



Configuring Failover

This chapter describes FWSM failover feature, which lets you configure two FWSMs so that one will take over operation if the other one fails. Failover is compatible with both routed and transparent firewall modes, and with single and multiple context modes.

This chapter includes the following sections:

- [Understanding Failover, page 13-1](#)
- [Configuring Failover, page 13-18](#)
- [Controlling and Monitoring Failover, page 13-38](#)

For sample failover configurations, see the “[Failover Example Configurations](#)” section on page B-18.

Understanding Failover

The failover configuration requires two identical FWSMs connected to each other through a dedicated failover link and, optionally, a state link. The health of the active interfaces and units is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

FWSM supports two failover configurations, Active/Active failover and Active/Standby failover. Each failover configuration has its own method for determining and performing failover.

With Active/Active failover, both units can pass network traffic. This lets you configure load balancing on your network. Active/Active failover is only available on units running in multiple context mode.

With Active/Standby failover, only one unit passes traffic while the other unit waits in a standby state. Active/Standby failover is available on units running in either single or multiple context mode.

Both failover configurations support stateful or stateless (regular) failover.

This section includes the following topics:

- [Failover System Requirements, page 13-2](#)
- [Failover and State Links, page 13-2](#)
- [Intra- and Inter-Chassis Module Placement, page 13-3](#)
- [Transparent Firewall Requirements, page 13-7](#)
- [Active/Standby and Active/Active Failover, page 13-8](#)
- [Regular and Stateful Failover, page 13-16](#)
- [Failover Health Monitoring, page 13-17](#)

Failover System Requirements

This section describes the software and license requirements for FWSMs in a failover configuration. This section contains the following topics:

- [Software Requirements, page 13-2](#)
- [License Requirements, page 13-2](#)

Software Requirements

The two units in a failover configuration must have the same major (first number) and minor (second number) software version. However, you can use different versions of the software during an upgrade process; for example, you can upgrade one unit from Version 3.1(1) to Version 3.1(2) and have failover remain active. We recommend upgrading both units to the same version to ensure long-term compatibility.

License Requirements

Both units must have the same license.

Failover and State Links

This section describes the failover and the state links, which are dedicated connections between the two units in a failover configuration. This section includes the following topics:

- [Failover Link, page 13-2](#)
- [State Link, page 13-3](#)

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit. The following information is communicated over the failover link:

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

**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.

The failover link uses a special VLAN interface that you do not configure as a normal networking interface; rather, it exists only for failover communications. This VLAN should only be used for the failover link (and optionally for the state link). Sharing the failover link VLAN with any other VLANs can cause intermittent traffic problems and ping and ARP failures. For inter-chassis failover, use dedicated interfaces on the switch for the failover link.

**Note**

If failover and the interface are configured but the VLAN is not downloaded from the switch and the failover is in disabled mode, then the FWSM doesn't send ARP requests by design because it could conflict with another FWSM currently running as Active in the system.

On systems running in multiple context mode, the failover link resides in the system context. This interface and the state link, if used, are the only interfaces that you can configure in the system context. All other interfaces are allocated to and configured from within security contexts.

**Note**

The IP address and MAC address for the failover link do not change at failover.

State Link

To use Stateful Failover, you must configure a state link to pass all state information. This link can be the same as the failover link, but we recommend that you assign a separate VLAN and IP address for the state link. The state traffic can be large, and performance is improved with separate links.

The state link interface is not configured as a normal networking interface; it exists only for Stateful Failover communications and, optionally, for the failover communication if you share the state and failover links.

In multiple context mode, the state link resides in the system context. This interface and the failover interface are the only interfaces in the system context. All other interfaces are allocated to and configured from within security contexts.

**Note**

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

**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.

Intra- and Inter-Chassis Module Placement

You can place the primary and secondary FWSMs within the same switch or in two separate switches. The following sections describe each option:

- [Intra-Chassis Failover, page 13-3](#)
- [Inter-Chassis Failover, page 13-5](#)

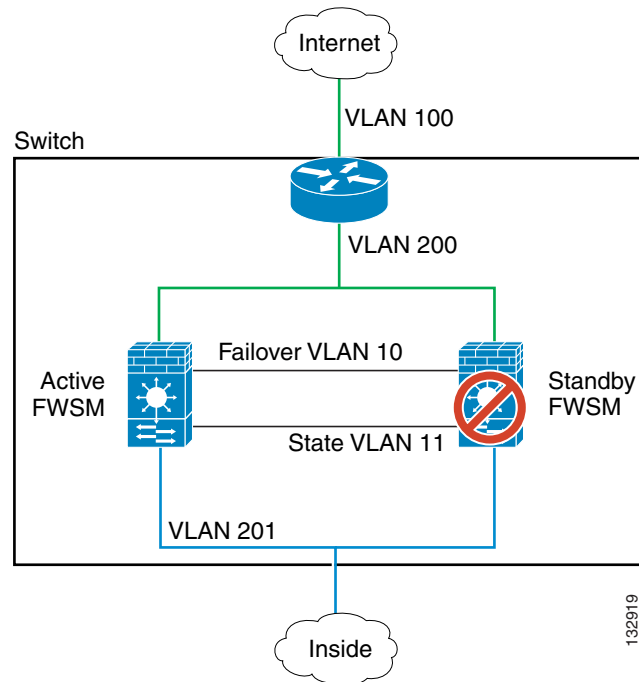
Intra-Chassis Failover

If you install the secondary FWSM in the same switch as the primary FWSM, you protect against module-level failure. To protect against switch-level failure, as well as module-level failure, see the “[Inter-Chassis Failover](#)” section on page 13-5.

Even though both FWSMs are assigned the same VLANs, only the active module takes part in networking. The standby module does not pass any traffic.

Figure 13-1 shows a typical intra-switch configuration.

Figure 13-1 Intra-Switch Failover



132919

Inter-Chassis Failover

To protect against switch-level failure, you can install the secondary FWSM in a separate switch. FWSM does not coordinate failover directly with the switch, but it works harmoniously with the switch failover operation. See the switch documentation to configure failover for the switch.

To accommodate the failover communications between FWSMs, we recommend that you configure a trunk port between the two switches that carries the failover and state VLANs. The trunk ensures that failover communication between the two units is subject to minimal failure risk.

For other VLANs, you must ensure that both switches have access to all firewall VLANs, and that monitored VLANs can successfully pass hello packets between both switches.

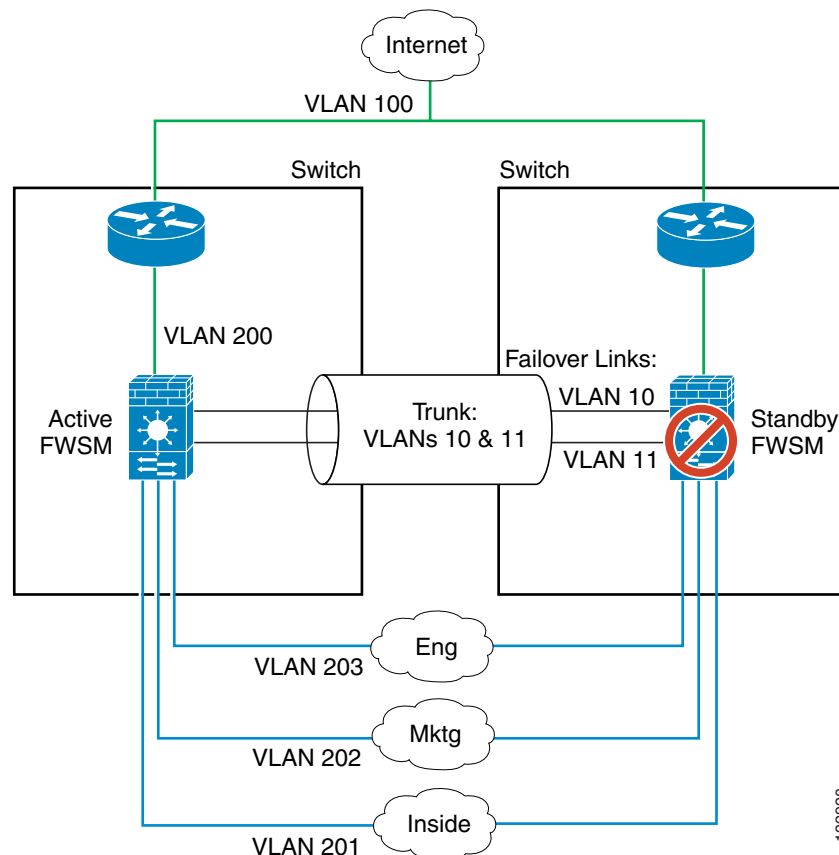
Figure 13-2 shows a typical switch and FWSM redundancy configuration. The trunk between the two switches carries the failover FWSM VLANs (VLANs 10 and 11).



