Configuring Active/Standby Failover

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

Information About Active/Standby Failover

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

Active/Standby Failover Overview

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

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

Primary/Secondary Status and Active/Standby Status

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

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

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

Device Initialization and Configuration Synchronization

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

The active unit is determined by the following:

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

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

Virtual MAC addresses guard against this disruption because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. In multiple context mode, the ASA generates virtual active and standby MAC addresses by default. See the "Information About MAC Addresses" section on page 5-11 for more information. In single context mode, you can manually configure virtual MAC addresses; see the "Configuring Virtual MAC Addresses" section on page 8-16 for more information.

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

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." During replication, commands entered on the active unit may not replicate properly to the standby unit, and commands entered on the standby unit may be overwritten by the configuration being replicated from the active unit. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

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

Note On the standby unit, the configuration exists only in running memory. To save the configuration to the flash memory on the standby unit, choose **File > Save Running Configuration to Flash**. Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

Command replication always flows from the active unit to the standby unit. As you apply your changes to the active unit in ASDM, the associated commands are sent across the failover link to the standby unit. You do not have to save the active configuration to flash memory to replicate the commands.

The following commands that 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 that 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**

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

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

Note Standby Failover does not replicate the following files and configuration components:

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

Failover Triggers

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

- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- You force a failover. (See the "Forcing Failover" section on page 8-17.)

Failover Actions

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

Table 8-1 shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 8-1 Failover Behavior

Failure Event	Policy	Active Action	Standby Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.

Table 8-1 Failover Behavior (continued)

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

Optional Active/Standby Failover Settings

You can configure the following Active/Standby failover options when you initially configuring failover or after failover has been configured:

- HTTP replication with Stateful Failover—Allows connections to be included in the state information replication.
- Interface monitoring—Allows you to monitor up to 250 interfaces on a unit and control which interfaces affect your failover.
- Interface health monitoring—Enables the ASA to detect and respond to interface failures more quickly.
- Failover criteria setup—Allows you to specify a specific number of interfaces or a percentage of monitored interfaces that must fail before failover occurs.
- Virtual MAC address configuration—Ensures that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit.

Licensing Requirements for Active/Standby Failover

The following table shows the licensing requirements for this feature:

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

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

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

Prerequisites for Active/Standby Failover

Active/Standby failover has the following prerequisites:

- Both units must be identical ASAs that are connected to each other through a dedicated failover link and, optionally, a Stateful Failover link.
- Both units must have the same software configuration and the proper license.
- Both units must be in the same mode (single or multiple, transparent or routed).

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- Supported in single and multiple context mode.
- For multiple context mode, perform all steps in the system execution space unless otherwise noted.

Firewall Mode Guidelines

- Supported in transparent and routed firewall mode.

IPv6 Guidelines

- IPv6 failover is supported.

Model Guidelines

- Stateful failover is not supported on the ASA 5505.

Additional Guidelines and Limitations

Configuring port security on the switch(es) connected to an ASA failover pair can cause communication problems when a failover event occurs. This is because if 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.

ASA failover replication fails if you try to make a configuration change in two or more contexts at the same time. The workaround is to make configuration changes on each unit sequentially.

The following guidelines and limitations apply for Active/Standby failover:

- To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.
- The standby IP addresses are used on the ASA that is currently the standby unit, and they must be in the same subnet as the active IP address on the corresponding interface on the active unit.
- If you change the console terminal pager settings on the active unit in a failover pair, the active console terminal pager settings change, but the standby unit settings do not. A default configuration issued on the active unit does affect behavior on the standby unit.
- When you enable interface monitoring, you can monitor up to 250 interfaces on a unit.
- By default, the ASA does not replicate HTTP session information when Stateful Failover is enabled. Because HTTP sessions are typically short-lived, and because HTTP clients typically retry failed connection attempts, not replicating HTTP sessions increases system performance without causing serious data or connection loss. The **failover replication http** command enables the stateful replication of HTTP sessions in a Stateful Failover environment, but it could have a negative impact upon system performance.
- AnyConnect images must be the same on both ASAs in a failover pair. If the failover pair has mismatched images when a hitless upgrade is performed, then the WebVPN 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.”

