interface.

failover link *if_name interface_id*

Example:

```
ciscoasa(config)# failover link folink gigabitethernet0/4
```

If you do not specify a separate interface, then the failover link is used for the statelink.

The *if_name* argument assigns a name to the interface.

The *interface_id* argument can be a physical interface, subinterface, or EtherChannel interface ID. On the Firepower 1010, the interface is a firewall interface ID; you cannot specify a switch port ID or VLAN ID.

Step 6 If you specified a separate state link, assign the active and standby IP addresses to the state link:

failover interface ip *state_if_name* {*ip_address mask* | *ipv6_address*/*prefix*} **standby** *ip_address*

Example:

```
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
```

Or:

```
ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby
2001:a0a:b00:a::a0a:b71
```

This 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. 169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.

The standby IP address must be in the same subnet as the active IP address.

Skip this step if you are sharing the state link.

Step 7 If you specified a separate state link, enable the state link.

interface *state_interface_id*

no shutdown

Example:

```
ciscoasa(config)# interface gigabitethernet 0/4
ciscoasa(config-if)# no shutdown
```

Skip this step if you are sharing the state link.

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

- (Preferred) Establish IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications:

  failover ipsec pre-shared-key [**0** | **8**] *key*

  Example:

  ```
  ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
  ```

  The *key* can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels.

  If you use a master passphrase (see Configure the Master Passphrase), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password.

  The **failover ipsec pre-shared-key** shows as ***** in **show running-config** output; this obscured key is not copyable.

  If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

  You cannot use both IPsec encryption and the legacy **failover key** encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase, you must first remove the failover key using the **no failover key** command before you configure IPsec encryption.

  Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.

- (Optional) Encrypt failover communication on the failover and state links:

  failover key [**0** | **8**] {**hex** *key* | *shared_secret*}

  Example:

  ```
  ciscoasa(config)# failover key johncr1cht0n
  ```

  Use a *shared_secret* from 1 to 63 characters or a 32-character **hex** *key*. For the *shared_secret*, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units.

  If you use a master passphrase (see Configure the Master Passphrase), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the shared secret or hex key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password.

  The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable.

  If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

Step 9 Enable failover:

failover

Step 10 Save the system configuration to flash memory:

write memory

Examples

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

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3

failover interface ip folink 172.27.48.0 255.255.255.254 standby 172.27.48.1
interface gigabitethernet 0/3
  no shutdown
failover link folink gigabitethernet0/3
failover ipsec pre-shared-key a3rynsun
failover
```

Configure the Secondary Unit for Active/Standby Failover

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

Before you begin

- You will not be able to enable failover if any interfaces are configured for DHCP.
- Do not configure a **nameif** for the failover and state links.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Procedure

Step 1 Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. See Configure the Primary Unit for Active/Standby Failover, on page 29.

For example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
INFO: Non-failover interface config is cleared on GigabitEthernet0/3 and its sub-interfaces
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
ciscoasa(config)# interface gigabitethernet 0/3
ciscoasa(config-ifc)# no shutdown
```

```
ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
ciscoasa(config)# failover
```

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

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

Configure Active/Active Failover

This section tells how to configure Active/Active failover.

Configure the Primary Unit for Active/Active Failover

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

Before you begin

- Enable multiple context mode according to Enable or Disable Multiple Context Mode.
- We recommend that you configure standby IP addresses for all interfaces except for the failover and state links according to Routed and Transparent Mode Interfaces. If you use a 31-bit subnet mask for point-to-point connections, do not configure a standby IP address. You will not be able to enable failover if any interfaces are configured for DHCP.
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Procedure

Step 1 Designate this unit as the primary unit:

failover lan unit primary

Step 2 Specify the interface to be used as the failover link:

failover lan interface *if_name interface_id*

Example:

```
ciscoasa(config)# failover lan interface folink gigabitethernet0/3
```

This interface cannot be used for any other purpose (except, optionally, the state link).

The *if_name* argument assigns a name to the interface.

The *interface_id* argument can be a physical interface, subinterface, or EtherChannel interface ID. For the Firepower 4100/9300, you can use any data-type interface.

Step 3 Assign the active and standby IP addresses to the failover link:

standby failover interface ip *if_name* {*ip_address mask* | *ipv6_address*/*prefix* } **standby** *ip_address*

Example:

```
ciscoasa(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2
```

Or:

```
ciscoasa(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby
2001:a0a:b00::a0a:b71
```

This address should be on an unused subnet. This subnet can be 31-bits (255.255.255.254) with only two IP addresses. 169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.