Note

FWSM failover is independent of the switch failover operation; however, FWSM works in any switch failover scenario.

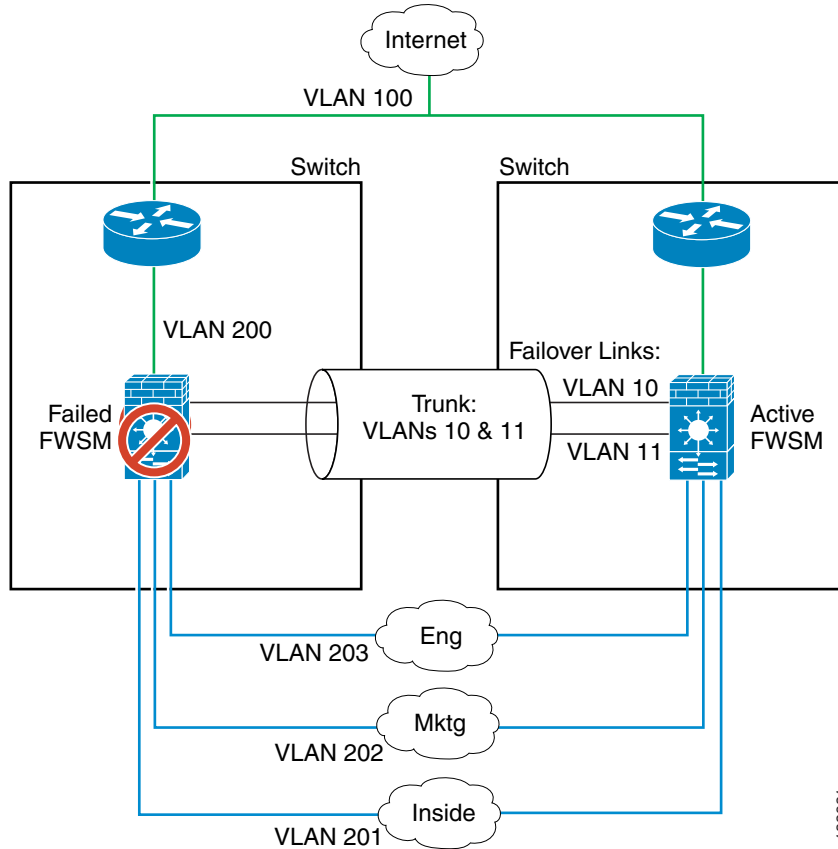
Figure 13-2 Normal Operation



132920

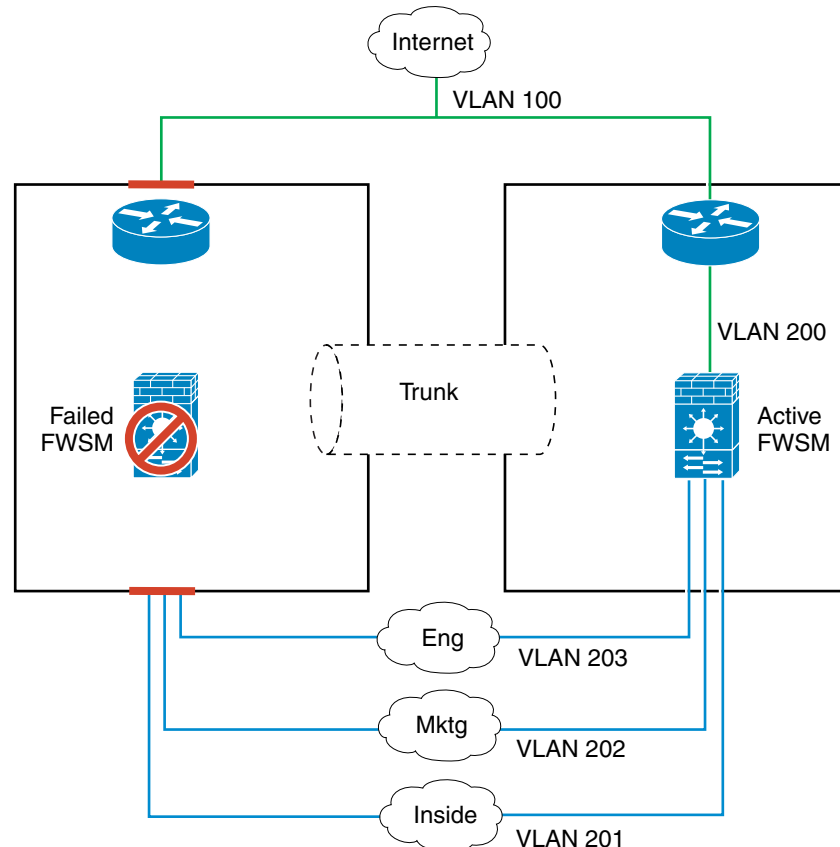
If the primary FWSM fails, then the secondary FWSM becomes active and successfully passes the firewall VLANs (Figure 13-3).

Figure 13-3 FWSM Failure



If the entire switch fails, as well as the FWSM (such as in a power failure), then both the switch and the FWSM fail over to their secondary units (Figure 13-4).

Figure 13-4 Switch Failure



Transparent Firewall Requirements

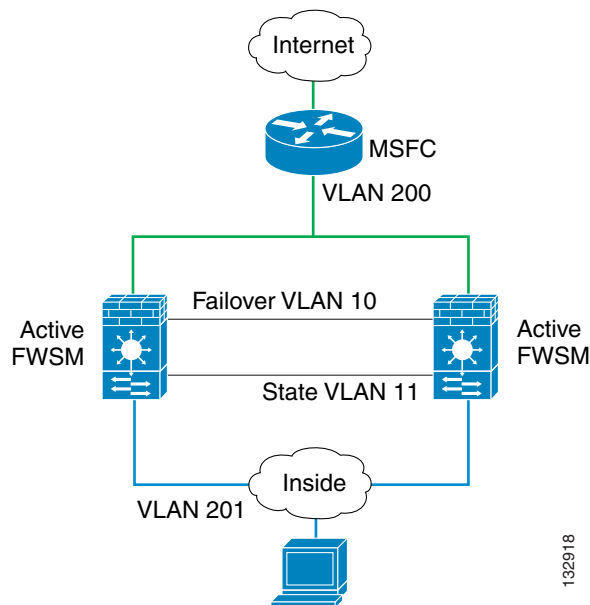
To avoid loops when you use failover in transparent mode, you must use switch software that supports BPDU forwarding, and you must configure the FWSM to allow BPDUs. See the [“Switch Hardware and Software Compatibility”](#) section on page A-1 for switch software versions that allow BPDUs automatically.

To allow BPDUs through the FWSM, configure an EtherType ACL and apply it to both interfaces according to the [“Adding an EtherType Access List”](#) section on page 10-8.

Loops can occur if both modules are active at the same time, such as when both modules are discovering the presence of the other module, or due to a bad failover link. Because the FWSMs bridge packets between the same two VLANs, loops can occur when inside packets destined for the outside get

endlessly replicated by both FWSMs (see [Figure 13-5](#)). The spanning tree protocol can break such loops if there is a timely exchange of BPDUs. To break the loop, BPDUs sent between VLAN 200 and VLAN 201 need to be bridged.

