CHAPTER 13

Configuring Active/Active Failover

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

Information About Active/Active Failover

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

Active/Active Failover Overview

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

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

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

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

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

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

Primary/Secondary Status and Active/Standby Status

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

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

  Note The ASA also provides load balancing, which is different from failover. Both failover and load balancing can exist on the same configuration. For information about load balancing, see the "Configuring Load Balancing" section on page 71-20.

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

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

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

- When both units boot at the same time, each failover group becomes active on its preferred unit after the configurations have been synchronized.

Device Initialization and Configuration Synchronization

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

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

When the replication starts, the ASA console on the unit sending the configuration 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 unit sending the configuration may not replicate properly to the peer unit, and commands entered on the unit receiving the configuration may be overwritten by the configuration being received. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

Note On the unit receiving the configuration, the configuration exists only in running memory. To save the configuration to the flash memory on both units, select **File > Save Running Configuration to flash** from the menu bar in the system execution space on the unit that has failover group 1 in the active state. Startup configurations saved on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts configuration files from the disk on the primary unit to an external server, and then copy them to disk on the secondary unit, where they become available when the unit reloads.

Command Replication

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

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

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

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

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

Table 13-1 lists the commands that are and are not replicated to the standby unit.

Table 13-1 Command Replication

Commands Replicated to the Standby Unit	Commands Not Replicated to the Standby Unit
All configuration commands except for **mode**, **firewall**, and **failover lan unit**	All forms of the **copy** command except for **copy running-config startup-config**
copy running-config startup-config	All forms of the **write** command except for **write memory**
delete	**debug**
mkdir	**failover lan unit**
rename	**firewall**
rmdir	**mode**
write memory	**show**

Failover Triggers

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

- The unit has a hardware failure.
- The unit has a power failure.
- The unit has a software failure.
- You force a failover. (See Forcing Failover, page 13-21.)

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

- Too many monitored interfaces in the group fail.
- You force a failover. (See Forcing Failover, page 13-21.)

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

See the "Failover Health Monitoring" section on page 11-18 for more information about interface and unit monitoring.

Failover Actions

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

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

Table 13-2 shows the failover action for each failure event. For each failure event, the policy (whether or not failover occurs), actions for the active failover group, and actions for the standby failover group are given.

Table 13-2 Failover Behavior for Active/Active Failover

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.
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.
Stateful Failover link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the failover interface as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Optional Active/Active Failover Settings

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

- Failover Group Preemption—Assigns a primary or secondary priority to a failover group to specify on which unit in the failover group becomes active when both units boot simultaneously.
- 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 security appliance 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/Active Failover

The following table shows the licensing requirements for this feature:

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

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

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

Prerequisites for Active/Active Failover

In Active/Active failover, both units must have the following:

- The same hardware model.
- The same number of interfaces.
- The same types of interfaces.

- The same software version, with the same major (first number) and minor (second number) version numbers. However you can use different versions of the software during an upgrade process; for example you can upgrade one unit from Version 7.0(1) to Version 7.9(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.
- The same software configuration.
- The same mode (multiple context mode).
- The proper license.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in multiple context mode only.

Firewall Mode Guidelines

Supported in routed and transparent firewall mode.

IPv6 Guidelines

IPv6 failover is supported.

Model Guidelines

Active/Active failover is not available on the Cisco ASA 5505.

Additional Guidelines and Limitations

No two interfaces in the same context should be configured in the same ASR group.

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 features are not supported for Active/Active failover:

- To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces.
- The standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.
- You can define a maximum number of two failover groups.
- Failover groups can only be added to the system context of devices that are configured for multiple context mode.
- You can create and remove failover groups only when failover is disabled.
- Entering the **failover group** command puts you in the failover group command mode. The **primary**, **secondary**, **preempt**, **replication http**, **interface-policy**, **mac address**, and **polltime interface** commands are available in the failover group configuration mode. Use the **exit** command to return to global configuration mode.

