



Configuring Redundancy

This chapter describes how to configure redundant supervisor engines and how to configure redundancy on Multilayer Switch Feature Cards (MSFCs) on the Catalyst 6000 family switches.

This chapter consists of these sections:

- [Understanding How Supervisor Engine Redundancy Works, page 22-2](#)
- [Configuring Redundant Supervisor Engines, page 22-3](#)
- [MSFC Redundancy, page 22-18](#)



Caution

Dual MSFCs in a single chassis are designed to be used in redundant mode only and must have identical configurations. See the [“MSFC Redundancy” section on page 22-18](#) for detailed information.

We do not support configurations where the MSFCs are not configured identically.



Note

Except where specifically differentiated, the information and procedures in this chapter apply to both Supervisor Engine 2 with Layer 3 Switching Engine II (Policy Feature Card 2 or PFC2) and Supervisor Engine 1 with Layer 3 Switching Engine II.



Note

The term *MSFC* is used to refer to the MSFC *and* MSFC2 except where specifically differentiated.

For more information about installing redundant Catalyst 6000 family supervisor engines, refer to the *Catalyst 6000 Family Module Installation Guide*. For syntax and usage information for the commands used in this chapter, refer to the *Catalyst 6000 Family Command Reference* publication.

Understanding How Supervisor Engine Redundancy Works

**Note**

Redundant supervisor engines must be of the same type with the same model feature card.

When you install two supervisor engines, the first supervisor engine to come online becomes the active module; the second supervisor engine goes into standby mode. All administrative and network management functions, such as SNMP, command-line interface (CLI) console, Telnet, Spanning Tree Protocol (STP), Cisco Discovery Protocol (CDP), and VLAN Trunk Protocol (VTP) are processed on the active supervisor engine.

On the standby supervisor engine, the console port is inactive, the module status shows as “standby,” and the status for the uplink ports is shown normally.

You must install redundant supervisor engines in slots 1 and 2 of the chassis. Redundant supervisor engines are hot swappable. The system continues to operate with the same configuration after switching over to the redundant supervisor engine.

**Note**

To allow you to control the booting of each supervisor engine separately, the configuration registers are not synchronized between the supervisor engines.

**Note**

The switchover time from the active to the standby supervisor engine does not include spanning tree convergence time.

At power-up, both supervisor engines run initial module-level diagnostics. Assuming both supervisor engines pass this level of diagnostics, the two supervisor engines communicate over the backplane, allowing them to cooperate during switching-bus diagnostics. The supervisor engine in slot 1 becomes active, and the supervisor engine in slot 2 enters standby mode. If the software versions of the two supervisor engines are different, or if the NVRAM configuration of the two supervisor engines is different, the active supervisor engine automatically downloads its software image and configuration to the standby supervisor engine.

If the background diagnostics on the active supervisor engine detect a major problem or an exception occurs, the active supervisor engine resets. The standby supervisor engine detects that the active supervisor engine is no longer running and becomes active. The standby supervisor engine can detect if the active supervisor engine is not functioning and can force a reset, if necessary. If the reset supervisor engine comes online again, it enters standby mode.

If you hot insert a second supervisor engine, the second module communicates with the active supervisor engine after completing its initial module-level diagnostics. Because the active supervisor engine is already switching traffic on the backplane, no switching-bus diagnostics are run for the second supervisor engine because running diagnostics can disrupt normal traffic. The second supervisor engine immediately enters standby mode. The active supervisor engine downloads the software image and configuration to the standby supervisor engine, if necessary.

The supervisor engines use two Flash images: the *boot image* and the *runtime image*. The boot image filename is specified in the BOOT environment variable, which is stored in NVRAM. The runtime image is the boot image that the ROM monitor uses to boot the supervisor engine. After the system boots, the runtime image resides in dynamic RAM (DRAM).

When you power up or reset a switch with redundant supervisor engines, synchronization occurs to ensure that the runtime and boot images on the standby supervisor engine are the same as the images on the active supervisor engine.

The supervisor engines can have different runtime and boot images. If the boot image and the runtime image are the same, and you change the BOOT environment variable or overwrite or destroy the current boot image on the Flash device that was used to boot the system, the runtime and boot images will differ. Whenever you reconfigure the boot image, the active supervisor engine synchronizes its current boot image with the standby supervisor engine.

The boot image is read directly into the Flash file system. You can perform operations (such as **copy**, **delete**, **undelete**, and so on) on files stored on Flash memory devices, and you can store the boot image of the active supervisor engine in the standby supervisor engine bootflash. For more information about using the Flash file system, see [Chapter 24, “Working With the Flash File System.”](#)

The supervisor engine has a Flash PC card (PCMCIA) slot (slot0) in addition to the onboard Flash memory; this slot can hold a Flash PC card that can store additional boot images.

**Note**

Throughout this publication, the term *Flash PC card* is used in place of the term *PCMCIA card*.

Because you can store multiple boot images, you must specify the name of the boot file image and the location of the image file in the Flash file system in order to boot and synchronize properly. For information about how to specify the name and location of the boot image, see [Chapter 23, “Modifying the Switch Boot Configuration.”](#)

In the synchronization process, the active supervisor engine checks the standby supervisor engine runtime image to make sure it matches its own runtime image. The active supervisor engine checks three conditions:

- If it needs to copy its boot image to the standby supervisor engine
- If the standby supervisor engine bootstring needs to be changed
- If the standby supervisor engine needs to be reset

The following section describes the conditions that can initiate Flash synchronization. For examples of how the system synchronizes the supervisor engine Flash images with various configurations, see the [“Supervisor Engine Synchronization Examples” section on page 22-14.](#)

Configuring Redundant Supervisor Engines

These sections describe how to configure redundant supervisor engines:

- [Synchronization Process Initiation, page 22-4](#)
- [Redundant Supervisor Engine Configuration Guidelines and Restrictions, page 22-4](#)
- [Verifying Standby Supervisor Engine Status, page 22-5](#)
- [Forcing a Switchover to the Standby Supervisor Engine, page 22-6](#)
- [High Availability, page 22-8](#)
- [Supervisor Engine Synchronization Examples, page 22-14](#)

Synchronization Process Initiation

These conditions initiate the synchronization of the runtime and boot images on the active and standby supervisor engines:

- Time stamp mismatch between the runtime images on the active and standby supervisor engines—The active supervisor engine synchronizes its runtime image with the standby supervisor engine if the time stamps of their respective runtime images differ when the system is booted or reset.
- Time stamp mismatch between the boot images on the active and standby supervisor engines—The active supervisor engine synchronizes its boot image with the standby supervisor engine if the time stamps of their respective boot images differ when the system is booted or reset, or if you change the BOOT environment variable.
- Current boot image overwritten—If you overwrite the current boot image stored on one of the Flash devices, the file system management module detects this event and initiates synchronization. The active supervisor engine copies its new boot image to the standby supervisor engine.
- BOOT environment variables changed—If you change the BOOT environment variables to specify a different default boot image, the active supervisor engine initiates boot-image synchronization. The NVRAM configuration module detects this event and calls the Flash synchronization function with the next probable boot filename by looking at the boot configuration parameter.
- Flash PC cards with same boot-image filename—If you change the Flash device on either the active or standby supervisor engine and the new Flash device contains a boot image that has the same name (but a different time stamp) as the boot image from the previous Flash device, the Flash file management module initiates synchronization.
- Current runtime image deleted—If you delete the current runtime image from the Flash device, the Flash file management module prompts you to verify that you want to delete the current runtime image. If you confirm the deletion, the Flash file management module initiates Flash synchronization and informs the NVRAM configuration module of the change. The NVRAM configuration module examines the BOOT environment variable to determine the next probable image to boot and calls the Flash synchronization function using the new image name.

Redundant Supervisor Engine Configuration Guidelines and Restrictions

These conditions and events can cause the synchronization of images between redundant supervisor engines to fail or to produce unexpected results:

- Downloading a new image to the active supervisor engine

When you download a new image to the active supervisor engine, it is copied to the file system (in bootflash or on a Flash PC card in the Flash PC card slot). Because you may or may not have configured this image as the boot image, the newly downloaded image is not copied to the standby supervisor engine automatically.

To initiate the synchronization function between the active and standby supervisor engines, you must configure this newly downloaded image as the boot image on the active supervisor engine. Synchronization occurs when you change the boot variable. To run the new image, you must reset the system.

- Unable to find the current runtime image

If the active supervisor engine is unable to find the current runtime image on any of the Flash devices, it signals an error condition. In this case, if the standby supervisor engine is inserted or reset, Flash synchronization does not occur. In addition, the STATUS LED on the standby supervisor engine turns red and the system generates a syslog error message.

- Active supervisor engine in slot 2

When the active supervisor engine is in slot 2, the standby supervisor engine is in slot 1. If you change the configuration to specify a new boot image and then reset the system, the supervisor engine in slot 1 becomes the active supervisor engine and loads its default boot image, canceling the configuration changes you have just made. To avoid this problem, the switch prompts you for Flash synchronization as soon as you change the boot file configuration.

Verifying Standby Supervisor Engine Status

You can verify the status of the standby supervisor engine using a number of CLI commands.



Note

The **show module** output provides information about installed daughter cards. The **show test** command provides information about onboard application-specific integrated circuits (ASICs).

To verify the status of the standby supervisor engine, perform one or more of these tasks:

Task	Command
Show the status of the standby supervisor engine.	show module [<i>mod</i>]
Show the state of the standby supervisor engine uplink ports.	show port [<i>mod</i> [/ <i>port</i>]]
Show diagnostic test results for the standby supervisor engine.	show test [<i>mod</i>]

This example shows how to check the status of the standby supervisor engine using the **show module** and **show test** commands:

```

Console> (enable) show module 2
Mod Slot Ports Module-Type           Model                Status
-----
2   2   2   1000BaseX Supervisor    WS-X6K-SUP1-2GE     ok

Mod Module-Name          Serial-Num
-----
2                        SAD02330231

Mod MAC-Address(es)      Hw    Fw    Sw
-----
2   00-e0-14-0e-f5-6c to 00-e0-14-0e-f5-6d 0.404  4.2(2038)  4.2(0.24)VAI50
   00-e0-14-0e-f5-6e to 00-e0-14-0e-f5-6f
   00-10-7b-bb-2b-00 to 00-10-7b-bb-2e-ff

Mod Sub-Type             Sub-Model            Sub-Serial  Sub-Hw
-----
2   L2 Switching Engine  WS-F6020             SAD02350211 0.101
Console> (enable)

```

```

Console> (enable) show test 2
Module 2 : 2-port 1000BaseX Supervisor
Network Management Processor (NMP) Status: (. = Pass, F = Fail, U = Unknown)
  ROM: .   Flash-EEPROM: .   Ser-EEPROM: .   NVRAM: .   EOBC Comm: .

Line Card Status for Module 1 : PASS

Port Status :
  Ports 1 2
  -----
  .   .

Line Card Diag Status for Module 2 (. = Pass, F = Fail, N = N/A)

Module 2
  Cafe II Status :
    NewLearnTest: .
    IndexLearnTest: .
    DontForwardTest: .
    DontLearnTest: .
    ConditionalLearnTest: .
    BadBpduTest: .
    TrapTest: .
  Loopback Status [Reported by Module 2] :
    Ports 1 2
    -----
    .   .

Console> (enable)

```

Forcing a Switchover to the Standby Supervisor Engine

You can force a switchover to the standby supervisor engine by resetting the active supervisor engine.



Note

Resetting the active supervisor engine disconnects any open Telnet sessions.

To force a switchover to the standby supervisor engine, perform this task in privileged mode:

Task	Command
Reset the active supervisor engine (where <i>mod</i> is the number of the active supervisor engine).	reset mod

In addition, you can also force a switchover to the standby supervisor engine by setting the CISCO-STACK-MIB moduleAction variable to reset(2) on the active supervisor engine. When the switchover occurs, the system sends a standard SNMP warm-start trap to the configured trap receivers.