The standby IP address must be in the same subnet as the active IP address.

Step 4 Enable the failover link:

interface *failover_interface_id*

no shutdown

Example:

```
ciscoasa(config)# interface gigabitethernet 0/3
ciscoasa(config-if)# no shutdown
```

Step 5 (Optional) If you want to use a separate interface for the state link, specify the interface.

failover link *if_name interface_id*

Example:

```
ciscoasa(config)# failover link statelink gigabitethernet0/4
```

We recommend that you use a separate state link from the failover link. If you do not specify a separate interface, then the failover link is used for the statelink.

The *if_name* argument assigns a name to the interface.

The *interface_id* argument can be a physical interface, subinterface, or EtherChannel interface ID.

Step 6 If you specified a separate state link, assign the active and standby IP addresses to the state link:

This 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. 169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets, and you cannot use them for the failover or state links.

The standby IP address must be in the same subnet as the active IP address.

Skip this step if you are sharing the state link.

failover interface ip state *if_name* {*ip_address mask* | *ipv6_address/prefix*} **standby** *ip_address*

Example:

```
ciscoasa(config)# failover interface ip statelink 172.27.49.1 255.255.255.0 standby
172.27.49.2
```

Or:

```
ciscoasa(config)# failover interface ip statelink 2001:a0a:b00:a::a0a:b70/64 standby
2001:a0a:b00:a::a0a:b71
```

Step 7 If you specified a separate state link, enable the state link:

interface *state_interface_id*

no shutdown

Example:

```
ciscoasa(config)# interface gigabitethernet 0/4
ciscoasa(config-if)# no shutdown
```

Skip this step if you are sharing the state link.

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

- (Preferred) Establish IPsec LAN-to-LAN tunnels on the failover and state links between the units to encrypt all failover communications:

  failover ipsec pre-shared-key [**0** | **8**] *key*

  ```
  ciscoasa(config)# failover ipsec pre-shared-key a3rynsun
  ```

  The *key* can be up to 128 characters in length. Identify the same key on both units. The key is used by IKEv2 to establish the tunnels.

  If you use a master passphrase (see Configure the Master Passphrase), then the key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password.

  The **failover ipsec pre-shared-key** shows as ***** in **show running-config** output; this obscured key is not copyable.

  If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

  You cannot use both IPsec encryption and the legacy **failover key** encryption. If you configure both methods, IPsec is used. However, if you use the master passphrase, you must first remove the failover key using the **no failover key** command before you configure IPsec encryption.

  Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license.

- (Optional) Encrypt failover communication on the failover and state links:

  failover key [**0** | **8**] {**hex** *key* | *shared_secret*}

```
ciscoasa(config)# failover key johncr1cht0n
```

Use a *shared_secret*, from 1 to 63 characters, or a 32-character **hex** *key*.

For the *shared_secret*, you can use any combination of numbers, letters, or punctuation. The shared secret or hex key is used to generate the encryption key. Identify the same key on both units.

If you use a master passphrase (see Configure the Master Passphrase), then the shared secret or hex key is encrypted in the configuration. If you are copying from the configuration (for example, from **more system:running-config** output), specify that the shared secret or hex key is encrypted by using the **8** keyword. **0** is used by default, specifying an unencrypted password.

The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable.

If you do not configure failover and state link encryption, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

Step 9 Create failover group 1:

failover group 1

primary

preempt [*delay*]

Example:

```
ciscoasa(config-fover-group)# failover group 1
ciscoasa(config-fover-group)# primary
ciscoasa(config-fover-group)# preempt 1200
```

Typically, you assign group 1 to the primary unit, and group 2 to the secondary unit. 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. The **preempt** command causes the failover group to become active on the designated unit automatically when that unit becomes available.

You can enter an optional *delay* value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit. Valid values are from 1 to 1200.

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.

If you manually fail over, the **preempt** command is ignored.

Step 10 Create failover group 2 and assign it to the secondary unit:

failover group 2

secondary

preempt [*delay*]

Example:

```
ciscoasa(config-fover-group)# failover group 2
ciscoasa(config-fover-group)# secondary
```

```
ciscoasa(config-fover-group)# preempt 1200
```

Step 11 Enter the context configuration mode for a given context, and assign the context to a failover group:

context *name*

join-failover-group {**1** | **2**}

Example:

```
ciscoasa(config)# context Eng
ciscoasa(config-ctx)# join-failover-group 2
```

Repeat this command for each context.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1; you cannot assign it to group 2.

