Configuring Failover

The Failover page provides access to failover settings for the selected security appliance. The available settings and the overall appearance of the Failover page may change slightly, depending upon the type of device selected, its mode of operation (routed or transparent), and its context mode (single or multiple).

In other words, how you configure failover depends upon both the operating mode and the security context of the security appliance.

Please note the following caveats when assigning an interface as a failover link:

- You can define the interface in the Add/Edit Interface dialog box, but do not configure it. In particular, **do not specify an interface Name**, as this parameter disqualifies the interface from being used as the failover link. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.
- IPv6 addresses are not supported for failover links.
- On an ASA 5505, an interface assigned as the backup for another interface cannot be used as a failover link (although no checking is performed to prevent this).
- Do not assign a PPPoE-enabled interface as a failover link. PPPoE and Failover should not be configured on the same device interface (although no checking is performed to prevent this).
- A failover interface cannot use the same IP address as another interface, especially the Management IP address (although no checking is performed to prevent this).

Note also that after you assign an interface as a failover link, the interface is listed on the Interfaces page, but you cannot edit or delete the interface from that page. The only exception is if you set a physical interface to be the stateful failover link—you can configure its speed and duplex.

This chapter contains the following topics:

Understanding Failover

Failover lets you configure two identical security appliances such that one will take over firewall operations if the other fails. Using a pair of security appliances, you can provide high system availability without operator intervention.

The linked security appliances communicate failover information over a dedicated link. This failover link can be either a LAN-based connection or, on PIX security appliances, a dedicated serial failover cable. The following information is communicated over the failover link:

- Current failover state (active or standby)
- "Hello" messages (also called "keep-alives")
- Network link status
- MAC address exchange
- Configuration replication
- Per-connection state information, in the case of Stateful failover

Caution All information sent over the failover link 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 user names, 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 failover communications with a failover key, particularly if you are using the security appliance to terminate VPN tunnels.

Cisco security appliances support two types of failover:

- **Active/Standby** – The *active* security appliance inspects all network traffic, while the *standby* security appliance remains idle until a failure occurs on the active appliance. Changes to the configuration of the active security appliance are transmitted over the failover link to the standby security appliance.

  When failover occurs, the standby security appliance becomes the active unit, and it assumes the IP and MAC addresses of the previously active unit. Because 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 operating in single- or multiple-context mode. In single-context mode, only Active/Standby failover is available, and all failover configuration is by means of the Failover page.

  Note When using Active/Standby failover, you must make all configuration changes on the active unit. The active unit automatically replicates the changes to the standby unit. The standby unit should not be imported or added to the Security Manager device list.

  Also, you must manually copy the authentication certificate from the active device to the standby device. See Additional Steps for an Active/Standby Failover Configuration, page 50-9 for additional information.

- **Active/Active** – Both security appliances inspect network traffic by alternating their roles—such that one is active and one is standby—on a per context basis. This means Active/Active failover is available only on security appliances operating in multiple-context mode.

  However, Active/Active failover is not required in multiple-context mode. That is, on a device operating in multiple-context mode, you can configure Active/Standby or Active/Active failover. In either case, you provide system-level failover settings in the system context, and context-level failover settings in the individual security contexts.

  See Active/Active Failover, page 50-3 for additional information about this topic.

In addition, failover can be stateless or stateful:

- **Stateless** – Also referred to as "regular" failover. With stateless failover, all active connections are dropped when failover occurs. Clients need to re-establish connections when the new active unit takes over.
- **Stateful** – The active unit in the failover pair continually passes per-connection state information to the standby unit. When failover occurs, the same connection information is available on the new active unit. Supported end-user applications are not required to reconnect to maintain the current communication session.

  See Stateful Failover, page 50-4 for more information.

Related Topics

- Chapter 50, "Configuring Failover"
- Basic Failover Configuration, page 50-5
- Failover Policies, page 50-10

Active/Active Failover

Active/Active failover is available only on security appliances operating in multiple-context mode. In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device.

The active and standby roles are assigned over the complete set of security contexts, more or less arbitrarily.

To enable Active/Active failover on the security appliance, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments 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 the unit on which 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.

Note To reliably manage security contexts in Active/Active failover mode, Cisco Security Manager requires an IP address for the management interface of each context so that it can communicate directly with the active security context of a failover pair.

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 is 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.

Note When bootstrapping the peer devices in an Active/Active Failover configuration, the bootstrap configurations are only applied to the system contexts of the respective failover peer devices.

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.

Stateful Failover

Note Stateful failover is not supported on the ASA 5505 appliance.

When stateful failover is enabled, the active unit in the failover pair continually updates the current connection-state information on the standby unit. When failover occurs, supported end-user applications are not required to reconnect to maintain the current communication session.

