



Using Failover

This chapter describes the Firewall Services Module (FWSM) failover feature, which allows a secondary FWSM to take over the functionality of a failed FWSM. This chapter includes the following sections:

- [Understanding Failover, page 15-1](#)
- [Configuring Failover, page 15-14](#)
- [Verifying the Failover Configuration, page 15-18](#)
- [Forcing Failover, page 15-22](#)
- [Disabling Failover, page 15-22](#)
- [Monitoring Failover, page 15-23](#)
- [Frequently Asked Failover Questions, page 15-23](#)
- [Failover Configuration Example, page 15-26](#)



Note

See the [“Configuring the Switch for Failover”](#) section on [page 2-11](#) to configure the switch for failover.

Understanding Failover

This section describes how failover works and includes the following sections:

- [Failover Overview, page 15-2](#)
- [Regular and Stateful Failover, page 15-2](#)
- [Failover and State Links, page 15-3](#)
- [Module Placement, page 15-4](#)
- [Transparent Firewall Requirements, page 15-9](#)
- [Primary/Secondary Status and Active/Standby Status, page 15-10](#)
- [Configuration Replication, page 15-10](#)
- [Failover Triggers, page 15-11](#)
- [Failover Actions, page 15-12](#)
- [Failover Monitoring, page 15-13](#)

Failover Overview

The failover feature lets you use a standby FWSM to take over the functionality of a failed FWSM. Failover is compatible with both routed and transparent firewall modes, and with single and multiple context modes.



Note

The two FWSMs must have the same major (first number) and minor (second number) software version, license, and operating modes (routed or transparent, single or multiple context). You can use different maintenance versions (third numbers) during an upgrade process; for example, you can upgrade one unit from 2.2(1) to 2.2(2) and failover is still active. However, we recommend upgrading both units to the same version to ensure long-term compatibility. We do not guarantee full compatibility for failover when the maintenance versions differ.

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 takes over the active unit IP addresses (or, for transparent firewall, the management IP address) and MAC address, and it begins passing traffic. The FWSM has one MAC address for all interfaces. The unit that was active and 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, failover is unnoticed by the rest of the network. However, the host switch needs to reassociate the new active and standby chassis slots with their corresponding MAC addresses. The FWSM helps this process by sending out gratuitous ARPs on all its VLAN interfaces. (See the [“Primary/Secondary Status and Active/Standby Status” section on page 15-10](#) section for more information about MAC addresses).

The standby unit can effectively take over as the active unit because it has the same configuration, and it is assigned the same VLANs from the switch.



Note

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

Regular and Stateful Failover

The FWSM supports two types of failover:

- Regular failover—When a failover occurs, all active connections are dropped and clients need to reestablish connections when the new active unit takes over.
- Stateful failover—During normal operation, the active unit continually passes per-connection stateful information (for each context) to the standby unit. The interval between stateful information updates is 10 seconds, but if you set the unit polltime to be greater than 10 seconds, then that interval is used.

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 data:

- NAT translation table
- TCP connection states
- UDP connection states (for connections lasting at least 15 seconds)
- HTTP connection states (Optional)
- H.323, SIP, and MGCP UDP media connections
- ARP table
- (Transparent firewall mode only) MAC address table

Failover and State Links

This section describes the failover link and, for stateful failover, the state link, and it includes the following topics:

- [Failover Link, page 15-3](#)
- [State Link, page 15-3](#)

Failover Link

The two units constantly communicate over a failover link to determine the operating status of each unit. Communications over the failover link include the following data:

- The unit state (active or standby).
- Hello messages (also sent on all other interfaces).
- Configuration synchronization between the two units. (See the “[Configuration Replication](#)” section on [page 15-10](#) section for more information.)

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).

For multiple context mode, the failover link resides in the system configuration. This interface (and the state link, if used) is the only configurable interface in the system configuration.



Note

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

State Link

To use stateful failover, 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.

In multiple context mode, the state link resides in the system configuration. This interface and the failover interface are the only interfaces in the system configuration.



Note

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

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 15-4](#)
- [Inter-Chassis Failover, page 15-4](#)

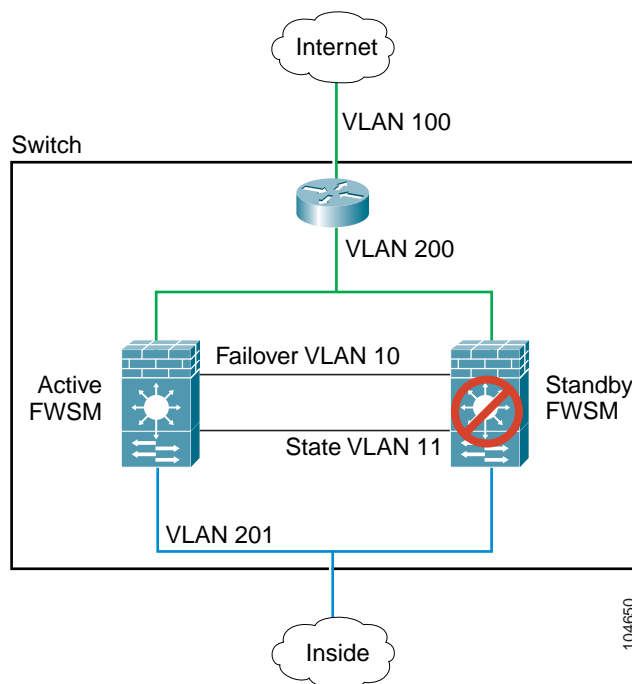
Intra-Chassis Failover

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 15-4.

