



CHAPTER 51

Configuring Active/Active Failover

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

- [Information About Active/Active Failover, page 51-1](#)
- [Licensing Requirements for Active/Active Failover, page 51-6](#)
- [Prerequisites for Active/Active Failover, page 51-7](#)
- [Guidelines and Limitations, page 51-7](#)
- [Configuring Active/Active Failover, page 51-8](#)
- [Remote Command Execution, page 51-21](#)
- [Monitoring Active/Active Failover, page 51-25](#)
- [Feature History for Active/Active Failover, page 51-25](#)

Information About Active/Active Failover

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

- [Active/Active Failover Overview, page 51-1](#)
- [Primary/Secondary Status and Active/Standby Status, page 51-2](#)
- [Device Initialization and Configuration Synchronization, page 51-3](#)
- [Command Replication, page 51-3](#)
- [Failover Triggers, page 51-4](#)
- [Failover Actions, page 51-5](#)

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 66-11.

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.

On the unit receiving the configuration, the configuration exists only in running memory. To save the configuration to flash memory after synchronization enter the **write memory all** command in the system execution space on the unit that has failover group 1 in the active state. The command is replicated to the peer unit, which proceeds to write its configuration to flash memory. Using the **all** keyword with this command causes the system and all context configurations to be saved.



Note

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:

- Commands 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.

- Commands 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.

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

Table 51-1 lists the commands that are and are not replicated to the standby unit.

Table 51-1 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

You can use the **write standby** command to resynchronize configurations that have become out of sync. For Active/Active failover, the **write standby** command behaves as follows:

- If you enter the **write standby** command in the system execution space, the system configuration and the configurations for all security contexts on the ASA are written to the peer unit. This includes configuration information for security contexts that are in the standby state. You must enter the command in the system execution space on the unit that has failover group 1 in the active state.



Note If there are security contexts in the active state on the peer unit, the **write standby** command causes active connections through those contexts to be terminated. Use the **failover active** command on the unit providing the configuration to make sure all contexts are active on that unit before entering the **write standby** command.

- If you enter the **write standby** command in a security context, only the configuration for the security context is written to the peer unit. You must enter the command in the security context on the unit where the security context appears in the active state.

Replicated commands are not saved to the flash memory when replicated to the peer unit. They are added to the running configuration. To save replicated commands to flash memory on both units, use the **write memory** or **copy running-config startup-config** command on the unit that you made the changes on. The command is replicated to the peer unit and cause the configuration to be saved to flash memory on the peer unit.

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 51-23](#).)

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 51-23](#).)

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 61-14](#) 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 51-2](#) 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 51-2 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.

Table 51-2 Failover Behavior for Active/Active Failover (continued)

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
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.

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.

ASA failover replication fails if you try to make a configuration change on 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 effect 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 group 1 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 Configuring Active/Active Failover, page 51-8](#)
- [Configuring the Primary Failover Unit, page 51-8](#)
- [Configuring the Secondary Failover Unit, page 51-11](#)

Task Flow for Configuring Active/Active Failover

To configure Active/Active Failover, perform the following steps:

-
- | | |
|---------------|--|
| Step 1 | Configure the primary unit, as shown in the “Configuring the Primary Failover Unit” section on page 51-8 . |
| Step 2 | Configure the secondary unit, as shown in the “Configuring the Secondary Failover Unit” section on page 51-11 . |
| Step 3 | (Optional) Configure optional Active/Active failover settings, as shown in the “Optional Active/Active Failover Settings” section on page 51-6 . |
-

Configuring the Primary Failover Unit

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


Restrictions



Do not configure an IP address for the Stateful Failover link if you are going to use a dedicated Stateful Failover interface. You use the **failover interface ip** command to configure a dedicated Stateful Failover interface in a later step.