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

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

Basic Failover Configuration

The following steps describe basic failover configuration. Please note the following caveats when assigning an interface as a failover link:

- You can define the interface in the Add/Edit Interface dialog box, but do not configure it. In particular, **do not specify an interface Name**, as this parameter disqualifies the interface from being used as the failover link.
- On an ASA 5505, an interface assigned as the backup for another interface cannot be used as a failover link (although no checking is performed to prevent this).
- Do not assign a PPPoE-enabled interface as a failover link. PPPoE and Failover should not be configured on the same device interface (although no checking is performed to prevent this).
- A Failover interface cannot use the same IP address as another interface, especially the Management IP address (although no checking is performed to prevent this).

Note When you save a failover configuration, it is applied to both the security appliance and the failover peer.

Before You Begin

Licenses installed on the device must allow failover configurations. On ASA 5505 and 5510 devices, this failover license is an optional license. You must install the failover license outside of Security Manager, using ASDM or the device CLI, and ensure that the **License Supports Failover** option is selected in the General page of the device properties (right-click the device and select **Device Properties**). If the license is installed when you add the device to the inventory, or you install the license and then rediscover device policies, Security Manager can identify the license and set this option appropriately.

If the option is selected and the license is not in fact installed, you will see deployment failures. If the option is not selected, Security Manager will not deploy the failover policy to the device even if you configure the policy.

Related Topics

- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26
- Understanding Failover, page 50-1
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

- Failover Policies, page 50-10

Step 1 Ensure Device View is your present application view; if necessary, click the **Device View** button on the toolbar.

Note For more information on using the Device View to configure device policies, see Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30.

Step 2 Select the appliance you want to configure.

Step 3 Expand the **Platform** entry in the Device Policy selector, then expand **Device Admin** and select **Failover**.

The Failover page is displayed.

Step 4 (PIX only) Choose the **Failover Method**: **Serial Cable** or **LAN Based**. If you choose Serial Cable, the LAN Failover settings are disabled; be sure the cable connecting the two devices is in place.

Step 5 Select **Enable Failover** to enable failover on this appliance.

Step 6 (Optional) Click the Settings button to open the Settings dialog box for the selected device. The contents of the Settings dialog box depend on the type of device, and whether it is operating in single or multiple mode—some options may not be available. Refer to the following sections:

- Settings Dialog Box, page 50-21 (ASA/PIX 7+)
- Advanced Settings Dialog Box, page 50-16 (FWSM)

Step 7 Click the **Bootstrap** button to open the Bootstrap configuration for LAN failover dialog box, which provides bootstrap configurations that can be applied to the primary and secondary devices in a LAN failover configuration. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.

Step 8 (Multiple-context devices only) In the Configuration section, select the failover mode: **Active/Active** or **Active/Standby**.

Step 9 (Optional) Follow these steps to configure an interface for **LAN Failover** communications between the two devices:

a. Assign a device **Interface** for LAN-based communications, and then press the Tab key on your keyboard to update the page.

On PIX and ASA devices, this drop-down list displays the interfaces defined on the device. You can type in a port ID (e.g., *gigabitethernet1*), or you can choose the port if you have already defined the interface.

On an FWSM, the Interface list is not populated with VLAN IDs; you must enter the numeric ID of the VLAN you wish to use.

Note In both cases, this cannot be a Named interface, nor can the interface be configured for PPPoE.

b. Provide a **Logical Name** for this failover interface.

c. Enter the **Active IP** address for failover communications.

d. Enter a **Standby IP** address for failover communications. The Standby IP address is used on the security appliance that is currently the standby unit.

e. Enter the **Subnet Mask** for both IP addresses. Both must be on the same subnet.

Step 10 (Optional) Follow these steps to enable and configure an interface for **Stateful Failover** communications between the two devices:

a. Assign a device **Interface** for update communications, and then press the Tab key on your keyboard to update the page.

You can type in a port ID (e.g., *gigabitethernet1*), or you can choose the port if you have already defined the interface; note that this cannot be a Named interface.

Note On an FWSM, this is a **VLAN** interface.

b. Provide a **Logical Name** for this interface.

c. Enter the **Active IP** address for connection updates.

d. Enter a **Standby IP** address for update communications.

e. Enter the **Subnet Mask** for both IP addresses. Both must be on the same subnet.

f. Select **Enable HTTP Replication** to preserve HTTP connection information.

Connection information is communicated to the standby unit for all TCP protocols except HTTP, because HTTP connections are generally short-lived. Select this option to maintain HTTP connections during failover.

Step 11 Provide a communications-encryption key: enter a **Shared Key** and then repeat it in the **Confirm** field. Be sure to enter the same key on both devices. (Not available on FWSM versions prior to 3.1)

The Shared Key can be any arbitrary string of up to 63 alphanumeric characters. If **HEX** is checked, the Shared Key is an arbitrary string of exactly 32 hexadecimal characters. (The HEX option is available only on PIX/ASA version 7.0.5 and later, and FWSM versions 3.1.3 and later.)