Figure 13-5 Potential Loops in Transparent Mode



132918

Active/Standby and Active/Active Failover

This section describes each failover configuration in detail. This section includes the following topics:

- [Active/Standby Failover, page 13-8](#)
- [Active/Active Failover, page 13-11](#)
- [Determining Which Type of Failover to Use, page 13-15](#)

Active/Standby Failover

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

- [Active/Standby Failover Overview, page 13-9](#)
- [Primary/Secondary Status and Active/Standby Status, page 13-9](#)
- [Device Initialization and Configuration Synchronization, page 13-9](#)
- [Command Replication, page 13-10](#)
- [Failover Triggers, page 13-10](#)
- [Failover Actions, page 13-11](#)

Active/Standby Failover Overview

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

**Note**

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

Primary/Secondary Status and Active/Standby Status

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

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

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

Device Initialization and Configuration Synchronization

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

The active unit is determined by the following:

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

**Note**

If the secondary unit boots without detecting the primary unit, it becomes the active unit. It uses its own MAC address for the active IP addresses. However, when the primary unit becomes available, the secondary unit changes the MAC address to that of the primary unit, which can cause an interruption in your network traffic.

When the configuration synchronization starts, the FWSM console on the active unit displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the FWSM console displays the message “End Configuration Replication to mate.” During the configuration synchronization, commands entered on the active unit may not replicate properly to the standby unit, and commands entered on the standby unit may be overwritten by the configuration being replicated from

the active unit. Avoid entering commands on either unit in the failover pair during the configuration replication process. Depending upon the size of the configuration, replication can take from a few seconds to several minutes.

If you enter the **write standby** command on the active unit, the standby unit clears its running configuration (except for the failover commands used to communicate with the active unit), and the active unit sends its entire configuration to the standby unit.

In multiple context mode, when you enter the **write standby** command in the system execution space, all contexts are replicated. If you enter the **write standby** command within a context, the command replicates only the context configuration.

On the standby unit, the replicated configuration exists only in running memory. To save the configuration to Flash memory after synchronization:

- In single context mode, enter the **write memory** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- In multiple context mode, enter the **write memory all** command on the active unit from the system execution space. This command saves the system configuration and all context configurations. The command is replicated to the standby unit, which proceeds to write its configurations to Flash memory. Contexts with startup configurations on external servers are accessible from either unit over the network and do not need to be saved separately for each unit. Alternatively, you can copy the contexts on disk from the active unit to an external server, and then copy them to disk on the standby unit, where they become available when the unit reloads.

Command Replication

As commands are entered on the active unit, they are sent across the failover link to the standby unit. Command replication always flows from the active unit to the standby unit. Replicated commands are stored in the running configuration of the standby unit. Saving the running configuration to the startup configuration on the active unit causes the running configuration to be saved to the startup configuration on the standby unit; however, you do not have to save the active configuration to Flash memory to replicate the commands.



Note

The RSA keys are not synchronized from the primary to the secondary unit in FWSM.



Note

The **mode** command is not replicated to the secondary unit.

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

Failover Triggers

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

- The unit has a hardware failure or a power failure.
- The unit has a software failure.
- Too many monitored interfaces fail.
- The **no failover active** command is entered on the active unit or the **failover active** command is entered on the standby unit.

Failover Actions

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

[Table 13-1](#) shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 13-1 Failover Behavior

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

Active/Active Failover

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

- [Active/Active Failover Overview](#), page 13-12
- [Primary/Secondary Status and Active/Standby Status](#), page 13-12
- [Device Initialization and Configuration Synchronization](#), page 13-12
- [Command Replication](#), page 13-13

- [Failover Triggers, page 13-14](#)
- [Failover Actions, page 13-14](#)

Active/Active Failover Overview

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

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

The failover group forms the base unit for failover in Active/Active failover. Interface failure monitoring, failover, and active/standby status are all attributes of a failover group rather than of the unit. The MAC address of the primary unit is used by all interfaces in the active contexts.

When an active failover group fails, it changes to the standby state while the associated standby failover group becomes active. The interfaces in the failover group that becomes active assume the MAC address and IP addresses of the interfaces in the failover group that failed. The interfaces in the failover group that is now in the standby state take over the standby MAC address and IP addresses.



Note

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

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

Primary/Secondary Status and Active/Standby Status

As in Active/Standby failover, one unit in an Active/Active failover pair is designated the primary unit, and the other unit the secondary unit. Unlike Active/Standby failover, this designation does not indicate which unit becomes active when both units start simultaneously. Instead, the primary/secondary designation determines which unit provides the running configuration to the pair and on which unit each failover group appears in the active state when both units start simultaneously.

Each failover group in the configuration is given a primary or secondary unit 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.



Note

FWSM does not provide load balancing services. Load balancing must be handled by a router passing traffic to FWSM.

Device Initialization and Configuration Synchronization

Configuration synchronization occurs when one or both units in a failover pair boot.

When a unit boots while the peer unit is not available, then both failover groups become active on the unit regardless of the primary or secondary designation for the failover groups and the unit. Configuration synchronization does not occur. Some reasons a peer unit may not be available are that the peer unit is powered down, the peer unit is in a failed state, or the failover link between the units has not been established.

When a unit boots while the peer unit is active (with both failover groups active on it), the booting unit contacts the active unit to obtain the running configuration. By default, the failover groups will remain active on the active unit regardless of the primary or secondary preference of each failover group and unit designation (unless configured with the **preempt** command). The failover groups remain active on the first unit until one of the following occurs:

- A failover condition causes the failover group to become active on the peer unit.
- You manually force a failover group to become active on the peer unit using the **no failover active** command.
- The **preempt** command forces the failover group to become active on its preferred unit when that unit becomes available.

When both units boot at the same time, the primary unit becomes the active unit. The secondary unit obtains the running configuration from the primary unit. Once the configuration has been synchronized, each failover group becomes active on its preferred unit.

Command Replication

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

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



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

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

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



Note The **mode** command does not get replicated.

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

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

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

Failover Triggers

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

- The unit has a hardware failure.
- The unit has a power failure.
- The unit has a software failure.
- The **no failover active** or the **failover active** command is entered in the system execution space.

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

- Too many monitored interfaces in the contexts that belong to the failover group fail.
- The **no failover active group** *group_id* command is entered.

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

See the “[Failover Health Monitoring](#)” section on page 13-17 for more information about interface and unit monitoring.

Failover Actions

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



Note

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

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

Table 13-2 Failover Behavior for Active/Active Failover

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
A unit experiences a power or software failure	Failover	Become standby Mark as failed	Become active Mark active as failed	When a unit in a failover pair fails, any active failover groups on that unit are marked as failed and become active on the peer unit.
Interface failure on active failover group above threshold	Failover	Mark active group as failed	Become active	None.

Table 13-2 Failover Behavior for Active/Active Failover (continued)

Failure Event	Policy	Active Group Action	Standby Group Action	Notes
Interface failure on standby failover group above threshold	No failover	No action	Mark standby group as failed	When the standby failover group is marked as failed, then the active failover group will not attempt to fail over, even if the interface failure threshold is surpassed.
Formerly active failover group recovers	No failover	No action	No action	Unless configured with the preempt command, the failover groups remain active on their current unit.
Failover link failed at startup	No failover	Become active	Become active	If the failover link is down at startup, both failover groups on both units will become active.
State link failed	No failover	No action	No action	State information will become out of date, and sessions will be terminated if a failover occurs.
Failover link failed during operation	No failover	n/a	n/a	Each unit marks the failover interface as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.

Determining Which Type of Failover to Use

The type of failover you choose depends upon your FWSM configuration and how you plan to use FWSM.

If you are running FWSM in single mode, then you can only use Active/Standby failover; Active/Active failover is only available to FWSMs running in multiple context mode. If you are running the FWSM in multiple context mode, then you can configure either Active/Active failover or Active/Standby failover.

If you are using an upstream router to provide load balancing, use Active/Active failover. If you do not want to provide load balancing, use either Active/Standby or Active/Active failover.