- The **failover polltime interface**, **failover interface-policy**, **failover replication http**, and **failover mac address** commands have no affect on Active/Active failover configurations. They are overridden by the following failover group configuration mode commands: **polltime interface**, **interface-policy**, **replication http**, and **mac address**.
- 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.
- VPN failover is unavailable. (It is available in Active/Standby failover configurations only.)

Configuring Active/Active Failover

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

This section includes the following topics:

- Task Flow for Using the High Availability and Scalability Wizard, page 13-8
- Configuring Failover (Without the Wizard), page 13-12
- Configuring Asymmetric Routing Groups in Multiple Context Mode, page 13-20

Task Flow for Using the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring Active/Active 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 13-9.

Step 2 Enter the IP address of the failover peer in 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 13-9.

Step 3 If the ASA or the failover peer are in single context mode, change them to multiple context mode in the Change Device to Multiple Mode screen. When you change the ASA to multiple context mode, it reboots. ASDM automatically reestablishes communication with the ASA when it has finished rebooting. See the "Change a Device to Multiple Mode" section on page 13-9.

Step 4 Assign security contexts to failover groups in the Context Configuration screen. You can add and delete contexts in this screen. See the "Security Context Configuration" section on page 13-10.

Step 5 Define the Failover Link in the Failover Link Configuration screen. See the "Failover Link Configuration" section on page 13-10.

Step 6 Define the Stateful Failover link in the State Link Configuration screen. See the "State Link Configuration" section on page 13-10.

Step 7 Add standby addresses to the ASA interfaces in the Standby Address Configuration screen. See the "Standby Address Configuration" section on page 13-11.

Step 8 Review your configuration in the Summary screen. See "Summary" section on page 13-11.

Step 9 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/Active 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)

Change a Device to Multiple Mode

The Change Device to Multiple Mode dialog box requires that the ASA be in multiple context mode. This dialog box lets you convert a ASA in single context mode to multiple context mode.

When you convert from single context mode to multiple context mode, the ASA creates the system configuration and the admin context from the current running configuration. The admin context configuration is stored in the admin.cfg file. The conversion process does not save the previous startup configuration, so if the startup configuration differed from the running configuration, those differences are lost.

Converting the ASA from single context mode to multiple context mode causes the ASA and its peer to reboot. However, the High Availability and Scalability Wizard restores connectivity with the newly created admin context and reports the status in the Devices Status field in this dialog box.

Note You must convert both the current ASA and its peer to multiple context mode before you can proceed.

To change the current ASA to multiple context mode, perform the following steps:

Step 1 Click **Change *device* To Multiple Context**, where *device* is the hostname of the ASA.

Step 2 Repeat this step for the peer ASA.

The status of the ASA appears during conversion to multiple context mode.

Security Context Configuration

The Security Context Configuration screen lets you assign security contexts to failover groups. It displays the name of currently configured security contexts, lets you add new ones, and change or remove existing ones as needed. In addition, it displays the failover group number to which the context is assigned and lets you change the failover group as needed. Although you can create security contexts in this screen, you cannot assign interfaces to those contexts or configure other properties for them. To configure context properties and assign interfaces to a context, choose **System > Security Contexts**.

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 (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

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, failover group, or context name to expand the list. Click the minus sign (-) by a device, failover group, or context name to collapse the list.

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)

Failover-Multiple Mode, Security Context

The fields displayed on the Failover pane in multiple context mode change depending upon whether the context is in transparent or routed firewall mode.

This section includes the following topics:

- Failover - Routed
- Failover - Transparent

Failover - Routed

Use this pane to define the standby IP address for each interface in the security context and to specify whether the status of the interface should be monitored.

Fields

- Interface table—Lists the interfaces on the ASA and identifies their active IP address, standby IP address, and monitoring status.
  - Interface Name column—Identifies the interface name.
  - Active IP column—Identifies the active IP address for this interface.
  - Standby IP column—Identifies the IP address of the corresponding interface on the standby failover unit.
  - Is Monitored column—Specifies whether this interface is monitored for failure.
- Edit—Displays the Edit Failover Interface Configuration dialog box for the selected interface.

Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to define the standby IP address for an interface and to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Active IP Address—Identifies the IP address for this interface. This field does not appear if an IP address has not been assigned to the interface.

- Subnet Mask/Prefix Length—Identifies the mask (for IPv4 addresses) or prefix (for IPv6 addresses) for this interface. This field does not appear if an IP address has not been assigned to the interface.
- Standby IP Address—Specifies the IP address of the corresponding interface on the standby failover unit. This field does not appear if an IP address has not been assigned to the interface.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. 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 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 - Transparent

Use this pane to define the standby IP address for the management interface for the security context and to specify whether the status of the interfaces on the security context should be monitored.

Fields

- Interface—Lists the interfaces for the security context and identifies their monitoring status.
  - Interface Name—Identifies the interface name.
  - Is Monitored—Specifies whether this interface is monitored for failure.
- Edit—Displays the Edit Failover Interface Configuration dialog box for the selected interface.
- Management IP Address—Identifies the active and standby management IP addresses for the security context.
  - Active—Identifies the management IP address for the active failover unit.
  - Standby—Specifies the management IP address for the standby failover unit.
- Management Netmask—Identifies the mask associated with the management address.

Edit Failover Interface Configuration

Use the Edit Failover Interface Configuration dialog box to specify whether the status of the interface should be monitored.

Fields

- Interface Name—Identifies the interface name.
- Monitor interface for failure—Specifies whether this interface is monitored for failure. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. 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 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-Multiple Mode, System

This pane includes tabs for configuring the system-level failover settings in the system context of an ASA in multiple context mode. In multiple mode, you can configure Active/Standby or Active/Active failover. Active/Active failover is automatically enabled when you create failover groups in the device manager. For both types of failover, you need to provide system-level failover settings in the system context, and context-level failover settings in the individual security contexts. For more information about configuring failover in general, see Chapter 11, "Information About Failover.".

Seethe following topics for more information:

- Failover > Setup Tab
- Failover > Criteria Tab
- Failover > Active/Active Tab
- Failover > MAC Addresses Tab

Failover > Setup Tab

Use this tab to enable failover on a ASA in multiple context mode. You also designate the failover link and the state link, if using Stateful Failover, on this tab.

Note During a successful failover event on the ASA, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The ASA does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

Fields

- Enable Failover—Checking this check box enables failover and lets you configure a standby ASA.

  Note The speed and duplex settings for an 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.

- Use 32 hexadecimal character key—Check this check box to enter a hexadecimal value for the encryption key in the Shared Key field. Uncheck this check box to enter an alphanumeric shared secret in the Shared Key field.
- Shared Key—Specifies the failover shared secret or key for encrypted and authenticated communications between failover pairs.

  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 you cleared the Use 32 hexadecimal character key check box, 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.

- LAN Failover—Contains the fields for configuring LAN Failover.
  - Interface—Specifies the interface used for failover communication. Failover requires a dedicated interface, however, you can use the same interface for Stateful Failover.

    Only unconfigured interfaces or subinterfaces that have not been assigned to a context 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 or assign that interface to a context.
  - Active IP—Specifies the IP address for the failover interface on the active unit. The IP address can be an IPv4 or an IPv6 address.
  - Subnet Mask/Prefix Length—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 on the primary and secondary unit.
  - Logical Name—Specifies the logical name of the interface used for failover communication.
  - Standby IP—Specifies the IP address used by the secondary unit to communicate with the primary unit. The IP address can be an IPv4 or an IPv6 address.
  - Preferred Role—Specifies whether the preferred role for this ASA is as the primary or secondary unit in a LAN failover.
- State Failover—Contains the fields for configuring Stateful Failover.
  - Interface—Specifies the interface used for failover communication. You can choose an unconfigured interface or subinterfaces or the LAN Failover interface.