Prerequisites

- Configure standby addresses for all IP addresses according to [Chapter 8, “Completing Interface Configuration \(Routed Mode\)”](#) or [Chapter 9, “Completing Interface Configuration \(Transparent Mode\)”](#).
- Complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	failover lan unit primary	Designates the unit as the primary unit.
Step 2	failover lan interface if_name phy_if Example: hostname(config)# failover lan interface folink GigabitEthernet0/3	Specifies the interface to be used as the failover interface. The <i>if_name</i> argument assigns a name to the interface specified by the <i>phy_if</i> argument. The <i>phy_if</i> argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the Stateful Failover link).
Step 3	failover interface ip if_name [ip_address mask standby ip_address ipv6_address/prefix standby ipv6_address] Example: hostname(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2 hostname(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71	Assigns the active and standby IP addresses to the failover link. You can assign either an IPv4 or an IPv6 address to the interface. You cannot assign both types of addresses to the failover link. The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask. The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.
Step 4	failover link if_name phy_if Example: hostname(config)# failover link folink GigabitEthernet0/2	(Optional) Specifies the interface to be used as the Stateful Failover link.  Note If the Stateful Failover link uses the failover link or a data interface, then you only need to supply the <i>if_name</i> argument. The <i>if_name</i> argument assigns a logical name to the interface specified by the <i>phy_if</i> argument. The <i>phy_if</i> argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the failover link).

Command	Purpose
<p>Step 5</p> <pre>failover interface ip if_name [ip_address mask standby ip_address ipv6_address/prefix standby ipv6_address]</pre> <p>Example:</p> <pre>hostname(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2</pre> <pre>hostname(config)# failover interface ip statelink 2001:a1a:b00::a0a:a70/64 standby 2001:a1a:b00::a0a:a71</pre>	<p>(Optional) Assigns an active and standby IP address to the Stateful Failover link. You can assign either an IPv4 or an IPv6 address to the interface. You cannot assign both types of addresses to the Stateful Failover link.</p> <p> Note If the Stateful Failover link uses the failover link or data interface, skip this step. You have already defined the active and standby IP addresses for the interface.</p> <hr/> <p>The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.</p> <p>The Stateful Failover link IP address and MAC address do not change at failover unless it uses a data interface. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.</p>
<p>Step 6</p> <pre>interface phy_if</pre> <pre>no shutdown</pre> <p>Example:</p> <pre>hostname(config)# interface GigabitEthernet 0/3 hostname(config-if)# no shutdown</pre>	<p>Enables the interface.</p> <p> Note If the Stateful Failover link uses the failover link or regular data interface, skip this step. You have already enabled the interface.</p>
<p>Step 7</p> <pre>failover group {1 2} primary secondary</pre> <p>Example:</p> <pre>hostname(config)# failover group 1 hostname(config-fover-group)# primary hostname(config-fover-group)# exit hostname(config)# failover group 2 hostname(config-fover-group)# secondary hostname(config-fover-group)# exit</pre>	<p>Configures the failover groups.</p> <p>You can have only two failover groups. The failover group command creates the specified failover group if it does not exist and enters the failover group configuration mode.</p> <p>For each failover group, specify whether the failover group has primary or secondary preference using the primary or secondary commands. You can assign the same preference to both failover groups. For traffic sharing configurations, you should assign each failover group a different unit preference.</p> <p>The exit command restores global configuration mode.</p> <p>The example assigns failover group 1 as the primary preference and failover group 2 as the secondary preference.</p>
<p>Step 8</p> <pre>context name</pre> <pre>join-failover-group {1 2}</pre> <p>Example:</p> <pre>hostname(config)# context Eng hostname(config-context)# join-failover-group 1 hostname(config-context) exit</pre>	<p>Assigns each user context to a failover group (in context configuration mode).</p> <p>Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.</p>

	Command	Purpose
Step 9	failover Example: hostname(config)# failover	Enables failover.
Step 10	copy running-config startup-config Example: hostname(config)# copy running-config startup-config	Saves the system configuration to flash memory.


