



# CHAPTER 12

## Failover

---

This section contains the following topics:

- [Understanding Failover](#)
- [Configuring Failover with the High Availability and Scalability Wizard](#)
- [Field Information for the Failover Panes](#)

## Understanding Failover

The Failover pane contains the settings for configuring failover on the security appliance. However, the Failover pane changes depending upon whether you are in multiple mode or single mode, and when you are in multiple mode, it changes based on the security context you are in.

Failover allows you to configure two security appliances so that one will take over operation if the other fails. Using a pair of security appliances, you can provide high availability with no operator intervention. The security appliance communicates failover information over a dedicated failover link. This failover link can be either a LAN-based connection or, on the PIX security appliance platform, a dedicated serial failover cable. The following information is communicated over the failover link:

- The failover state (active or standby).
- Hello messages (keep-alives).
- Network link status.
- MAC address exchange.
- Configuration replication.



**Caution**

All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any usernames, passwords and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communication with a failover key if you are using the security appliance to terminate VPN tunnels.

The security appliance supports two types of failover, Active/Standby and Active/Active. Additionally, failover can be stateful or stateless. For more information about the types of failover, see the following topics:

- [Active/Standby Failover](#)

- [Active/Active Failover](#)
- [Stateless \(Regular\) Failover](#)
- [Stateful Failover](#)

## Active/Standby Failover

In an Active/Standby configuration, the active security appliance handles all network traffic passing through the failover pair. The standby security appliance does not handle network traffic until a failure occurs on the active security appliance. Whenever the configuration of the active security appliance changes, it sends configuration information over the failover link to the standby security appliance.

When a failover occurs, the standby security appliance becomes the active unit. It assumes the IP and MAC addresses of the previously active unit. Because the other devices on the network do not see any changes in the IP or MAC addresses, ARP entries do not change or time out anywhere on the network.

Active/Standby failover is available to security appliances in single mode or multiple mode.

## Active/Active Failover

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

To enable Active/Active failover on the security appliance, you need to create failover groups. If you enable failover without creating failover groups, you are enabling Active/Standby failover. A failover group is simply a logical group of one or more security contexts. You can create two failover groups on the security appliance. You should create the failover groups on the unit that will have failover group 1 in the active state. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

As in Active/Standby failover, each unit in an Active/Active failover pair is given a primary or secondary designation. Unlike Active/Standby failover, this designation does not indicate which unit is active when both units start simultaneously. Each failover group in the configuration is given a primary or secondary role preference. This preference determines on which unit in the failover pair the contexts in the failover group appear in the active state when both units start simultaneously. 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.

Initial configuration synchronization occurs when one or both units start. This synchronization occurs as follows:

- When both units start simultaneously, the configuration is synchronized from the primary unit to the secondary unit.
- When one unit starts while the other unit is already active, the unit that is starting up receives the configuration from the already active unit.

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 will cause the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

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, 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.

---

## Stateless (Regular) Failover

Stateless failover is also referred to as regular failover. In stateless failover, all active connections are dropped when a failover occurs. Clients need to reestablish connections when the new active unit takes over.

## Stateful Failover

**Note**

---

Stateful Failover is not supported on the ASA 5505 series adaptive security appliance.

---

When Stateful Failover is enabled, the active unit in the failover pair continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

**Note**

---

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

---

To use Stateful Failover, you must configure a state link to pass all state information to the standby unit. If you are using a LAN failover connection rather than the serial failover interface (available on the PIX security appliance platform only), you can use the same interface for the state link as the failover link. However, we recommend that you use a dedicated interface for passing state information the standby unit.

The following information is passed to the standby unit when Stateful Failover is enabled:

- NAT translation table.
- TCP connection table (except for HTTP), including the timeout connection.
- HTTP connection states (if HTTP replication is enabled).
- H.323, SIP, and MGCP UDP media connections.
- The system clock.

- The ISAKMP and IPsec SA table.

The following information is not copied to the standby unit when Stateful Failover is enabled:

- HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- The ARP table.
- Routing tables.