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

[Figure 15-1](#) shows a typical intra-switch configuration.

Figure 15-1 Intra-Switch Failover



Inter-Chassis Failover

To protect against switch-level failure, you can install the secondary FWSM in a separate switch. The 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 the FWSMs, you must configure a trunk port between the two switches that carries all the FWSM VLANs. Because this trunk also accommodates FWSM traffic when a module fails, this trunk should be at least as large as the maximum amount of

traffic you expect to be inspected by the FWSM. The FWSM has an internal 6-Gbps EtherChannel to the switch, so if the FWSM runs at full capacity, the trunk between the two devices needs to include at least six 1-Gbps interfaces. EtherChannel aggregates the bandwidth of up to eight compatibly configured ports into a single logical link. (See the “[Adding a Trunk Between a Primary Switch and Secondary Switch](#)” section on page 2-12 for more information.)

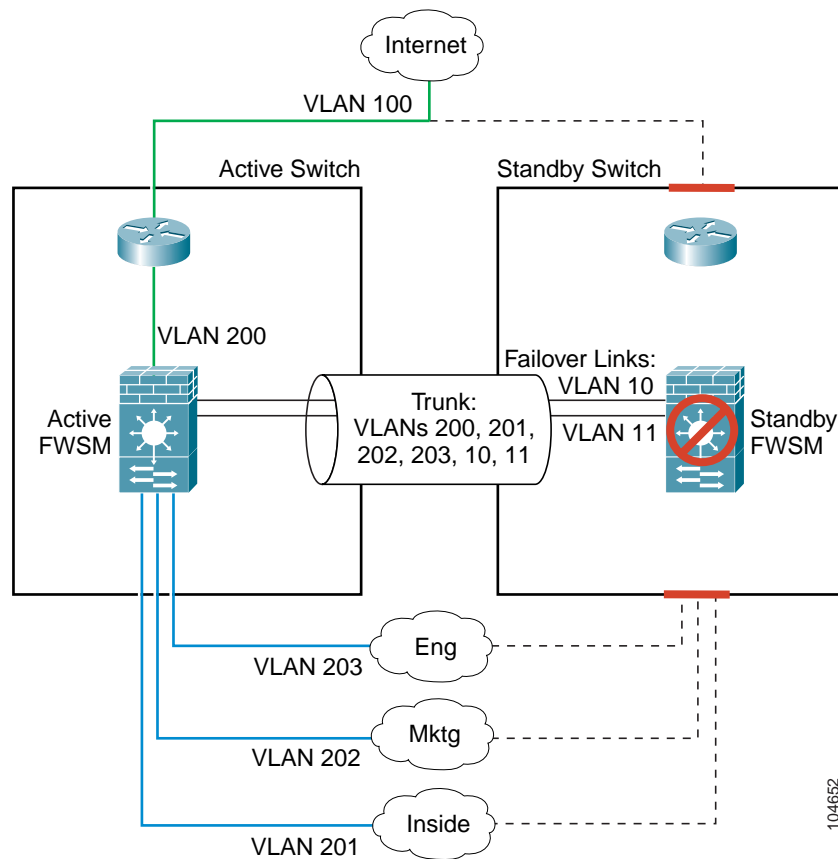
Figure 15-2 shows a typical switch and FWSM redundancy configuration. The Spanning Tree algorithm ensures that the VLANs pass through only one switch, which also contains the active FWSM. The trunk between the two switches carries all FWSM VLANs, including the failover and state links (VLANs 10 and 11).



Note

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

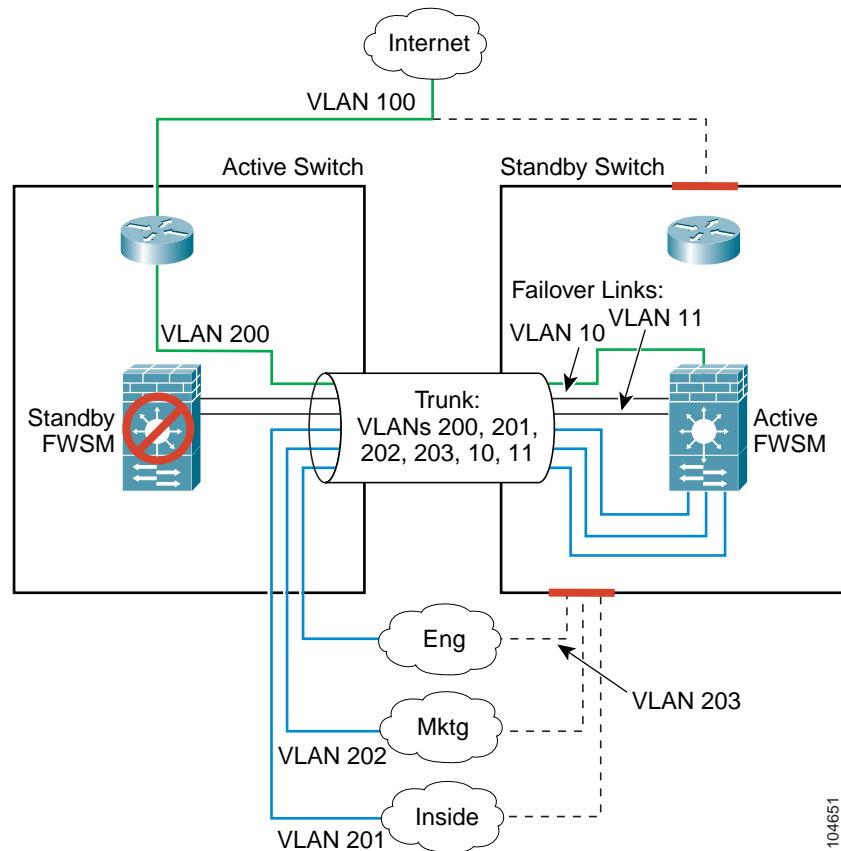
Figure 15-2 Normal Operation with Standby Units