Configuring the Secondary Failover Unit

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

Detailed Steps

To configure the secondary failover unit, perform the following steps:

	Command	Purpose
Step 1	failover lan interface <i>if_name</i> <i>phy_if</i> Example: hostname(config)# failover lan interface folink GigabitEthernet0/3	Specifies the interface to be used as the failover interface. The <i>if_name</i> argument assigns a name to the interface specified by the <i>phy_if</i> argument. The <i>phy_if</i> argument can be the physical port name, such as Ethernet1, or a previously created subinterface, such as Ethernet0/2.3. This interface should not be used for any other purpose (except, optionally, the Stateful Failover link).
Step 2	failover interface ip <i>if_name</i> [<i>ip_address</i> <i>mask</i> standby <i>ip_address</i> <i>ipv6_address/prefix</i> standby <i>ipv6_address</i>] Example: hostname(config)# failover interface ip folink 172.27.48.1 255.255.255.0 standby 172.27.48.2 hostname(config)# failover interface ip folink 2001:a0a:b00::a0a:b70/64 standby 2001:a0a:b00::a0a:b71	Assigns the active and standby IP addresses to the failover link. You can assign either an IPv4 or an IPv6 address to the interface. You cannot assign both types of addresses to the failover link. The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask. The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.
Step 3	interface <i>phy_if</i> no shutdown Example: hostname(config-if)# interface GigabitEthernet0/3	Enables the interface.

	Command	Purpose
Step 4	<code>failover lan unit secondary</code> Example: <code>hostname(config)# failover lan unit secondary</code>	(Optional) Designates this unit as the secondary unit:  Note This step is optional because, by default, units are designated as secondary unless previously configured.
Step 5	<code>failover</code> Example: <code>hostname(config)# failover</code>	Enables failover. After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the active unit console.
Step 6	<code>copy running-config startup-config</code> Example: <code>hostname(config)# copy running-config startup-config</code>	Saves the configuration to flash memory. Enter the command after the running configuration has completed replication.
Step 7	<code>no failover active group group_id</code> Example: <code>hostname(config)# no failover active group 1</code>	If necessary, force any failover group that is active on the primary to the active state on the secondary unit. To force a failover group to become active on the secondary unit, enter this command in the system execution space on the primary unit. The <i>group_id</i> argument specifies the group you want to become active on the secondary unit.

Configuring Optional Active/Active Failover Settings

The following optional Active/Active failover settings can be configured when you are initially configuring failover or after you have already established failover. Unless otherwise noted, the commands should be entered on the unit that has failover group 1 in the active state.

This section includes the following topics:


- [Configuring Failover Group Preemption, page 51-12](#)
- [Enabling HTTP Replication with Stateful Failover, page 51-14](#)
- [Disabling and Enabling Interface Monitoring, page 51-14](#)
- [Configuring Interface Health Monitoring, page 51-15](#)
- [Configuring Failover Criteria, page 51-16](#)
- [Configuring Virtual MAC Addresses, page 51-16](#)
- [Configuring Support for Asymmetrically Routed Packets, page 51-18](#)

Configuring Failover Group Preemption

Assigning a primary or secondary priority to a failover group specifies which unit the failover group becomes active on when both units boot simultaneously. However, if one unit boots before the other, then both failover groups become active on that unit. When the other unit comes online, any failover groups that have the unit as a priority do not become active on that unit unless manually forced over, unless a

failover occurs, or unless the failover group is configured with the **preempt** command. The **preempt** command causes a failover group to become active on the designated unit automatically when that unit becomes available.

To configure preemption for the specified failover group, enter the following commands:

	Command	Purpose
Step 1	failover group {1 2}	Specifies the failover group.
	Example: hostname(config)# failover group 1	
Step 2	preempt [delay]	Causes the failover group to become active on the designated unit. You can enter an optional <i>delay</i> value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit. Valid values are from 1 to 1200.
	Example: hostname(config-fover-group)# preempt 1200	
		 Note If Stateful Failover is enabled, the preemption is delayed until the connections are replicated from the unit on which the failover group is currently active.