This example shows the console output on the active supervisor engine when you force a switchover from the active to the standby supervisor engine:

```

Console> (enable) reset 1
This command will force a switch-over to the standby Supervisor module.
Do you want to continue (y/n) [n]? y
Console> (enable) 12/07/1998,17:04:39:SYS-5:Module 1 reset from Console//

System Bootstrap, Version 3.1(2)
Copyright (c) 1994-1997 by cisco Systems, Inc.

System Bootstrap, Version 3.1(2)

```

```

Copyright (c) 1994-1997 by cisco Systems, Inc.
Presto processor with 32768 Kbytes of main memory

Autoboot executing command: "boot bootflash:cat6000-sup.5-4-1a.bin"
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
Uncompressing file: #####

System Power On Diagnostics
NVRAM Size .. .....512KB
ID Prom Test .....Passed
DPRAM Size .....16KB
DPRAM Data 0x55 Test .....Passed
DPRAM Data 0xaa Test .....Passed
DPRAM Address Test .....Passed
Clearing DPRAM .....Done
System DRAM Memory Size .....32MB
DRAM Data 0x55 Test .....Passed
DRAM Data 0xaa Test .....Passed
DRAM Address Test .....Passed
Clearing DRAM .....Done
EARLII .....Present
EARLII RAM Test .....Passed
EARL Serial Prom Test .....Passed
Level2 Cache .....Present
Level2 Cache test.....Passed

Boot image: bootflash:cat6000-sup.5-4-1a.bin
Downloading epld sram device please wait ...
Programming successful for Altera 10K50 SRAM EPLD
This module is now in standby mode.
Console is disabled for standby supervisor

```

This example shows the console output on the standby supervisor engine when you force a switchover from the active to the standby supervisor engine:

```

Cisco Systems Console

Enter password:
12/07/1998,17:04:43:MLS-5:Multilayer switching is enabled
12/07/1998,17:04:43:MLS-5:Netflow Data Export disabled
12/07/1998,17:04:44:SYS-5:Module 2 is online
12/07/1998,17:04:45:SYS-5:Module 5 is online
12/07/1998,17:04:45:SYS-5:Module 7 is online
12/07/1998,17:04:45:SYS-5:Module 3 is online
12/07/1998,17:04:52:MLS-5:Route Processor 172.20.52.6 added
12/07/1998,17:05:10:SYS-5:Module 8 is online
12/07/1998,17:05:14:SYS-5:Module 9 is online
12/07/1998,17:05:22:SYS-5:Module 4 is online
12/07/1998,17:06:13:SYS-5:Module 1 is in standby mode
Supervisor image synchronization process will start in 10 seconds
12/07/1998,17:06:37:SYS-5:Ports on standby supervisor(Module 1) are UP
12/07/1998,17:06:41:SYS-5:Active supervisor is synchronizing the NMP image.
12/07/1998,17:06:44:SYS-5:The active supervisor has synchronized the NMP image.

Console>

```

High Availability

High availability allows you to minimize the switchover time from the active supervisor engine to the standby supervisor engine if the active supervisor engine fails.

Prior to this feature, fast switchover ensured that a switchover to the standby supervisor engine happened quickly. However, with fast switchover, because the state of the switch features before the switchover was unknown, you had to reinitialize and restart all the switch features when the standby supervisor engine assumed the active role.

High availability removes this limitation; high availability allows the active supervisor engine to communicate with the standby supervisor engine, keeping feature protocol states synchronized. Synchronization between the supervisor engines allows the standby supervisor engine to take over in the event of a failure.

In addition, high availability provides a *versioning* option that allows you to run different software images on the active and standby supervisor engines.

These features are discussed in these sections:

- [High-Availability Overview, page 22-8](#)
- [High-Availability Supported Features, page 22-9](#)
- [Versioning Overview, page 22-10](#)
- [CLI Commands, page 22-11](#)
- [Loading a Different \(but Compatible\) Image on the Standby Supervisor Engine, page 22-13](#)

High-Availability Overview

For high availability, a system database is maintained on the active supervisor engine and updates are sent to the standby supervisor engine for any change of data in the system database. The active supervisor engine communicates and updates the standby supervisor engine when any state changes occur, ensuring that the standby supervisor engine knows the current protocol state of supported features. The standby supervisor engine knows the current protocol states for all modules, ports, and VLANs; the protocols can initialize with this state information and start running immediately.

The active supervisor engine controls the system bus (backplane), sends and receives packets to and from the network, and controls all modules. Protocols run on the active supervisor engine only.

The standby supervisor engine is isolated from the system bus and does not switch packets. But it *does* receive packets from the switching bus to learn and populate its Layer 2 forwarding table for Layer 2-switched flows. The standby supervisor engine also receives packets from the switching bus to learn and populate the Multilayer Switching (MLS) table for Layer 3-switched flows. The standby supervisor engine does not participate in forwarding any packets and does not communicate with any modules.

If you enable high availability when the standby supervisor engine is running, image version compatibility is checked and if found compatible, the database synchronization is started. High availability compatible features continue from the saved states on the standby supervisor engine after a switchover.

When you disable high availability, the database synchronization is not done and all features must restart on the standby supervisor engine after a switchover.

If you change high availability from enabled to disabled, synchronization from the active supervisor engine is stopped and the standby supervisor engine discards all current synchronization data.

If you change high availability from disabled to enabled, synchronization from the active to standby supervisor engine is started (provided the standby supervisor engine is present and its image version is compatible).

NVRAM synchronization occurs irrespective of high availability being enabled or disabled (provided there are compatible NVRAM versions on the two supervisor engines).

If you do not install a standby supervisor engine during system bootup, the active supervisor engine detects this and the database updates are not queued for synchronization. Similarly, when you reset or remove the standby supervisor engine, the synchronization updates are not queued and any pending updates in the synchronization queue are discarded. When you hot insert or restart a second supervisor engine that becomes the standby supervisor engine, the active supervisor engine downloads the entire system database to the standby supervisor engine. Only after this global synchronization is completed, the active supervisor engine queues and synchronizes the individual updates to the standby supervisor engine.

**Note**

When you hot insert or restart a second supervisor engine, it might take a few minutes for the global synchronization to complete.

High-Availability Supported Features

**Note**

MLS flows are preserved from the active supervisor engine to the standby supervisor engine.

**Note**

High availability does not preserve routing table entries on the active MSFC because high availability is not run on the MSFC IOS software. However, you can configure both MSFCs on the active and standby supervisor engines with the same configuration to preserve routing table entries across the active and standby MSFCs. You can then configure HSRP on the MSFCs to provide automatic routing backup. See the “[MSFC Redundancy](#)” section on page 22-18 for detailed information.

High availability for the Catalyst 6000 family switch is classified into three categories (see [Table 22-1](#)):

- **Supported features**—High availability is fully supported; the feature’s database is synchronized from the active supervisor engine to the standby supervisor engine.
- **Compatible features**—High availability is not supported; the feature’s database is not synchronized from the active supervisor engine to the standby supervisor engine. However, the feature can be enabled (operational) with high availability.
- **Incompatible features**—High availability is not supported; the feature’s database is not synchronized from the active supervisor engine to the standby supervisor engine. The feature cannot be enabled if high availability is enabled and similarly, high availability cannot be enabled if the feature is enabled.

**Note**

Timers and statistics are not synchronized from the active to the standby supervisor engine.

Table 22-1 High Availability Feature Support

Supported Features	Compatible Features	Incompatible Features
CEF	ASLB	Dynamic VLAN
COPS-DS	CDP	GVRP
COPS-PR	GMRP	Port security
DTP	IGMP snooping	Protocol filtering
EtherChannel	RMON	
IOS ACLs	RSVP	
MLS	SNMP	
PAgP	Telnet sessions	
QoS	UplinkFast	
SPAN	VTP pruning	
STP		
Trunking		
UDLD		
VACLs		
VTP		

Versioning Overview

When you enable high-availability versioning, you can have two different but compatible images on the active and standby supervisor engines. The active supervisor engine exchanges image version information with the standby supervisor engine and determines whether the images are compatible for enabling high availability. If the active and standby supervisor engines are not running compatible image versions, you cannot enable high availability.

Image versioning is supported in supervisor engine software release 5.4(1)CSX and later releases. With versioning enabled, high availability is fully supported with the active and standby supervisor engines running different images as long as the images are compatible. The *only* fully compatible images are as follows:

- 5.5(3) and 5.5(4)
- 6.1(3) and 6.1(4)

Images that are compatible with all modules except Gigabit Ethernet switching modules are as follows:

- 5.4(3) and 5.4(4)
- 5.5(3) and 5.5(5)
- 5.5(4) and 5.5(5)

Images that are compatible with Gigabit Ethernet switching modules but not compatible with 10/100BASE-T modules are as follows:

- 5.5(6a) and 5.5(7)



Note

Attempting to run incompatible image versions could result in configuration loss.

**Note**

When you install two supervisor engines, the first supervisor engine to come online becomes the active module; the second supervisor engine goes into standby mode. If two supervisor engines are installed in your system, at power up the supervisor engine in slot 1 becomes active, and the supervisor engine in slot 2 enters standby mode. If the software versions of the two supervisor engines are different, or if the NVRAM configuration of the two supervisor engines is different, and if you do not enable versioning, the active supervisor engine automatically downloads its software image and configuration to the standby supervisor engine.

CLI Commands

This section describes the CLI commands for high availability and versioning.

Enabling or Disabling High Availability

High availability is disabled by default. To enable or disable high availability, perform this task in privileged mode:

Task	Command
Enable or disable high availability.	set system highavailability {enable disable}

This example shows how to enable high availability:

```
Console> (enable) set system highavailability enable
System high availability enabled.
Console> (enable)
```

This example shows how to disable high availability:

```
Console> (enable) set system highavailability disable
System high availability disabled.
Console> (enable)
```

Enabling or Disabling High-Availability Versioning

High-availability versioning is disabled by default. To enable or disable high-availability versioning, perform this task in privileged mode:

Task	Command
Enable or disable high-availability versioning.	set system highavailability versioning {enable disable}

This example shows how to enable high-availability versioning:

```
Console> (enable) set system highavailability versioning enable
Image versioning enabled.
Console> (enable)
```

This example shows how to disable high-availability versioning:

```
Console> (enable) set system highavailability versioning disable
Image versioning disabled.
Console> (enable)
```

Showing High-Availability Settings and Operational Status

The **show system highavailability** command displays the following:

- High-availability setting (enabled or disabled)
- Versioning setting (enabled or disabled)
- High-availability operational status (based on whether the standby supervisor engine is present and operational). The operational status field displays one of the following:
 - OFF (high-availability-not-enabled): The high availability option in NVRAM is disabled.
 - OFF (standby-supervisor-not-present): The standby supervisor engine is not installed.
 - OFF (standby-supervisor-image-incompatible): The standby supervisor engine is running a different image than the active supervisor engine and it is not version compatible (the versioning option in NVRAM is enabled). No synchronization is done (even a configuration change in NVRAM on the active supervisor engine cannot be propagated to the standby supervisor engine because of the version incompatibility).
 - OFF (standby-supervisor-image-nvram-only-compat): The standby supervisor engine is running a different image than the active supervisor engine (versioning option in NVRAM is enabled) and the image is only NVRAM compatible (that is, a configuration change in NVRAM on the active supervisor engine is propagated to the standby supervisor engine). However, high availability cannot be supported.
 - OFF (standby-supervisor-not-operational-yet): The standby supervisor engine is detected but is not operational (not online yet).
 - OFF (high-availability-not-operational-yet): The standby supervisor engine is operational (online), but high availability is not operational yet (when the system is booted from reset, it takes a few minutes before high availability is operational).
 - ON: High availability is operational. The active supervisor engine's features have started queuing their state changes for synchronizing to the standby supervisor engine.