[Table 13-3](#) provides a comparison of some of the features supported by each type of failover configuration.

Table 13-3 Failover Configuration Feature Support

Feature	Active/Active	Active/Standby
Single Context Mode	No	Yes
Multiple Context Mode	Yes	Yes
Load Balancing Network Configurations	Yes	No
Unit Failover	Yes	Yes
Failover of Groups of Contexts	Yes	No
Failover of Individual Contexts	No	No

Regular and Stateful Failover

FWSM supports two types of failover, regular and stateful. This section includes the following topics:

- [Regular Failover, page 13-16](#)
- [Stateful Failover, page 13-16](#)

Regular Failover

When a failover occurs, all active connections are dropped. Clients need to reestablish connections when the new active unit takes over.

Stateful Failover

When Stateful Failover is enabled, the active unit 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.

The state information passed to the standby unit includes the following:

- NAT translation table.
- TCP connection states.
- UDP connection states.
- The ARP table.
- The Layer 2 bridge table (when running in transparent firewall mode).
- The HTTP connection states (if HTTP replication is enabled).
- The ISAKMP and IPsec SA table.
- GTP PDP connection database.

The information that is not passed to the standby unit when Stateful Failover is enabled includes the following:

- The HTTP connection table (unless HTTP replication is enabled).
- The user authentication (uauth) table.
- The routing tables.
- Multicast traffic information.

**Note**

If failover occurs during an active Cisco IP SoftPhone session, the call will remain active because the call session state information is replicated to the standby unit. When the call is terminated, the IP SoftPhone client will lose connection with the CallManager. This occurs because there is no session information for the CTIQBE hangup message on the standby unit. When the IP SoftPhone client does not receive a response back from the CallManager within a certain time period, it considers the CallManager unreachable and unregisters itself.

**Note**

Because transparent FWSM relies on a Layer 2 MAC table for forwarding, the connection entry for a pair of hosts might still be active when the MAC table entries for one or both hosts have timed out due to inactivity. In such a situation, if a failover event occurs before either host sends another packet to re-populate the MAC address table, the peer FWSM is not able to generate switch CAM table refresh packets for the given endpoints. Therefore, if the CAM table entries on the switch for the given hosts are still active and point to the formerly active unit, traffic is incorrectly switched to the standby FWSM and dropped there (if the idle connection starts passing traffic again after the failover event and before the CAM table entries age out on the switch).

Failover Health Monitoring

FWSM monitors each unit for overall health and for interface health. See the following sections for more information about how FWSM performs tests to determine the state of each unit:

- [Unit Health Monitoring, page 13-17](#)
- [Interface Monitoring, page 13-17](#)

Unit Health Monitoring

FWSM determines the health of the other unit by monitoring the failover link. When a unit does not receive hello messages on the failover link, then the unit sends an ARP request on all interfaces, including the failover interface. FWSM retries a user-configurable number of times. The action FWSM takes depends on the response from the other unit. See the following possible actions:

- If FWSM receives a response on any interface, then it does not fail over.
- If FWSM does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.
- If FWSM does not receive a response on the failover link only, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.

**Note**

If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Interface Monitoring

You can monitor up to 250 interfaces divided between all contexts. You can configure one context to monitor a shared interface (because the interface is shared, all contexts benefit from the monitoring).

When a unit does not receive hello messages on a monitored interface, it runs the following tests:

1. **Link Up/Down test**—A test of the interface status. If the Link Up/Down test indicates that the interface is operational, then FWSM performs network tests. The purpose of these tests is to generate network traffic to determine which (if either) unit has failed. At the start of each test, each unit clears its received packet count for its interfaces. At the conclusion of each test, each unit looks to see if it has received any traffic. If it has, the interface is considered operational. If one unit receives traffic for a test and the other unit does not, the unit that received no traffic is considered failed. If neither unit has received traffic, then the next test is used.

2. Network Activity test—A received network activity test. The unit counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops. If no traffic is received, the ARP test begins.
3. ARP test—A reading of the unit ARP cache for the 2 most recently acquired entries. One at a time, the unit sends ARP requests to these machines, attempting to stimulate network traffic. After each request, the unit counts all received traffic for up to 5 seconds. If traffic is received, the interface is considered operational. If no traffic is received, an ARP request is sent to the next machine. If at the end of the list no traffic has been received, the ping test begins.
4. Broadcast Ping test—A ping test that consists of sending out a broadcast ping request. The unit then counts all received packets for up to 5 seconds. If any packets are received at any time during this interval, the interface is considered operational and testing stops.

If all network tests fail for an interface, but this interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed. If the threshold for failed interfaces is met, then a failover occurs. If the other unit interface also fails all the network tests, then both interfaces go into the “Unknown” state and do not count towards the failover limit.

An interface becomes operational again if it receives any traffic. A failed FWSM returns to standby mode if the interface failure threshold is no longer met.


Note

If a failed unit does not recover and you believe it should not be failed, you can reset the state by entering the **failover reset** command. If the failover condition persists, however, the unit will fail again.

Configuring Failover

This section describes how to configure failover and includes the following topics:

- [Using Active/Standby Failover, page 13-18](#)
- [Using Active/Active Failover, page 13-23](#)
- [Configuring Failover Communication Authentication/Encryption, page 13-28](#)
- [Verifying the Failover Configuration, page 13-29](#)

Using Active/Standby Failover

This section provides step-by-step procedures for configuring Active/Standby failover. This section includes the following topics:

- [Prerequisites and General Information, page 13-18](#)
- [Primary/Secondary Status and Active/Standby Status, page 13-19](#)
- [Configuring Active/Standby Failover, page 13-19](#)
- [Configuring Optional Active/Standby Failover Settings, page 13-22](#)

See the “Failover Example Configurations” section on page B-18 for examples of typical failover configurations.

Prerequisites and General Information

Before you begin, verify the following:

- Both units have the proper license.
- If the primary unit is in single context mode, the secondary unit must also be in single context mode and also be in the same firewall mode as the primary unit.
- If the primary unit is in multiple context mode, the secondary unit must also be in multiple context mode. You do not need configure the firewall mode of the security contexts on the secondary unit because the failover and state links reside in the system context. The secondary unit obtains the security context configuration from the primary unit.

Before you begin, be aware of the following:

- The **mode** command does not get replicated to the secondary unit.

Primary/Secondary Status and Active/Standby Status

The main difference between the two units in a failover pair is related to which unit is active and which unit is standby, namely which IP addresses are used and which unit actively passes traffic. However, a few differences also exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary.

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC address is always coupled with the active IP addresses.
- By default, the MAC address used for the active FWSM comes from the Burned-in MAC address of the primary FWSM.
- Under certain circumstances, MAC addresses used for the active FWSMs are changed, such as in the following cases:
 - Case 1—The primary FWSM in a failover pair is replaced with a new FWSM
 - Case 2—The secondary FWSM boots and becomes active because it did not detect the primary FWSM.

In Case 1 above, if the primary FWSM is replaced, then as soon as it becomes part of the failover set, the secondary/active FWSM changes the MAC addresses to those of the new primary FWSM.

In Case 2 above, if the secondary FWSM boots without knowing the Burned-in MAC address of the primary FWSM, then it uses its own Burned-in MAC address until it hears from the primary, at which time it swaps the MAC addresses.

Any time the secondary/active FWSM applies new MAC addresses, it sends out gratuitous ARPs for the interface IP addresses but not for the other IP addresses that it owns. These other IP addresses consist of global IP addresses in static and global statements. Therefore, if the Burned-in MAC address of the secondary/active FWSM changes, you must clear the ARP table on the devices that Layer 2 adjacent to the FWSM. Otherwise, the ARP entries for the global IP addresses on those devices will be old and invalid.

Configuring Active/Standby Failover

This section describes how to configure Active/Standby failover. You must configure the secondary unit to recognize the failover link before the secondary unit can obtain the running configuration from the primary unit.