Example

The following example configures failover group 1 with the primary unit as the higher priority and failover group 2 with the secondary unit as the higher priority. Both failover groups are configured with the **preempt** command with a wait time of 100 seconds, so the groups will automatically become active on their preferred unit 100 seconds after the units become available.

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# mac-address e1 0000.a000.a011 0000.a000.a012
hostname(config-fover-group)# exit
hostname(config)#
```

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information.

You can use the **replication http** command to cause a failover group to replicate HTTP state information when Stateful Failover is enabled.

	Command	Purpose
Step 1	failover group {1 2}	Specifies the failover group.
	Example: hostname(config)# failover group 1	
Step 2	replication http	Enables HTTP state replication for the specified failover group.
	Example: hostname(config-fover-group)# replication http	This command affects only the failover group in which it was configured. To enable HTTP state replication for both failover groups you must enter this command in each group. This command should be entered in the system execution space.

Example

The following example shows a possible configuration for a failover group:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# replication http
hostname(config-fover-group)# exit
```

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 security appliance 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.

In Active/Active failover, this command is only valid within a context.

To enable or disable interface monitoring for specific interfaces, enter one of the following commands:

Do one of the following:

no monitor-interface *if_name*

Disables health monitoring for an interface.

Example:

```
hostname/context (config)#
no monitor-interface 1
```

monitor-interface *if_name*

Enables health monitoring for an interface.

Example:

```
hostname/context (config)#
monitor-interface 1
```

Example

The following example enables monitoring on an interface named “inside”:

```
hostname(config)# monitor-interface inside
hostname(config)#
```

Configuring Interface Health Monitoring

The ASA sends hello packets out of each data interface to monitor interface 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.

To change the default interface poll time, perform the following steps:

	Command	Purpose
Step 1	failover group {1 2}	Specifies the failover group.
	Example: hostname(config)# failover group 1	
Step 2	polltime interface <i>seconds</i>	Specifies the data interface poll and hold times in the Active/Active failover configuration.
	Example: hostname(config-fover-group)# polltime interface seconds	Valid values for the poll time are from 1 to 15 seconds or, if the optional msec keyword is used, from 500 to 999 milliseconds. The hold time determines how long it takes from the time a hello packet is missed to when the interface is marked as failed. Valid values for the hold time are from 5 to 75 seconds. You cannot enter a hold time that is less than 5 times the poll time.

Example

The following partial example shows a possible configuration for a failover group. The interface poll time is set to 500 milliseconds and the hold time to 5 seconds for data interfaces in failover group 1.

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# polltime interface msec 500 holdtime 5
hostname(config-fover-group)# exit
hostname(config)#
```

Configuring Failover Criteria

By default, if a single interface fails, failover occurs. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs. The failover criteria is specified on a failover group basis.

To change the default failover criteria for the specified failover group, perform the following steps:

	Command	Purpose
Step 1	failover group {1 2}	Specifies the failover group.
	Example: hostname(config)# failover group 1	
Step 2	interface-policy num[%]	Specifies the policy for failover when monitoring detects an interface failure.
	Example: hostname(config-fover-group)# interface-policy 225	When specifying a specific number of interfaces, the <i>num</i> argument can be from 1 to 250. When specifying a percentage of interfaces, the <i>num</i> argument can be from 1 to 100.

The following partial example shows a possible configuration for a failover group:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# interface-policy 25%
hostname(config-fover-group)# exit
hostname(config)#
```

Configuring Virtual MAC Addresses

Active/Active failover uses virtual MAC addresses on all interfaces. If you do not specify the virtual MAC addresses, then they are computed 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 for all failover groups.

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

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

To configure specific active and standby MAC addresses for an interface, perform the following steps.