To show the high-availability configuration and operational states, perform this task:

Task	Command
Show high-availability configuration and operational states.	show system highavailability

This example shows how to disable high availability and versioning:

```
Console> (enable) show system highavailability
Highavailability: disabled
Highavailability versioning: disabled
Highavailability Operational-status: OFF (high-availability-not-enabled)
Console> (enable)
```

This example shows how to enable high availability:

```
Console> (enable) set system highavailability enable
System high availability enabled.
Console> (enable)

Console> (enable) show system highavailability
Highavailability: enabled
Highavailability versioning: disabled
Highavailability Operational-status: ON
Console> (enable)
```

Loading a Different (but Compatible) Image on the Standby Supervisor Engine

Use this procedure to load a new image on the standby supervisor engine that is different from the image on the active supervisor engine. From the active supervisor engine console port, perform these steps (active supervisor engine is in slot 1):

Step 1 Enable high availability versioning.

```
Console> (enable) set system highavailability enable
System high availability enabled.
Console> (enable)
```

Step 2 Download the new image to the active supervisor engine bootflash.

```
Console> (enable) copy tftp:image2.bin bootflash
IP address or name of remote host []? 172.20.52.3

8763532 bytes available on device bootflash, proceed (y/n) [n]? y

... display text truncated
Console> (enable)
```

Step 3 Copy the new image to the standby supervisor engine bootflash.

```
Console> (enable) copy bootflash:image2.bin 2/bootflash:

5786532 bytes available on device bootflash, proceed (y/n) [n]? y

... display text truncated
Console> (enable)
```

Step 4 Modify the BOOT environment variable so the standby supervisor engine boots the new image.

```
Console> (enable) set boot system flash bootflash:image2.bin prepend 2
BOOT variable = bootflash:image2.bin,1;slot0:image1.bin,1
Console> (enable)
```

Step 5 To boot the new image, reset the standby supervisor engine.

```
Console> (enable) reset 2
This command will reset the system.
Do you want to continue (y/n) [n]? y

... display text truncated
Console> (enable)
```

Supervisor Engine Synchronization Examples

These sections contain examples that show what happens when the synchronization function encounters certain conditions:

- [Synchronizing the Runtime Image with the Bootstring, page 22-14](#)
- [Synchronizing the Boot Images on the Active and Standby Supervisor Engines, page 22-16](#)



Note

In the following examples, the number **1** following the filename in the bootstring (for example, **bootflash:f1,1**) indicates the number of Trivial File Transfer Protocol (TFTP) boot retries that are attempted. However, the supervisor engine does not support TFTP booting. The number is included in these examples to be consistent with Cisco IOS conventions.



Note

These examples are not intended to cover every possible condition.

Synchronizing the Runtime Image with the Bootstring

This section contains four examples in which the active supervisor engine runtime image is synchronized with the standby supervisor engine.

Example 1: Runtime image not synchronized

The configuration for example 1 is as follows:

- The active supervisor engine configuration is as follows (if the image in the standby supervisor engine is identical to the image in the active supervisor engine, the output is the same):
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1`
 - Bootflash: `f1`
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine.
- The expected results are as follows:
 - The active supervisor engine **f1** image is not copied to the standby supervisor engine.
 - The standby supervisor engine bootstring is not modified.
 - The standby supervisor engine is not reset.

Example 2: File copied, bootstring changed, standby supervisor engine reset

The configuration for example 2 is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1`
 - Bootflash: `f1`

- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f2`
 - Boot string: `bootflash:f2,1`
 - Bootflash: `f2`
- The time stamp for **f1** on the active supervisor engine is not the same as **f2** on the standby supervisor engine.
- The expected results are as follows:
 - The active supervisor engine copies **f1** to the standby supervisor engine and renames the file **RTSYNC_f1**.
 - The standby supervisor engine bootflash is modified to the following: **f2, RTSYNC_f1**.
 - The standby supervisor engine bootstring is modified to the following: **bootflash:RTSYNC_f1,1;f2,1;**
 - The standby supervisor engine is reset.

Example 3: File not copied, bootstring changed, standby supervisor engine reset

The configuration for example 3 is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1`
 - Bootflash: `f1`
- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f2`
 - Boot string: `bootflash:f2,1`
 - Bootflash: `f1, f2`
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine but is not the same as **f2** on the standby supervisor engine.
- The expected results are as follows:
 - The active supervisor engine runtime image is synchronized to the standby supervisor engine.
 - The active supervisor engine **f1** image is not copied to the standby supervisor engine.
 - The standby supervisor engine boot string is modified to the following: **f1,1;f2,1;**
 - The standby supervisor engine is reset.

Example 4: Oldest bootflash file deleted, bootflash squeezed

The configuration for example 4 is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1`
 - Bootflash: `f1`

- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f2`
 - Boot string: `bootflash:f2,1;`
 - Bootflash: `f2, f3, f4` (less than 1 MB left on device)
- The time stamp for **f1** on the active supervisor engine is not the same as **f2** on the standby supervisor engine. The **f2** time stamp is older than **f3**, and the **f3** time stamp is older than **f4**.
- The expected results are as follows:
 - The active supervisor engine runtime image is synchronized with the standby supervisor engine.
 - The active supervisor engine attempts to copy its **f1** image to the standby supervisor engine.
 - Because there is not enough space on the standby supervisor engine bootflash, the redundant synchronization function finds the oldest file, deletes it, and squeezes bootflash.
 - The active supervisor engine copies the **f1** image to the standby supervisor engine and renames it **RTSYNC_f1**.
 - The standby supervisor engine bootflash is modified to the following: **f3, f4, RTSYNC_f1**.
 - The standby supervisor engine boot string is modified to the following: **RTSYNC_f1,1;f2,1;**
 - The standby supervisor engine is reset.

Synchronizing the Boot Images on the Active and Standby Supervisor Engines

This section contains four examples in which the bootstrings on the active and standby supervisor engines are synchronized.

Example 1: Unable to allocate the boot image

The configuration for this example is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1`
- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1`
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine.
- The system attempts to modify the active supervisor engine bootstring to the following: **f2,1;**
- The expected results are as follows:
 - The active supervisor engine is unable to allocate **f2**, causing the synchronization to fail.
 - An error is recorded in syslog.

- The active supervisor engine **f1** image is not copied to the standby supervisor engine.
- The standby supervisor engine bootstring is not modified.
- The standby supervisor engine is not reset.

Example 2: File copied, bootflash modified, standby supervisor engine not reset

The configuration for this example is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1,f2`
- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash:
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine.
- You modify the active supervisor engine bootstring to the following: **f2,1;**
- The expected results are as follows:
 - The active supervisor engine copies its **f2** image to the standby supervisor engine and renames it **BTSYNC_f2**.
 - The standby supervisor engine bootflash is modified to the following: **f1, BTSYNC_f2**.
 - The standby supervisor engine bootstring is modified to the following: **bootflash:BTSYNC_f2,1;f1,1;**
 - The standby supervisor engine is not reset.

Example 3: File not copied, bootstring modified, standby supervisor engine not reset

The configuration for this example is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1,f2`
- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1,f2`
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine; the time stamp for **f2** on the active supervisor engine is the same as **f2** on the standby supervisor engine.
- The active supervisor engine bootstring is modified to the following: **f2,1; f1,1;**

- The expected results are as follows:
 - The active supervisor engine **f1** image is not copied to the standby supervisor engine.
 - The standby supervisor engine bootstring is modified to the following:
bootflash:f2,1;bootflash:f1,1;
 - The standby supervisor engine is not reset.

Example 4: File copied, oldest file deleted, bootflash squeezed, bootstring modified, standby supervisor engine not reset

The configuration for this example is as follows:

- The active supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f1, f2`
- The standby supervisor engine configuration is as follows:
 - Runtime image: `bootflash:f1`
 - Boot string: `bootflash:f1,1;`
 - Bootflash: `f0, f1, f3` (less than 1 MB left on device)
- The time stamp for **f1** on the active supervisor engine is the same as **f1** on the standby supervisor engine. The time stamp for **f0** is older than **f1**, and the time stamp for **f1** is older than **f3**.
- The active supervisor engine bootstring is modified to the following: **bootflash:f2,1;bootflash:f1,1;**
- The expected results are as follows:
 - The active supervisor engine attempts to copy its **f2** image to the standby supervisor engine.
 - Because there is not enough space available on the standby supervisor engine bootflash, the redundant synchronization function finds the oldest file (**f0**), deletes it, and squeezes bootflash.
 - The active supervisor engine copies its **f2** image to the standby supervisor engine and renames it **BTSYNC_f2**.
 - The standby supervisor engine bootflash is modified to the following: **f1, f3, BTSYNC_f2**.
 - The standby supervisor engine boot string is modified to the following:
bootflash:BTSYNC_f2,1;bootflash:f1,1;

MSFC Redundancy

MSFC redundancy is described in these sections:

- [Dual MSFC Redundancy, page 22-19](#)
- [Single Router Mode Redundancy, page 22-41](#)
- [Manual-Mode MSFC Redundancy, page 22-45](#)

Dual MSFC Redundancy

**Caution**

You must configure both MSFCs identically. [Table 22-2 on page 22-20](#) summarizes the identical requirements and the exceptions for Layer 3 redundancy for a single switch chassis.

We do not support configurations where the MSFCs are not configured identically.

These sections describe how to configure MSFC redundancy:

- [Hardware and Software Requirements, page 22-19](#)
- [Layer 3 Redundancy for a Single Chassis, page 22-19](#)
- [Routing Protocol Peering, page 22-20](#)
- [Access Control List Configuration, page 22-22](#)
- [Dual MSFC Operational Model for Redundancy and Load Sharing, page 22-22](#)
- [Understanding Failure Scenarios, page 22-24](#)

Hardware and Software Requirements

To configure Layer 3 redundancy, you must have at least one of the following configurations:

- A single chassis with two identical supervisor engine daughter card configurations:
 - Supervisor Engine 1 with Policy Feature Card (PFC) and MSFC or MSFC2 (both supervisor engines must have the same type of MSFC)
 - Supervisor Engine 2 with PFC2 and MSFC2
- Two chassis with a supervisor engine in each—You must have at least one supervisor engine in each chassis. Each supervisor engine must be equipped with a PFC and an MSFC.

**Note**

Each MSFC must be running the same release of Cisco IOS software.

Layer 3 Redundancy for a Single Chassis

In a single Catalyst 6000 family chassis, you can have redundant supervisor engines, each with an MSFC. You can configure HSRP on the MSFCs to provide transparent default gateway redundancy for IP hosts in the network. HSRP configuration can coexist with IPX and AppleTalk configuration on the same interfaces.

If one MSFC fails, HSRP allows one MSFC (router) to assume the function automatically of the other. Combined with the high-availability feature of supervisor engine software release 5.4(1), this configuration provides an added level of redundancy for your network.

**Caution**

You *must* configure both MSFCs identically. [Table 22-2](#) summarizes the identical requirements and the exceptions for Layer 3 redundancy for a single switch chassis.

Table 22-2 Single Chassis Layer 3-Redundancy Requirements

Identical Requirements— Global and Interface Levels	Exceptions—Interface Level	Exceptions—Global Level
<ul style="list-style-type: none"> • Both MSFCs <i>must</i> have the following: <ul style="list-style-type: none"> – Same routing protocols – Same static routes – Same default routes – Same policy routes – Same VLAN interfaces – Same IOS ACLs^{1, 2} • All interfaces <i>must</i> have the same administrative status 	<ul style="list-style-type: none"> • HSRP standby commands • IP address commands³ • IPX network³ 	<ul style="list-style-type: none"> • IP default-gateway • IPX internal-network • IPX default-route

1. Dynamic and reflexive ACLs, which are based on actual data flow, may be programmed by either MSFC.
2. In addition to defining the same ACLs on both MSFCs, you must also apply the ACLs to the same VLAN interfaces, in the same direction, on both MSFCs.
3. The IP or IPX addresses do not have to be identical on both MSFCs, but there *must* be an IP or IPX address configured on both MSFCs.

For information on specifying alternate configurations for the interface and global level exceptions listed in Table 22-2, see the “[alt Keyword Usage](#)” section on page 22-33.

Redundant supervisor engines *must* have identical hardware (MSFC and PFC). See the “[Hardware and Software Requirements](#)” section on page 22-19 for more information.