Step 12 Enable failover:

failover

Step 13 Save the system configuration to flash memory:

write memory

Examples

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

```
failover lan unit primary
failover lan interface folink gigabitethernet0/3

failover interface ip folink 172.27.48.0 255.255.255.254 standby 172.27.48.1
interface gigabitethernet 0/3
  no shutdown
failover link statelink gigabitethernet0/4

failover interface ip statelink 172.27.48.2 255.255.255.254 standby 172.27.48.3
interface gigabitethernet 0/4
  no shutdown
failover group 1
  primary
  preempt
failover group 2
  secondary
  preempt
context admin
  join-failover-group 1
failover ipsec pre-shared-key a3rynsun
failover
```

Configure the Secondary Unit for Active/Active Failover

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

Before you begin

- Enable multiple context mode according to Enable or Disable Multiple Context Mode.
- You will not be able to enable failover if any interfaces are configured for DHCP.
- Do not configure a **nameif** for the failover and state links.
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Procedure

Step 1 Re-enter the exact same commands as on the primary unit *except* for the **failover lan unit primary** command. You can optionally replace it with the **failover lan unit secondary** command, but it is not necessary because **secondary** is the default setting. You also do not need to enter the **failover group** and **join-failover-group** commands, as they are replicated from the primary unit. See Configure the Primary Unit for Active/Active Failover, on page 34.

For example:

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

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

ciscoasa(config)# **write memory**

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

failover active group 2

Configure Optional Failover Parameters

You can customize failover settings as desired.

Configure Failover Criteria and Other Settings

See Defaults for Failover, on page 29 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.

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 Change the unit poll and hold times:

failover polltime [**unit**] [**msec**] *poll_time* [**holdtime** [**msec**] *time*]

Example:

```
ciscoasa(config)# failover polltime unit msec 200 holdtime msec 800
```

The **polltime** range is between 1 and 15 seconds or between 200 and 999 milliseconds. The **holdtime** range is between 1and 45 seconds or between 800 and 999 milliseconds.You cannot enter a holdtime value that is less than 3 times the unit poll time. With a faster poll time, the ASA can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested.

If a unit does not hear hello packet on the failover communication interface for one polling period, additional testing occurs through the remaining interfaces. If there is still no response from the peer unit during the hold time, the unit is considered failed and, if the failed unit is the active unit, the standby unit takes over as the active unit.

In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.

Step 2 Configure BFD for unit health monitoring.

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.

a) Define a BFD template to be used for failover health detection:

bfd-template single-hop *template_name*

bfd interval min-tx *milliseconds***min-rx** *milliseconds* **multiplier** *multiplier_value*

Example:

```
ciscoasa(config)# bfd template single-hop failover-temp
ciscoasa(config-bfd)# bfd interval min-tx 50 min-rx 50 multiplier 3
```

The **min-tx** specifies the rate at which BFD control packets are sent to the failover peer. The range is 50 to 999 milliseconds. The **min-rx** specifies the rate at which BFD control packets are expected to be received from the failover peer. The range is 50 to 999 milliseconds. The **multiplier** specifies the number of consecutive BFD control packets that must be missed from a failover peer before BFD declares that the peer is unavailable. The range is 3 to 50.

You can also configure echo and authentication for this template; see Create the BFD Template.

b) Enable BFD for health monitoring:

failover health-check bfd *template_name*

Example:

```
ciscoasa(config)# failover health-check bfd failover-temp
```

Step 3 Change the interface link state poll time:

failover polltime link-state msec *poll_time*

Example:

```
ciscoasa(config)# failover polltime link-state msec 300
```

The range is between 300 and 799 milliseconds. 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.

In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.

Step 4 Set the session replication rate in connections per second:

failover replication rate *conns*

Example:

```
ciscoasa(config)# failover replication rate 20000
```

The minimum and maximum rate is determined by your model. The default is the maximum rate. In Active/Active mode, you set this rate for the system; you cannot set this rate per failover group.