Detailed Steps

	Command	Purpose
Step 1	failover group {1 2}	Specifies the failover group.
	Example: hostname(config)# failover group 1	
Step 2	mac address <i>phy_if</i> <i>active_mac</i> <i>standby_mac</i>	Specifies the virtual MAC addresses for the active and standby units.
	Example: hostname(config-fover-group)# mac address gigabitethernet1/0 0000.a000.a011 0000.a000.a012	The <i>phy_if</i> argument is the physical name of the interface, such as GigabitEthernet1/0. The <i>active_mac</i> and <i>standby_mac</i> arguments are MAC addresses in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The <i>active_mac</i> address is associated with the active IP address for the interface, and the <i>standby_mac</i> is associated with the standby IP address for the interface.

Example

The following partial example shows a possible configuration for a failover group:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# preempt 100
hostname(config-fover-group)# mac address gigabitethernet1/0 0000.a000.a011 0000.a000.a012
hostname(config-fover-group)# exit
hostname(config)#
```

Configuring Support for Asymmetrically Routed Packets

When running in Active/Active failover, a unit may receive a return packet for a connection that originated through its peer unit. Because the ASA that receives the packet does not have any connection information for the packet, the packet is dropped. This most commonly occurs when the two ASAs in an Active/Active failover pair are connected to different service providers and the outbound connection does not use a NAT address.

You can prevent the return packets from being dropped using the **asr-group** command on interfaces where this is likely to occur. When an interface configured with the **asr-group** command receives a packet for which it has no session information, it checks the session information for the other interfaces that are in the same group. If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

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

**Note**

Using the **asr-group** command to configure asymmetric routing support is more secure than using the **static** command with the **nailed** option.

The **asr-group** command does not provide asymmetric routing; it restores asymmetrically routed packets to the correct interface.

Prerequisites

You must have to following configured for asymmetric routing support to function properly:

- Active/Active Failover
- Stateful Failover—Passes state information for sessions on interfaces in the active failover group to the standby failover group.
- Replication HTTP—HTTP session state information is not passed to the standby failover group, and therefore is not present on the standby interface. For the ASA to be able to re-route asymmetrically routed HTTP packets, you need to replicate the HTTP state information.

You can configure the **asr-group** command on an interface without having failover configured, but it does not have any effect until Stateful Failover is enabled.

Detailed Steps

To configure support for asymmetrically routed packets, perform the following steps:

- Step 1** Configure Active/Active Stateful Failover for the failover pair. See the [“Configuring Active/Active Failover”](#) section on page 51-8.
- Step 2** For each interface that you want to participate in asymmetric routing support, enter the following command. You must enter the command on the unit where the context is in the active state so that the command is replicated to the standby failover group. For more information about command replication, see [Command Replication](#), page 51-3.


```
hostname/ctx(config)# interface phy_if
hostname/ctx(config-if)# asr-group num
```

Valid values for *num* range from 1 to 32. You need to enter the command for each interface that participates in the asymmetric routing group. You can view the number of ASR packets transmitted, received, or dropped by an interface using the **show interface detail** command. You can have more than one ASR group configured on the ASA, but only one per interface. Only members of the same ASR group are checked for session information.

Example

Figure 51-1 shows an example of using the **asr-group** command for asymmetric routing support.

Figure 51-1 ASR Example



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

Example 51-1 Primary Unit System Configuration

```
hostname primary
interface GigabitEthernet0/1
description LAN/STATE Failover Interface
interface GigabitEthernet0/2
no shutdown
interface GigabitEthernet0/3
no shutdown
interface GigabitEthernet0/4
no shutdown
interface GigabitEthernet0/5
no shutdown
failover
```

```

failover lan unit primary
failover lan interface folink GigabitEthernet0/1
failover link folink
failover interface ip folink 10.0.4.1 255.255.255.0 standby 10.0.4.11
failover group 1
primary
failover group 2
secondary
admin-context admin
context admin
description admin
allocate-interface GigabitEthernet0/2
allocate-interface GigabitEthernet0/3
config-url flash:/admin.cfg
join-failover-group 1
context ctx1
description context 1
allocate-interface GigabitEthernet0/4
allocate-interface GigabitEthernet0/5
config-url flash:/ctx1.cfg
join-failover-group 2

```

Example 51-2 admin Context Configuration