    If you choose the LAN Failover interface, the interface needs enough capacity to handle both the LAN Failover and Stateful Failover traffic. Also, 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.

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

  - Active IP—Specifies the IP address for the Stateful Failover interface on the primary unit. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.
  - Subnet Mask/Prefix Length—Specifies the mask (IPv4 address) or prefix (IPv6 address) for the Stateful Failover interfaces on the primary and secondary units. This field is dimmed if the LAN Failover interface or Use Named option is selected in the Interface drop-down list.
  - Logical Name—Specifies the logical interface used for failover communication. 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.
  - Standby IP—Specifies the IP address used by the secondary unit to communicate with the primary unit. This field is dimmed if the LAN Failover interface or Use Named option is chosen from the Interface drop-down list.

  - Enable HTTP replication—Checking this check box enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.

Failover > Criteria Tab

Use this 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.

Note If you are configuring Active/Active failover, you do not use this tab to define the interface policy; instead, you define the interface policy for each failover group using the Failover > Active/Active Tab. With Active/Active failover, the interface policy settings defined for each failover group override the settings on this tab. If you disable Active/Active failover, then the settings on this tab are used.

Fields

- Interface Policy—Contains the fields for defining the policy for failover when monitoring detects an interface failure.
  - Number of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the value you set with this command, then the ASA fails over. The range is between 1 and 250 failures.
  - Percentage of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the percentage you set with this command, then the ASA fails over.
- Failover Poll Times—Contains the fields for defining how often hello messages are sent on the failover link, and, optionally, how long to wait before testing the peer for failure if no hello messages are received.
  - 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.
  - 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.

Failover > Active/Active Tab

Use this tab to enable Active/Active failover on the ASA by defining failover groups. In an Active/Active failover configuration, both ASAs pass network traffic. Active/Active failover is only available to ASAs in multiple mode.

A failover group is simply a logical group of security contexts. You can create two failover groups on the ASA. You must create the failover groups on the active unit in the failover pair. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

Note During a successful failover event on the ASA, the interfaces are brought down, roles are switched (IP addresses and MAC addresses are swapped), and the interfaces are brought up again. However, the process is transparent to users. The ASA does not send link-down messages or system log messages to notify users that interfaces were taken down during failover (or link-up messages for interfaces brought up by the failover process).

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

Fields

- Failover Groups—Lists the failover groups currently defined on the ASA.
  - Group Number—Specifies the failover group number. This number is used when assigning contexts to failover groups.
  - Preferred Role—Specifies the unit in the failover pair, primary or secondary, on which the failover group appears in the active state when both units start up simultaneously or when the preempt option is specified. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.
  - Preempt Enabled—Specifies whether the unit that is the preferred failover device for this failover group should become the active unit after rebooting.
  - Preempt Delay—Specifies the number of seconds that the preferred failover device should wait after rebooting before taking over as the active unit for this failover group. The range is between 0 and 1200 seconds.
  - Interface Policy—Specifies either the number of monitored interface failures or the percentage of failures that are allowed before the group fails over. The range is between 1 and 250 failures or 1 and 100 percent.
  - Interface Poll Time—Specifies the amount of time between polls among interfaces. The range is between 1 and 15 seconds.
  - Replicate HTTP—Identifies whether Stateful Failover should copy active HTTP sessions to the standby firewall for this failover group. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link. This setting overrides the HTTP replication setting on the Setup tab.
- Add—Displays the Add Failover Group dialog box. This button is only enabled if less than 2 failover groups exist. See Add/Edit Failover Group for more information.
- Edit—Displays the Edit Failover Group dialog box for the selected failover group. See Add/Edit Failover Group for more information.
- Delete—Removes the currently selected failover group from the failover groups table. This button is only enabled if the last failover group in the list is selected.

Add/Edit Failover Group

Use the Add/Edit Failover Group dialog box to define failover groups for an Active/Active failover configuration.