Note This step is optional, but we strongly recommend encrypting failover communications.

Step 12 To specify a failover reconnect timeout value for asymmetrically routed sessions, enter a length of time in the **Timeout** field, in the form hh:mm:ss (the minutes and seconds values are optional). If the field is blank (the default), or contains a zero, reconnections are prevented. Setting this value to -1 disables the timeout, allowing connections to reconnect after any amount of time.

Step 13 (Optional) You can configure Bidirectional Forwarding Detection (BFD), to communicate with a failover pair and this can be used to monitor the health of the failover unit. Create or select a **BFD template** from the **Health-Check Monitoring** section.

Note This is applicable only for Firepower failover devices running ASA 9.7.1 and above.

Tip BFD failover commands are supported only in the Active/Standby mode.
In a multi-context device, BFD failover commands are supported only in the system context.
BFD Failover commands are not supported in the transparent mode.

Step 14 (FWSM only) – Configured interfaces are listed in the Interface Configuration table. To edit the failover configuration for a listed interface, select it and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23.

Adding A Security Context to Failover Group 2

To add a new security context to an existing failover group 2, you must save the new context configuration to a deployment file and then manually add it to the appropriate device. Otherwise, until the first successful deployment, Security Manager will attempt to communicate with the new context through the device's Admin context. This will fail since group 2 cannot be reached through the Admin context (unless both group 1 and 2 are active on the same device).

The following steps outline creating a new security context and adding it to failover group 2.

1. Create the new security context.

   Be sure to define: context Name, Configuration URL, assign an Interface, choose Failover Group 2, and provide a Management IP Address. See Managing Security Contexts, page 59-7 for more information.

2. Save and submit these changes.

3. Provide the following context-configuration information, saving each change as you go:

   - On the Credentials page of the Device Properties window for the new context, provide Username and Password. See Viewing or Changing Device Properties, page 3-40 for additional information.
   - On the context's Interfaces page, edit the assigned interface, providing a Name, IP address and Subnet Mask. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for additional information.
   - On the context's Failover Page (ASA/PIX 7.0+), page 50-17, edit the interface configuration to provide a Standby IP Address.
   - On the HTTP Page, page 49-2, check Enable HTTP Server and then define HTTP access.
   - On the Credentials page, provide the Username and Password to be used when contacting the context. See Configuring Device Credentials, page 48-17 for additional information.

4. Choose **Deploy** from the Configuration Manager's File menu. Submit your changes, and then in the Deploy Saved Changes dialog box, be sure only this new context is selected, and then click Edit Deploy method. In the Edit Deploy Method dialog box, change the Method to File and then specify the Destination and a file name. Click OK to close the Edit Deploy Method dialog box, and then click Deploy the Deploy Saved Changes dialog box.

   The context configuration is saved to the specified file. See Deploying to a File, page 8-11 for more information about this step.

5. After uploading the configuration file to the device, use the CLI to enable HTTP access for the context. For example:

```
ciscoasa/group2(config-if)#  int g3/0
ciscoasa/group2(config-if)# nameif man
ciscoasa/group2(config-if)# security-level 100
ciscoasa/group2(config-if)# ip add 203.0.113.176 255.255.254.0 st 203.0.113.177
ciscoasa/group2(config-if)# exit
ciscoasa/group2(config)# http serv ena
ciscoasa/group2(config)# http 0.0.0.0 0.0.0.0 man
ciscoasa/group2(config)# username cisco pass cisco
ciscoasa/group2(config)#wr
```

Following this process, any new changes to the context can be successfully deployed to the context with Security Manager (attempts to reach the context will not go through the Admin context's management IP address).

Alternative

Another approach to this issue is to add the new context to failover group 1 first, and then perform the configuration via Security Manager. However, in order to then move this context to failover group 2, both groups (1 and 2) must be active on the same device. Otherwise, this error will be reported:

```
"join-failover-group 2
ERROR: Command requires failover-group 2 and 1 to be in the same state or no nameif comand for all interfaces
in this context"
```

Additional Steps for an Active/Standby Failover Configuration

Cisco Security Manager lets you authenticate a PIX/ASA/FWSM device by validating the certificate installed on the device. When configuring firewalls in an active/standby failover configuration, you must manually copy the certificate from the active device to the standby device so that Security Manager can communicate with the standby device after a failover occurs.

The following procedures describe how to export or display the identity certificate, CA certificate, and keys for a security appliances in your network using ASDM, and then import that information onto a standby device using ASDM.

- Exporting the Certificate to a File or PKCS12 data, page 50-9
- Importing the Certificate onto the Standby Device, page 50-9

Exporting the Certificate to a File or PKCS12 data

To export a trustpoint configuration, follow these steps using ASDM:

Step 1 Go to Configuration > Features > Device Administration > Certificate > Trustpoint > Export.