Configuring Active/Standby Failover

This section describes how to configure Active/Standby failover. This section includes the following topics:

Using the High Availability and Scalability Wizard

The High Availability and Scalability Wizard guides you through a step-by-step process of creating either an Active/Active failover configuration, an Active/Standby failover configuration, or a VPN Cluster Load Balancing configuration.

As you go through the wizard, screens appear according to the type of failover that you are configuring and the hardware platform that you are using.

This section includes the following topics:

Task Flow for Using the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Standby failover using the High Availability and Scalability Wizard. Each step in the procedure corresponds to a wizard screen. Click **Next** after completing each step, except for the last one, before proceeding to the next step. Each step also includes a reference to additional information that you may need to complete the step.

Step 1 Start the wizard. See the "Start the Wizard" section on page 8-8.

Step 2 Enter the IP address of the failover peer on the Failover Peer Connectivity and Compatibility Check screen. Click **Test Compatibility**. You cannot move to the next screen until all compatibility tests have been passed. See the "Failover Peer Connectivity and Compatibility Check" section on page 8-9.

Step 3 Define the Failover Link in the Failover Link Configuration screen. See the "Failover Link Configuration" section on page 8-9.

Step 4 (Not available on the ASA 5505 ASA) Define the Stateful Failover link in the State Link Configuration screen. See the "State Link Configuration" section on page 8-10.

Step 5 Add standby addresses to the ASA interfaces in the Standby Address Configuration screen. See the "Standby Address Configuration" section on page 8-10.

Step 6 Review your configuration in the Summary screen. If necessary, click **Back** to go to a previous screen and make changes. See the "Summary" section on page 8-11.

Step 7 Click **Finish**. The failover configuration is sent to the ASA and to the failover peer.

Start the Wizard

Step 1 Choose **Wizards > High Availability and Scalability**.

Step 2 In the Configuration Type screen, click **Configure Active/Standby failover**.

Failover Peer Connectivity and Compatibility Check

The Failover Peer Connectivity and Compatibility Check screen lets you verify that the selected failover peer is reachable and compatible with the current unit. If any of the connectivity and compatibility tests fail, you must correct the problem before you can proceed with the wizard.

To check failover peer connectivity and compatibility, perform the following steps:

Step 1 Enter the IP address of the peer unit. This address does not have to be the failover link address, but it must be an interface that has ASDM access enabled on it. The field accepts both IPv4 and IPv6 addresses.

Step 2 Click **Next** to perform the following connectivity and compatibility tests:

- Connectivity test from this ASDM to the peer unit
- Connectivity test from this firewall device to the peer firewall device
- Hardware compatibility test for the platform
- Software version compatibility
- Failover license compatibility
- Firewall mode compatibility (routed or transparent)
- Context mode compatibility (single or multiple)

Failover Link Configuration

The Failover Link Configuration screen appears *only* if you are configuring LAN-based failover.

To configure LAN-based failover, perform the following steps:

Step 1 Choose the LAN interface to use for failover communication from the drop-down list.

Step 2 Enter a name for the interface.

Step 3 Enter the IP address used for the failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 4 Enter the IP address used for the failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 5 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Step 6 (For ASA 5505 only) Choose the switch port from the drop-down list, which includes the current VLAN assigned to each switch port and any name associated with the VLAN. Because a default VLAN exists for every switch port, do not choose VLAN 1 for the inside port, because one less inside port will be available for another use.

Note To provide sufficient bandwidth for failover, do not use trunks or PoE for failover.

Step 7 (Optional) Enter the secret key used to encrypt failover communication. If you leave this field blank, failover communication, including any passwords or keys in the configuration that are sent during command replication, will be in clear text.

State Link Configuration

Note The State Link Configuration screen does not appear on the ASA 5505.

The State Link Configuration screen lets you enable and disable Stateful Failover, and configure Stateful Failover link properties.

To enable Stateful Failover, perform the following steps:

Step 1 To pass state information across the LAN-based failover link, click **Use the LAN link as the State Link**.

Step 2 To disable Stateful Failover, click **Disable Stateful Failover**.