Note

For MSFC and MSFC2 memory requirements, refer to the Release Notes for MSFC publication: <http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/index.htm>

Routing Protocol Peering

In a redundant supervisor engine and dual MSFC configuration, one supervisor engine is fully operational (active) and the other supervisor engine is in standby mode; however, both MSFCs are operational (in terms of programming the PFC on the active supervisor engine) and act as independent routers.



Note

PFC: With the PFC, MLS entries can be associated with either MSFC (based on which MSFC routed the first packet). Only the PFC on the active supervisor engine switches the packets.

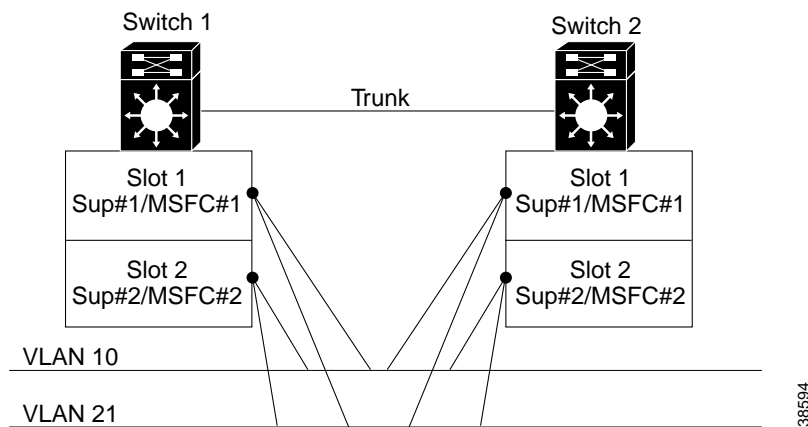


Note

PFC2: With PFC2, only the designated MSFC programs the forwarding information base (FIB) the adjacency table, Cisco IOS software, and policy routing ACLs on the active supervisor engine. If you configure static routes or policy routing, you must have the *identical* configuration on both MSFCs. If you have a static route on the nondesignated MSFC that is not on the designated MSFC, that route *will not* be programmed in the PFC2.

Both MSFCs are operational from a routing protocol peering perspective. For example, if you have two MSFCs in a single Catalyst 6000 family switch chassis, each configured with interface VLAN 10 and VLAN 21, the MSFCs are peered to each other over these VLANs. Combined with a dual chassis and dual MSFC design for the same VLANs, each MSFC has 6 peers: its peer in the same chassis as well as the 2 MSFCs in the second chassis (3 in VLAN 10 and 3 in VLAN 21). See [Figure 22-1](#).

Figure 22-1 Dual Chassis and Dual MSFC Peering



Although the MSFCs (from a peering perspective) act as independent routers, the two MSFCs in the chassis operate at the same time, have the same interfaces, and run the same routing protocols.

If you combine high availability on the supervisor engines with HSRP on the MSFCs, you have the following Layer 2 and Layer 3 redundancy mechanisms:

- Layer 2 redundancy for the supervisor engines (one active and one in standby)—If the active supervisor engine fails (the MSFC installed on it will also fail), both Layer 2 and Layer 3 functions roll over to the redundant supervisor engine and MSFC combination.
- Layer 3 redundancy and load sharing for the two MSFCs—If one MSFC fails, the other MSFC takes over almost immediately (using HSRP) without any Layer 2 disruption (the active supervisor engine continues to forward Layer 2 traffic).

The Layer 3 entries programmed by the failed MSFC on the active supervisor engine are used until they gracefully age out and are replaced by the Layer 3 entries populated by the newly active MSFC. Aging takes 4 minutes and allows the newly active MSFC to repopulate the MLS entries using its XTAG value, while concurrently hardware-switching flows yet to be aged. In addition, this process prevents a newly active MSFC from being overwhelmed with initial flow traffic.



Note

Each MSFC has its own XTAG value to identify itself as the MLS Route Processor. MSFC #1 (on the active supervisor engine) has an XTAG of 1, and MSFC #2 (on the standby supervisor engine) has an XTAG of 2.

Only Supervisor Engine 1 uses the XTAG values; XTAG values are not used on Supervisor Engine 2.



Caution

For same-chassis Layer 3 redundancy to function as expected, the configuration on each MSFC *must* be the same (see [Table 22-2 on page 22-20](#)).

**Note**

Table 22-2 lists configuration exceptions. For example, in Figure 22-1, there are 4 MSFCs on VLAN 10; therefore, each MSFC has different IP addresses and HSRP priorities.

Access Control List Configuration

If you use Cisco IOS access control lists (ACLs) on the MSFC, you *must* configure the ACLs on both MSFCs identically, globally, and at the interface level. Only the designated MSFC (the MSFC to come online first, or the MSFC that has been online the longest) programs the PFC with ACL information.

The active supervisor engine's PFC multilayer switches packets (CEF [Cisco Express Forwarding] for PFC2) after consulting with its ACL ASIC to determine whether a packet is forwarded or not, depending on the IOS ACL configured. If a designated MSFC fails, the new designated MSFC must reprogram the PFC for static ACLs. For consistent results, both MSFCs *must* have identical ACL configurations, including static ACLs.

**Note**

In addition to defining the same ACLs on both MSFCs, you must also apply the ACLs to the same VLAN interfaces on both MSFCs.

**Note**

Dynamic and reflexive ACLs, which are based on actual data flow, may be programmed by either MSFC.

**Note**

PFC: For detailed information on hardware and software handling of IOS ACLs with the PFC, see the “[Hardware and Software Handling of Cisco IOS ACLs with PFC](#)” section on page 16-10.

**Note**

PFC2: For detailed information on hardware and software handling of IOS ACLs with the PFC2, see the “[Hardware and Software Handling of Cisco IOS ACLs with PFC2](#)” section on page 16-12.

To determine the status of the designated MSFC, enter the **show fm features** or the **show redundancy** command:

```
Router-15# show redundancy
Designated Router: 1 Non-designated Router:2
Redundancy Status: non-designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled
```

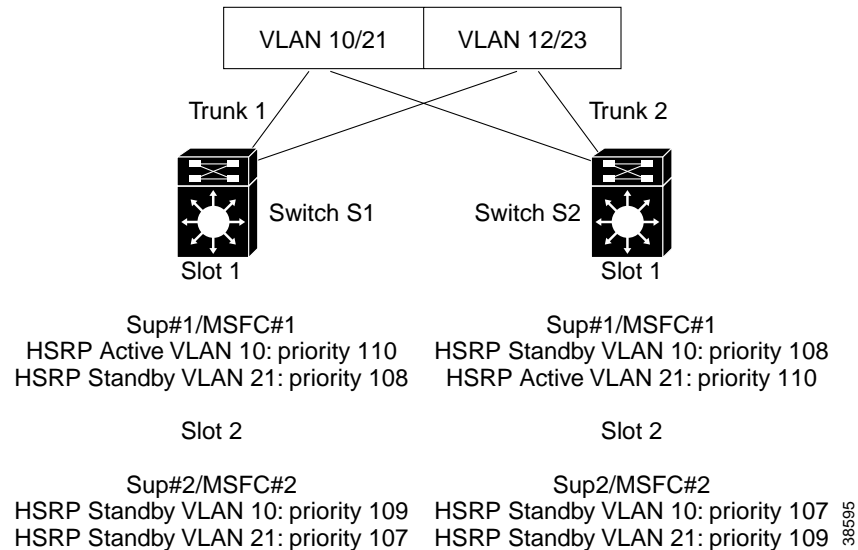
```
Router-16# show redundancy
Designated Router: 1 Non-designated Router:2
Redundancy Status: designated
Config Sync AdminStatus : enabled
Config sync RuntimeStatus: enabled
```

Dual MSFC Operational Model for Redundancy and Load Sharing

Figure 22-2 shows a typical access and distribution layer building block with multiple VLANs in an access layer switch. Because there is no Layer 2 loop, HSRP is used for convergence and load sharing. Switches S1 and S2 have a supervisor engine with an MSFC in slot 1 (Sup #1/MSFC #1) and in slot 2

(Sup #2/MSFC #2). Sup #1 is active and Sup #2 is in standby mode in both switches. High availability is enabled on the supervisor engines. The supervisor engines automatically perform image and configuration synchronization; you must manually synchronize the images and configurations on the MSFCs.

Figure 22-2 Dual MSFC Operational Model for Redundancy and Load Sharing—VLANs 10 and 21



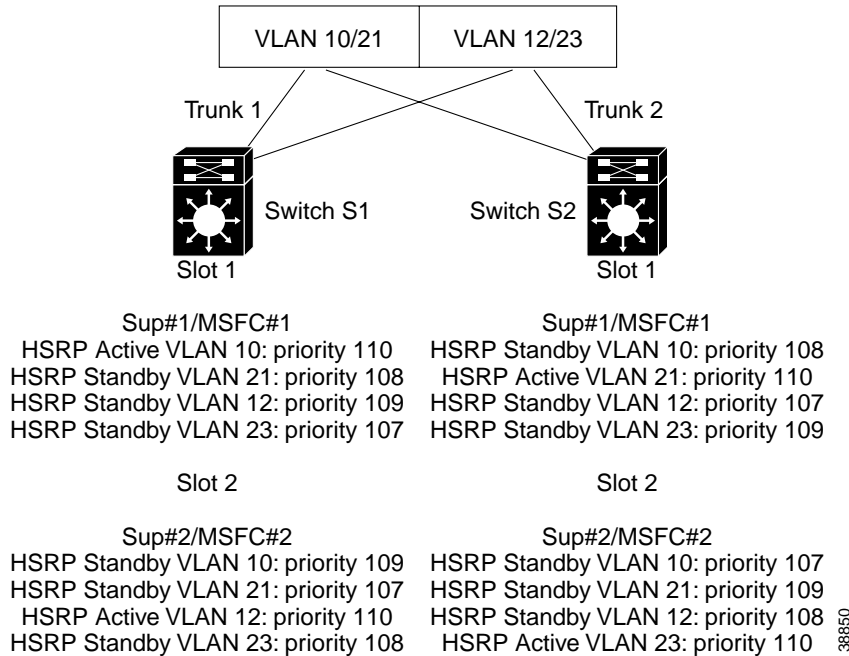
In [Figure 22-2](#), you should configure redundancy and load sharing as follows:

- VLAN 10 (even-numbered VLANs)—Configure MSFC #1 in Switch S1 as the primary HSRP router (priority 110) and configure MSFC #2 as the standby router (priority 109).
- VLAN 21 (odd-numbered VLANs)—Configure MSFC #1 in Switch S2 as the primary HSRP router (priority 110) and configure MSFC #2 as the standby router (priority 109).

Load sharing is achieved by having the even-numbered VLANs routed by Switch S1 and the odd-numbered VLANs by Switch S2. In a complete switch failure, the remaining switch would service both even and odd VLANs.

You can achieve further load sharing by using MSFC #2 in Switch S1 as the primary HSRP router for VLAN 12 and MSFC #2 as the primary HSRP router in Switch S2 for VLAN 23 (see [Figure 22-3](#)).

Figure 22-3 Dual MSFC Operational Model for Redundancy and Load Sharing—VLANs 10, 12, 21, and 23



Only the active HSRP router for a VLAN will respond with the HSRP MAC address for ARP requests to the HSRP IP address. The active HSRP router will in turn ARP for the end stations' MAC address and populate its ARP cache. By using both MSFCs in a single chassis to share HSRP duties for even VLANs, you can share the control plane ARP traffic. In an MSFC failure, only the ARP entries on the affected VLAN would need to be relearned.

The tradeoff for this level of redundancy and load sharing is the added complexity of keeping track of the even and odd VLANs on the MSFCs within a Catalyst 6000 family switch chassis.

MLS entries are created for packets arriving at the HSRP MAC addresses as well as those arriving with the router's real MAC addresses. HSRP is used for unicast traffic first-hop redundancy; for traffic received through another router attached to VLAN 10, for example, the actual MAC address of Sup #1/MSFC #1 is used.