## Configuring Failover with the High Availability and Scalability Wizard

The High Availability and Scalability Wizard steps you through the process of creating an Active/Active failover configuration, and Active/Standby failover configuration, or a VPN Cluster Load Balancing configuration.

See the following topics for information about using the High Availability and Scalability Wizard:

- [Accessing and Using the High Availability and Scalability Wizard](#)
- [Configuring Active/Active Failover with the High Availability and Scalability Wizard](#)
- [Configuring Active/Standby Failover with the High Availability and Scalability Wizard](#)
- [Configuring VPN Load Balancing with the High Availability and Scalability Wizard](#)
- [Field Information for the High Availability and Scalability Wizard](#)

## Accessing and Using the High Availability and Scalability Wizard

To open the High Availability and Scalability Wizard, choose **Wizards > High Availability and Scalability Wizard** from the ASDM menu bar. The first screen of the wizard appears.

To move to the next screen of the wizard, click the **Next** button. You must complete the mandatory field of each screen before you can move to the next screen.

To move to a previous screen of the wizard, click the **Back** button. If information filled in on later screens of the wizard is not affected by the change you make to an earlier screen, that information remains on the screen as you move forward through the wizard again. You do not need to reenter it.

To leave the wizard at any time without saving any changes, click **Cancel**.

To send your configuration to the security appliance at the end of the wizard, click **Finish**.

## Configuring Active/Active Failover with 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 with a wizard screen. Click **Next** after completing each step, except for the last step, before moving to the next step. Each step also contains a reference to additional information that you may need to complete the step.

- 
- Step 1** Choose **Configure Active/Active** failover on the Choose the type of failover configuration screen.

- See [Choose the Type of Failover Configuration](#) for more information about this screen.
- Step 2** Enter the IP address of the failover peer on the Check Failover Peer Connectivity and Compatibility screen. Click **Test Compatibility**. You will not be able to move to the next screen until all compatibility tests are passed.
- See [Check Failover Peer Connectivity and Compatibility](#) for more information about this screen.
- Step 3** If the security appliance or the failover peer are in single context mode, change them to multiple context mode on the Change Device to Multiple Mode screen. When you change the security appliance to multiple context mode, it will reboot. ASDM automatically reestablishes communication with the security appliance when it has finished rebooting.
- See [Change Device to Multiple Mode](#) for more information about this screen.
- Step 4** (PIX 500 series security appliance only) Select cable-based or LAN-based failover on the Select Failover Communication Media screen.
- See [Select Failover Communication Media](#) for more information about this screen.
- Step 5** Assign security contexts to failover groups on the Context Configuration screen. You can add and delete contexts on this screen.
- See [Security Context Configuration](#) for more information about this screen.
- Step 6** Define the Failover Link on the Failover Link Configuration screen.
- See [Failover Link Configuration](#) for more information about this screen.
- Step 7** (Not available on the ASA 5505 security appliance) Define the Stateful Failover link on the State Link Configuration screen.
- See [State Link Configuration](#) for more information about this screen.
- Step 8** Add standby addresses to the security appliance interfaces on the Standby Address Configuration screen.
- See [Standby Address Configuration](#) for more information about this screen.
- Step 9** Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.
- See [Summary](#) for more information about this screen.
- Step 10** Click **Finish**.
- The failover configuration is sent to the security appliance and to the failover peer.
- 

## Configuring Active/Standby Failover with the High Availability and Scalability Wizard

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

---

- Step 1** Choose **Configure Active/Standby** failover on the Choose the type of failover configuration screen. Click next.
- See [Choose the Type of Failover Configuration](#) for more information about this screen.