Step 3 To configure an unused interface as the Stateful Failover interface, click **Configure another interface for Stateful failover**.

Step 4 Choose the interface to use for Stateful Failover communication from the drop-down list.

Step 5 Enter the name for the Stateful Failover interface.

Step 6 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the active state. This field accepts an IPv4 or IPv6 address.

Step 7 Enter the IP address for the Stateful Failover link on the unit that has failover group 1 in the standby state. This field accepts an IPv4 or IPv6 address.

Step 8 Enter or choose a subnet mask (IPv4 addresses or a prefix (IPv6 Addresses) for the Active IP and Standby IP addresses.

Standby Address Configuration

Use the Standby Address Configuration screen to assign standby IP addresses to the interface on the ASA. The interfaces currently configured on the failover devices appear. The interfaces are grouped by context, and the contexts are grouped by failover group.

To assign standby IP addresses to the interface on the ASA, perform the following steps:

Step 1 Click the plus sign (+) by a device name to display the interfaces on that device. Click the minus sign (-) by a device name to hide the interfaces on that device.

Step 2 Double-click the **Active IP** field to edit or add an active IP address. Changes to this field also appear in the Standby IP field for the corresponding interface on the failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 3 Double-click the **Standby IP** field to edit or add a standby IP address. Changes to this field also appear in the Active IP field for the corresponding interface on the failover peer unit. This field accepts IPv4 or IPv6 addresses.

Step 4 Check the **Is Monitored** check box to enable health monitoring for that interface. Uncheck the check box to disable health monitoring. By default, health monitoring of physical interfaces is enabled, and health monitoring of virtual interfaces is disabled.

Step 5 Choose the asynchronous group ID from the drop-down list. This setting is only available for physical interface. For virtual interfaces, this field displays "None."

Summary

The Summary screen displays the results of the configuration steps that you performed in the previous wizard screens.

Verify your settings and click **Finish** to send your configuration to the device. If you are configuring failover, the configuration is also sent to the failover peer. If you need to change a setting, click **Back** to return to the screen that you want to change. Make the change, and click **Next** until you return to the Summary screen.

Configuring Failover (Without the Wizard)

Follow these steps to configure Active/Standby failover on both units.

The speed and duplex settings for the failover interface cannot be changed when Failover is enabled. To change these settings for the failover interface, you must configure them in the Configuration > Interfaces pane before enabling failover.

Step 1 Choose the **Configuration > Device Management > Failover > Setup** tab.

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

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

Step 3 To encrypt the failover link, do the following:

a. (Optional) Check the **Use 32 hexadecimal character key** check box to enter a hexadecimal value for the encryption key in the Shared Key field.

b. Enter the encryption key in the Shared Key field.

If you checked the Use 32 hexadecimal character key check box, then enter a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

If the Use 32 hexadecimal character key check box is unchecked, then enter an alphanumeric shared secret. The shared secret can be from 1 to 63 characters. Valid character are any combination of numbers, letters, or punctuation. The shared secret is used to generate the encryption key.

Step 4 Select the interface to use for the failover link from the Interface list. Failover requires a dedicated interface, however you can share the interface with Stateful Failover.

Only unconfigured interfaces or subinterfaces are displayed in this list and can be selected as the LAN Failover interface. Once you specify an interface as the LAN Failover interface, you cannot edit that interface in the Configuration > Interfaces pane.

Step 5 Specify the logical name of the interface used for failover communication in the Logical Name field.

Step 6 Specify the active IP address for the interface in the Active IP field. The IP address can be either an IPv4 or an IPv6 address. You cannot configure both types of addresses on the failover link interface.

Step 7 Depending upon the type of address specified for the Active IP, enter a subnet mask (IPv4 addresses) or a prefix length (IPv6 address) for the failover interface in the Subnet Mask/Prefix Length field. The name of the field changes depending upon the type of address specified in the Active IP field.

Step 8 Specify the IP address used by the secondary unit to communicate with the primary unit in the Standby IP field. The IP address can be an IPv4 or an IPv6 address.

Step 9 Select **Primary** or **Secondary** in the Preferred Role field to specify whether the preferred role for this ASA is as the primary or secondary unit.