Step 2 Fill in the Trustpoint Name, Encryption Passphrase, and Confirm Passphrase fields. For information on these fields, click Help.

Step 3 Select a method for exporting the trustpoint configuration.

- Export to a File—Type the filename or browse for the file.
- Display the trustpoint configuration in PKCS12 format—Display the entire trustpoint configuration in a text box and then copy it for importing. For more information, click Help.

Step 4 Click **Export**.

Importing the Certificate onto the Standby Device

To import a trustpoint configuration, follow these steps using ASDM:

Step 1 Go to Configuration > Features > Device Administration > Certificate > Trustpoint > Import.

Step 2 Fill in the Trustpoint Name, Decryption Passphrase, and Confirm Passphrase fields. For information on these fields, click Help. The decryption passphrase is the same as the encryption passphrase used when the trustpoint configuration was exported.

Step 3 Select a method for importing the trustpoint configuration.

- Import from a File—Type the filename or browse for the file.
- Enter the trustpoint configuration in PKCS12 format—Paste the entire trustpoint configuration from the exported source into a text box. For more information, click Help.

Failover Policies

This section lists the pages that describe configuring failover on various types of security appliances; the pages are organized by device type.

PIX 6.x Firewalls

- Failover Page (PIX 6.3), page 50-10
  - Edit Failover Interface Configuration Dialog Box (PIX 6.3), page 50-12
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Firewall Services Modules

- Failover Page (FWSM), page 50-13
  - Advanced Settings Dialog Box, page 50-16
  - Add/Edit Interface MAC Address Dialog Box, page 50-23
  - Edit Failover Interface Configuration Dialog Box, page 50-23
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Adaptive Security Appliances and PIX 7.0 Firewalls

- Failover Page (ASA/PIX 7.0+), page 50-17
  - Settings Dialog Box, page 50-21
  - Edit Failover Group Dialog Box, page 50-25
  - Edit Failover Interface Configuration Dialog Box, page 50-23
  - Add/Edit Interface MAC Address Dialog Box, page 50-23
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Failover Page (PIX 6.3)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure failover settings for a PIX 6.3.x Firewall.

Navigation Path

Select a PIX 6.3.x device in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10

Field Reference

Table 50-1 Failover Page (PIX 6.3)

Element	Description
Failover	
Failover Method	Choose the type of failover link: **Serial Cable** or **LAN Based**. If you choose Serial Cable, ensure the physical cable is connected to both devices.
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key type, flash memory, and RAM. On PIX devices with LAN Based chosen as the Failover Method, you must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.
Failover Poll Time	Specify the amount of time between hello messages among units. Values can range from 3 to 15 seconds; default is 15.
LAN-Based Failover These fields are available when LAN Based is the chosen Failover Method.	
Interface	Choose the interface to be used for LAN-based failover. If "Not Selected" is chosen, LAN-based failover is disabled.
Shared Key Confirm	Used to encrypt communications between the primary and standby devices. Value can be any alphanumeric string. Re-enter the Shared Key in the Confirm field.
Stateful Failover (Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
Interface	Choose the interface to be used for Stateful Failover. If "Not Selected" is chosen, Stateful Failover is disabled. **Note** You must choose a fast LAN link from the list (for example, 100full, 1000full, or 1000sxfull).
Enable HTTP Replication	When selected, active HTTP sessions are copied to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Interface Configuration The table lists all available named interfaces. To define a Standby IP address and Active and Standby MAC addresses for an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box (PIX 6.3), page 50-12.	

Edit Failover Interface Configuration Dialog Box (PIX 6.3)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Edit Failover Interface Configuration dialog box to configure failover interfaces for the selected PIX 6.3.x device.

Note The failover interface cannot be configured for PPPoE.

Navigation Path

You can access the Edit Failover Interface Configuration dialog box from the Interface Configuration table on the Failover Page (PIX 6.3), page 50-10.

Related Topics

- Failover Policies, page 50-10

Field Reference

Table 50-2 Edit Failover Interface Configuration Dialog Box (PIX 6.3)

Element	Description
Interface	The name of the interface; read-only.
Active IP Address	Displays the IP address of the active interface. This address is used by the standby device to communicate with the active device. The address must be on the same network as the system IP address. The active IP address of this interface; read-only. This address is used by the standby device to communicate with the active device. This field is blank if an IP address has not been assigned to the interface. **Tip** You can use this IP address with the ping tool to check the status of the active device.
Netmask	The subnet mask for the active IP address; read-only. This field is blank if an IP address has not been assigned to the interface.
Standby IP Address	Specify the IP address of the corresponding interface on the standby failover unit. This address is used by the active device to communicate with the standby device. The address must be on the same network as the system IP address. This field does not appear if an IP address has not been assigned to the interface. **Tip** You can use this IP address with the ping tool to check the status of the standby device.
Failover MAC Addresses	
These parameters let you define virtual MAC addresses for a physical interface that is configured for failover; these addresses are optional.	