- Step 2** Enter the IP address of the failover peer on the Check Failover Peer Connectivity and Compatibility screen. Click **Test Compatibility**. You will not be able to move to the next screen until all compatibility tests are passed.
- See [Check Failover Peer Connectivity and Compatibility](#) for more information about this screen.
- Step 3** (PIX 500 series security appliance only) Select cable-based or LAN-based failover on the Select Failover Communication Media screen.
- See [Select Failover Communication Media](#) for more information about this screen.
- Step 4** Define the Failover Link on the Failover Link Configuration screen.
- See [Failover Link Configuration](#) for more information about this screen.
- Step 5** (Not available on the ASA 5505 security appliance) Define the Stateful Failover link on the State Link Configuration screen.
- See [State Link Configuration](#) for more information about this screen.
- Step 6** Add standby addresses to the security appliance interfaces on the Standby Address Configuration screen.
- See [Standby Address Configuration](#) for more information about this screen.
- Step 7** Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.
- See [Summary](#) for more information about this screen.
- Step 8** Click **Finish**.
- The failover configuration is sent to the security appliance and to the failover peer.
- 

## Configuring VPN Load Balancing with the High Availability and Scalability Wizard

The following procedure provides a high-level overview for configuring VPN cluster load balancing using the High Availability and Scalability Wizard. Each step in the procedure corresponds with a wizard screen. Click **Next** after completing each step, except for the last step, before moving to the next step. Each step also contains a reference to additional information that you may need to complete the step.

- Step 1** Choose **Configure VPN Cluster Load Balancing** failover on the Choose the type of failover configuration screen.
- See [Choose the Type of Failover Configuration](#) for more information about this screen.
- Step 2** Configure the VPN load balancing settings on the VPN Cluster Load Balancing Configuration screen.
- See [VPN Cluster Load Balancing Configuration](#) for more information about this screen.
- Step 3** Review your configuration on the Summary screen. If necessary, use the Back button to go to a previous screen and make changes.
- See [Summary](#) for more information about this screen.
- Step 4** Click **Finish**.
- The failover configuration is sent to the security appliance and to the failover peer.
-

## Field Information for the High Availability and Scalability Wizard

The following dialogs are available in the High Availability and Scalability Wizard. You will not see every dialog box when you run through the wizard; each dialog box appears depending on the type of failover you are configuring and the hardware platform you are configuring it on.

- [Choose the Type of Failover Configuration](#)
- [Check Failover Peer Connectivity and Compatibility](#)
- [Change Device to Multiple Mode](#)
- [Security Context Configuration](#)
- [Failover Link Configuration](#)
- [State Link Configuration](#)
- [Standby Address Configuration](#)
- [VPN Cluster Load Balancing Configuration](#)
- [Summary](#)

### Choose the Type of Failover Configuration

The Choose the Type of Failover Configuration screen lets you select the type of failover to configure.

#### Fields

The Choose the Type of Failover Configuration displays the following informational fields. These are useful for determining the failover capabilities of the security appliance.

- **Hardware Model**—(*Display only*) Displays the security appliance model number.
- **No. of Interfaces**—(*Display only*) Displays the number of interfaces available on the security appliance.
- **No. of Modules**—(*Display only*) Displays the number of modules installed on the security appliance.
- **Software Version**—(*Display only*) Displays the version of the platform software on the security appliance.
- **Failover License**—(*Display only*) Displays the type of failover license installed on the device. You may need to purchase an upgraded license to configure failover.
- **Firewall Mode**—(*Display only*) Displays the firewall mode (routed or transparent) and the context mode (single or multiple).

Choose the type of failover configuration you are configuring:

- **Configure Active/Active Failover**—Configures the security appliance for Active/Active failover.
- **Configure Active/Standby Failover**—Configures the security appliance for Active/Standby failover.
- **Configure VPN Cluster Load Balancing**—Configures the security appliance to participate in VPN load balancing as part of a cluster.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
<b>Routed</b>	<b>Transparent</b>	<b>Single</b>	<b>Context</b>	<b>System</b>
•	•	•	—	•

## Check Failover Peer Connectivity and Compatibility

The Check Failover Peer Connectivity and Compatibility 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.

### Fields

- Peer IP Address—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.
- Test Compatibility—Click this button 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
  - Software version compatibility
  - Failover license compatibility
  - Firewall mode compatibility

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
<b>Routed</b>	<b>Transparent</b>	<b>Single</b>	<b>Context</b>	<b>System</b>
•	•	•	—	•

## Change Device to Multiple Mode

The Change Device to Multiple Mode dialog box appears for Active/Active failover configuration only. Active/Active failover requires the security appliance to be in multiple context mode. This dialog box lets you convert a security appliance in single context mode to multiple context mode.