Step 5 Disable the ability to make any configuration changes directly on the standby unit or context:

failover standby config-lock

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

Step 6 (Active/Active mode only) Specify the failover group you want to customize:

failover group {**1** | **2**}

Example:

```
ciscoasa(config)# failover group 1
ciscoasa(config-fover-group)#
```

Step 7 Enable HTTP state replication:

- For Active/Standby mode:

  failover replication http

- For Active/Active mode:

  replication http

To allow HTTP connections to be included in the state information replication, you need to enable HTTP replication. We recommend enabling HTTP state replication.

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 8 Set the threshold for failover when interfaces fail:

- For Active/Standby mode:

  failover interface-policy *num* [%]

  Example:

  ```
  ciscoasa (config)# failover interface-policy 20%
  ```

- For Active/Active mode:

  interface-policy *num* [%]

  Example:

  ```
  ciscoasa(config-fover-group)# interface-policy 20%
  ```

By default, one interface failure causes failover.

When specifying a specific number of interfaces, the *num* argument can be from 1 to 1025.

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

Step 9 Change the interface poll and hold times:

- For Active/Standby mode:

  failover polltime interface [**msec**] *polltime* [**holdtime** *time*]

  Example:

  ```
  ciscoasa(config)# failover polltime interface msec 500 holdtime 5
  ```

- For Active/Active mode:

  polltime interface [**msec**] *polltime* [**holdtime***time*]

  Example:

```
ciscoasa(config-fover-group)# polltime interface msec 500 holdtime 5
```

- *polltime*—Sets how long to wait between sending a hello packet to the peer. Valid values for the polltime are from 1 to 15 seconds or, if the optional **msec** keyword is used, from 500 to 999 milliseconds. The default is 5 seconds.
- **holdtime***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.

Step 10 Configure the virtual MAC address for an interface:

- For Active/Standby mode:

  failover mac address *phy_if active_mac standby_mac*

  Example:

```
ciscoasa(config)# failover mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8
```

- For Active/Active mode:

  mac address *phy_if active_mac standby_mac*

  Example:

```
ciscoasa(config-fover-group)# mac address gigabitethernet0/2 00a0.c969.87c8 00a0.c918.95d8
```

The *phy_if* argument is the physical name of the interface, such as gigabitethernet0/1.

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

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

You can also set the MAC address using other commands or methods, but we recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.

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

Step 11 (Active/Active mode only) Repeat this procedure for the other failover group.

Configure Interface Monitoring

By default, monitoring is enabled on all physical interfaces, or for the Firepower 1010, all VLAN interfaces. Firepower 1010 switch ports are not elegible for interface monitoring.

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).

Before you begin

In multiple context mode, configure interfaces within each context.

Procedure

Enable or disable health monitoring for an interface:

[**no**] **monitor-interface** {*if_name*}

Example:

```
ciscoasa(config)# monitor-interface inside
ciscoasa(config)# no monitor-interface eng1
```

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.

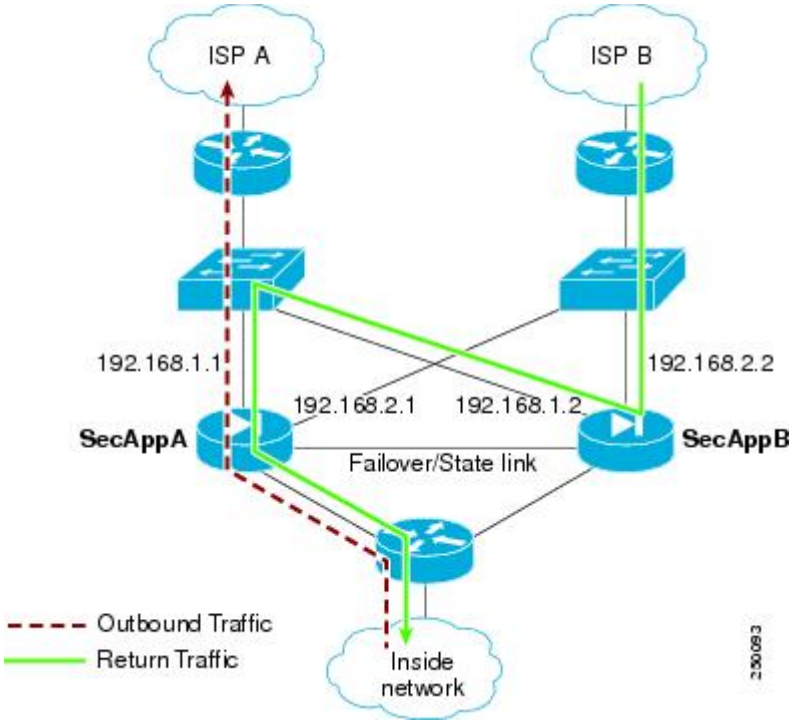
You can prevent the return packets from being dropped by allowing asymmetrically routed packets. To do so, you assign the similar interfaces on each ASA to the same ASR group. For example, both ASAs connect to the inside network on the inside interface, but connect to separate ISPs on the outside interface. 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 6: 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, specify the interface for which you want to allow asymmetrically routed packets:

interface *phy_if*

Example:

```
primary/admin(config)# interface gigabitethernet 0/0
```

Step 2 Set the ASR group number for the interface:

asr-group *num*

Example:

```
primary/admin(config-ifc)# asr-group 1
```

Valid values for *num* range from 1 to 32.

Step 3 On the secondary unit, specify the similar interface for which you want to allow asymmetrically routed packets:

interface *phy_if*

Example:

```
secondary/ctx1(config)# interface gigabitethernet 0/1
```

Step 4 Set the ASR group number for the interface to match the primary unit interface:

asr-group *num*

Example:

```
secondary/ctx1(config-ifc)# asr-group 1
```

Examples

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

Primary Unit System Configuration

```
interface GigabitEthernet0/1
  description LAN/STATE Failover Interface
interface GigabitEthernet0/2
  no shutdown
interface GigabitEthernet0/3
  no shutdown
interface GigabitEthernet0/4
  no shutdown
interface GigabitEthernet0/5
  no shutdown
failover
failover lan unit primary
failover lan interface folink GigabitEthernet0/1
failover link folink
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
  primary
failover group 2
  secondary
admin-context SecAppA
context admin
  allocate-interface GigabitEthernet0/2
  allocate-interface GigabitEthernet0/3
  config-url flash:/admin.cfg
  join-failover-group 1
context SecAppB
  allocate-interface GigabitEthernet0/4
  allocate-interface GigabitEthernet0/5
  config-url flash:/ctx1.cfg
  join-failover-group 2
```

SecAppA Context Configuration

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

SecAppB Context Configuration

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

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.

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 Force a failover when entered on the *standby* unit. The standby unit becomes the active unit.

If you specify the **group** *group_id*, then this command forces a failover when entered on the *standby* unit for the specified Active/Active failover group. The standby unit becomes the active unit for the failover group.

- For Active/Standby mode on the standby unit:

  failover active

- For Active/Active mode on the standby unit:

  failover active [**group** *group_id*]

  Example:

  ```
  standby# failover active group 1
  ```

Step 2 Force a failover when entered on the *active* unit. The active unit becomes the standby unit.

If you specify the **group** *group_id*, then this command forces a failover when entered on the *active* unit for the specified failover group. The active unit becomes the standby unit for the failover group.

- For Active/Standby mode on the active unit:

  no failover active

- For Active/Active mode on the active unit:

  no failover active [**group** *group_id*]

  Example:

```
active# no failover active group 1
```

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 48. 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 Disable failover:

no failover

Step 2 To completely disable failover, save the configuration and reload:

write memory

reload

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 Restore a failed unit to an unfailed state:

- For Active/Standby mode:

  failover reset

- For Active/Active mode:

  failover reset [**group** *group_id*]

  Example:

```
ciscoasa(config)# failover reset group 1
```

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

If you specify the **group** *group_id*, this command restores a failed Active/Active failover group to an unfailed state.

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

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

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

Replicated commands are stored in the running configuration.



Note If the maximum interface limit is reached, the number of VPN tunnels on the active and standby units may be inconsistent. In this situation run the**write standby**command on the active unit to sync the configuration.

Test the Failover Functionality

To test failover functionality, perform the following procedure.

Procedure

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

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

Active/Standby mode:

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

Active/Active mode:

ciscoasa(config)# **no failover active group** *group_id*

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

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

Step 5 When you are finished, you can restore the unit to active status by enter the following command on the newly active unit:

Active/Standby mode:

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

Active/Active mode:

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

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

Remote Command Execution

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

Send a Command

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

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

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

Procedure

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

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

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

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

Commands that cause a command mode change do not change the prompt for the current session. You must use the **show failover exec** command to display the command mode the command is executed in. See Change Command Modes, page 7-53 for more information.