This section includes the following topics:

- [Configuring the Primary Unit, page 13-20](#)

- [Configuring the Secondary Unit, page 13-21](#)

Configuring the Primary Unit

Follow these steps to configure the primary unit in an Active/Standby failover configuration. These steps provide the minimum configuration needed to enable failover on the primary unit. For multiple context mode, all steps are performed in the system execution space unless otherwise noted.

To configure the primary unit in an Active/Standby failover pair, perform the following steps:

- Step 1** If you have not done so already, configure the active and standby IP addresses for each interface (routed mode) or for the management address (transparent mode). The standby IP address is used on the FWSM that is currently the standby unit. It must be in the same subnet as the active IP address.



Note Do not configure an IP address for the failover link or for the state link (if you are going to use Stateful Failover).

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```



Note In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context.

- Step 2** Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

- Step 3** Define the failover interface.

- a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name vlan vlan
```

The *if_name* argument assigns a name to the interface specified by the *vlan* argument.

- b. Assign the active and standby IP address to the failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The failover link IP address and MAC address do not change at failover. The active IP address for the failover link always stays with the primary unit, while the standby IP address stays with the secondary unit.

- Step 4** (Optional) To enable Stateful Failover, configure the state link. The state link must be configured on an unused interface.

- a. Specify the interface to be used as state link:

```
hostname(config)# failover link if_name [vlan vlan]
```



Note If the state link uses the failover link, then you only need to supply the *if_name* argument.

The *if_name* argument assigns a logical name to the interface specified by the *vlan* argument. This interface should not be used for any other purpose except, optionally, the failover link.

- b. Assign an active and standby IP address to the state link.



Note If the state link uses the failover link, skip this step. You have already defined the failover link active and standby IP addresses.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The state link IP address and MAC address do not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- Step 5** To enable monitoring on an interface, enter the following command:

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

The maximum number of interfaces to monitor on the FWSM (divided between all contexts) is 250.



Note In multiple context mode, you must configure interface monitoring from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config)#`, where *context* is the name of the current context.

- Step 6** Enable failover:

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

- Step 7** Save the configuration:

```
hostname(config)# write memory
```



Note In multiple context mode, enter **write memory all** in the system execution space to save all context configurations.

Configuring the Secondary Unit

The only configuration required on the secondary unit is for the failover interface. The secondary unit requires these commands to initially communicate with the primary unit. After the primary unit sends its configuration to the secondary unit, the only permanent difference between the two configurations is the **failover lan unit** command, which identifies each unit as primary or secondary.

For multiple context mode, all steps are performed in the system execution space unless noted otherwise.

To configure the secondary unit, perform the following steps:

- Step 1** Define the failover interface. Use the same settings as you used for the primary unit.

- a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name vlan vlan
```

The *if_name* argument assigns a name to the interface specified by the *vlan* argument.

- b. Assign the active and standby IP address to the failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```



Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface on the primary unit.

- Step 2** (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```



Note This step is optional because by default units are designated as secondary unless previously configured.

- Step 3** Enable failover:

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

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages “Beginning configuration replication: Sending to mate” and “End Configuration Replication to mate” appear on the active unit console.

- Step 4** After the running configuration has completed replication, save the configuration to Flash memory:

```
hostname(config)# write memory
```

Configuring Optional Active/Standby Failover Settings

You can configure the following optional Active/Standby failover setting when you are initially configuring failover or after failover has already been configured. Unless otherwise noted, the commands should be entered on the active unit.

This section includes the following topics:

- [Enabling HTTP Replication with Stateful Failover, page 13-22](#)
- [Configuring Interface and Unit Poll Times, page 13-23](#)
- [Configuring Failover Criteria, page 13-23](#)

Enabling HTTP Replication with Stateful Failover

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

Enter the following command in global configuration mode to enable HTTP state replication when Stateful Failover is enabled:

```
hostname(config)# failover replication http
```

Configuring Interface and Unit Poll Times

FWSM monitors both unit and interface health for failover. You can configure the amount of time between hello messages when monitoring interface and unit health. Decreasing the poll time allows an interface or unit failure to be detected more quickly, but consumes more system resources.

To change the interface poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime interface seconds
```

To change the unit poll time, enter the following command in global configuration mode:

```
hostname(config)# failover polltime seconds
```

To change the unit hold time, enter the following command in global configuration mode:

```
hostname(config)# failover holdtime seconds
```

The defaults are as follows:

- The interface **poll time** is 15 seconds.
- The unit **poll time** is 1 second.
- The **holdtime** time is 3 times the **poll time** (with a minimum value of 3 seconds) if you specify a **poll time** but do not specify a hold time with the **holdtime** keyword. If you specify a hold time using the **holdtime** keyword, it must be at least 3 times the **poll time**. If you enter the **clear configure failover** command, the hold time is 15 seconds.



Note

You cannot enter a holdtime value that is less than 3 times the unit poll time. With a faster poll time, the FWSM can detect failure and trigger failover faster. However, faster detection can cause unnecessary switchovers when the network is temporarily congested.

Configuring Failover Criteria

By default, failure of 50% of monitored interfaces causes failover. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs.

To change the default failover criteria, enter the following command in global configuration mode:

```
hostname(config)# failover interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Using Active/Active Failover

This section describes how to configure Active/Active failover.

This section includes the following topics:

- [Prerequisites, page 13-24](#)
- [Configuring Active/Active Failover, page 13-24](#)
- [Configuring Optional Active/Active Failover Settings, page 13-27](#)

See the “Failover Example Configurations” section on page B-18 for examples of typical failover configurations.

Prerequisites

Before you begin, verify the following:

- Both units have the proper license.
- Both units are in multiple context mode. You do not need configure the firewall mode of the security contexts on the secondary unit because the failover and state links reside in the system context. The secondary unit obtains the security context configuration from the primary unit.



Note

The **mode** command does not get replicated to the secondary unit.

Configuring Active/Active Failover

This section describes how to configure Active/Active failover. You must configure the secondary unit to recognize the failover link before the secondary unit can obtain the running configuration from the primary unit.

This section includes the following topics:

- [Configure the Primary Unit, page 13-24](#)
- [Configure the Secondary Unit, page 13-26](#)

Configure the Primary Unit

To configure the primary unit in an Active/Active failover configuration, perform the following steps:

- Step 1** If you have not done so already, configure the active and standby IP addresses for each interface (routed mode) or for the management address (transparent mode). The standby IP address is used on the FWSM that is currently the standby unit. It must be in the same subnet as the active IP address.



Note

Do not configure an IP address for the failover link or for the state link (if you are going to use Stateful Failover).

```
hostname(config-if)# ip address active_addr netmask standby standby_addr
```



Note

In multiple context mode, you must configure the interface addresses from within each context. Use the **changeto context** command to switch between contexts. The command prompt changes to `hostname/context(config-if)#`, where *context* is the name of the current context.

- Step 2** Configure the basic failover parameters in the system execution space.

- a. Designate the unit as the primary unit:

```
hostname(config)# failover lan unit primary
```

- b. Specify the failover link:

```
hostname(config)# failover lan interface if_name vlan vlan
```

The *if_name* argument assigns a logical name to the interface specified by the *vlan* argument. This interface should not be used for any other purpose (except, optionally, the state link).

- c. Specify the failover link active and standby IP addresses:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby IP address subnet mask. The failover link IP address does not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- Step 3** (Optional) To enable Stateful Failover, configure the state link. The state link must be configured on an unused interface.

- a. Specify the interface to be used as state link:

```
hostname(config)# failover link if_name [vlan vlan]
```

The *if_name* argument assigns a logical name to the interface specified by the *vlan* argument. This interface should not be used for any other purpose (except, optionally, the failover link).



Note If the state link uses the failover link, then you only need to supply the *if_name* argument.

- b. Assign an active and standby IP address to the state link.



Note If the state link uses the failover link, skip this step. You have already defined the failover link active and standby IP addresses.

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

The state link IP address does not change at failover. The active IP address always stays with the primary unit, while the standby IP address stays with the secondary unit.

- Step 4** Configure the failover groups. You can have at most two failover groups. The **failover group** command creates the specified failover group if it does not exist and enters the failover group configuration mode.