```

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

```

Example 51-3 ctx1 Context Configuration

```

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

```

Figure 51-1 shows the ASR support working as follows:

1. An outbound session passes through ASA SecAppA. It exits interface outsideISP-A (192.168.1.1).
2. Because of asymmetric routing configured somewhere upstream, the return traffic comes back through the interface outsideISP-B (192.168.2.2) on ASA SecAppB.
3. Normally the return traffic would be dropped because there is no session information for the traffic on interface 192.168.2.2. However, the interface is configured with the command **asr-group 1**. The unit looks for the session on any other interface configured with the same ASR group ID.

4. The session information is found on interface outsideISP-A (192.168.1.2), which is in the standby state on the unit SecAppB. Stateful Failover replicated the session information from SecAppA to SecAppB.
5. Instead of being dropped, the layer 2 header is rewritten with information for interface 192.168.1.1 and the traffic is redirected out of the interface 192.168.1.2, where it can then return through the interface on the unit from which it originated (192.168.1.1 on SecAppA). This forwarding continues as needed until the session ends.

Remote Command Execution

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

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

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

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

To send a command to a failover peer, perform the following steps:

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

If you are in single context mode, skip to the next step.

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

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

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

Commands that cause a command mode change do not change the prompt for the current session. You must use the **show failover exec** command to display the command mode the command is executed in. See [Changing Command Modes, page 51-22](#), for more information.

Changing Command Modes

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

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

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

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

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

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

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

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

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

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

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

Security Considerations

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

Limitations of Remote Command Execution

When you use remote commands you face the following limitations:

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

Controlling Failover

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

- [Forcing Failover, page 51-23](#)
- [Disabling Failover, page 51-24](#)
- [Restoring a Failed Unit or Failover Group, page 51-24](#)

Forcing Failover

Enter the following command in the system execution space of the unit where the failover group is in the standby state:

```
hostname# failover active group group_id
```

Or, enter the following command in the system execution space of the unit where the failover group is in the active state:

```
hostname# no failover active group group_id
```

Entering the following command in the system execution space causes all failover groups to become active:

```
hostname# failover active
```

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, enter the following command:

```
hostname(config)# no failover
```

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, enter the following command:

```
hostname(config)# failover reset
```

To restore a failed Active/Active failover group to an unfailed state, enter the following command:

```
hostname(config)# failover reset group group_id
```

Testing the Failover Functionality

To test failover functionality, perform the following steps:

-
- Step 1** Test that your active unit or failover group is passing traffic as expected by using FTP (for example) to send a file between hosts on different interfaces.
 - Step 2** Force a failover to the standby unit by entering the following command on the unit where the failover group containing the interface connecting your hosts is active:


```
hostname(config)# no failover active group group_id
```
 - Step 3** Use FTP to send another file between the same two hosts.
 - Step 4** If the test was not successful, enter the **show failover** command to check the failover status.

- Step 5** When you are finished, you can restore the unit or failover group to active status by enter the following command on the unit where the failover group containing the interface connecting your hosts is active:

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

Monitoring Active/Active Failover

To monitor Active/Active Failover, perform one of the following tasks. Commands are entered in the system execution space unless otherwise noted.

Command	Purpose
<code>show failover</code>	Displays information about the failover state of the unit.
<code>show failover group</code>	Displays information about the failover state of the failover group. The information displayed is similar to that of the <code>show failover</code> command but limited to the specified group.
<code>show monitor-interface</code>	Displays information about the monitored interface. Enter this command within a security context.
<code>show running-config failover</code>	Displays the failover commands in the running configuration.

For more information about the output of the monitoring commands, see the *Cisco ASA 5500 Series Command Reference*.

Feature History for Active/Active Failover

[Table 51-3](#) lists each feature change and the platform release in which it was implemented.

Table 51-3 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 commands: failover interface ip , show failover , ipv6 address , show monitor-interface .