When you convert from single context mode to multiple context mode, the security appliance 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 security appliance from single context mode to multiple context mode causes the security appliance 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.

You need to convert both the current security appliance and the peer security appliance to multiple context mode before you can proceed.

#### Fields

- Change *device* To Multiple Context—Causes the security appliance to change to multiple context mode. *device* is the hostname of the security appliance.
- Change *device* (peer) To Multiple Context—Causes the peer unit to change to multiple context mode. *device* is the hostname of the security appliance.
- Device Status—(*Display only*) Displays the status of the security appliance while converting to multiple context mode.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

## Select Failover Communication Media

The Select Failover Communication Media appears only on PIX 500 series security appliances. This screen lets you select between using a failover cable or LAN-based connection for the failover link.

#### Fields

- Use Failover Cable—Choose this option to use a dedicated failover cable for failover communication.
- Use LAN-based connection—Choose this option to use a network connection for failover communication.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

## Security Context Configuration

The Security Context Configuration screen appears for Active/Active configuration only. The Security Context Configuration screen lets you assign security contexts to failover groups. It displays the security contexts currently configured on the device and lets you add new ones or remove existing ones as needed.

Although you can create security contexts on this screen, you cannot assign interfaces to those contexts or configure any other properties for them. To configure context properties and assign interfaces to a context, you need to use the System > Security Contexts pane.

#### Fields

- **Name**—Displays the name of the security context. To change the name, click the name and type a new name.
- **Failover Group**—Displays the failover group the context is assigned to. To change the failover group for a security context, click the failover group and select the new failover group number from the drop-down list.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

## Failover Link Configuration

The Failover Link Configuration screen only appears if you are configuring LAN-based failover; it does not appear if you are configuring a PIX 500 series security appliance for cable-based failover.

#### Fields

- **LAN Interface**—Choose the interface to use for failover communication from the drop-down list.
- **Logical Name**—Type a name for the interface.
- **Active IP**—Type the IP address used for the failover link on the unit that has failover group 1 in the active state.
- **Standby IP**—Type the IP address used for the failover link on the unit that has failover group 1 in the standby state.
- **Subnet Mask**—Type or select a subnet mask for the Active IP and Standby IP addresses.
- **Secret Key**—(Optional) Enter the key used to encrypt failover communication. If this field is left blank, failover communication, including any passwords or keys in the configuration sent during command replication, is in clear text.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

## State Link Configuration

The State Link Configuration screen does not appear in the wizard for ASDM running on the ASA 5505 platform.

The State Link Configuration lets you enable Stateful Failover and configure the Stateful Failover link properties.

### Fields

- Use the LAN link as the State Link—Choose this option to pass state information across the LAN-based failover link. This option is not available on PIX 500 series security appliances configured for cable-based failover.
- Disable Stateful Failover—Choose this option to disable Stateful Failover.
- Configure another interface for Stateful failover—Choose this option to configure an unused interface as the Stateful Failover interface.
  - State Interface—Choose the interface you want to use for Stateful Failover communication from the drop-down list.
  - Logical Name—Type the name for the Stateful Failover interface.
  - Active IP—Type the IP address for the Stateful Failover link on the unit that has failover group 1 in the active state.
  - Standby IP—Type the IP address for the Stateful Failover link on the unit that has failover group 1 in the standby state.
  - Subnet Mask—Type or select a subnet mask for the Active IP and Standby IP addresses.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

## Standby Address Configuration

Use the Standby Address Configuration screen to assign standby addresses to the interface on the security appliance.

### Fields

- Device/Interface—(Active/Standby failover) Displays the interfaces configured on the failover units. Click the plus sign (+) by a device name to displays the interfaces on that device. Click the minus sign (-) by a device name to hides the interfaces on that device.
- Device/Group/Context/Interface—(Active/Active failover) Displays the interfaces configured on the failover unit. The interfaces are grouped by context and the contexts are grouped by failover group. 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.