Step 10 (Optional) Configure the Stateful Failover link by doing the following:

Note Stateful Failover is not available on the ASA 5505 platform. This area does not appear on ASDM running on an ASA 5505 ASA.

a. Specifies the interface used for state communication. You can choose an unconfigured interface or subinterface, the LAN Failover interface, or the Use Named option.

Note We recommend that you use two separate, dedicated interfaces for the LAN Failover interface and the Stateful Failover interface.

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

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

If you choose the Use Named option, the Logical Name field becomes a drop-down list of named interfaces. Choose the interface from this list. The Active IP, Subnet Mask/Prefix Length, and Standby IP values do not need to be specified. The values specified for the interface are used. Be sure to specify a standby IP address for the selected interface on the Interfaces tab.

Note Because Stateful Failover can generate a large amount of traffic, performance for both Stateful Failover and regular traffic can suffer when you use a named interface.

b. Specify the IP address for the Stateful Failover interface in the Active IP field. The IP address can be either an IPv4 or an IPv6 address. You cannot configure both types of addresses on the failover link interface. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.

c. Specify the mask (IPv4 address) or prefix (IPv6 address) for the Stateful Failover interface in the Subnet Mask/Prefix Length. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.

d. Specify the interface name used for failover communication in the Logical Name field. If you chose the Use Named option in the Interface drop-down list, this field displays a list of named interfaces. This field is dimmed if the LAN Failover interface is chosen from the Interface drop-down list.

e. Specify the IP address used by the secondary unit to communicate with the primary unit in the Standby IP field. The IP address can be an IPv4 or an IPv6 address. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.

f. (Optional) Enable HTTP replication by checking the **Enable HTTP Replication** check box. This enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected in the event of a failover.

Step 11 Click **Apply**.

The configuration is saved to the devices in the failover pair.

Configuring Interface Standby Addresses

Configuring standby IP address in ASDM changes depending upon the mode in which the unit is operating. This section includes the following topics:

Configuring Interface Standby Addresses in Routed Firewall Mode

To configure a standby address for each interface on the ASA, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. The IP address for each interface appears in the Active IP Address column. If configured, the standby IP address for the interface appears in the Standby IP address column. The failover interface and Stateful failover interface do not display IP address; you cannot change those address from this tab.

Step 2 For each interface that does not have a standby IP address, double-click the Standby IP Address field and do one of the following:

- Click the **...** button and select an IP address from the list.
- Type an IP address into the field. The address can be an IPv4 or an IPv6 address.

You can also specify whether or not the interface is monitored from this tab. For more information about configuring interface monitoring, see the "Disabling and Enabling Interface Monitoring" section on page 8-14.

Configuring the Management Interface Standby Address in Transparent Firewall Mode

If you are in multiple context mode, you must perform this procedure in each context.

To configure the management interface standby address on the ASA, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. Only the Management interface shows an IP address.

Step 2 For the Management interface that does not have a standby IP address, double-click the Standby IP Address field and do one of the following:

- Click the **...** button and select an IP address from the list.

- Type an IP address into the field. The address can be an IPv4 or an IPv6 address.

You can also specify whether or not the interface is monitored from this tab. For more information about configuring interface monitoring, see "Disabling and Enabling Interface Monitoring" section on page 8-14.

Configuring Optional Active/Standby Failover Settings

This section includes the following topics:

You can configure the optional Active/Standby failover settings when initially configuring the primary unit in a failover pair or on the active unit in the failover pair after the initial configuration.

Disabling and Enabling Interface Monitoring

You can control which interfaces affect your failover policy by disabling the monitoring of specific interfaces and enabling the monitoring of others. This feature enables you to exclude interfaces attached to less critical networks from affecting your failover policy.

You can monitor up to 250 interfaces on a unit. By default, monitoring physical interfaces is enabled and monitoring subinterfaces is disabled.

Hello messages are exchanged during every interface poll frequency time period between the ASA failover pair. The failover interface poll time is 3 to 15 seconds. For example, if the poll time is set to 5 seconds, testing begins on an interface if 5 consecutive hellos are not heard on that interface (25 seconds).