For each failover group, you need to specify whether the failover group has primary or secondary preference using the **primary** or **secondary** command. You can assign the same preference to both failover groups. For load balancing configurations, you should assign each failover group a different unit preference.

The following example assigns failover group 1 a primary preference and failover group 2 a secondary preference:

```
hostname(config)# failover group 1
hostname(config-fover-group)# primary
hostname(config-fover-group)# exit
hostname(config)# failover group 2
hostname(config-fover-group)# secondary
hostname(config-fover-group)# exit
```

- Step 5** Assign each context to a failover group using the **join-failover-group** command in context configuration mode.

Any unassigned contexts are automatically assigned to failover group 1. The admin context is always a member of failover group 1.

Enter the following commands to assign each context to a failover group:

```
hostname(config)# context context_name
hostname(config-context)# join-failover-group {1 | 2}
```

Step 6 Enable failover:

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

Step 7 To enable monitoring on an interface, change to the context and enter the following command:

```
hostname(config)# changeto context context_name
hostname(config)# monitor-interface interface_name
```

The maximum number of interfaces to monitor on the FWSM (divided between all contexts) is 250.

Configure the Secondary Unit

You need to configure the secondary unit to recognize the failover link. This allows the secondary unit to communicate with and receive the running configuration from the primary unit.

To configure the secondary unit in an Active/Active failover configuration, perform the following steps:

Step 1 Define the failover interface. Use the same settings as you used for the primary unit.

- a. Specify the interface to be used as the failover interface:

```
hostname(config)# failover lan interface if_name vlan vlan
```

The *if_name* argument assigns a logical name to the interface specified by the *vlan* argument.

- b. Assign the active and standby IP address to the failover link:

```
hostname(config)# failover interface ip if_name ip_addr mask standby ip_addr
```



Note Enter this command exactly as you entered it on the primary unit when you configured the failover interface.

The standby IP address must be in the same subnet as the active IP address. You do not need to identify the standby address subnet mask.

Step 2 (Optional) Designate this unit as the secondary unit:

```
hostname(config)# failover lan unit secondary
```



Note This step is optional because by default units are designated as secondary unless previously configured otherwise.

Step 3 Enable failover:

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

After you enable failover, the active unit sends the configuration in running memory to the standby unit. As the configuration synchronizes, the messages `Beginning configuration replication: Sending to mate` and `End Configuration Replication to mate` appear on the active unit console.

Step 4 After the running configuration has completed replication, enter the following command to save the configuration to Flash memory:

```
hostname(config)# write memory
```

- Step 5** If necessary, force any failover group that is active on the primary to the active state on the secondary unit. To force a failover group to become active on the secondary unit, enter the following command in the system execution space on the primary unit:

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

The *group_id* argument specifies the group you want to become active on the secondary unit.

Configuring Optional Active/Active Failover Settings

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

This section includes the following topics:

- [Configuring Failover Group Preemption, page 13-27](#)
- [Enabling HTTP Replication with Stateful Failover, page 13-27](#)
- [Configuring Interface and Unit Poll Times, page 13-28](#)
- [Configuring Failover Criteria, page 13-28](#)

Configuring Failover Group Preemption

Assigning a primary or secondary priority to a failover group specifies which unit the failover group becomes active on when both units boot simultaneously. However, if one unit boots before the other, then both failover groups become active on that unit. When the other unit comes online, any failover groups that have the unit as a priority do not become active on that unit unless manually forced over, a failover occurs, or the failover group is configured with the **preempt** command. The **preempt** command causes a failover group to become active on the designated unit automatically when that unit becomes available.

Enter the following commands to configure preemption for the specified failover group:

```
hostname(config)# failover group {1 | 2}  
hostname(config-fover-group)# preempt [delay]
```

You can enter an optional *delay* value, which specifies the number of seconds the failover group remains active on the current unit before automatically becoming active on the designated unit.

Enabling HTTP Replication with Stateful Failover

To allow HTTP connections to be included in the state information, you need to enable HTTP replication. Because HTTP connections are typically short-lived, and because HTTP clients typically retry failed connection attempts, HTTP connections are not automatically included in the replicated state information. You can use the **replication http** command to cause a failover group to replicate HTTP state information when Stateful Failover is enabled.

To enable HTTP state replication for a failover group, enter the following command. This command only affects the failover group in which it was configured. To enable HTTP state replication for both failover groups, you must enter this command in each group. This command should be entered in the system execution space.

```
hostname(config)# failover group {1 | 2}  
hostname(config-fover-group)# replication http
```

Configuring Interface and Unit Poll Times

You can configure the amount of time between hello messages when monitoring the health of the interfaces in a failover group. Decreasing the interface poll time allows failover to occur faster when an interface fails, but consumes more system resources.

To change the default interface poll time, enter the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# polltime interface seconds
```

The unit poll time specifies the amount of time between hello messages sent across the failover link to determine the health of the peer unit. Decreasing the unit poll time allows a failed unit to be detected faster, but consumes more system resources. To change the unit poll time, enter the following command in global configuration mode of the system execution space:

```
hostname(config)# failover polltime seconds
```

Configuring Failover Criteria

By default, failure of 50% of monitored interfaces causes failover. You can specify a specific number of interfaces or a percentage of monitored interfaces that must fail before a failover occurs. The failover criteria is specified on a failover group basis.

To change the default failover criteria for the specified failover group, enter the following commands:

```
hostname(config)# failover group {1 | 2}
hostname(config-fover-group)# interface-policy num[%]
```

When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

Configuring Failover Communication Authentication/Encryption

You can encrypt and authenticate the communication between failover peers by specifying a shared secret or hexadecimal 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 FWSM 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 FWSM to terminate VPN tunnels.

Enter the following command on the active unit of an Active/Standby failover pair or on the unit that has failover group 1 in the active state of an Active/Active failover pair:

```
hostname(config)# failover key {secret | hex key}
```

The *secret* argument specifies a shared secret that is used to generate the encryption key. It can be from 1 to 63 characters. The characters can be any combination of numbers, letters, or punctuation. The *hex key* argument specifies a hexadecimal encryption key. The key must be 32 hexadecimal characters (0-9, a-f).

**Note**

To prevent the failover key from being replicated to the peer unit in clear text for an existing failover configuration, disable failover on the active unit (or in the system execution space on the unit that has failover group 1 in the active state), enter the failover key on both units, and then reenable failover. When failover is reenabled, the failover communication will be encrypted with the key.

For new failover configurations, the **failover key** command should be part of the initial failover pair configuration.

Verifying the Failover Configuration

This section describes how to verify your failover configuration. This section includes the following topics:

- [Viewing Failover Status, page 13-29](#)
- [Viewing Monitored Interfaces, page 13-37](#)
- [Viewing the Failover Configuration, page 13-37](#)
- [Testing the Failover Functionality, page 13-37](#)

Viewing Failover Status

This section describes how to view the failover status. On each unit you can verify the failover status by entering the **show failover** command. The information displayed depends upon whether you are using Active/Standby or Active/Active failover.

This section includes the following topics:

- [Viewing Failover Status for Active/Standby, page 13-29](#)
- [Viewing Failover Status for Active/Active, page 13-33](#)

Viewing Failover Status for Active/Standby

The following is sample output from the **show failover** command for Active/Standby failover. [Table 13-4](#) provides descriptions for the information shown.