The path that the traffic takes after failover depends on which device fails as follows:

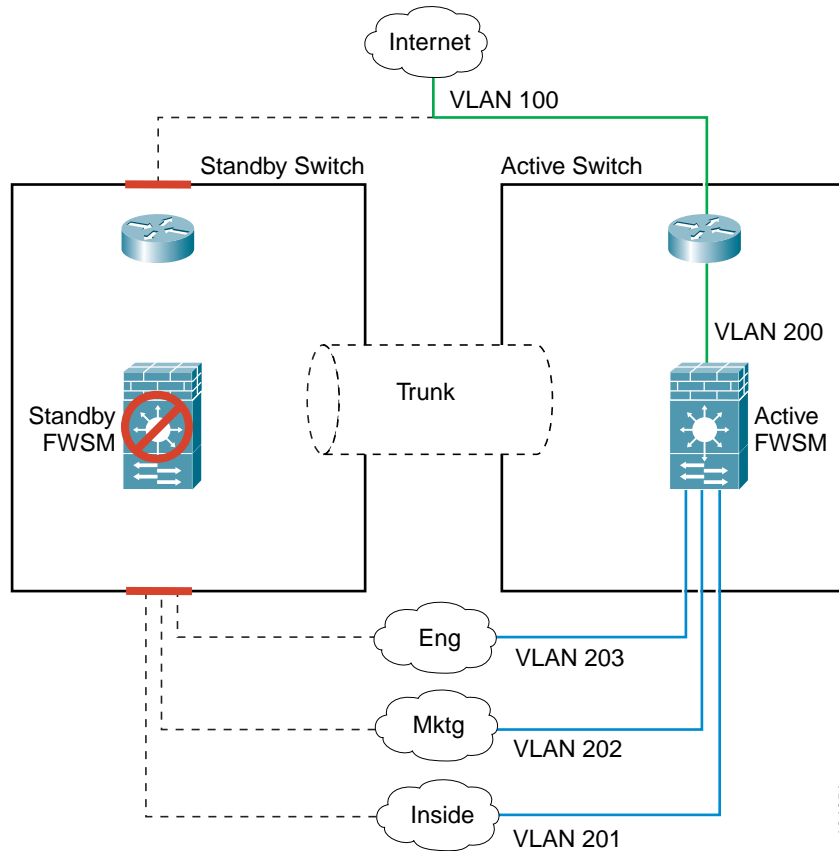
- **FWSM failure only**—If the primary FWSM fails, then the secondary FWSM becomes active. However, if the primary switch is still active, all VLAN traffic destined for the FWSM continues to enter the primary switch. The secondary (now active) FWSM receives and sends all traffic over the trunk (Figure 15-3).

Figure 15-3 FWSM Failure Only



- Switch/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 15-4).