Table 50-2 Edit Failover Interface Configuration Dialog Box (PIX 6.3) (continued)

Element	Description
Active MAC Address	Specify a MAC address for the active interface in hexadecimal format (for example, 0123.4567.89ab).
Standby MAC Address	Specify a MAC address for the standby interface in hexadecimal format (for example, 0123.4567.89ab).

Failover Page (FWSM)

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure basic failover settings for the selected Firewall Services Module (FWSM).

Navigation Path

To access this feature, select a FWSM in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9
- Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Field Reference

Table 50-3 Failover Page (FWSM)

Element	Description
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key, flash memory, and RAM. You must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Settings button	Click to display the Advanced Settings Dialog Box, page 50-16, used to define when failover should occur.
Configuration	
This section is presented only for FWSM 3.1.1+ devices operating in multiple-context mode.	

Table 50-3 Failover Page (FWSM) (continued)

Element	Description
Active/Active	In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device. To enable Active/Active failover on the device, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments 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. See Add/Edit Security Context Dialog Box (FWSM), page 59-8 for information about assigning a context to a failover group.
Active/Standby	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.
LAN Failover	
VLAN	Enter the numeric ID of the VLAN interface you are using for the failover link; for example, 11. This list is not automatically populated with VLAN IDs—you must highlight “Not Selected” and type the desired VLAN ID number; press your keyboard’s Tab key to activate the related fields. When configured for failover, the interface is directly connected to the standby device.
Logical Name	Enter a logical name for the failover VLAN interface.
Active IP Address	Specify the active IP address for this interface.
Standby IP Address	Specify a standby IP address for this interface. 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.
Subnet Mask	Enter the Subnet Netmask for the Active and Standby IP addresses.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.

Table 50-3 Failover Page (FWSM) (continued)

Element	Description
Stateful Failover	
(Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
VLAN	Enter the numeric ID of the VLAN interface you are using for the failover link; for example, 12. This list is not automatically populated with VLAN IDs—you must highlight "Not Selected" and type the desired VLAN ID number; press your keyboard's Tab key to activate the related fields. When configured for failover, the interface is directly connected to the standby device.
Logical Name	Enter a logical name for the Stateful failover VLAN interface.
Active IP Address	Specify the active IP address for this interface.
Standby IP Address	Specify a standby IP address for this interface. 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.
Subnet Mask	Enter the Subnet Netmask for the Active and Standby IP addresses.
Enable HTTP Replication	When selected, allows stateful failover to copy active HTTP sessions to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Shared Key (FWSM 3.1.1+ only) The options in this section let you encrypt the communications between the active and standby devices by providing a shared encryption key. ⚠ **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 this device is used to terminate VPN tunnels, this information includes any user names, passwords and shared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communications with a shared key.	
Shared Key Confirm	Enter any string of characters up to 63 numbers, letters and punctuation characters. This string is used to generate the encryption key. Re-enter this string the **Confirm** field. If you select **HEX**, the entry in the Shared Key and Confirm fields must be exactly 32 hexadecimal characters (0-9, a-f).
Interface Configuration This table is presented on the Failover page for devices operating in single-context mode, or for individual security contexts only. The table lists all available named interfaces. To enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23. Select or deselect **Monitor this interface for failure**.	

Advanced Settings Dialog Box

The Advanced Settings dialog box lets you configure additional failover settings for the selected FWSM.

Note The following reference table describes all fields that can be presented in the Advanced Settings dialog box. The fields actually presented depend on operating mode (routed or transparent) and whether the device is hosting single or multiple contexts.

Navigation Path

You can access the Advanced Settings dialog box by clicking the Settings button on the Failover Page (FWSM), page 50-13.

Related Topics

- Failover Policies, page 50-10

Field Reference

Table 50-4 Advanced Settings Dialog Box

Element	Description
Interface Policy	
Select a failed-interfaces option and provide an appropriate value.	
Number of failed interfaces	When the number of failed monitored interfaces exceeds this value, the security appliance fails over. Valid values range from 1 to 250.
Percentage of failed interfaces	When the number of failed monitored interfaces exceeds this percentage, the security appliance fails over.
Failover Poll Time	
These fields define how often hello messages are sent on the failover link, and 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 between failover units. Enter a value between 1 and 15 seconds, or if **msec** is checked, between 500 and 999 milliseconds.
Unit Hold Time	The amount of time to wait for a hello message on the failover link, after which the unit begins testing for peer failure. Enter a value between 3 and 45 seconds. This value must be at least three times the Unit Failover value.
Monitored Interface	The amount of time between polls among interfaces. Enter a value between 3 and 15 seconds.
MAC Address Mapping	
In Active/Standby mode, this table lists interface-virtual MAC address mappings. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, which are described in Using Tables, page 1-48. To add or edit interface mappings, click the Add Row or Edit Row button to open the Add/Edit Interface MAC Address Dialog Box, page 50-23.	