```
hostname# show failover

Failover On
Failover unit Primary
Failover LAN Interface: fover Vlan 100(up)
Unit Poll frequency 1 seconds, holdtime 3 seconds
Interface Poll frequency 15 seconds
Interface Policy 1
Monitored Interfaces 2 of 250 maximum
failover replication http
Last Failover at: 22:44:03 UTC Dec 8 2004
  This host: Primary - Active
    Active time: 13434 (sec)
    Interface inside (10.130.9.3): Normal
    Interface outside (10.132.9.3): Normal
  Other host: Secondary - Standby Ready
    Active time: 0 (sec)
    Interface inside (10.130.9.4): Normal
    Interface outside (10.132.9.4): Normal
```

```

Stateful Failover Logical Update Statistics
Link : fover Vlan100 (up)
Stateful Obj   xmit      xerr      rcv      rerr
General       1950      0         1733     0
sys cmd       1733      0         1733     0
up time       0         0         0         0
RPC services  0         0         0         0
TCP conn      6         0         0         0
UDP conn      0         0         0         0
ARP tbl       106      0         0         0
Xlate_Timeout 0         0         0         0
VPN IKE upd   15        0         0         0
VPN IPSEC upd 90        0         0         0
VPN CTCP upd  0         0         0         0
VPN SDI upd   0         0         0         0
VPN DHCP upd  0         0         0         0

Logical Update Queue Information
          Cur      Max      Total
Recv Q:  0         2       1733
Xmit Q:   0         2      15225

```

In multiple context mode, using the **show failover** command in a security context displays the failover information for that context. The information is similar to the information shown when using the command in single context mode. Instead of showing the active/standby status of the unit, it displays the active/standby status of the context. [Table 13-4](#) provides descriptions for the information shown.

```

Failover On
Last Failover at: 04:03:11 UTC Jan 4 2003
  This context: Negotiation
    Active time: 1222 (sec)
    Interface outside (192.168.5.121): Normal
    Interface inside (192.168.0.1): Normal
  Peer context: Not Detected
    Active time: 0 (sec)
    Interface outside (192.168.5.131): Normal
    Interface inside (192.168.0.11): Normal

```

```

Stateful Failover Logical Update Statistics
Status: Configured.
Stateful Obj   xmit      xerr      rcv      rerr
RPC services  0         0         0         0
TCP conn      99        0         0         0
UDP conn      0         0         0         0
ARP tbl       22        0         0         0
Xlate_Timeout 0         0         0         0
GTP PDP       0         0         0         0
GTP PDMCB     0         0         0         0

```

Table 13-4 Show Failover Display Description

Field	Options
Failover	<ul style="list-style-type: none"> • On • Off
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the name of the failover link.

Table 13-4 Show Failover Display Description (continued)

Field	Options
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	<i>n</i> seconds The number of seconds you set with the failover polltime interface command. The default is 15 seconds.
Interface Policy	Displays the number or percentage of interfaces that must fail to trigger failover.
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
failover replication http	Displays if HTTP state replication is enabled for Stateful Failover.
Last Failover at:	The date and time of the last failover in the following form: <i>hh:mm:ss UTC DayName Month Day yyyy</i> UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host:	For each host, the display shows the following information.
Other host:	
Primary or Secondary	<ul style="list-style-type: none"> Active Standby
Active time:	<i>n</i> (sec) The amount of time the unit has been active. This time is cumulative, so the standby unit, if it was active in the past, will also show a value.
Interface <i>name</i> (<i>n.n.n.n</i>):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: <ul style="list-style-type: none"> Failed—The interface has failed. No Link—The interface line protocol is down. Normal—The interface is working correctly. Link Down—The interface has been administratively shut down. Unknown—FWSM cannot determine the status of the interface. Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	<ul style="list-style-type: none"> <i>interface_name</i>—The interface used for the Stateful Failover link. Unconfigured—You are not using Stateful Failover. up—The interface is up and functioning. down—The interface is either administratively shutdown or is physically down. failed—The interface has failed and is not passing stateful data.

Table 13-4 Show Failover Display Description (continued)

Field	Options
Stateful Obj	<p>For each field type, the following statistics are shown. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit.</p> <ul style="list-style-type: none"> • xmit—Number of transmitted packets to the other unit. • xerr—Number of errors that occurred while transmitting packets to the other unit. • rcv—Number of received packets. • rerr—Number of errors that occurred while receiving packets from the other unit.
General	Sum of all stateful objects.
sys cmd	Logical update system commands; for example, LOGIN and Stay Alive.
up time	Up time, which the active unit passes to the standby unit.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPCB	GTP PDPCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	<p>For each field type, the following statistics are used:</p> <ul style="list-style-type: none"> • Cur—Current number of packets • Max—Maximum number of packets • Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Viewing Failover Status for Active/Active

The following is sample output from the **show failover** command for Active/Active failover. [Table 13-5](#) provides descriptions for the information shown.

```

hostname# show failover

Failover On
Failover unit Primary
Failover LAN Interface: fover Vlan 100 (up)
Unit Poll frequency 1 seconds, holdtime 15 seconds
Interface Poll frequency 4 seconds
Interface Policy 1
Monitored Interfaces 8 of 250 maximum
failover replication http
Group 1 last failover at: 13:40:18 UTC Dec 9 2004
Group 2 last failover at: 13:40:06 UTC Dec 9 2004

This host:      Primary
Group 1         State:          Active
                Active time:    2896 (sec)
Group 2         State:          Standby Ready
                Active time:    0 (sec)

admin Interface outside (10.132.8.5): Normal
admin Interface third (10.132.9.5): Normal
admin Interface inside (10.130.8.5): Normal
admin Interface fourth (10.130.9.5): Normal
ctx1 Interface outside (10.1.1.1): Normal
ctx1 Interface inside (10.2.2.1): Normal
ctx2 Interface outside (10.3.3.2): Normal
ctx2 Interface inside (10.4.4.2): Normal

Other host:     Secondary
Group 1         State:          Standby Ready
                Active time:    190 (sec)
Group 2         State:          Active
                Active time:    3322 (sec)

admin Interface outside (10.132.8.6): Normal
admin Interface third (10.132.9.6): Normal
admin Interface inside (10.130.8.6): Normal
admin Interface fourth (10.130.9.6): Normal
ctx1 Interface outside (10.1.1.2): Normal
ctx1 Interface inside (10.2.2.2): Normal
ctx2 Interface outside (10.3.3.1): Normal
ctx2 Interface inside (10.4.4.1): Normal

Stateful Failover Logical Update Statistics
Link : fover Vlan100 (up)
Stateful Obj   xmit      xerr      rcv        rerr
General        1973      0          1895       0
sys cmd        380       0          380        0
up time        0         0          0          0
RPC services   0         0          0          0
TCP conn       1435     0          1450       0
UDP conn       0         0          0          0
ARP tbl        124      0          65         0
Xlate_Timeout  0         0          0          0
VPN IKE upd    15        0          0          0
VPN IPSEC upd  90        0          0          0
VPN CTCP upd   0         0          0          0
VPN SDI upd    0         0          0          0
VPN DHCP upd   0         0          0          0

```

```

Logical Update Queue Information
          Cur      Max      Total
Recv Q:    0       1      1895
Xmit Q:    0       0      1940

```

The following is sample output from the **show failover group** command for Active/Active failover. The information displayed is similar to that of the **show failover** command, but limited to the specified group. [Table 13-5](#) provides descriptions for the information shown.

```

hostname# show failover group 1

Last Failover at: 04:09:59 UTC Jan 4 2005

This host:   Secondary
            State:      Active
            Active time: 186 (sec)

            admin Interface outside (192.168.5.121): Normal
            admin Interface inside (192.168.0.1): Normal

Other host:  Primary
            State:      Standby
            Active time: 0 (sec)

            admin Interface outside (192.168.5.131): Normal
            admin Interface inside (192.168.0.11): Normal