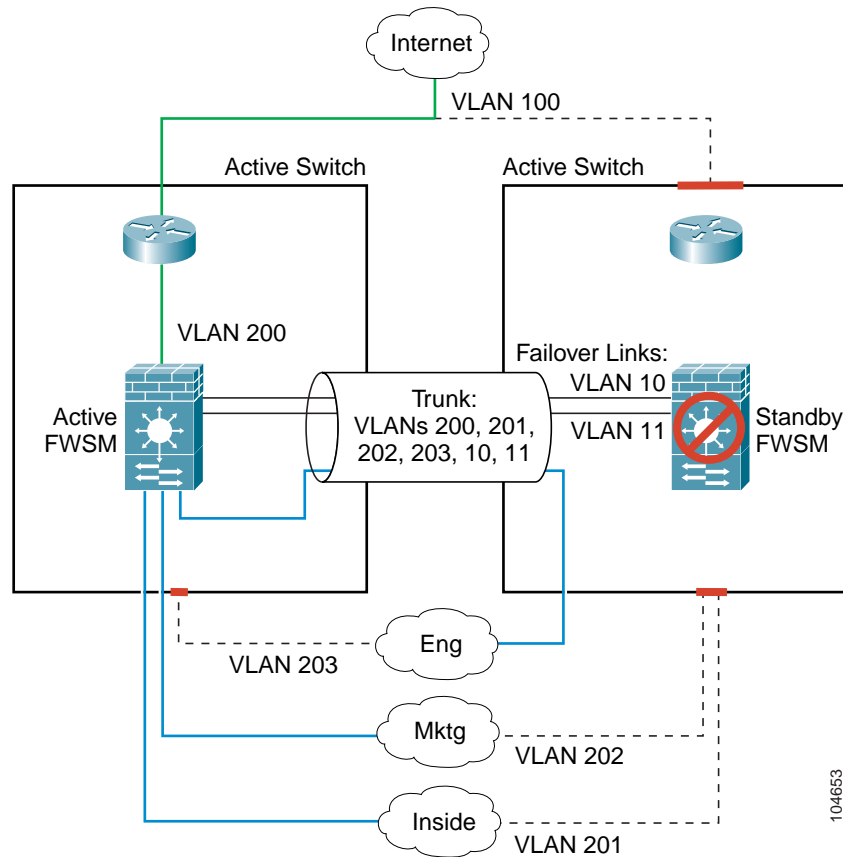
Figure 15-4 Switch/FWSM Failure



104654

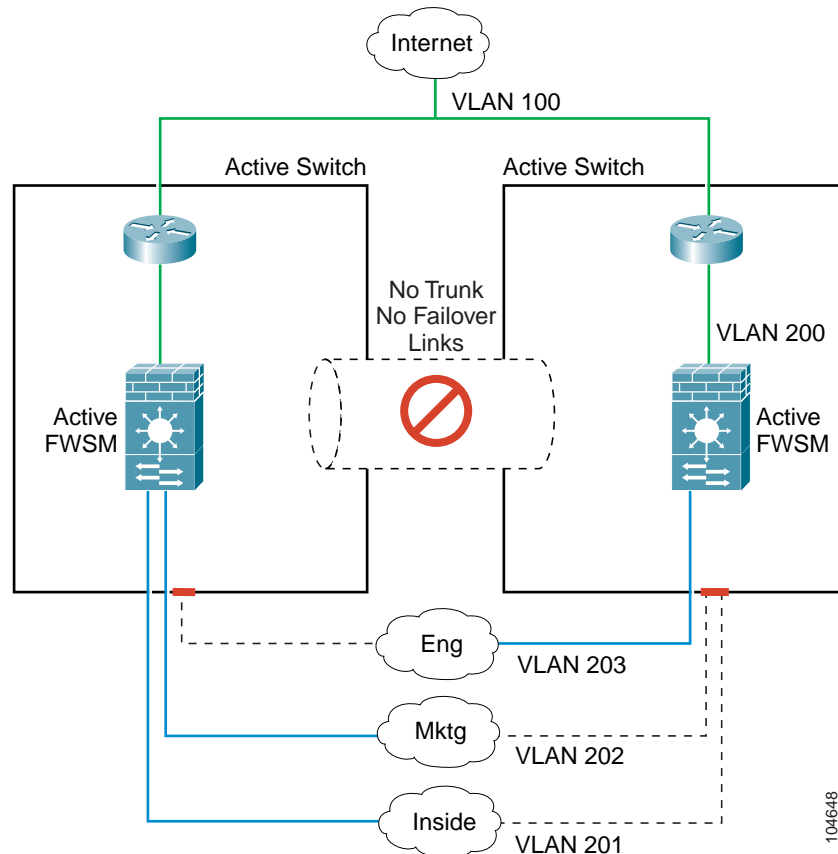
- Partial switch failure—If one or more interfaces on the switch fails, both switches would be partially active, but only one FWSM is active. The FWSM, which operates independently of the switch, has no reason to fail over because the active FWSM receives FWSM traffic from the secondary switch over the trunk (Figure 15-5).

Figure 15-5 Partial Switch Failure



- Trunk failure—If the trunk between the switches fails, all communication between the FWSMs terminates, which results in both FWSMs becoming active. Spanning Tree prevents any loops, however, and traffic is handled successfully by one or both FWSMs until you resolve the trunk issue (Figure 15-6).

Figure 15-6 Trunk Failure



104648

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 [“Chassis System Requirements” section on page 1-2](#) 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 Control List” section on page 10-16](#).

Loops can occur if both units are active at the same time, such as when both units are discovering each other’s presence, or due to a bad failover link as described in the [“Basic Failover Questions” section on page 15-24](#). Because the FWSM units 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 15-7](#)). 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 15-7 Potential Loops in Transparent Mode



Primary/Secondary Status and Active/Standby Status

The main differences 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.

Configuration Replication

The two FWSM units share almost the identical configuration. The configuration can be the same because it includes both the active IP addresses and the standby IP addresses. When a unit is active, it uses the active IP addresses; when a unit is standby, it uses the standby IP addresses.



Note

Because the configuration is the same on both units, the host names, usernames, and passwords are also the same.

The only difference in the configuration is the primary and secondary designation, although you must also pre-configure the failover link on the secondary unit before the units can communicate. All other configuration is automatically replicated from the active to the standby unit.

The active unit sends the configuration in running memory to the standby unit. On the standby unit, the configuration exists only in running memory. You can optionally save the configuration to Flash memory so that when you reboot the standby unit when the active unit is unavailable, the standby unit can become the active unit. To save the configuration to Flash memory after replication:

- For single context mode, enter the **copy running-config startup-config** command on the active unit. The command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- For multiple context mode, enter the **copy running-config startup-config** command on the active unit from the system execution space and within each context on disk. The command is replicated to the standby unit, which proceeds to write its configuration 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. (See the [“Downloading a Text Configuration”](#) section on page 16-6 for more information.)

Configuration replication from the active unit to the standby unit occurs in the following circumstances:

- When the standby unit completes its initial startup, it clears its running configuration (except for the failover commands that must be pre-configured and are not replicated), and the active unit sends its entire configuration to the standby unit.
- As commands are entered on the active unit, they are sent across the failover link to the standby unit. You do not have to save the active configuration to Flash memory to replicate the commands.
- If you enter the **write standby** command on the active unit, the standby unit clears its running configuration (except for the failover commands that must be pre-configured and are not replicated), and the active unit sends its entire configuration to the standby unit.

For 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.

**Note**

Changes made on the standby unit are not replicated to the active unit. If you enter a command on the standby unit, the 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.

When the replication starts, the FWSM console displays the message “Beginning configuration replication: Sending to mate,” and when it is complete, the FWSM displays the message “End Configuration Replication to mate.” During the replication, information cannot be entered on the FWSM terminal. Depending on the size of the configuration, replication can take several minutes.

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.