Fields

- Preferred Role—Specifies the unit in the failover pair, primary or secondary, on which the failover group appears in the active state. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.
- Preempt after booting with optional delay of—Checking this check box causes the unit that is the preferred failover device for a failover group to become the active unit after rebooting. Checking this check box also enables the Preempt after booting with optional delay of field in which you can specify a period of time that the device should wait before becoming the active unit.
- Preempt after booting with optional delay of—Specifies the number of seconds that a unit should wait after rebooting before taking over as the active unit for any failover groups for which it is the preferred failover device. The range is between 0 and 1200 seconds.
- Interface Policy—Contains the fields for defining the policy for failover when monitoring detects an interface failure. These settings override any interface policy settings on the Criteria tab.
  - Number of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the value you set with this command, then the ASA fails over. The range is between 1 and 250 failures.
  - Percentage of failed interfaces that triggers failover—When the number of failed monitored interfaces exceeds the percentage you set with this command, then the ASA fails over.
- Poll time interval for monitored interfaces—The amount of time between polls among interfaces. The range is between 1 and 15 seconds.
- Enable HTTP replication—Checking this check box enables Stateful Failover to copy active HTTP sessions to the standby firewall. If you do not allow HTTP replication, then HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link. This setting overrides the HTTP replication setting on the Setup tab.
- MAC Addresses—Lists physical interfaces on the ASA for which an active and standby virtual MAC address has been configured.
  - Physical Interface—Displays the physical interface for which failover virtual MAC addresses are configured.
  - Active MAC Address—Displays the MAC address for the interface and failover group on the unit where the failover group is active.
  - Standby MAC Address—Displays the MAC address for the interface and failover group on the unit where the failover group is in the standby state.
- Add—Displays the Add Interface MAC Address dialog box. You cannot assign virtual MAC addresses to the LAN failover and Stateful Failover interfaces. See Add/Edit Interface MAC Address for more information.
- Edit—Displays the Edit Interface MAC Address dialog box for the selected interface. See Add/Edit Interface MAC Address for more information.
- Delete—Removes the currently selected interface from the MAC addresses table. There is no confirmation or undo.

Add/Edit Interface MAC Address

Use the Add/Edit Interface MAC Address dialog box to define the active and standby virtual MAC addresses for the interfaces in a failover group. If you do not specify a virtual MAC address for an interface, the interface is given a default virtual MAC address as follows:

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

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

These MAC addresses override the physical MAC addresses for the interface.

Fields

- Physical Interface—Specifies the physical interface for which you are defining failover virtual MAC addresses. Because the MAC addresses do not change for the LAN failover and Stateful Failover interfaces during failover, you cannot choose these interfaces.
- MAC Addresses—Contains the fields for specifying the active and standby virtual MAC addresses for the interface.
  - Active Interface—Specifies the MAC address for the interface and failover group on the unit where the failover group is active. Each interface may have up to two MAC addresses, one for each failover group, which override the physical MAC address. Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).
  - Standby Interface—Specifies the MAC address for the interface and failover group on the unit where the failover group is in the standby state. Each interface may have up to two MAC addresses, one for each failover group, which override the physical MAC address. Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Failover > MAC Addresses Tab

The MAC Addresses tab lets you configure the virtual MAC addresses for the interfaces in an Active/Standby failover pair.

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, then the failover pair uses the burned-in NIC address as the MAC address.

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

In Active/Active failover, the MAC addresses configured on this tab are not in effect. Instead, the MAC addresses defined in the failover groups are used.

Fields

- MAC Addresses—Lists physical interfaces on the ASA for which an active and standby virtual MAC address has been configured.
  - Physical Interface—Identifies the physical interface for which failover virtual MAC addresses are configured.
  - Active MAC Address—Identifies the MAC address on the active ASA (usually primary).
  - Standby MAC Address—Identifies the MAC address on the standby ASA (usually secondary).
- Add—Displays the Add/Edit Interface MAC Address dialog box.
- Edit—Displays the Add/Edit Interface MAC Address dialog box for the selected interface.
- Delete—Removes the currently selected interface from the MAC addresses table. There is no confirmation or undo.