- **Active IP**—Double-click this 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 peer unit.
- **Standby IP**—Double-click this 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 peer unit.
- **Is Monitored**—Check this check box to enable health monitoring for that interface. Uncheck the check box to disable the health monitoring. By default, health monitoring of physical interfaces is enabled and health monitoring of virtual interfaces is disabled.
- **ASR Group**—Select the asynchronous group ID from the drop-down list. This setting is only available for physical interface. For virtual interfaces this field displays “None”.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			<b>Multiple</b>	
<b>Routed</b>	<b>Transparent</b>	<b>Single</b>	<b>Context</b>	<b>System</b>
•	•	•	—	•

## VPN Cluster Load Balancing Configuration

If you have a remote-client configuration in which you are using two or more security appliances connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called *load balancing*. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. It makes efficient use of system resources and provides increased performance and availability.

Use the VPN Cluster Load Balancing Configuration screen to set parameters necessary for this device to participate in a load balancing cluster.



### Note

VPN load balancing runs only on security appliance models ASA 5520 and higher.

VPN load balancing requires an active 3DES/AES license. The security appliance checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the security appliance prevents the enabling of load balancing and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

Enabling load balancing involves:

- Configuring the load-balancing cluster by establishing a common virtual cluster IP address, UDP port (if necessary), and IPsec shared secret for the cluster. These values are identical for every device in the cluster.
- Configuring the participating device by enabling load balancing on the device and defining device-specific properties. These values vary from device to device.

**Note**

Load balancing is effective only on remote sessions initiated with the Cisco VPN Client (Release 3.0 and later), the Cisco VPN 3002 Hardware Client (Release 3.5 and later), or the ASA 5505 operating as an Easy VPN Client. All other clients, including LAN-to-LAN connections, can connect to a security appliance on which load balancing is enabled, but they cannot participate in load balancing.

To implement load balancing, you group together logically two or more devices on the same private LAN-to-LAN network into a *virtual cluster*.

**Fields**

- **VPN Load Balancing**—Configures virtual cluster device parameters.
  - **Cluster IP Address**—Specifies the single IP address that represents the entire virtual cluster. Choose an IP address that is within the public subnet address range shared by all the security appliances in the virtual cluster.
  - **Cluster UDP Port**—Specifies the UDP port for the virtual cluster in which this device is participating. The default value is 9023. If another application is using this port, enter the UDP destination port number you want to use for load balancing.
  - **Enable IPsec Encryption**—Enables or disables IPsec encryption. If you select this check box, you must also specify and verify a shared secret. The security appliances in the virtual cluster communicate via LAN-to-LAN tunnels using IPsec. To ensure that all load-balancing information communicated between the devices is encrypted, select this check box.

**Note**

When using encryption, you must have previously configured the load-balancing inside interface. If that interface is not enabled on the load-balancing inside interface, you get an error message when you try to configure cluster encryption.

If the load-balancing inside interface is enabled when you configured cluster encryption, but is disabled before you configure the participation of the device in the virtual cluster, you get an error message when you select the Participate in Load Balancing Cluster check box, and encryption is not enabled for the cluster.

- **Shared Secret Key**—Specifies the shared secret to between IPsec peers when you enable IPsec encryption. The value you enter in the box appears as consecutive asterisk characters.
- **Priority Of This Device**—Specifies the priority assigned to this device within the cluster. The range is from 1 to 10. The priority indicates the likelihood of this device becoming the virtual cluster master, either at start-up or when an existing master fails. The higher you set the priority (for example, 10), the more likely this device becomes the virtual cluster master.
- **Public Interface Of This Device**—Specifies the name or IP address of the public interface for this device.
- **Private Interface Of This Device**—Specifies the name or IP address of the private interface for this device.