Because the FWSM can have a large number of interfaces, it cannot monitor every interface. Rather, you configure the FWSM to monitor a subset of interfaces. The FWSM fails over when a certain number of monitored interfaces fails; you configure the failure threshold to be an absolute value or a percentage of the total number of monitored interfaces.

See the “[Failover Monitoring](#)” section on page 15-13 for more information about when a unit or interface is considered to be failed.

Failover Actions

Table 15-1 shows the failover action for each failure event.

Table 15-1 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.
Trunk failure in inter-switch setup	No failover	No action	Become active	Both units become active if all communication between the modules is terminated, as in the case of a trunk failure. Neither unit receives hello messages, and both units assume the active role.

Failover Monitoring

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

- [Unit Health Monitoring, page 15-13](#)
- [Interface Monitoring, page 15-13](#)

Unit Health Monitoring

The 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. The FWSM retries a user-configurable number of times. The action the FWSM takes depends on the response from the other unit. See the following possible actions:

- If the FWSM receives a response on any interface, then it does not fail over.
- If the 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 the 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 should monitor important interfaces, for example, you might configure one context to monitor a shared VLAN (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 VLAN status. If the Link Up/Down test indicates that the VLAN is operational, then the 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:

- [Configuring the Primary Unit, page 15-14](#)
- [Configuring the Secondary Unit, page 15-17](#)

Configuring the Primary Unit

Follow these steps to configure the primary unit. For multiple context mode, all steps are performed in the system execution space unless otherwise noted.

**Note**

At any time during the procedure, you can enter the **show failover** command to see the failover status. See the [“Using the Show Failover Command” section on page 15-18](#) section for detailed information.

- Step 1** To configure the VLAN interface you are using for the failover link, enter the following command. For multiple context mode, enter this command in the system execution space:

```
primary(config)# failover lan interface interface_name vlan vlan
```

Note this setting because this command is the same on the secondary unit.

This VLAN should not be used for any other purpose (except, optionally, the state link) or be assigned to any switch ports. This VLAN does need to be assigned to the FWSM by the switch.

Do not assign an ACL to this interface; failover traffic is allowed automatically, and other traffic is denied.

- Step 2** To set the IP address of the failover interface, enter the following command:

```
primary(config)# failover interface ip failover_interface ip_address mask standby ip_address
```

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.

Note this setting because this command is the same on the secondary unit.

The failover 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 3** (Stateful failover only) To configure the VLAN interface you are using for the state link, enter the following command:

```
primary(config)# failover link interface_name [vlan vlan]
```

This VLAN should not be used for any other purpose (except, optionally, the failover link) or be assigned to any switch ports. This VLAN does need to be assigned to the FWSM by the switch.

If the interface is the same as the failover interface, you do not need to identify the VLAN.

Do not assign an ACL to this interface; failover traffic is allowed automatically, and other traffic is denied.

- Step 4** (Stateful failover only) To set the IP address of the state interface, enter the following command:

```
primary(config)# failover interface ip state_interface ip_address mask standby ip_address
```

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** (Stateful failover only—Optional), To allow HTTP connections to be included in the state information, enter the following command:

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

If you do not allow HTTP replication, then HTTP connections are disconnected at failover. HTTP connections are brief and frequent, and the state information, although updated constantly, might not include the latest HTTP states at failover. For this reason, you might want to disable HTTP replication to reduce the amount of traffic on the state link.

- Step 6** To set the threshold for monitored interface failure, enter the following command:

```
primary(config)# failover interface-policy number[%]
```

When the number of failed monitored interfaces meets the value you set with this command, then the FWSM fails over. You can set the following arguments:

- *number*—An absolute value.
- *number%*—A percentage of all monitored interfaces.

- Step 7** To set this FWSM as the primary unit, enter the following command:

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

**Note**

This command is the only configuration difference between the primary and secondary units, although you need to set other **failover** commands on the secondary unit before the FWSM can replicate the active configuration.

- Step 8** (Optional) To set 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, enter the following command:

```
primary(config)# failover polltime [unit] [msec] number [holdtime seconds]
```

See the following arguments:

- **polltime unit [msec] number**—The amount of time between hello messages. Set the time in seconds between 1 and 15. The default is 1 second. If you specify **msec**, you can set the time between 500 and 999 milliseconds.
- **holdtime number**—Sets the time during which a unit must receive a hello message on the failover link, or else the unit begins the testing process for peer failure. Set the time in seconds between 15 and 45. The default is the greater of 15 seconds or 3 times the polltime. You cannot enter a value that is less than 3 times the polltime.

For example, if the polltime is 1 second, then a 15 second holdtime means 15 hello messages are missed before the unit is tested for failure.



Note

The interval between stateful information updates is 10 seconds, but if you set the polltime to be greater than 10, then that interval is used.

- Step 9** (Optional) To set the time in seconds between hello messages on monitored interfaces, enter the following command:

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

If the interface does not receive five consecutive hello messages, the FWSM begins the testing process for interface failure. See the [“Failover Monitoring” section on page 15-13](#) for more information.

The *seconds* is an integer between 3 and 15. The default is 15 seconds, which means an interface receives no reply for 75 seconds (5 times the polling interval) before the interface is tested for failure.

- Step 10** To enable failover, enter the following command:

```
primary(config)# failover
```

- Step 11** (Multiple context mode only) To save the system configuration to Flash memory, enter the following command:

```
primary(config)# copy running-config startup-config
```

- Step 12** (Multiple context mode only) To change to a context to configure the standby IP addresses (if you have not already done so) and to configure the interface monitoring, enter the following command:

```
primary(config)# changeto context name
```

- Step 13** If you have not done so already, set the standby IP address for each interface (routed mode) or for the management IP address (transparent mode) by entering the command appropriate for your firewall mode.