Monitored failover interfaces can have the following status:

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

To enable or disable health monitoring for specific interfaces on units in single configuration mode, enter one of the following commands. Alternately, for units in multiple configuration mode, you must enter the commands within each security context.

To disable or enable monitoring of an interface, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

A list of configured interfaces appears. The Monitored column displays whether or not an interface is monitored as part of your failover criteria. If it is monitored, a check appears in the Monitored check box.

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

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

Configuring Failover Criteria

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

Use the Configuration > Device Management > High Availability > Criteria tab to define criteria for failover, such as how many interfaces must fail and how long to wait between polls. The hold time specifies the interval to wait without receiving a response to a poll before unit failover.

For information about configuring the hold and poll times, see Configuring the Unit and Interface Health Poll Times, page 8-15.

To configure the interface policy, perform the following steps:

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Criteria** tab.

Step 2 In the Interface Policy area, do one of the following:

- To define a specific number of interfaces that must fail to trigger failover, enter a number from 1 to 250 in the Number of failed interfaces field. When the number of failed monitored interfaces exceeds the value you specify, the ASA fails over.
- To define a percentage of configured interfaces that must fail to trigger failover, enter a percentage in the Percentage of failed interfaces field. When the number of failed monitored interfaces exceeds the percentage you set, the ASA fails over.

Step 3 Click **Apply**.

Configuring the Unit and Interface Health Poll Times

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

Decreasing the poll and hold times enables the ASA to detect and respond to interface failures more quickly but may consume more system resources. Increasing the poll and hold times prevents the ASA from failing over on networks with higher latency.

Step 1 Choose the **Configuration > Device Management > High Availability > Failover > Criteria** tab.

Step 2 To configure the interface poll and hold times, change the following values in the Failover Poll Times area:

- Monitored Interfaces—The amount of time between polls among interfaces. The range is between 1and 15 seconds or 500 to 999 milliseconds.
- Interface Hold Time—Sets the time during which a data interface must receive a hello message on the data interface, after which the peer is declared failed. Valid values are from 5 to 75 seconds.

Step 3 To configure the unit poll and hold times, change the following values in the Failover Poll Times area:

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

Step 4 Click **Apply**.

Configuring Virtual MAC Addresses

The Configuration > Device Management > High Availability > MAC Addresses tab displays the virtual MAC addresses for the interfaces in an Active/Standby failover pair.

Note This tab is not available on the ASA 5505 platform.

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

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

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

To configure the virtual MAC address for an interface, perform the following steps:

Step 1 Open the **Configuration > Device Management > High Availability > Failover > MAC Addresses** tab.

Step 2 To edit an existing virtual MAC address entry, double-click the row for the interface whose MAC addresses you want to change. To add a new virtual MAC address entry, click **Add**.

The Add/Edit Interface MAC Address dialog box appears.

Step 3 Type the new MAC address for the active interface in the Active MAC Address field.

Step 4 Type the new MAC address for the standby interface in the Standby MAC Address field.

Step 5 Click **OK**.

Step 6 To delete a virtual MAC address entry, perform the following steps:

a. Click the interface to select the table row.

b. Click **Delete**.

c. Click **OK**.

Controlling Failover

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

Forcing Failover

To force the standby unit to become active, perform the following steps:

Step 1 Choose **Monitoring > Properties > Failover > Status**.

Step 2 Click one of the following buttons:

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

Disabling Failover

To disable failover, perform the following steps:

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

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

Restoring a Failed Unit

To restore a failed unit to an unfailed state, perform the following steps:

Step 1 Choose **Monitoring > Properties > Failover > Status**.

Step 2 Click **Reset Failover**.

Note

Monitoring Active/Standby Failover

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

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

Feature History for Active/Standby Failover

Table 8-2 lists the release history for this feature.

Table 8-2 Feature History for Optional Active/Standby Failover Settings

Feature Name	Releases	Feature Information
This feature was introduced.	7.0	This feature was introduced.
IPv6 support for failover added.	8.2(2)	We modified the following screens: Configuration > Device Management > High Availability > Failover > Setup Configuration > Device Management > High Availability > Failover > Interfaces