**Note**

If the devices in the virtual cluster are powered up at different times, the first device to be powered up assumes the role of virtual cluster master. Because every virtual cluster requires a master, each device in the virtual cluster checks when it is powered-up to ensure that the cluster has a virtual master. If none exists, that device takes on the role. Devices powered up and added to the cluster later become secondary devices. If all the devices in the virtual cluster are powered up simultaneously, the device with the highest priority setting becomes the virtual cluster master. If two or more devices in the virtual cluster are powered up simultaneously, and both have the highest priority setting, the one with the lowest IP address becomes the virtual cluster master.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

**Summary**

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

**Fields**

The configuration appears in the center of the screen. 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** until you reach the screen where you need to make the change. Make the change and click **Next** until you return to the Summary screen.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	•

**Field Information for the Failover Panes**

What displays on the failover pane depends upon the mode you are in (single or multiple context mode) and whether you are in the system execution space or in a security context.

This section contains the following topics:

- [Failover - Single Mode](#)
- [Failover-Multiple Mode, Security Context](#)

- [Failover-Multiple Mode, System](#)

## Failover - Single Mode

The Failover pane contains the tabs where you can configure Active/Standby failover in single context mode. For more information about failover, see [Understanding Failover](#). For more information about configuring the settings on each tab of the Failover pane, see the following information. Note that the Interfaces tabs changes based on whether you are in routed firewall mode or transparent firewall mode.

- [Failover: Setup](#)
- [Failover: Interfaces \(Routed Firewall Mode\)](#)
- [Failover: Interfaces \(Transparent Firewall Mode\)](#)
- [Failover: Criteria](#)
- [Failover: MAC Addresses](#)

### Failover: Setup

Use this tab to enable failover on the security appliance. You also designate the failover link and the state link, if using Stateful Failover, on this tab.

For more information about configuring failover in general, see [Understanding Failover](#).

#### Fields

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



#### Note

---

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

---

ASDM displays a dialog box asking if you want to configure the peer unit when you enable failover. This dialog box also appears when the Preferred Role setting or, on the PIX security appliance platform, the Enable LAN rather than serial cable failover setting changes.

- **Peer IP Address**—Enter an IP address on the peer unit that ASDM can connect to. This field appears on the Do you want to configure the failover peer firewall dialog box.
- **Use 32 hexadecimal character key**—Check this check box to enter a hexadecimal value for the encryption key in the Shared Key box. Uncheck this check box to enter an alphanumeric shared secret in the Shared Key box.
- **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 unchecked 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.

- Enable LAN rather than serial cable failover—(PIX security appliance platform only) Check this check box to enable LAN Failover. Uncheck this check box to use the dedicated serial cable as the failover link.
- 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 share the interface with Stateful Failover.  
Only unconfigured interfaces or subinterfaces are displayed in this list and can be selected as the LAN Failover interface. Once you specify an interface as the LAN Failover interface, you cannot edit that interface in the Configuration > Interfaces pane.
  - Active IP—Specifies the IP address for the failover interface on the active unit.
  - Subnet Mask—Specifies the mask 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
  - Preferred Role—Specifies whether the preferred role for this security appliance is as the primary or secondary unit in a LAN failover.
- State Failover—Contains the fields for configuring Stateful Failover.




---

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

---

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




---

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

---

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

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

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




---

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

---

- 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 selected in the Interface drop-down list.

- Subnet Mask—Specifies the mask 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 selected 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 selected in 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 selected in the Interface drop-down list.
- Enable HTTP replication—Selecting 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## Failover: Interfaces (Routed Firewall Mode)

Use this tab to define the standby IP address for each interface on the security appliance and to specify whether the status of the interface should be monitored.

For more information about configuring failover in general, see [Understanding Failover](#).

### Fields

- Interface—Lists the interfaces on the security appliance 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 \(Routed Firewall Mode\)](#) dialog box for the selected interface.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

**For More Information**

For more information about failover in general, see [Understanding Failover](#).

**Edit Failover Interface Configuration (Routed Firewall Mode)**

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—Identifies the mask 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. 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.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

**For More Information**

For more information about failover in general, see [Understanding Failover](#).