Stateful Failover Logical Update Statistics
Status: Configured.
RPC services    0          0          0          0
TCP conn        33         0          0          0
UDP conn        0          0          0          0
ARP tbl         12         0          0          0
Xlate_Timeout   0          0          0          0
GTP PDP         0          0          0          0
GTP PDMCB       0          0          0          0

```

Table 13-5 Show Failover Display Description

Field	Options
Failover	<ul style="list-style-type: none"> • On • Off
Failover Unit	Primary or Secondary.
Failover LAN Interface	Displays the name of the failover link.
Unit Poll frequency	Displays the number of seconds between hello messages sent to the peer unit and the number of seconds during which the unit must receive a hello message on the failover link before declaring the peer failed.
Interface Poll frequency	<p><i>n</i> seconds</p> <p>The number of seconds you set with the failover polltime interface command. The default is 15 seconds.</p>
Interface Policy	Displays the number or percentage of interfaces that must fail before triggering failover.

Table 13-5 Show Failover Display Description (continued)

Field	Options
Monitored Interfaces	Displays the number of interfaces monitored out of the maximum possible.
Group 1 Last Failover at: Group 2 Last Failover at:	The date and time of the last failover for each group in the following form: <i>hh:mm:ss UTC DayName Month Day yyyy</i> UTC (Coordinated Universal Time) is equivalent to GMT (Greenwich Mean Time).
This host: Other host:	For each host, the display shows the following information.
Role	Primary or Secondary
System State	<ul style="list-style-type: none"> Active or Standby Ready Active Time in seconds
Group 1 State Group 2 State	<ul style="list-style-type: none"> Active or Standby Ready Active Time in seconds
<i>context</i> Interface <i>name</i> (<i>n.n.n.n</i>):	For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions: <ul style="list-style-type: none"> Failed—The interface has failed. No link—The interface line protocol is down. Normal—The interface is working correctly. Link Down—The interface has been administratively shut down. Unknown—FWSM cannot determine the status of the interface. Waiting—Monitoring of the network interface on the other unit has not yet started.
Stateful Failover Logical Update Statistics	The following fields relate to the Stateful Failover feature. If the Link field shows an interface name, the Stateful Failover statistics are shown.
Link	<ul style="list-style-type: none"> <i>interface_name</i>—The interface used for the Stateful Failover link. Unconfigured—You are not using Stateful Failover. up—The interface is up and functioning. down—The interface is either administratively shutdown or is physically down. failed—The interface has failed and is not passing stateful data.

Table 13-5 Show Failover Display Description (continued)

Field	Options
Stateful Obj	For each field type, the following statistics are used. They are counters for the number of state information packets sent between the two units; the fields do not necessarily show active connections through the unit. <ul style="list-style-type: none"> xmit—Number of transmitted packets to the other unit xerr—Number of errors that occurred while transmitting packets to the other unit rcv—Number of received packets rerr—Number of errors that occurred while receiving packets from the other unit
General	Sum of all stateful objects.
sys cmd	Logical update system commands; for example, LOGIN and Stay Alive.
up time	Up time, which the active unit passes to the standby unit.
RPC services	Remote Procedure Call connection information.
TCP conn	TCP connection information.
UDP conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
L2BRIDGE tbl	Layer 2 bridge table information (transparent firewall mode only).
Xlate_Timeout	Indicates connection translation timeout information.
VPN IKE upd	IKE connection information.
VPN IPSEC upd	IPSec connection information.
VPN CTCP upd	cTCP tunnel connection information.
VPN SDI upd	SDI AAA connection information.
VPN DHCP upd	Tunneled DHCP connection information.
GTP PDP	GTP PDP update information. This information appears only if inspect GTP is enabled.
GTP PDPMCB	GTP PDPMCB update information. This information appears only if inspect GTP is enabled.
Logical Update Queue Information	For each field type, the following statistics are used: <ul style="list-style-type: none"> Cur—Current number of packets Max—Maximum number of packets Total—Total number of packets
Recv Q	The status of the receive queue.
Xmit Q	The status of the transmit queue.

Viewing Monitored Interfaces

To view the status of monitored interfaces, enter the following command. In single context mode, enter this command in global configuration mode. In multiple context mode, enter this command within a context.

```
primary/context(config)# show monitor-interface
```

For example:

```
hostname/context(config)# show monitor-interface
  This host: Primary - Active
    Interface outside (192.168.1.2): Normal
    Interface inside (10.1.1.91): Normal
  Other host: Secondary - Standby
    Interface outside (192.168.1.3): Normal
    Interface inside (10.1.1.100): Normal
```

Viewing the Failover Configuration

To view the failover commands in the running configuration, enter the following command:

```
hostname(config)# show running-config failover
```

All of the failover commands are displayed. On units running multiple context mode, enter this command in the system execution space. Entering **show running-config all failover** displays the failover commands in the running configuration and includes commands for which you have not changed the default value.

Testing the Failover Functionality

To test failover functionality, perform the following steps:

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


```
hostname(config)# no failover active
```
 - For Active/Active failover, enter the following command on the unit where failover group containing the interface connecting your hosts is active:


```
hostname(config)# no failover active group group_id
```
- Step 3** Use FTP to send another file between the same two hosts.
- Step 4** If the test was not successful, enter the **show failover** command to check the failover status.
- Step 5** When you are finished, you can restore the unit or failover group to active status by enter the following command:
- For Active/Standby failover, enter the following command on the active unit:


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

- For Active/Active failover, enter the following command on the unit where the failover group containing the interface connecting your hosts is active:

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

Controlling and Monitoring Failover

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

- [Forcing Failover, page 13-38](#)
- [Disabling Failover, page 13-39](#)
- [Disabling Configuration Synchronization, page 13-39](#)
- [Restoring a Failed Unit or Failover Group, page 13-39](#)
- [Monitoring Failover, page 13-39](#)

Forcing Failover

To force the standby unit or failover group to become active, enter one of the following commands:

- For Active/Standby failover:

Enter the following command on the standby unit:

```
hostname# failover active
```

Or enter the following command on the active unit:

```
hostname# no failover active
```

- For Active/Active failover:

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

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

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

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

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

```
hostname# failover active
```

Disabling Failover

To disable failover, enter the following command:

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

Disabling failover on an Active/Standby pair causes the active and standby state of each unit to be maintained until you restart. For example, the standby unit remains in standby mode so that both units do not start passing traffic. To make the standby unit active (even with failover disabled), see the “[Forcing Failover](#)” section on page 13-38.

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

Disabling Configuration Synchronization

Management applications may lose connectivity when upgrading the FWSM with complex configurations. This can result in incomplete configuration files being applied to the standby FWSM. You can disable the automatic configuration synchronization in order to avoid incomplete configurations being applied to the standby FWSM. You need to disable configuration synchronization when upgrading a software image or changing the configuration on the active FWSM to verify that the configuration files are complete before the configuration is synchronized with the standby FWSM configuration. After you verify that the configuration is complete, reenable configuration synchronization.

To disable configuration synchronization, enter this command:

```
hostname(config)# failover suspend-config-sync
```

To reenable configuration synchronization, use the **no** form of the this command.

Restoring a Failed Unit or Failover Group

To restore a failed unit to an unfailed state, enter the following command:

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

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

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

Restoring a failed unit or group to an unfailed state does not automatically make it active; restored units or groups remain in the standby state until made active by failover (forced or natural). An exception is a failover group configured with the **preempt** command. If previously active, a failover group will become active if it is configured with the **preempt** command and if the unit on which it failed is its preferred unit.

Monitoring Failover

When a failover occurs, both FWSMs send out system messages. This section includes the following topics:

- [Failover System Messages, page 13-40](#)
- [Debug Messages, page 13-40](#)
- [SNMP, page 13-40](#)

Failover System Messages

FWSM issues a number of system messages related to failover at priority level 2, which indicates a critical condition. To view these messages, see the *Catalyst 6500 Series Switch and Cisco 7600 Series Router Firewall Services Module Logging Configuration and System Log Messages* to enable logging and to see descriptions of the system messages.

**Note**

During switchover, failover will logically shut down and then bring up interfaces, generating system log messages 411001 and 411002. This is normal activity.

Debug Messages

To see debug messages, enter the **debug fover** command. See the *Catalyst 6500 Series Switch and Cisco 7600 Series Router Firewall Services Module Command Reference* for more information.

**Note**

Because debugging output is assigned high priority in the CPU process, it can drastically affect system performance. For this reason, use the **debug fover** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco TAC.

SNMP

To receive SNMP syslog traps for failover, configure the SNMP agent to send SNMP traps to SNMP management stations, define a syslog host, and compile the Cisco syslog MIB into your SNMP management station. See the **snmp-server** and **logging** commands in the *Catalyst 6500 Series Switch and Cisco 7600 Series Router Firewall Services Module Command Reference* for more information.