- For routed mode, enter the following command for each interface:

```
primary/contexta(config)# ip address interface_name ip_address mask standby ip_address
```

- For transparent mode, enter the following command:

```
primary/contexta(config)# ip address ip_address mask standby ip_address
```

The standby IP address is used on the FWSM that is currently the standby unit.

To add the standby address, reenter the **ip address** command for each interface (or management IP address) and add the **standby ip_address** option.

This IP address must be in the same subnet as the active IP address. You do not identify the subnet mask. To check the current IP address settings, enter the **show ip address** command.

- Step 14** To enable monitoring on an interface, enter the following command:

```
primary/contexta(config)# monitor-interface interface_name
```

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

- Step 15** To save the configuration for the context (in multiple context mode) or for the single mode FWSM, enter the following command:

```
primary/contexta(config)# copy running-config startup-config
```

- Step 16** (Multiple context mode only) Repeat [Step 12](#) through [Step 15](#) for each context.

See the [“Failover Configuration Example”](#) section on page 15-26 for a typical failover configuration.

Configuring the Secondary Unit

The only configuration required for 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.



Note

At any time during the procedure, you can enter the **show failover** command to see the failover status. See the [“Using the Show Failover Command”](#) section on page 15-18 section for detailed information.

To configure the secondary unit, follow these steps:

- Step 1** If required, and if you have not already done so, enter the activation key to enable the same number of contexts as are licensed on the primary unit by entering the following command:

```
secondary(config)# activation-key key
```

- Step 2** If you have not already done so, set the context mode to match the primary unit by entering the following command:

```
secondary(config)# mode {single | multiple}
```

The FWSM reboots.

- Step 3** To configure the VLAN interface you are using for the failover link, enter the following command:

```
secondary(config)# failover lan interface interface_name vlan vlan
```

Use the same setting as the primary unit.

- Step 4** To set the IP address of the failover interface, enter the following command:

```
secondary(config)# failover interface ip interface_name ip_address mask standby ip_address
```

Use the same setting as the primary unit.

- Step 5** (Optional) To set this FWSM as the secondary unit, enter the following command:

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

The default is secondary.


Note

This command is the only configuration difference between the primary and secondary units.

Step 6 To enable failover, enter the following command:

```
secondary(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 7 To save the configuration to Flash memory, enter the following command:

```
secondary(config)# copy running-config startup-config
```

See the [“Failover Configuration Example” section on page 15-26](#) for a typical failover configuration.

Verifying the Failover Configuration

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

- [Using the Show Failover Command, page 15-18](#)
- [Viewing Monitored Interfaces, page 15-21](#)
- [Testing the Failover Functionality, page 15-22](#)

See the [“Monitoring Failover” section on page 15-23](#) section for other troubleshooting tools.

Using the Show Failover Command

On each unit, verify the failover status by entering the following command in the system execution space:

```
primary(config)# show failover
```

This command shows the following information:

- The failover status, either on or off
- The active unit
- The IP addresses assigned for the active and standby units
- The failover link information
- The interface policy
- The stateful failover statistics

See the following sample **show failover** command output. A description of each field follows.

```

FWSM(config)# show failover
Failover On
Failover unit Primary
Failover LAN Interface fover Vlan 150
Unit Poll frequency 15 seconds
Interface Poll frequency 15 seconds
Interface Policy 50%
Monitored Interfaces 249 of 250 maximum
Last Failover at: 10:58:08 Apr 15 2004
  This host: Primary - Active
    Active time: 2232 (sec)
    admin Interface inside (10.6.8.91): Normal
    admin Interface outside (70.1.1.2): Normal
  Other host: Secondary - Standby
    Active time: 0 (sec)
    admin Interface inside (10.6.8.100): Normal
    admin Interface outside (70.1.1.3): Normal

Stateful Failover Logical Update Statistics
Link : 4th
Stateful Obj   xmit      xerr      rcv       rerr
General        0          0          0          0
sys cmd        0          0          0          0
up time        0          0          0          0
xlate          0          0          0          0
tcp conn       0          0          0          0
udp conn       0          0          0          0
ARP tbl        0          0          0          0
RIP Tbl        0          0          0          0