Table 50-4 Advanced Settings Dialog Box (continued)

Element	Description
Failover Groups	
	In Active/Active mode, this table lists both failover groups. To edit failover parameters for either group, select it in the list and click the Edit Row button to open the Edit Failover Group Dialog Box, page 50-25.
Bridge Group Configuration	
	In single-context transparent mode, this table lists all currently defined bridge groups (see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26). To add a standby IP address to a bridge group, select it in the list and click the Edit Row button to open the Edit Failover Bridge Group Configuration Dialog Box, page 50-17.

Edit Failover Bridge Group Configuration Dialog Box

Use this dialog box to add a standby IP address to a failover bridge group.

Navigation Path

You can access the Edit Failover Bridge Group Configuration dialog box as follows:

- On the Failover page presented for an individual security context in transparent mode on an ASA.
- From the Bridge Group Configuration table in the Advanced Settings Dialog Box, page 50-16 presented by an FWSM in transparent mode.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13

Field Reference

Table 50-5 Edit Failover Bridge Group Configuration Dialog Box

Element	Description
Name	Identifies the bridge group; not editable.
IP Address	Identifies the IP address assigned to the bridge group; not editable.
Network Mask	Identifies the subnet mask for the IP Address; not editable.
Standby Address	Enter the IP address of the standby bridge group; this address must be on the same subnet as the primary address.

Failover Page (ASA/PIX 7.0+)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure basic failover settings for ASA and PIX 7.0+ security devices

Note The features and options presented on the Failover page vary according to type of device selected, operating system version, firewall mode (routed or transparent), and security contexts (single or multiple). Thus, some of the elements described in the following table may not appear on the Failover page for your currently selected device.

Navigation Path

Select an ASA or PIX 7.0+ in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

Field Reference

Table 50-6 Failover Page (ASA/PIX 7.0+)

Element	Description
Failover Method	Choose the type of failover link: **Serial Cable** or **LAN Based**. If you choose Serial Cable, ensure the physical cable is connected to both devices. **Note** This option is available only on PIX devices.
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key type, flash memory, and RAM. On PIX devices with LAN Based chosen as the Failover Method, and on all ASAs, you must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.
Settings button	Click to display the Settings Dialog Box, page 50-21, used to define when failover should occur.
Timeout	The failover Timeout specifies the amount of time after a system boots or becomes active that "nailed" sessions are accepted; used in conjunction with static translation rules (see Static Rules Tab, page 24-26 for more information). Enter a value in this field to specify the failover reconnect timeout value for asymmetrically routed sessions. The value is in the form hh:mm:ss (hours:minutes:seconds); both minutes and seconds are optional. Valid values for the number of hours are -1 to 1193; the default value is 0, which means connections cannot be re-established. Setting this value to -1 disables the timeout, allowing reconnections after any amount of time.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Configuration	
This section is presented only for devices operating in multiple-context mode.	
Active/Active	In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device. To enable Active/Active failover on the security appliance, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments 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. See Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 for information about assigning a context to a failover group.
Active/Standby	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.
LAN Failover	
Interface	Choose the interface to use as the failover link; all interfaces available on the device are listed. When configured for failover, the interface is directly connected to the standby device. **Note** You can choose an EtherChannel interface as the failover link. As with any other type of interface assigned as a failover link, the EtherChannel interface cannot be named, and none of the EtherChannel's member interfaces can be named. Further, while being used as an active failover link, changes to the interface configuration are not allowed. Refer to Configuring EtherChannels, page 46-9 for more information.
Logical Name	Enter a logical name for the failover interface.
Active IP Address	Specify the active IP address for this interface.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Standby IP Address	Specify a standby IP address for this interface. 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.
Subnet Mask	Enter the Subnet Netmask for the active and standby IP addresses.
Stateful Failover (Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
Interface	Choose the interface to use for the stateful failover link; all interfaces available on the device are listed. **Note** You can choose an EtherChannel interface as the stateful failover link. As with any other type of interface assigned as a failover link, the EtherChannel interface cannot be named, and none of the EtherChannel's member interfaces can be named. Further, while being used as an active failover link, changes to the interface configuration are not allowed. Refer to Configuring EtherChannels, page 46-9 for more information.
Logical Name	Enter the logical name of the interface on the active firewall device to communicate with standby device for failover. When configured for stateful failover, the interface is directly connected to the standby device.
Active IP Address	Specify the IP address of the active interface.
Standby IP Address	Specify the IP address of the standby interface.
Subnet Mask	Enter the Subnet Netmask for the active and standby IP addresses.
Enable HTTP Replication	When selected, active HTTP sessions are copied to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Key The options in this section let you encrypt the communications between the active and standby devices. Select the type and provide a string of characters to produce the shared encryption key. **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 this device is used to terminate VPN tunnels, this information includes any user names, passwords and shared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communications with a shared key.	
Any string HEX	If you select **Any string**, the entry in the Shared Key field can be any combination of up to 63 numbers, letters and punctuation characters. This string is used to generate the encryption key. If you select **HEX**, the entry in the Shared Key and Confirm fields must be exactly 32 hexadecimal characters (0-9, a-f). This string is used as the encryption key.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Shared Key Confirm	Enter any string of characters appropriate to the selected key type: Any string or HEX. Re-enter the string the **Confirm** field.