Add/Edit Interface MAC Address

Use the Add/Edit Interface MAC Address dialog box to define the active and standby virtual MAC addresses for an interface.

Fields

- Physical Interface—Specifies the physical interface for which you are defining failover virtual MAC addresses. Because the MAC addresses do not change for the LAN failover and Stateful Failover interfaces during failover, you cannot choose these interfaces.
- MAC Addresses—Contains the fields for specifying the active and standby virtual MAC addresses for the interface.
  - Active Interface—Specifies the MAC address of the interface on the active ASA (usually primary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).
  - Standby Interface—Specifies the MAC address of the interface on the standby ASA (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

Configuring Asymmetric Routing Groups in Multiple Context Mode

Note To configure asymmetric routing (ASR) groups, you must be in the admin context and it must be active.

To configure ASR groups, perform the following steps:

Step 1 Choose **Configuration > Device Management > Routing > ASR Groups**.

Step 2 Specify the ASR group IDs of the configured interfaces from the ASR Group ID drop-down list. The maximum number of groups that you can assign to one interface is eight. If other contexts have assigned interfaces, to a group, you may assign fewer groups to this context,

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

Controlling Failover

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

- Forcing Failover, page 13-21
- Disabling Failover, page 13-21
- Restoring a Failed Unit or Failover Group, page 13-22

Forcing Failover

To force failover at the unit level, follow these steps:

Step 1 Open **System > Monitoring > Failover > System**.

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 the active unit.

To force failover at the failover group level, follow these steps:

Step 1 Open **System > Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to control.

Step 2 Cick one of the following buttons:

- Click **Make Active** to make the the failover group active on the security appliance.
- Click **Make Standby** to make the failover group active on the other security appliance.

Disabling Failover

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

To disable failover, perform the following steps:

Step 1 Open the **System > Configuration> Device Setup > High Availability > Failover > Setup** tab.

Step 2 Clear the **Enable Failover** checkbox.

Restoring a Failed Unit or Failover Group

Restoring a failed unit or failover group moves the unit or failover group from the failed state to the standby state; it does not automatically make the failover group or unit active. Restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with 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.

To restore a failed unit to an unfailed state, follow these steps:

Step 1 Open **System > Monitoring > Failover > System**.

Step 2 Click **Reset Failover**. Clicking this button on the active unit resets the standby unit.

To restore a failed failover group to an unfailed state, follow these steps:

Step 1 Open Open **System > Monitoring > Failover > Failover Group #**, where # is the number of the failover group you want to restore.

Step 2 Click **Reset Failover**.

Monitoring Active/Active Failover

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

- System, page 13-22
- Failover Group 1 and Failover Group 2, page 13-23

System

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

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

Fields

Failover state of the system—*Display only*. Displays the failover state of the ASA. The information shown is the same output you would receive from the **show failover** command. Refer to *Cisco ASA 5500 Series Command Reference* for more information about the displayed output.

The following actions are available on the System pane:

- Make Active—Click this button to make the ASA the active unit in an active/standby configuration. In an active/active configuration, clicking this button causes both failover groups to become active on the ASA.

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

Failover Group 1 and Failover Group 2

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

Fields

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

You can perform the following actions from this pane:

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

Feature History for Active/Active Failover

Table 13-3 lists the release history for this feature.

Table 13-3 Feature History for Active/Active Failover

Feature Name	Releases	Feature Information
Active/Active failover	7.0	In an Active/Active failover configuration, both ASAs can pass network traffic. We introduced this feature and the relevant commands.
IPv6 Support in failover	8.2(2)	We modified the following screens: Configuration > Device Managment > High Availability > Failover > Setup Configuration > Device Managment > High Availability > Failover > Interfaces