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

```

Table 15-2 describes the **show failover** output.

Table 15-2 Show Failover Display Description

Field	Options
Failover	<ul style="list-style-type: none"> On Off
Failover Unit	<ul style="list-style-type: none"> Primary Secondary
Failover LAN Interface	<p>Shows the interface name and VLAN for the failover link: <i>interface_name</i> vlan <i>number</i></p> <p>If you have not configured the failover interface, the display shows: Not configured</p>
Unit Poll frequency	<p><i>n</i> seconds</p> <p>The number of seconds you set with the failover poll unit command. The default is 15 seconds.</p>
Interface Poll frequency	<p><i>n</i> seconds</p> <p>The number of seconds you set with the failover poll interface command. The default is 15 seconds.</p>

Table 15-2 Show Failover Display Description (continued)

Field	Options
Interface Policy	<i>n</i> [%] The threshold for interface failure that you set with the failover interface-policy command. The default is 50%.
Monitored Interfaces	<i>n</i> of 250 maximum The number of interfaces you are monitoring.
Last Failover	The last time a failover occurred.
This host: Other host:	For each host, the display shows the following information.
Primary or Secondary	<ul style="list-style-type: none"> • Active—The unit is in active mode. • Standby—The unit is in standby mode, • Disabled—The unit has failover disabled, or the failover link is not configured. • Listen—The unit is attempting to discover an active unit by listening for polling messages. • Learn—The unit detected an active unit, and is not synchronizing the configuration before going to standby mode. • Failed—The unit is failed.
Active time:	<i>n</i> (sec) The amount of time the unit has been in the active state. This time is cumulative, so the standby unit, if it was active in the past, will also show a value.
[<i>context_name</i>] Interface <i>name</i> (<i>n.n.n.n</i>):	<p>For each interface, the display shows the IP address currently being used on each unit, as well as one of the following conditions:</p> <ul style="list-style-type: none"> • Failed—The interface has failed. • Link Down—The interface line protocol is down. • Normal—The interface is working correctly. • No Link—The interface has been administratively shut down. • Unknown—The FWSM cannot determine the status of the interface. • (Waiting)—The interface has not yet received any polling messages from the other unit. • Testing—The interface is being tested. <p>In multiple context mode, the context name appears before each interface.</p>
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.

Table 15-2 Show Failover Display Description (continued)

Field	Options
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.
Stateful Obj	<p>For each field type, the following statistics are used:</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.
xlate	Translation information.
tcp conn	TCP connection information.
udp conn	Dynamic UDP connection information.
ARP tbl	Dynamic ARP table information.
RIP Tbl	Dynamic router table information.
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 Monitored Interfaces

To view the status of monitored interfaces, (from within the context) enter the following command:

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

For example:

```
primary/contexta(config)# show monitor-interface
This host: Primary - Active
  Interface outside (88.1.1.2): Normal
  Interface inside (10.6.8.91): Normal
Other host: Secondary - Standby
  Interface outside (88.1.1.3): Normal
  Interface inside (10.6.8.100): Normal
```

Testing the Failover Functionality

Follow these steps to ensure that failover works:

-
- Step 1** Test that your primary (active) unit 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:
- ```
primary(config)# no failover active
```
- 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 leave the secondary unit as active or force the primary unit to be active again by entering the following command:
- ```
primary(config)# failover active
```
-

Forcing Failover

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

- Enter this command on the active unit to failover to the standby unit:


```
primary(config)# no failover active
```
- Enter this command on the standby unit to force it to become active:


```
secondary(config)# failover active
```

Disabling Failover

When you disable failover, the active and standby state of each unit is 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 above.

To disable failover, enter the following command:

```
primary(config)# no failover
```

This command is not replicated to the standby unit so you must disable failover of the standby unit separately.

To verify that failover is off, enter the **show failover** command:

```
primary(config)# show failover
Failover Off
...
```

Monitoring Failover

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

- [Failover System Messages, page 15-23](#)
- [SNMP, page 15-23](#)
- [Debug Messages, page 15-23](#)

Failover System Messages

The 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 System Messages Guide* to enable logging and to see descriptions of the system messages.

SNMP

To receive SNMP syslog traps for failover, see the “[Using SNMP](#)” section on page 17-1 for more information.

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.

Frequently Asked Failover Questions

This section contains frequently asked questions about the failover features and includes the following topics:

- [Configuration Replication Questions, page 15-23](#)
- [Basic Failover Questions, page 15-24](#)
- [Stateful Failover Questions, page 15-25](#)

Configuration Replication Questions

See the following questions and answers for configuration replication:

- Does configuration replication save the configuration to Flash memory on the standby unit?
No, the configuration is only in running memory.
- How can both units be configured the same without manually entering the configuration twice?
Commands entered on the active unit are automatically replicated to the standby unit.

- What happens if I enter commands on the standby unit?
You will see an error message telling you that the configurations are out of sync. However, the command is still applied.
If you enter individual commands on the active unit that are replicated to the standby unit, your alterations on the standby unit are preserved.
If you use the **write standby** command on the active unit, it will erase any new commands you entered on the standby unit.
- What happens if I enter the **copy running-config startup-config** command on the active unit?
The **copy running-config startup-config** command is replicated to the standby unit, which proceeds to write its configuration to Flash memory.
- What happens if the configuration in Flash memory on the secondary unit differs from the configuration on the primary unit?
After startup, the primary unit sends its configuration to the secondary unit and erases the running configuration on the secondary unit. However, the secondary unit startup configuration remains unaltered in Flash memory.
- How can I view the running configuration and the Flash memory configuration?
 - **show running**—Shows the running configuration. You can also enter **write terminal**.
 - **show config**—Shows the configuration in Flash memory.
- Are contexts that are on disk saved to disk on the standby unit?
No, all contexts are loaded into running memory only. The startup configuration for a context continues to reside on the primary unit disk. You can copy the context configurations to the standby unit so that if the standby unit starts up and needs to be the active unit, it can load the contexts from disk.
- Can I fail over a context, but not the entire module?
No, you can only have one active FWSM.

Basic Failover Questions

See the following questions and answers for basic failover:

- Which unit becomes active if you restart both units?
The primary unit.
- What happens if the active unit has a power failure?
After hello messages are not acknowledged, the standby unit becomes active.
- What happens when the formerly active unit comes online again?
No failover occurs. It remains in standby mode.
- How long does it take to detect a failure?
 - Network errors are detected within three consecutive polling intervals (by default, 15 second intervals). The polling interval is user-configurable using the **failover poll interface** command.
 - Failover communication errors are detected within a user-configurable number of seconds (the default is 15). The polling time is user-configurable using the **failover poll unit** command.

- What maintenance is required?

Syslog messages are generated when any errors or switches occur. Evaluate the failed unit and fix or replace it.

- Is it possible to have both FWSM units become active at the same time?

Yes, in the following circumstances:

- Both units have configurations in Flash memory
- Both units have failover enabled
- The failover link is down at startup

or

In an inter-switch failover scenario, the trunk between the switches fails

- What prevents the standby unit from passing traffic?

The FWSM failover feature ensures that only traffic aimed to the standby unit (hello messages, Telnet if enabled) is successful, while traffic aimed through the unit is dropped.

Stateful Failover Questions

See the following questions and answers for stateful failover:

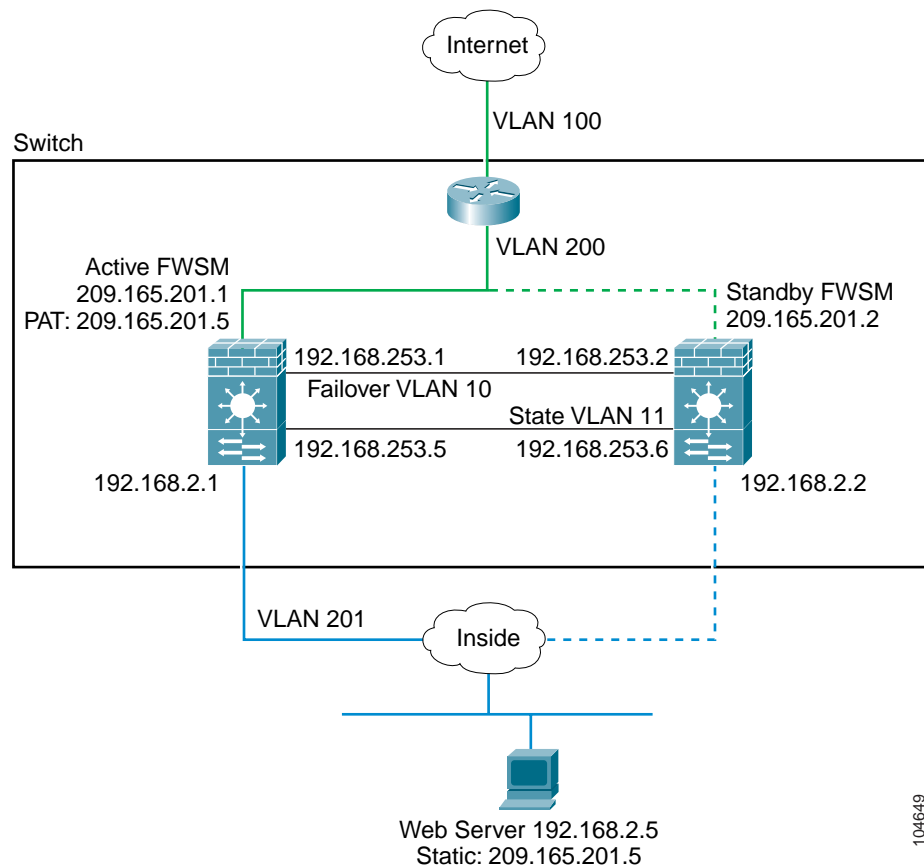
- What information is not replicated to the standby FWSM on stateful failover?
 - The user authentication (uauth) table.
 - The ISAKMP and IPSec SA table (for management access only).
 - The ARP table.
 - Routing information.
 - Other UDP connections.
- Can I share the state link interface with the failover link?

Yes, however, we recommend that you use a separate interface.

Failover Configuration Example

Figure 15-8 shows the network diagram for a failover configuration within a switch. The only difference between the configuration of inter-switch and intra-switch failover is on the switch; the configuration on the FWSM is the same.

Figure 15-8 Failover Scenario



104649

[Example 15-1](#) lists the typical commands in a failover configuration. This example shows how to configure multiple context mode and shows one context, the admin context. For single context mode, simply combine the two configurations, and remove the **admin-context** command and the **context** commands.

Example 15-1 Failover Configuration

System Configuration:

```
hostname FW5M
enable password farscape
password crichton
admin-context adminctxt
context adminctxt
    allocate-interface vlan200
    allocate-interface vlan201
    config-url disk:/adminctxt.cfg
failover lan interface faillink vlan 10
failover link statelink vlan 11
failover lan unit primary
failover interface ip faillink 192.168.253.1 255.255.255.252 standby 192.168.253.2
failover interface ip statelink 192.168.253.5 255.255.255.252 standby 192.168.253.6
failover interface-policy 1
failover replication http
failover
```

Context Configuration:

```
nameif vlan200 outside security0
nameif vlan201 inside security100
enable password aeryn
password rygel
telnet 192.168.2.45 255.255.255.255 [A host on the context network, not shown]
ip address outside 209.165.201.1 255.255.255.224 standby 209.165.201.2
ip address inside 192.168.2.1 255.255.255.0 standby 192.168.2.2
monitor-interface inside
monitor-interface outside
global (outside) 1 209.165.201.3 netmask 255.255.255.224
nat (inside) 1 0.0.0.0 0.0.0.0 0 0
static (inside,outside) 209.165.201.5 192.168.2.5 netmask 255.255.255.255 0 0
access-list acl_out permit tcp any 209.165.201.5 eq 80
access-group acl_out in interface outside
route outside 0 0 209.165.201.4 1
```

[Example 15-2](#) shows the configuration for the secondary unit.

Example 15-2 Failover Configuration: Secondary Unit

```
failover lan interface faillink vlan 10
failover lan unit secondary
failover interface ip faillink 192.168.253.1 255.255.255.252 standby 192.168.253.2
failover
```