Interface Configuration
(in some instances, labeled Monitor Interface Configuration)

This table is presented on the Failover page for ASA 8.4.1+ devices operating in single-context, transparent mode, and for individual contexts on PIX/ASA devices. Otherwise, it appears in the Settings Dialog Box, page 50-21.

The table lists all available named interfaces. To enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23. Select or deselect **Monitor this interface for failure**.

Settings Dialog Box

The Settings dialog box lets you define criteria for when failover should occur on the selected ASA or PIX 7.x appliance.

Navigation Path

You can access the Settings dialog box by clicking the Settings button on the Failover Page (ASA/PIX 7.0+), page 50-17.

Note The following reference table presents all fields that can be presented in the Settings dialog box. The fields actually presented depend on operating mode (routed or transparent) and whether the device is hosting single or multiple contexts.

Related Topics

- Failover Policies, page 50-10
- Edit Failover Interface Configuration Dialog Box, page 50-23
- Add/Edit Interface MAC Address Dialog Box, page 50-23
- Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Field Reference

Table 50-7 Settings Dialog Box

Element	Description
Interface Policy	
Number of failed interfaces	When the number of failed monitored interfaces exceeds this value, the security appliance fails over. The range is between 1 and 250 failures.
Percentage of failed interfaces	When the number of failed monitored interfaces exceeds this percentage, the security appliance fails over.

Table 50-7 Settings Dialog Box (continued)

Element	Description
Failover Poll Time	
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 if the **Change units to msec** option is checked.
Unit Hold Time	Sets the time during which a unit must receive a hello message on the failover link, or the unit begins the testing process for peer failure. The range is between 3 and 45 seconds, or between 800 and 999 milliseconds if the **msec** option is checked. You cannot enter a value that is less than three times the Unit Failover value.
Monitored Interface	The amount of time between polls among interfaces. The range is between 3 and 15 seconds, or between 500 and 999 milliseconds if the msec option is checked.
Interface Hold Time	Sets the time during which a data interface must receive a hello message, after which the peer is declared failed. Valid values are from 5 to 75 seconds. This value must be at least five times the Unit Failover value.
Link State Interval	Sets the interval after which each ASA in a failover pair checks the link state of its interfaces. By default the link state interval value is 500msec. You can customize the polltime; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. Valid range is between 300 and 799 milliseconds. **Note** The Link State Interval is available for ASA 9.7.1 and higher.
Failover Groups In Active/Active mode, this table lists both failover groups. To edit failover parameters for either group, select it in the list and click the Edit Row button to open the Edit Failover Group Dialog Box, page 50-25.	
MAC Address Mapping In Active/Standby mode, this table lists interface-virtual MAC address mappings. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, which are described in Using Tables, page 1-48. To add or edit interface mappings, click the Add Row or Edit Row button to open the Add/Edit Interface MAC Address Dialog Box, page 50-23.	
Monitor Interface Configuration In single-context mode, this table lists all available named interfaces. To define a Standby IP address for, and enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23.	
Management IP Address In single-context transparent mode, this section presents the management IP address and netmask defined for the device (on the Management IP Page, page 47-10); you cannot change these values.	
Standby	Enter the management IP address of the standby unit; this address must be on the same subnet as the primary address.

Add/Edit Interface MAC Address Dialog Box

The Add/Edit Interface MAC Address dialog box lets you define virtual MAC addresses for a physical interface on ASA, FWSM 3.x and PIX 7.x security appliances that are configured for failover (not available on ASA 5505 devices).

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

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

Navigation Path

You can open the Add/Edit Interface MAC Address dialog box from the Settings Dialog Box, page 50-21.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Edit Failover Group Dialog Box, page 50-25

Field Reference

Table 50-8 Add/Edit Interface MAC Address Dialog Box

Element	Description
Physical Interface	Choose the physical interface on which failover virtual MAC addresses are to be configured.
MAC Address	
Active Interface	Enter a virtual MAC address for the active interface in hexadecimal format (for example, 0023.4567.89ab).
Standby Interface	Enter a virtual MAC address for the standby interface in hexadecimal format (for example, 0023.4567.89ab).

Edit Failover Interface Configuration Dialog Box

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

Note A failover interface cannot be configured for PPPoE.