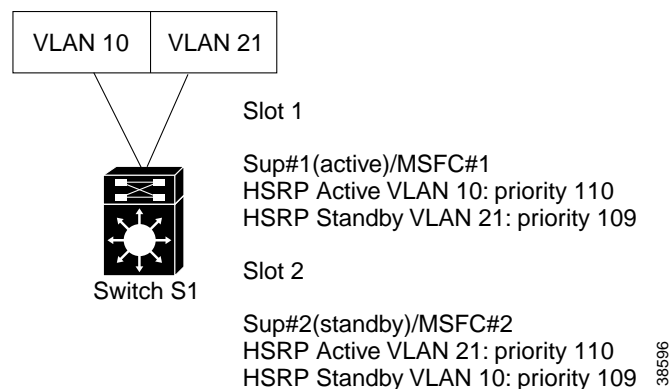
Understanding Failure Scenarios

The five examples in this section describe possible failure scenarios within a single chassis with dual supervisor engines and dual MSFCs (see [Figure 22-4](#)) when you enable high availability. The designated MSFC refers to the MSFC that is used to program the ACL ASIC for static ACLs.

**Note**

While the examples are specific to the PFC, the failover scenarios for the PFC2/MSFC2 would be similar for handling ACLs and CEF table entries. On a Supervisor Engine 2, the designated MSFC2 programs many of the ASICs on the PFC2 including building the CEF table. In a designated MSFC2 HSRP failover to the nondesignated MSFC2, the PFC2 continues to function with the CEF table programmed by the previously designated MSFC2. Similar to the process with the MLS cache in a Supervisor Engine 1/MSFC configuration, the newly designated MSFC2 eventually reprograms the CEF table with its own entries and the old entries age out.

Figure 22-4 Single Chassis with Dual Supervisor Engines and Dual MSFCs



Failure Case 1: Designated MSFC #1 Fails

This sequence occurs when the designated MSFC #1 fails:

1. MLS entries for MSFC #1 gracefully age out of the Sup #1 Layer-3 cache, while MSFC #2 takes temporary ownership of these MLS entries using its XTAG value.
2. MLS entries for MSFC #2 are not affected.
3. MSFC #2 removes all dynamic and reflexive ACLs programmed in hardware by MSFC #1.
4. MSFC #2 reprograms the static ACLs in the Sup #1 ACL ASIC because it is now the designated MSFC.

Failure Case 2: Nondesignated MSFC #2 Fails

This sequence occurs when the nondesignated MSFC #2 fails:

1. MLS entries for MSFC #2 gracefully age out of the Sup #1 Layer 3 cache, while MSFC #1 takes temporary ownership of these MLS entries using its XTAG value.
2. MLS entries from MSFC #1 are not affected.
3. MSFC #1 removes all dynamic and reflexive ACLs programmed in hardware by MSFC #2.
4. MSFC #1 remains the designated MSFC.

Failure Case 3: Active Sup #1 Fails

This sequence occurs when the active supervisor engine (Sup #1) fails:

1. Because the Layer 3 state is maintained, MLS entries of MSFC #1 gracefully age out of the Sup #2 Layer 3 cache while MSFC #2 takes temporary ownership of these MLS entries using its XTAG value.
2. The standby supervisor engine maintains the Layer 2 state so that there is no Layer 2 convergence time.
3. MSFC #2 removes all dynamic and reflexive ACLs programmed in hardware by MSFC #1.
4. MSFC #2 reprograms the static ACLs in the Sup #2 ACL ASIC. MSFC #2 is now the designated MSFC.

Failure Case 4: Standby Sup #2 Fails

This sequence occurs when the standby supervisor engine (Sup #2) fails:

1. MLS entries for MSFC #2 gracefully age out of the Sup #1 Layer 3 cache while MSFC #1 takes temporary ownership of these MLS entries using its XTAG value.
2. MLS entries from MSFC #1 are not affected.
3. MSFC #1 removes all dynamic and reflexive ACLs programmed in hardware by MSFC #2. MSFC #1 remains the designated MSFC.

Failure Case 5: New or Previously Failed Supervisor Comes Back Online

This sequence occurs when the previously failed supervisor engine (Sup #2) comes online:

1. Sup #1 continues to be the active supervisor engine.
2. Sup #2 synchronizes its image and configuration with Sup #1 (unless high-availability versioning is enabled).
3. MSFC #2 (on Sup #2) comes up. If the HSRP preempt for VLAN 21 is configured, then MSFC #2 becomes HSRP active. The MLS entries for MSFC #1 are purged and then relearned via MSFC #2.
4. MSFC #1 remains the designated MSFC for the static ACLs.

Configuring Redundancy with HSRP

Although the supervisor engine software high-availability feature maintains the protocol state between redundant supervisor engines, you need to configure HSRP for failover between redundant MSFCs. HSRP is used to provide the first-hop, unicast redundancy. You can configure one or more HSRP groups on MSFC VLAN interfaces to provide automatic routing backup for your network. Each VLAN interface in an HSRP group shares a virtual IP address and MAC address. You can configure end stations and other devices to use the HSRP address as the default gateway so that if one router interface fails, service is not interrupted to those devices.

The interface with the highest HSRP priority is the active interface for that HSRP group.



Note

PFC2: The PFC2 supports a maximum of 16 unique HSRP group numbers. You can use the same HSRP group numbers in different VLANs. If you configure more than 16 HSRP groups, this restriction prevents use of the VLAN number as the HSRP group number.

**Note**

PFC2: Identically numbered HSRP groups use the same virtual MAC address, which might cause errors if you configure bridging on the MSFC.

The **standby use-bia** option should not be used in an HSRP configuration. MLS entries are not created when you use the **standby use-bia** option. When the **standby use-bia** option is configured, if an HSRP active interface goes up and down, there will be no router CAM address for the standby VLAN interface and without the router CAM entry, no shortcuts are created. This problem is independent of any MSFC Cisco IOS release. (This problem is documented in caveat CSCdz17169.)

To configure HSRP on an MSFC VLAN interface, perform this task in interface configuration mode:

	Task	Command
Step 1	Enable HSRP and specify the HSRP IP address. If you do not specify a <i>group_number</i> , group 0 is used. To assist in troubleshooting, configure the group number to match the VLAN number.	Router(config-if)# standby [<i>group_number</i>] ip [<i>ip_address</i>]
Step 2	Specify the priority for the HSRP interface. Increase the priority of at least one interface in the HSRP group (the default is 100). The interface with the highest priority becomes active for that HSRP group.	Router(config-if)# standby [<i>group_number</i>] priority <i>priority</i>
Step 3	Configure the interface to preempt the current active HSRP interface and become active if the interface priority is higher than the priority of the current active interface.	Router(config-if)# standby [<i>group_number</i>] preempt [delay <i>delay</i>]
Step 4	(Optional) Set the HSRP hello timer and holdtime timer for the interface. The default values are 3 (hello) and 10 (holdtime). All interfaces in the HSRP group should use the same timer values.	Router(config-if)# standby [<i>group_number</i>] timers <i>hellotime</i> <i>holdtime</i>
Step 5	(Optional) Specify a clear-text HSRP authentication string for the interface. All interfaces in the HSRP group should use the same authentication string.	Router(config-if)# standby [<i>group_number</i>] authentication <i>string</i>

This example shows how to configure an interface as part of HSRP group 100:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan100
Router(config-if)# standby 100 ip 172.20.100.10
Router(config-if)# standby 100 priority 110
Router(config-if)# standby 100 preempt
Router(config-if)# standby 100 timers 5 15
Router(config-if)# standby 100 authentication Secret
Router(config-if)# ^Z
Router#
```

Configuration Examples

This section describes three configuration options for achieving redundancy:

- [Example 1—Two Chassis with One Supervisor Engine and One MSFC Each, page 22-28](#)
- [Example 2—Single Chassis with Dual Supervisor Engines and MSFCs, page 22-29](#)
- [Example 3—Double Chassis with Dual Supervisor Engines and MSFCs, page 22-30](#)

For the following examples, the designated MSFC is on the active supervisor engine. To determine the status of the designated MSFC, enter the **show fm features** or the **show redundancy** command. This example shows that Router-16 is the designated MSFC:

```
Router-15# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: non-designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled


Router-16# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: designated
Config Sync AdminStatus : enabled
Config sync RuntimeStatus: enabled
```

Example 1—Two Chassis with One Supervisor Engine and One MSFC Each

In the example in [Figure 22-5](#), high availability cannot be configured on the supervisor engines but HSRP can be configured on the MSFCs.

Figure 22-5 Two Chassis with One Supervisor Engine and One MSFC Each



This example shows how to configure HSRP on the MSFC in Switch S1:

```
Console> (enable) switch console 15
Trying Router-15...
Connected to Router-15.
Type ^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 110
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
```

```

Router(config-if)# standby 21 priority 109
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^C^C^C

```

This example shows how to configure HSRP on the MSFC in Switch S2:

```

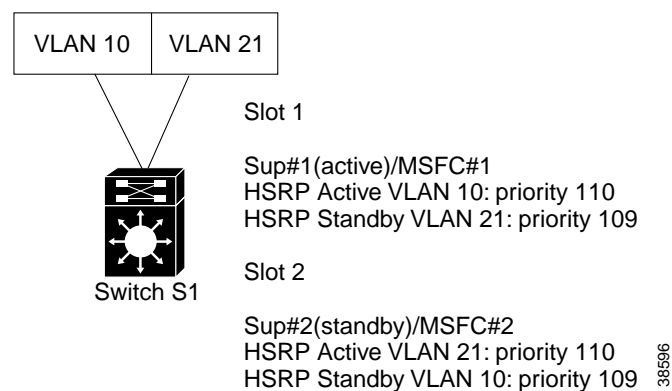
Console> (enable) switch console 15
Trying Router-15...
Connected to Router-15.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 109
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 110
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^C^C^C

```

Example 2—Single Chassis with Dual Supervisor Engines and MSFCs

In the example in [Figure 22-6](#), high availability is configured on the supervisor engines, and HSRP is configured on the MSFCs.

Figure 22-6 Single Chassis with Redundant Supervisors and MSFCs



This example shows how to configure HSRP on the MSFC in Switch S1:

```

Console> (enable) switch console 15
Trying Router-15...
Connected to Router-15.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10

```

```

Router(config-if)# standby 10 priority 110
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 109
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^c^c^c

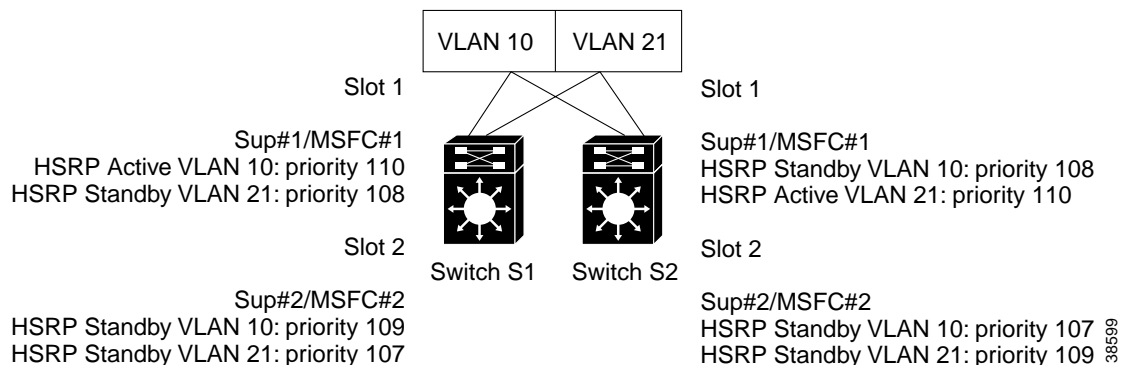
Console> (enable) switch console 16
Trying Router-16...
Connected to Router-16.
Type ^c^c^c to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 109
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 110
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^c^c^c

```

Example 3—Double Chassis with Dual Supervisor Engines and MSFCs

Figure 22-7 shows two Catalyst 6000 family switches (S1 and S2), each with a supervisor engine and MSFC in slot 1 (Sup #1/MSFC #1) and slot 2 (Sup #2/MSFC #2). Because there is no Layer-2 loop, HSRP is used for convergence and load sharing. In both switches, Sup #1 is the active supervisor engine, and Sup #2 is the standby supervisor engine.

Figure 22-7 Dual MSFC Operational Model for Redundancy and Load Sharing



38699

This example shows how to configure HSRP on the MSFC in Switch S1:

```
Console> (enable) switch console 15
Trying Router-15...
Connected to Router-15.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 110
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 108
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^C^C^C

Console> (enable) switch console 16
Trying Router-16...
Connected to Router-16.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 109
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 107
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^C^C^C
```

This example shows how to configure HSRP on the MSFC in Switch S2:

```
Console> (enable) switch console 15
Trying Router-15...
Connected to Router-15.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 108
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 110
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
```

```

Router(config-if)# ^Z
Router# ^C^C^C

Console> (enable) switch console 16
Trying Router-16...
Connected to Router-16.
Type ^C^C^C to switch back...
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# interface vlan10
Router(config-if)# standby 10 ip 172.20.100.10
Router(config-if)# standby 10 priority 107
Router(config-if)# standby 10 preempt
Router(config-if)# standby 10 timers 5 15
Router(config-if)# standby 10 authentication Secret
Router(config-if)# interface vlan21
Router(config-if)# standby 21 ip 192.20.100.21
Router(config-if)# standby 21 priority 109
Router(config-if)# standby 21 preempt
Router(config-if)# standby 21 timers 5 15
Router(config-if)# standby 21 authentication Secret
Router(config-if)# ^Z
Router# ^C^C^C

```

MSFC Configuration Synchronization Overview

MSFC high availability allows for automatic synchronization of the startup configuration and running configuration between the designated MSFC (the MSFC to come online first, or the MSFC that has been online the longest) and the nondesignated MSFC. High-availability redundancy is disabled by default.



Caution

Configuration synchronization is only supported for IP and IPX configurations. Before enabling synchronization, you must ensure that both MSFCs have identical configurations for all protocols. If you are using AppleTalk, DECnet, VINES or any other routing, you must manually ensure that identical configurations are on both MSFCs for all protocols.

To determine the status of the designated MSFC, enter the **show fm features** or the **show redundancy** command:

```

Router-15# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: non-designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled

Router-16# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: designated
Config Sync AdminStatus : enabled
Config sync RuntimeStatus: enabled

```

High-availability redundancy provides startup and running configuration synchronization.

When you enable high-availability redundancy, the startup configuration of both MSFCs is updated when you enter either of these commands on the designated MSFC:

- **write mem**
- **copy source startup-config**

When you enable high-availability redundancy, every configuration command executed on the designated MSFC is sent to the nondesignated MSFC. Also, the running configuration synchronization is updated when you enter the **copy source running-config** command on the designated MSFC.

These sections provide information about MSFC configuration synchronization:

- [Configuration Synchronization States, page 22-33](#)
- [alt Keyword Usage, page 22-33](#)

Configuration Synchronization States

The two states for the configuration synchronization are as follows:

- Config Sync AdminStatus—signifies what the user has configured for this feature at that moment
- Config Sync RuntimeStatus—enabled only when the following occurs:
 - The Config Sync AdminStatus is enabled on both designated and nondesignated MSFCs
 - The designated and nondesignated MSFCs are running compatible images

When you enable the Config Sync RuntimeStatus, the following occurs:

- No configuration mode is available on the CLI of the nondesignated MSFC; EXEC mode is available
- The **alt** keyword is available and required (see the “[alt Keyword Usage](#)” section on page 22-33 for more information on the **alt** keyword)
- The running and startup configurations are synchronized

When the Config Sync RuntimeStatus is in disabled mode, the following occurs:

- Configuration mode is available on the CLI of both MSFCs
- The **alt** keyword is available but optional
- The running and startup configurations are not synchronized

Various configuration and operation cases are covered in the “[High-Availability Redundancy Configuration Examples](#)” section on page 22-35.

alt Keyword Usage

When you enable the Config Sync RuntimeStatus, the configuration mode on the nondesignated MSFC is disabled; only the EXEC mode is still available. Configuration of both MSFCs is made through the console or a Telnet session on the designated MSFC.

To configure both MSFCs from a single console, enter the **alt** keyword to specify an alternate configuration. When specifying the alternate configuration, the configuration specified before the **alt** keyword relates to the MSFC on the supervisor engine in slot 1 of the switch; the configuration specified after the **alt** keyword relates to the MSFC on the supervisor engine in slot 2.



Note

The **alt** keyword is required when Config Sync AdminStatus is enabled.

[Table 22-3](#) shows the interface and global configuration commands that contain the **alt** keyword.

Table 22-3 Interface and Global Configuration Commands Containing the alt Keyword

Interface Configuration Commands	Global Configuration Commands
<ul style="list-style-type: none"> • [no] standby [<i>group_number</i>] ip [<i>ip_address</i>] [secondary] alt [no] standby [<i>group_number</i>] ip [<i>ip_address</i>] [secondary] • [no] standby [<i>group_number</i>] priority <i>priority</i> [preempt [<i>delay delay</i>]] alt [no] standby [<i>group_number</i>] priority <i>priority</i> [preempt [<i>delay delay</i>]] • [no] ip address <i>ip_address mask</i> [secondary] alt [no] ip address <i>ip_address mask</i> [secondary] • [no] ipx network <i>network</i> [encapsulation encapsulation_type] [secondary] [alt [no] ipx network <i>network</i> [encapsulation encapsulation_type] [secondary]] 	<ul style="list-style-type: none"> • [no] hostname <i>hostname alt hostname hostname</i> • [no] ip default-gateway <i>ip_address alt</i> [no] ip default-gateway <i>ip_address</i> • router bgp <i>autonomous_system bgp router-id ip_address</i> [alt <i>ip_address</i>] • router ospf <i>process_id router-id ip_address</i> [alt <i>ip_address</i>]

This example shows how the **alt** keyword is used when entering the **ip address** command:

```
Router-1(config-if)# ip address 1.2.3.4 255.255.255.0 alt ip address 1.2.3.5 255.255.255.0
```

Enabling or Disabling Configuration Synchronization

To enable high-availability redundancy, perform this task in privileged mode:

	Task	Command
Step 1	Enable redundancy.	redundancy
Step 2	Enable high availability.	high-availability
Step 3	Enable or disable configuration synchronization.	[no] config-sync

This example shows how to enable high-availability redundancy and configuration synchronization (Router-15 is the designated MSFC):

```
Console>(enable) session 15
Trying Router-15...
Connected to Router-15.
Escape character is '^'.

Router-15> enable
Router-15# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router-15(config)# redundancy
Router-15(config-r)# high-availability
Router-15(config-r-ha)# config-sync
Router-15(config-r-ha)# end
```



Note

When you enable high-availability redundancy, the configuration mode is disabled on the nondesignated MSFC; only the EXEC mode is available.

In this example, Router-16 is the nondesignated MSFC; high-availability redundancy and configuration synchronization are enabled:

```

Console>(enable) session 16
Trying Router-16...
Connected to Router-16.
Escape character is '^]'.

Router-16> enable
Router-16# configure terminal
Config mode is disabled on non-designated Router, please configure from designated Router

```

High-Availability Redundancy Configuration Examples

This section discusses different scenarios for enabling high availability and configuration synchronization:

- [Scenario 1: Enabling Configuration Synchronization on Both MSFCs, page 22-35](#)
- [Scenario 2: Disabling Configuration Synchronization on the Designated MSFC, page 22-38](#)
- [Scenario 3: Designated MSFC Comes Up, page 22-39](#)
- [Scenario 4: Nondesignated MSFC Comes Up, page 22-39](#)
- [Scenario 5: Designated MSFC Goes Down, page 22-40](#)

Scenario 1: Enabling Configuration Synchronization on Both MSFCs

This scenario assumes both MSFCs are up.

When you enable configuration synchronization on both MSFCs, the IP addresses on all the interfaces are checked first. If an IP address is specified for the designated MSFC but not specified for the nondesignated MSFC, a message is displayed indicating the first interface for which the alternate IP address was not specified.

After checking IP addresses, the HSRP addresses are checked; if an HSRP address is specified for the designated MSFC but not specified for the nondesignated MSFC, a message is displayed indicating the first interface for which the alternate HSRP (standby) address was not specified.

After checking the HSRP addresses, the IPX network address is checked.

The designated MSFC is configured first. This example shows a missing alternate configuration for the VLAN 1 interface:

```

Router-16# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router-16(config)# redundancy
Router-16(config-r)# high-availability
Router-16(config-r-ha)# config-sync

Alternate IP address missing for Vlan1
The alternate configuration is missing. The auto-config sync can not be enabled

```



Note

When specifying the alternate IP configuration, the configuration specified before the **alt** keyword relates to the MSFC on the supervisor engine in slot 1 of the switch; the configuration specified after the **alt** keyword relates to the MSFC on the supervisor engine in slot 2. See the [“alt Keyword Usage” section on page 22-33](#) for more information.

This example shows how to specify the alternate configuration for VLAN 1:

```
Router-16(config)# interface vlan 1
Router-16(config-if)# ip address 70.0.70.4 255.255.0.0 alt ip address 70.0.70.5
255.255.0.0
Router-16(config-if)# exit
```

This example shows that high-availability redundancy is accepted:

```
Router-16(config)# redundancy
Router-16(config-r)# high-availability
Router-16(config-r-ha)# config-sync
Router-16(config-r-ha)# end
Router-16#
00:03:31: %SYS-5-CONFIG_I: Configured from console by console
```

Because the Config Sync AdminStatus on the nondesignated MSFC is disabled, the Config Sync RuntimeStatus on the designated MSFC will remain in disabled mode. The following message is displayed on the designated MSFC:

```
00:17:05: %RUNCFGSYNC-6-SYNCEVENT:
Non-Designated Router is now online
High-Availability Redundancy Feature is not enabled on the Non-Designated Router
```

This example shows how to enable the configuration synchronization feature on the nondesignated MSFC:

```
Router-151(config)# redundancy
Router-151(config-r)# high-availability
Router-151(config-r-ha)# config-sync
Router-151(config-r-ha)# end
Router-151#
00:03:31: %SYS-5-CONFIG_I: Configured from console by console
```



Note

When you enable high-availability redundancy, the configuration mode is disabled on the console of the nondesignated MSFC; only the EXEC mode is available.

The following message, acknowledging that the high-availability redundancy is enabled, and that the configuration mode will be automatically exited, is displayed on the nondesignated MSFC:

```
00:18:57: %RUNCFGSYNC-6-SYNCEVENT:
The High-Availability Redundancy Feature is enabled
The config mode is no longer accessible

Router-15#

00:19:41: %RUNCFGSYNC-6-SYNCEVENT:
Non-Designated Router is now online
Running Configuration Synchronization will begin in 1 minute
```

A one-minute timer will start, allowing for stabilization of the nondesignated MSFC. When the timer expires, a snapshot of the current running configuration is sent to the nondesignated MSFC. This message is displayed before the running configuration is synchronized:

```
00:20:41: %RUNCFGSYNC-6-SYNCEVENT:
Syncing Running Configuration to the Non-Designated Router

00:20:41: %RUNCFGSYNC-6-SYNCEVENT:
Syncing Startup Configuration to the Non-Designated Router
```

These examples show that the designated MSFC and nondesignated MSFC have the same running configuration after synchronization:

```
<designated MSFC>
Router-16# show running-config
Building configuration...

Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Router-15 alt hostname Router-16
!
boot bootldr bootflash:c6msfc-boot-mz.120-7.XE1
!
ip subnet-zero
!
ip cef
redundancy
  high-availability
  config-sync
cns event-service server
!
!
!
interface Vlan1
  ip address 70.0.70.4 255.255.0.0 alt ip address 70.0.70.5 255.255.0.0
!
interface Vlan10
ip address 192.10.10.1 255.255.255.0 alt ip address 192.10.10.2 255.255.255.0
  no ip redirects
  shutdown
  standby ip 192.20.20.1 alt standby ip 192.20.20.1
!
ip classless
ip route 223.255.254.0 255.255.255.0 70.0.100.0
no ip http server
!
!
!
line con 0
  transport input none
line vty 0 4
  login
  transport input lat pad mop telnet rlogin udptn nasi
!
end
```

```
<nondesignated MSFC>
Router-15# show running-config
Building configuration...

Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
```

```

hostname Router1 alt hostname Router2
!
boot bootldr bootflash:c6msfc-boot-mz.120-7.XE1
!
ip subnet-zero
!
ip cef
redundancy
  high-availability
  config-sync
cns event-service server
!
!
!
interface Vlan1
  ip address 70.0.70.4 255.255.0.0 alt ip address 70.0.70.5 255.255.0.0
!
interface Vlan10
  ip address 192.10.10.1 255.255.255.0 alt ip address 192.10.10.2 255.255.255.0
  no ip redirects
  shutdown
  standby ip 192.20.20.1 alt standby ip 192.20.20.1
!
ip classless
ip route 223.255.254.0 255.255.255.0 70.0.100.0
no ip http server
!
!
!
line con 0
  transport input none
line vty 0 4
  login
  transport input lat pad mop telnet rlogin udptn nasi
!
end

```

Scenario 2: Disabling Configuration Synchronization on the Designated MSFC

In this scenario, configuration synchronization is enabled. These examples show how to disable configuration synchronization:

```

Router-16# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router2(config)# redundancy
Router2(config-r)# high-availability
Router2(config-r-ha)# no config-sync

```

When configuration synchronization is disabled, the following message is displayed on the nondesignated MSFC:

```

00:13:00: %RUNCFGSYNC-6-SYNCEVENT:
The High-Availability Redundancy Feature is now disabled
The config mode is now accessible

```

Configuration mode is available on the CLI of the designated and nondesignated MSFC.

Scenario 3: Designated MSFC Comes Up

In this scenario, Config Sync AdminStatus is enabled. The designated MSFC validates the alternate configuration, allowing configuration synchronization to occur when the nondesignated MSFC comes up.

Because the nondesignated MSFC is not up yet, Config Sync RuntimeStatus is disabled, and there is no configuration synchronization. See the [“Scenario 4: Nondesignated MSFC Comes Up”](#) section on page 22-39 for information on the nondesignated MSFC.

This example shows that Router-16 is the designated MSFC, Config Sync AdminStatus is enabled, and Config Sync RuntimeStatus is disabled:

```
Router-16# show redundancy
Designated Router: 1 Non-designated Router:0

Redundancy Status: designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: disabled
```

Scenario 4: Nondesignated MSFC Comes Up

Config Sync AdminStatus is Enabled

In this scenario, the nondesignated MSFC notifies the designated MSFC that it is up and Config Sync AdminStatus is enabled. The designated MSFC requests the nondesignated MSFC to enable Config Sync RuntimeStatus. The nondesignated MSFC enables Config Sync RuntimeStatus.

The following message is displayed on the nondesignated MSFC:

```
00:00:07: %RUNCFGSYNC-6-SYNCEVENT:
The High-Availability Redundancy Feature is enabled
The config mode is no longer accessible

00:00:51: %RUNCFGSYNC-6-SYNCEVENT:
Non-Designated Router is now online
Running Configuration Synchronization will begin in 1 minute
```

A one-minute timer will start, allowing the nondesignated MSFC to stabilize. When the timer expires, a snapshot of the current running configuration is sent to the nondesignated MSFC. The following message is displayed before synchronizing the running configuration:

```
00:01:51: %RUNCFGSYNC-6-SYNCEVENT:
Syncing Running Configuration to the Non-Designated Router
```

Config Sync AdminStatus is Disabled

In this scenario, the nondesignated MSFC notifies the designated MSFC that it is up. Because the Config Sync AdminStatus is disabled on the nondesignated MSFC, the designated MSFC displays the following message indicating that high-availability redundancy needs to be enabled on the nondesignated MSFC:

```
Router-16#
Non-Designated Router came up.
High-Availability Redundancy Feature is not enabled on the Non-Designated Router
```

This example shows how to enable the high-redundancy availability feature on the nondesignated MSFC:

```
Router-15# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router-15(config)# redundancy
Router-15(config-r)# high-availability
Router-15(config-r-ha)# config-sync
```

```

Router-15(config-r-ha)#
00:03:47: %SYS-5-CONFIG_I: Configured from console by console
00:03:47: %RUNCFGSYNC-6-SYNCEVENT:
The High-Availability Redundancy Feature is enabled
The config mode is no longer accessible

00:00:51: %RUNCFGSYNC-6-SYNCEVENT:
Non-Designated Router is now online
Running Configuration Synchronization will begin in 1 minute

```

A one-minute timer will start, allowing the nondesignated MSFC to stabilize. When the timer expires, a snapshot of the current running configuration is sent to the nondesignated MSFC. This message is displayed before synchronizing the running configuration:

```

00:01:51: %RUNCFGSYNC-6-SYNCEVENT:
Syncing Running Configuration to the Non-Designated Router

```

These examples show that Config Sync AdminStatus and RuntimeStatus are enabled on the designated and nondesignated MSFCs:

```

Router-15# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: non-designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled

```

```

Router-16# show redundancy
Designated Router: 1 Non-designated Router:2

Redundancy Status: designated
Config Sync AdminStatus : enabled
Config sync RuntimeStatus: enabled

```

Scenario 5: Designated MSFC Goes Down

In this scenario, the nondesignated MSFC will become the designated MSFC. Configuration synchronization is disabled, and the configuration mode on the CLI is now available.

When the previously designated MSFC comes back up, it will become the nondesignated MSFC; see the [“Scenario 4: Nondesignated MSFC Comes Up”](#) section on page 22-39.

Single Router Mode Redundancy

These sections describe how to configure single router mode (SRM) redundancy:

- [Hardware and Software Requirements, page 22-41](#)
- [Configuration Guidelines, page 22-42](#)
- [Configuring Single Router Mode Redundancy, page 22-42](#)
- [Upgrading Images with Single Router Mode Enabled, page 22-44](#)
- [Getting Out of Single Router Mode, page 22-45](#)

SRM redundancy is an alternative to internally redundant (dual) MSFC2 configurations where both MSFC2s are active at the same time. In SRM redundancy, only the designated router is visible to the network at any given time. The nondesignated router is booted up completely and participates in configuration synchronization which is automatically enabled when entering SRM. All configuration following the “alt” keyword is ignored in SRM. Due to this, the nondesignated router’s configuration is exactly the same as the designated router but its interfaces are kept in a line down state and are not visible to the network. Processes, such as routing protocols, are created on the nondesignated router and the designated router, but all nondesignated router interfaces are in a line down state; they do not send or receive updates from the network.

When the designated router fails, the nondesignated router changes its state from a nondesignated router to a designated router and its interface state changes to link up. It builds up its routing table while the existing supervisor engine switch processor entries are used to forward Layer 3 traffic. After the newly designated router builds its routing table, the entries in the switch processor are updated.

Hardware and Software Requirements

To configure SRM redundancy, you must have the following hardware and software:

- A single chassis with two identical supervisor engine daughter card configurations:
 - Supervisor Engine 2 with PFC2 and MSFC2
 - Supervisor Engine 1 with PFC and MSFC or MSFC2



Note Cisco IOS Release 12.1(8a)E4 provides initial support for single router mode (SRM) redundancy with Supervisor Engine 1 and MSFC.

When using Supervisor Engine 1 with the MSFC or MSFC2 for SRM redundancy, be aware that failover to the second MSFC is not stateful for multicast MLS. When the primary MSFC fails, all multicast MLS entries are removed and are then recreated and reinstalled in the hardware by the newly active MSFC.

- Supervisor engine software release 6.3(1) or later releases
- Cisco IOS Release 12.1(8a)E2 or later releases

Configuration Guidelines

Use these guidelines when configuring SRM redundancy:

- SRM redundancy requires that both the designated router and nondesignated router run the same Cisco IOS image.
- SRM redundancy requires that a Cisco IOS image is present in the bootflash of both the designated router and nondesignated router.
- With SRM redundancy, the nondesignated router cannot connect to external networks.
- With SRM redundancy, we do not recommend booting from an external network with the designated router. Booting from the network could severely degrade SRM functionality.
- With SRM redundancy, the designated router can reach external networks and copy commands such as **copy tftp:** can be used without any restrictions.
- For SRM to work properly, high availability must be enabled on the supervisor engine.
- When using authentication methods to control access to the switch such as RADIUS or TACACS+, you need to configure a fallback option to login in with a local username and password if you want to be able to access the nondesignated router through the **switch console** or **session** commands.

See [Chapter 21, “Configuring Switch Access Using AAA”](#) for information on configuring the fallback option.

Configuring Single Router Mode Redundancy

To configure SRM redundancy, perform these steps:



Caution

Before going from dual router mode to SRM redundancy, we recommend that you use the **copy running-config** command on the MSFCs to save the non-SRM configuration to bootflash. When going to SRM redundancy, the alternative configuration (the configuration following the **alt** keyword) is lost. Therefore, before enabling SRM redundancy, save the dual router mode configuration to bootflash by entering the following command on both MSFCs:
copy running-config bootflash:nosrm_dual_router_config.

See the [“Getting Out of Single Router Mode”](#) section on page 22-45 for additional information.



Note

This procedure assumes that the designated router is the MSFC2 in slot 1 and the nondesignated router is the MSFC2 in slot 2; the active supervisor engine is in slot 1 and the standby supervisor engine is in slot 2.

-
- Step 1** Enter the **show version** command to ensure that both supervisor engines are running supervisor engine software release 6.3(1) or later releases.
- Step 2** Enter the **set system highavailability enable** command to enable high availability on the active supervisor engine. Enter the **show system highavailability** command to verify that high availability is enabled.
- Step 3** If you have a console connection, enter the **switch console** command to access the designated router. If connected through a Telnet session, enter the **session mod** command to access the designated router.

- Step 4** Copy the Cisco IOS Release 12.1(8a)E2 or later image to the bootflash of the designated router and nondesignated router.
- Step 5** Set the boot image and configuration register on the designated router and nondesignated router to boot the new image on a reload:
- For the designated router, enter **boot system flash bootflash:image_name** and ensure that this image is the first in the boot list. Clear any existing “boot system” commands that appear in the running configuration (**show running-config**) using the **no** form of the **boot system** command.
- For the nondesignated router, set the configuration register to auto boot by entering **config-register 0x102**.
- Step 6** Enter the **reload** command to reload the designated router and nondesignated router.



Note If you already have SRM-capable Cisco IOS images loaded, you do not need to perform Step 6.

- Step 7** Disable configuration synchronization (**config-sync**) on the designated router using the **no** form of the command. Enter the **write memory** command. This lets you have access to configuration mode on both designated and nondesignated routers.
- Step 8** Enable SRM on the designated router first, and then enable SRM on the nondesignated router as follows:

```
Router(config)#redundancy
Router(config-r)#high-availability
Router(config-r-ha)#single-router-mode
```

- Step 9** Enter the **write memory** command on the designated router to ensure that the nondesignated router’s start-up configuration has SRM enabled.
- Step 10** Enter the **show startup-config** command on the nondesignated router to ensure that the nondesignated router has the following configuration statements:

```
redundancy
 high-availability
 single-router-mode
```

- Step 11** Enter the **show redundancy** command on the designated router and nondesignated router to ensure that both have the following configuration statement:

```
Single Router Mode RuntimeStatus: enabled
```

If not, repeat Steps 9 and 10 allowing sufficient time between steps.

- Step 12** Enter the **reload** command to reload the nondesignated router. When asked whether the configuration should be saved, enter **no**.

This display summarizes the above configuration commands used on the designated router and nondesignated router to enable SRM redundancy:

Time	Designated Router	Nondesignated Router
----	---	----
t0:	conf t->red->hi->no config-sync	
t1:		conf t->red->hi->no config-sync
t2:	conf t->red->hi->single-router-mode	
t3:		conf t->red->hi->single-router-m
t4:	write mem	
t5:		reload

Upgrading Images with Single Router Mode Enabled

This section describes how to upgrade the Cisco IOS image on the active and standby MSFC when SRM is running. The new image name is c6msfc2-jsv-mz.9E. The standby MSFC cannot load an image using TFTP, but it can load an image from the supervisor engine Flash PC card (sup-slot0:).



Note

This procedure impacts data traffic. We recommend that it be performed during a scheduled maintenance window.

To upgrade the images, perform these steps:

-
- Step 1** On the active supervisor engine, enter the **copy tftp sup-slot0:** command and follow the prompts to load the new (c6msfc2-jsv-mz.9E) image onto the supervisor engine Flash PC card.
- Step 2** If you have a console connection, enter the **switch console** command to access the active MSFC. If you are connected through a Telnet session, enter the **session mod** command to access the active MSFC.
- Step 3** On the active MSFC, copy the new image from the supervisor engine Flash PC card to the MSFC bootflash as follows:
- ```
copy sup-slot0:c6msfc2-jsv-mz.9E bootflash:c6msfc2-jsv-mz.9E
```
- Step 4** Access the standby MSFC by entering the **switch supervisor** command and then the **switch console** command on the active supervisor engine.



### Note

The standby MSFC does not appear in the **show module** command display that is issued from the active supervisor engine.

- 
- Step 5** On the standby MSFC, copy the new image from the supervisor engine Flash PC card to the MSFC bootflash as follows:
- ```
copy sup-slot0:c6msfc2-jsv-mz.9E bootflash:c6msfc2-jsv-mz.9E
```
- Step 6** On the active MSFC, specify that the new image is booted when the MSFC is reloaded as follows:
- ```
boot system flash bootflash:c6msfc2-jsv-mz.9E
```
- Step 7** On the active MSFC, enter the **write memory** command to ensure that the standby MSFC start-up configuration gets the boot information.
- Step 8** Enter the **reload** command to reload the standby MSFC.
- Step 9** Enter the **show redundancy** command on the active and standby MSFCs to ensure that both have the following configuration statement:
- ```
Single Router Mode RuntimeStatus: enabled
```
- Step 10** Enter the **reload** command to reload the active MSFC.
- Both MSFCs are now running the c6msfc2-jsv-mz.9E image.
-

Getting Out of Single Router Mode

**Note**

If you saved a copy of the running configuration used in dual router mode before configuring SRM redundancy, you do not need to use the procedure in this section. To get out of SRM redundancy and back to dual router mode, enter the following command on both MSFCs:
copy bootflash:nosrm_dual_router_config startup-config. After the configurations are copied, reload the MSFCs using the **reload** command.

To get out of SRM, perform these steps:

-
- Step 1** On the designated router, disable SRM using the **no** form of the command as follows:
- ```
Router(config)#redundancy
Router(config-r)#high-availability
Router(config-r-ha)#no single-router-mode
```
- Step 2** Enter the **write memory** command on the designated router and nondesignated router.
- Step 3** Enter the **show startup-config** command on the designated and nondesignated routers to ensure that “single-router mode” is not in the startup configuration.
- Step 4** Enter the **reload** command to reload the designated router and nondesignated router.  
SRM is now disabled on the designated router and nondesignated router.
- 

## Manual-Mode MSFC Redundancy

**Note**

Manual-mode MSFC redundancy will be supported until December, 2002, due to the release of supervisor engine software release 6.3(1), which contains the feature SRM. Cisco recommends using SRM rather than manual-mode MSFC redundancy to attain automatic Layer-3 failover capabilities in addition to unlimited support of the feature.

These sections describe how to configure redundant MSFCs with one MSFC active and the other MSFC in ROM-monitor mode:

- [Hardware and Software Requirements, page 22-46](#)
- [Guidelines for Configuring Manual-Mode MSFC Redundancy, page 22-46](#)
- [Accessing the Standby MSFC, page 22-47](#)
- [Manually Booting the MSFC, page 22-47](#)
- [Setting the MSFC Configuration Register, page 22-47](#)
- [MSFC Recovery Procedures, page 22-48](#)

## Hardware and Software Requirements

To configure Layer 3 redundancy, you must have at least one of the following configurations:

- A single chassis with two identical supervisor engine daughtercard configurations:
  - Supervisor Engine 1 with Policy Feature Card (PFC) and MSFC or MSFC2 (both supervisor engines must have the same type of MSFC)
  - Supervisor Engine 2 with PFC2 and MSFC2
- Two chassis with a supervisor engine in each—You must have at least one supervisor engine in each chassis. Each supervisor engine must be equipped with a PFC and an MSFC.
- Manual-mode MSFC redundancy requires the following software:
  - Supervisor engine software release 6.1(3) or later releases and MSFC IOS Release 12.1(7)E or later releases
  - Supervisor engine software release 5.5.8 or later releases and MSFC IOS Release 12.1(7a)E1 or later releases



Note

---

Each MSFC must be running the same release of Cisco IOS software.

---

## Guidelines for Configuring Manual-Mode MSFC Redundancy

Follow these guidelines to configure manual-mode MSFC redundancy:

- Because the MSFC switchover is manual, we recommend that you have this feature only in environments where externally redundant routers are present and where either HSRP is used or some form of gateway discovery is implemented on hosts.
- Ensure that the configuration register on the active MSFC (MSFC-15) is set to 0x2102 and that the configuration register on the MSFC in ROM-monitor mode (MSFC-16) is set to 0x0. This setting prevents both MSFCs from becoming active at the same time and allows the active MSFC to come online after a reset. See the [“Setting the MSFC Configuration Register”](#) section on page 22-47 for details on setting the configuration register.



Note

---

Setting both MSFCs to 0x0 is a supported option but requires user intervention in the event the switch is reset.

---

- To conserve IP address space and reduce the overall Layer 3 complexity, ensure that configuration synchronization is disabled on both MSFCs and that all “alt” addresses are removed. If alt addresses are used, IP address space is not conserved and in cases where link-level peering is present (such as BGP), the Layer 3 complexity is increased.
- When the MSFC in ROM-monitor mode is brought up during a maintenance window, ensure that it has the exact same configuration as the active MSFC. Follow the configuration guidelines in [Table 22-2](#) on page 22-20.
- During manual-mode MSFC redundancy, high availability should be enabled on the supervisor engine to keep Layer 2 downtime to a minimum when doing an MSFC switchover. Since high availability is not compatible with protocol filtering, port security, DVLAN, or GVRP, we recommend that you disable these features when using manual-mode MSFC redundancy.

- Ensure that the console port on both supervisor engines is accessible to operations personnel (out-of-band access through terminal server or modem).

The procedures in this section use the **switch console** command to access the MSFC from the supervisor engine. The **switch console** command is not supported on Telnet sessions.

## Accessing the Standby MSFC

To access the standby MSFC, enter the **switch supervisor** command followed by the **switch console** command.



Note

The standby MSFC does not appear in the **show module** command display issued from the active supervisor engine.

## Manually Booting the MSFC

If the configuration register on both MSFCs is set to 0x0, then MSFC manual mode requires that the MSFC be manually booted each time the switch is reset. To manually boot the MSFC, perform these steps:

- 
- Step 1** Enter the **switch console** command to gain access to the MSFC ROMMON prompt.
  - Step 2** Enter the **boot bootflash:image** command.
  - Step 3** Once the MSFC has booted, enter **^C^C^C** at the Router> prompt to return to the supervisor engine prompt. Now you may enter the **session** command to access the MSFC.
- 

## Setting the MSFC Configuration Register

For manual-mode MSFC redundancy, set the configuration registers as follows:

- 
- Step 1** From Cisco IOS configuration mode on the active MSFC (MSFC-15), perform the following:
 

```
Router(config)#config-register 0x2102
Router(config)#
```
  - Step 2** From Cisco IOS configuration mode on the MSFC in ROM-monitor mode (MSFC-16), perform the following:
 

```
Router(config)#config-register 0x0
Router(config)#
```
- 



Note

We recommend that **boot** system commands in both MSFC configurations point to a valid image on bootflash and that you do not set the configuration registers to ignore these boot commands.

## MSFC Recovery Procedures

This section describes how to recover from temporary or permanent MSFC failures.

A temporary failure of the active MSFC results in the MSFC simply rebooting because the configuration register is set to 0x2102.

A suspected permanent failure of the active MSFC first needs to be verified. Do this by entering the **reset 15** command from the active supervisor engine's console port and see if the active MSFC reboots without problems. If it does not, you have the following two options to switch over to the standby MSFC.

### Option 1: If You Have Physical Access to the Switch

If you have physical access to the switch, use this option. You can remove the active supervisor engine with the problematic MSFC, so the redundant supervisor engine will take over. From the redundant supervisor engine's physical console port, perform these steps:

- 
- Step 1 Enter the **switch console** command.
  - Step 2 From the ROMMON prompt, enter the **boot bootflash:image** command.
  - Step 3 After the standby MSFC has booted, from Cisco IOS configuration mode enter the **config-register 0x2102** command to ensure the MSFC will boot when the switch is reset.
- 

### Option 2: If You Have Remote Access Only to the Switch

If you only have remote access to the switch, use this option. From the active supervisor engine with the problematic MSFC, perform these steps:




---

**Note** If the problematic MSFC is on the standby supervisor engine, enter the **switch supervisor** command.

---

- 
- Step 1 Enter the **switch console** command.
  - Step 2 Send a Break signal to get into the problematic MSFC's ROMMON (the break will work if the MSFC is continually rebooting). You need to time the break so that it is issued after the system bootstrap message, but before the main Cisco IOS image is decompressed (see the two arrows in the following display output):

```
System Bootstrap, Version 12.0(3)XE, RELEASE SOFTWARE
Copyright (c) 1998 by cisco Systems, Inc.
Cat6k-MSFC platform with 131072 Kbytes of main memory <===== ISSUE BREAK AFTER THIS
POINT
```

```
Self decompressing the image :
#####
[OK]
```

```
<=====BUT BEFORE THIS POINT
```

```
Self decompressing the image :
#####
#####
#####
```

```
#####
#####
#####
#####
#####
[OK]
```

- Step 3** At the ROMMON prompt, enter the **confreg** command:
- Enter **y** at the “do you wish to change the configuration? y/n [n]:” prompt
  - Press **Enter** to accept the default for all questions until you reach this prompt: “change the boot characteristics? y/n [n]:”
  - Enter **y**
  - Enter **0** to select the “0 = ROM Monitor” option at the next prompt
  - Review the Configuration Summary to ensure the following value: boot: the ROM Monitor
  - You are again prompted with: “do you wish to change the configuration? y/n [n]:”
  - Enter **n**
  - You are returned to the ROMMON prompt
- Step 4** Enter the **reset** command and verify that the MSFC boots into ROMMON. This step ensures that this MSFC and the active MSFC will not boot concurrently.
- Step 5** Enter **^C^C^C** to return to the supervisor engine prompt.
- Step 6** Ensure that high availability has synchronized the supervisor engine state by entering the **show system highavailability** command and verifying that high availability “Operational-status” is ON.
- Step 7** Enter the **switch supervisor** command.
- Step 8** Enter the **switch console** command.
- Step 9** From the standby MSFC’s ROMMON prompt, perform step 3 above but in step 3d, select option 2 “boot system” as follows:
- ```
change the boot characteristics? y/n [n]: y
enter to boot:
 0 = ROM Monitor
 1 = the boot helper image
 2-15 = boot system
 [2]: 2 <=====

Configuration Summary
enabled are:
load rom after netboot fails
console baud: 9600
boot: the ROM Monitor

do you wish to change the configuration? y/n [n]: n

You must reset or power cycle for new config to take effect
rommon 2 >
```
- Step 10** Enter the **reset** command at the ROMMON prompt to boot the system.
- Step 11** After the MSFC has booted from the IOS configuration mode on the newly active MSFC’s console port, enter the **config-register 0x2102** command.

Step 12 Enter `^C^C^C` to return to the supervisor engine prompt.