## Failover: Interfaces (Transparent Firewall Mode)

Use this tab to define the standby management IP address and to specify whether the status of the interfaces on the security appliance should be monitored.

### Fields

- Interface—Lists the interfaces on the security appliance and identifies their monitoring status.
  - Interface Name column—Identifies the interface name.
  - Is Monitored column—Specifies whether this interface is monitored for failure.
- Edit—Displays the [Edit Failover Interface Configuration \(Transparent Firewall Mode\)](#) dialog box for the selected interface.
- Management IP Address—Identifies the active and standby management IP addresses for the security appliance or for a context in transparent firewall mode.
  - Active—Identifies the active management IP address.
  - Standby—Specifies the management IP address on the standby failover unit.
- Management Netmask—Identifies the mask associated with the active and standby management IP addresses.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
—	•	•	—	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## Edit Failover Interface Configuration (Transparent Firewall Mode)

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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
—	•	•	—	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## Failover: Criteria

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.

### 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 security appliance 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 security appliance 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 1 and 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 1 and 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## Failover: MAC Addresses

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



#### Note

This tab is not available on the ASA 5505 platform.

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

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

### Fields

- **MAC Addresses**—Lists physical interfaces on the security appliance for which an active and standby virtual MAC address has been configured.
  - **Physical Interface column**—Identifies the physical interface for which failover virtual MAC addresses are configured.
  - **Active MAC Address column**—Identifies the MAC address of the active security appliance (usually primary).
  - **Standby MAC Address column**—Identifies the MAC address of the standby security appliance (usually secondary).
- **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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	—	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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 security appliance (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 security appliance (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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 contains 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 security appliance 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.

#### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	—	•	—

#### For More Information

For more information about failover in general, see [Understanding Failover](#).

### 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—Identifies the mask 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	—	•	—

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
—	•	—	•	—

**For More Information**

For more information about failover in general, see [Understanding Failover](#).

**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.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
—	•	—	•	—

**For More Information**

For more information about failover in general, see [Understanding Failover](#).

## Failover-Multiple Mode, System

This pane includes tabs for configuring the system-level failover settings in the system context of a security appliance 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 [Understanding Failover](#).

See the 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 security appliance in multiple context mode. You also designate the failover link and the state link, if using Stateful Failover, on this tab.

#### Fields

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



---

**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 characters are any combination of numbers, letters, or punctuation. The shared secret is used to generate the encryption key.

- **Enable LAN rather than serial cable failover**—(PIX security appliance platform only) Check this check box to enable LAN failover. Uncheck this check box to use the dedicated serial link as the failover link.
- **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.
- Subnet Mask—Specifies the mask for the failover interface on the active unit.
- Logical Name—Specifies the logical name for the failover interface.
- Standby IP—Specifies the IP address of the standby unit.
- Preferred Role—Specifies whether the preferred role for this security appliance 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 subinterface 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 active unit.
- Subnet Mask—Specifies the mask for the Stateful Failover interface on the active unit.
- Logical Name—Specifies the logical name for the Stateful Failover interface.
- Standby IP—Specifies the IP address of the standby unit.
- 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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 security appliance 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 security appliance 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 1 and 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 1 and 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## Failover > Active/Active Tab

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

A failover group is simply a logical group of security contexts. You can create two failover groups on the security appliance. 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**

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 security appliance.
  - 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.

**Modes**

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

**For More Information**

For more information about failover in general, see [Understanding Failover](#).

**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 security appliance 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 security appliance 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 security appliance 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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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\_id01.
- Standby unit default MAC address: 00a0.c9: $physical\_port\_number$ .failover\_group\_id02.



### 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).

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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 security appliance 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 security appliance (usually primary).

- Standby MAC Address—Identifies the MAC address on the standby security appliance (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.

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

### For More Information

For more information about failover in general, see [Understanding Failover](#).

## 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 security appliance (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 security appliance (usually secondary). Enter the MAC address in hexadecimal format (for example, 0123.4567.89AB).

### Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	—	•

**For More Information**

For more information about failover in general, see [Understanding Failover](#).