Navigation Path

You can access the Edit Failover Interface Configuration dialog box from the Settings Dialog Box, page 50-21 (ASA/PIX 7.0+), Advanced Settings Dialog Box, page 50-16 (FWSM), and from the Failover page itself for ASA 8.4.1+ devices operating in single-context transparent mode, and for individual ASA/PIX security contexts.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13
- Edit Failover Group Dialog Box, page 50-25

Field Reference

Table 50-9 Edit Failover Interface Configuration Dialog Box

Element	Description
Interface Name	The name of the interface; read-only.
Active IP Address	The active IP address of this interface; read-only. This field is blank if an IP address has not been assigned to the interface; for example if DHCP is enabled on the interface.
Mask	The subnet mask for the active IP address; read-only. This field is blank if an IP address has not been assigned to the interface; for example, if DHCP is enabled on the interface.
Standby IP Address	Specify 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 this interface for failure	Specifies whether this interface is monitored for failure: check this box to enable monitoring. 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. 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 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.

Table 50-9 Edit Failover Interface Configuration Dialog Box (continued)

Element	Description
ASR Group Number	If this interface is part of an asymmetric routing group, provide its ASR group number. Valid values for ASR group numbers are 1 through 32. Stateful failover must be enabled for asymmetric routing support to function properly between units in failover configurations.

Edit Failover Group Dialog Box

Use the Edit Failover Group dialog box to configure failover parameters for groups of security contexts in an Active/Active failover configuration. See Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9, or Add/Edit Security Context Dialog Box (FWSM), page 59-8, for information about assigning a context to a failover group.

Navigation Path

You can access the Add Failover Group dialog box from the PIX/ASA Settings Dialog Box, page 50-21, or the FWSM Advanced Settings Dialog Box, page 50-16.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13

Field Reference

Table 50-10 Edit Failover Group Dialog Box

Element	Description
Preferred Role	Specifies the unit in the failover pair, primary or secondary, on which this failover group appears in the active state when both units start up simultaneously, or when the Preempt option is selected. Choose **Primary** or **Secondary**. You can have both failover groups in the active state on a single unit in the pair; however, a more typical configuration is to assign each failover group a different role to make each one active on a different unit, balancing the traffic across the devices.
Poll time interval for monitored interfaces	Specify the amount of time between polling of monitored interfaces. Valid values are from 3 to 15 seconds (or 500 to 999 milliseconds if **msec** is checked).
Hold Time	Specify the time period within which the group must receive a hello message, after which the other group is declared failed. Valid values are from 5 to 75 seconds.
Preempt after Reboot	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. Valid values are from 0 to 1200 seconds.

Table 50-10 Edit Failover Group Dialog Box (continued)

Element	Description
Enable HTTP Replication	Indicates whether active HTTP sessions are copied to the standby device for this failover group as part of Stateful failover. If you do not allow HTTP replication, 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 Failover page.
Failover Criteria	Select a failed-interfaces criterion for this group and specify the appropriate value: • **Number of failed interfaces** – When this number of interfaces have failed, failover is triggered. Valid values are 1 to 250. • **Percentage of failed interfaces** – When this percentage of the total number of interfaces have failed, failover is triggered. Valid values are 1 to 100.
MAC Address Mapping This table displays interfaces to which active and standby MAC addresses are mapped.	

Failover Page (Security Context)

The Failover page for individual ASA and PIX 7.0+ security contexts presents the **Interface Configuration** table, which lists all available named interfaces.

You can select an interface in the table and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23, where you can specify a standby IP address and an ASR group number, and enable or disable monitoring of the interface.

For individual transparent-mode contexts on ASA 8.4.1+ devices, the Failover page also presents the **Bridge Group Configuration** table, which lists all currently defined failover bridge groups.

You can select an entry in the table and click the Edit Row button to open the Edit Failover Bridge Group Configuration Dialog Box, page 50-17, where you can specify a standby IP address for the selected bridge group.

Navigation Path

Select a security context in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10
- About Bridging on Firewall Devices, page 47-1

Bootstrap Configuration for LAN Failover Dialog Box

The Bootstrap Configuration for LAN Failover dialog box provides you with bootstrap configuration that can be applied to the primary and secondary devices in a LAN failover configuration.

Navigation Path

You can access the Bootstrap Configuration for LAN Failover dialog box from the Failover page. For more information about the Failover page, see:

- Failover Page (PIX 6.3), page 50-10
- Failover Page (FWSM), page 50-13
- Failover Page (ASA/PIX 7.0+), page 50-17

Related Topics

- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

Field Reference

Table 50-11 Bootstrap Configuration for LAN Failover Dialog Box

Element	Description
Primary	Contains the bootstrap configuration for the primary device. Open a console connection to the primary device and then paste this configuration to activate failover on the device.
Secondary	Contains the bootstrap configuration for the secondary device. After the primary device becomes active, open a console connection to the secondary device and then paste this configuration to activate failover on the device.

Note For Active/Active Failover, the bootstrap configurations are only applied to the system contexts of the respective failover peer devices.