Change Command Modes

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

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

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

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

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

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

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

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

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

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

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

Security Considerations

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

Limitations of Remote Command Execution

When you use remote commands, you face the following limitations:

- If you upgrade one unit using the zero-downtime upgrade procedure and not the other, both units must be running software that supports the **failover exec** command.
- Command completion and context help is not available for the commands in the *cmd_string* argument.
- In multiple context mode, you can only send commands to the peer context on the peer unit. To send commands to a different context, you must first change to that context on the unit to which you are logged in.
- You cannot use the following commands with the **failover exec** command:

- **changeto**
- **debug** (**undebug**)

- If the standby unit is in the failed state, it can still receive commands from the **failover exec** command if the failure is due to a service card failure; otherwise, the remote command execution will fail.
- You cannot use the **failover exec** command to switch from privileged EXEC mode to global configuration mode on the failover peer. For example, if the current unit is in privileged EXEC mode, and you enter **failover exec mate configure terminal**, the **show failover exec mate** output will show that the failover exec session is in global configuration mode. However, entering configuration commands for the peer unit using **failover exec** will fail until you enter global configuration mode on the current unit.
- You cannot enter recursive failover exec commands, such as the **failover exec mate failover exec mate** command.
- Commands that require user input or confirmation must use the **noconfirm** option. For example, to reload the mate, enter:

  failover exec mate reload noconfirm

Monitoring Failover

Monitoring Failover is a system capability that enables you to view 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

To monitor failover status, enter one of the following commands:

- **show failover**

  Displays information about the failover state of the unit.

- **show failover group**

  Displays information about the failover state of the failover group. The information displayed is similar to that of the **show failover** command but limited to the specified group.

- **show monitor-interface**

  Displays information about the monitored interface.

- **show running-config failover**

  Displays the failover commands in the running configuration.

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.
Support for a hex value for the failover key	7.0(4)	You can now specify a hex value for failover link encryption. We modified the following command: **failover key hex**.
Support for the master passphrase for the failover key	8.3(1)	The failover key now supports the master passphrase, which encrypts the shared key in the running and startup configuration. If you are copying the shared secret from one ASA to another, for example from the **more system:running-config** command, you can successfully copy and paste the encrypted shared key. **Note** The **failover key** shared secret shows as ***** in **show running-config** output; this obscured key is not copyable. We modified the following command: **failover key** [**0** | **8**].
IPv6 support for failover added.	8.2(2)	We modified the following commands: **failover interface ip**, **show failover**, **ipv6 address**, **show monitor-interface**.

Feature Name	Releases	Feature Information
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 (the **failover key** command), you can now use an IPsec LAN-to-LAN tunnel for failover and state link encryption. **Note** Failover LAN-to-LAN tunnels do not count against the IPsec (Other VPN) license. We introduced or modified the following commands: **failover ipsec pre-shared-key**, **show vpn-sessiondb**.
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 command: **monitor-interface service-module**
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 introduced the following command: **failover standby config-lock**
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 commands: **failover lan interface, failover link**
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 modified the following command: **show local-host**

Feature Name	Releases	Feature Information
Improved sync time for dynamic ACLs from Secure Client when using Active/Standby failover	9.6(2)	When you use Secure Client 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 commands.
Stateful failover for Secure Client connections in multiple context mode	9.6(2)	Stateful failover is now supported for Secure Client connections in multiple context mode. We did not modify any commands.
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 introduced the following command: **failover polltime link-state**
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 introduced the following command: **failover health-check bfd**
Disable failover delay	9.15(1)	When you use bridge groups or IPv6 DAD, when a failover occurs the new active unit waits up to 3000 ms for the standby unit to finish networking tasks and transition to the standby state. Then the active unit can start passing traffic. To avoid this delay, you can disable the waiting time, and the active unit will start passing traffic before the standby unit transitions. New/Modified commands: **failover wait-disable**
Config-Sync Optimization feature for faster HA peering	9.18(1)	The Config-Sync Optimization feature enables comparing the configuration of the joining unit and the active unit by exchanging config-hash values. If the hash computed on both active and joining units match, the joining unit skips full config-sync and rejoin the HA. This feature enables faster HA peering and reduces maintenance window and upgrade time.
Heartbeat module redundancy	9.20(1)	Introduced an additional heartbeat module in the data plane of ASA high availability. This heartbeat module helps to avoid false failovers or split-brain scenarios that can happen due to traffic congestion in the control plain or CPU overload.