



Configuring High Availability

This chapter describes how to configure high availability, and describes the switchover processes.

- [Finding Feature Information, on page 1](#)
- [Feature History for High Availability, on page 1](#)
- [Configuring High Availability, on page 2](#)

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Feature History for High Availability

This table lists the New and Changed features.

Table 1: New and Changed Features

| Feature Name | Release | Feature Information |
|--------------------------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standby Supervisor's mgmt0 Status | 9.2(1) | <p>A syslog was introduced to alert the user if the standby supervisor's Ethernet management port is disconnected or down before performing an ISSU or system switchover.</p> <p>The show interface mgmt number standby command was introduced to display the status of the supervisor's mgmt0 link when issued from the active supervisor.</p> <p>The system switchover bypass-standby-mgmt0 command was introduced to skip checking for the status of the standby supervisor's mgmt0 link during a system switchover.</p> |
| Internal CRC Detection and Isolation | 8.5(1) | Internal CRC detection and error logging without isolation is enabled by default. |
| Standby Supervisor's mgmt0 Link | 8.4(2) | The standby supervisor's management Ethernet link on Cisco MDS Director switches is brought up when the supervisor reaches the standby state. |
| Internal CRC Detection and Isolation | 8.4(2) | <p>Added an option to log internal CRC errors without taking any action.</p> <p>The following command was modified:</p> <p>hardware fabric crc [threshold count] [log-only]</p> |

Configuring High Availability

This chapter describes how to configure high availability, and describes the switchover processes.

About High Availability

Process restartability provides the high availability functionality in Cisco MDS 9000 Series switches. This process ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This process is able to restore its state prior to the failure and continues executing from the failure point going forward.

From Cisco MDS NX-OS Release 8.4(2), the standby supervisor's management Ethernet link on Cisco MDS Director switches is brought up when the supervisor reaches the standby state. This will help prevent the port in the adjacent Ethernet switch from being detected as continuously down and potentially decommissioned.

From Cisco MDS NX-OS Release 9.2(1), NX-OS checks and prints a syslog to alert the user if the standby supervisor's Ethernet management link is disconnected or down before performing an In-Service Software

Upgrade (ISSU), In-Service Software Downgrade (ISSD), or system switchover. You can also use the **show interface mgmt number standby** command to display the status of the standby supervisor's mgmt0 link when issued from the active supervisor. Use the **system switchover bypass-standby-mgmt0** command to skip checking for the status of the standby supervisor's mgmt0 link during a system switchover. For information on system messages, see the [Cisco MDS 9000 Family and Nexus 7000 Series NX-OS System Messages Reference](#).

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because control traffic is not impacted.
- It does not disrupt data traffic because the switching modules are not impacted.
- Switching modules are not reset.



Note Switchover is not allowed if **auto-copy** is in progress.

Switchover Processes

Switchovers occur by one of the following two processes:

- The active supervisor module fails and the standby supervisor module automatically takes over.
- You manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a switchover process has started another switchover process cannot be started on the same switch until a stable standby supervisor module is available.



Caution If the standby supervisor module is not in a stable state (ha-standby), a switchover is not performed.

Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module automatically synchronizes its image with the running image on the active supervisor module.



Note The image a supervisor module is booted up from cannot be deleted from bootflash. This is to ensure that the new standby supervisor module is able to synchronize during the process.

Manual Switchover Guidelines

Be aware of the following guidelines when performing a manual switchover:

- When you manually initiate a switchover, system messages indicate the presence of two supervisor modules.
- A switchover can only be performed when two supervisor modules are functioning in the switch.

- The modules in the chassis are functioning as designed.

Manually Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, use the **system switchover** command. After you enter this command, another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

To ensure that an HA switchover is possible, enter the **show system redundancy status** command or the **show module** command. If the command output displays the HA standby state for the standby supervisor module, then the switchover is possible. See [Verifying Switchover Possibilities](#) for more information.

Verifying Switchover Possibilities

This section describes how to verify the status of the switch and the modules before a manual switchover.

- Use the **show interface mgmt number standby** command to verify that the standby supervisor's mgmt0 link is up.
- Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.
- Use the **show module** command to verify the status (and presence) of a module at any time. A sample output of the **show module** command follows:

```
switch# show module
Mod Ports Module-Type Model Status
-----
2 8 IP Storage Services Module DS-X9308-SMIP ok
5 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 active *
6 0 Supervisor/Fabric-1 DS-X9530-SF1-K9 ha-standby
8 0 Caching Services Module DS-X9560-SMAP ok
9 32 1/2 Gbps FC Module DS-X9032 ok
Mod MAC-Address(es) Serial-Num
-----
2 00-05-30-00-9d-d2 to 00-05-30-00-9d-de JAB064605a2
5 00-05-30-00-64-be to 00-05-30-00-64-c2 JAB06350B1R
6 00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd JAB06350B1R
8 00-05-30-01-37-7a to 00-05-30-01-37-fe JAB072705ja
9 00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 JAB06280ae9
* this terminal session
```

The Status column in the output should display an OK status for switching modules and an active or HA-standby status for supervisor modules. If the status is either OK or active, you can continue with your configuration.

- Use the **show boot auto-copy** command to verify the configuration of the auto-copy feature and if an auto-copy to the standby supervisor module is in progress. Sample outputs of the **show boot auto-copy** command follow:

```
switch# show boot auto-copy
Auto-copy feature is enabled
switch# show boot auto-copy list
No file currently being auto-copied
```

Configuring Internal CRC Detection and Isolation



Note This functionality is disabled by default.

To configure internal CRC detection and isolation, perform these steps:

-
- Step 1** Enter configuration mode:
 switch# **configure terminal**
- Step 2** Enable internal CRC detection, isolation, and error logging:
 switch(config)# **hardware fabric crc [threshold count]**
 Or
 Enable internal CRC detection and error logging without isolation in Cisco MDS NX-OS Release 8.4(2) and later releases:
 switch(config)# **hardware fabric crc [threshold count] log-only**
 From Cisco MDS NX-OS Release 8.5(1), internal CRC detection and error logging without isolation is enabled by default.
 The error rate is measured over a sequential 24-hour window, where the error count is reset to 0 at the start of each window. The threshold range is 1–100. The default threshold is 3 when the threshold is not specified.
- Step 3** (Optional) Disable internal CRC detection, isolation, and error logging:
 switch(config)# **no hardware fabric crc**
- Step 4** Save the configuration change:
 switch(config)# **copy running-config startup-config**
-

Default Settings for Internal CRC Detection and Isolation

The table below lists the default settings for interface parameters.

Table 2: Default Settings for Internal CRC Detection and Isolation

| Parameters | Default |
|-----------------------------|----------|
| Internal CRC Error Handling | Disabled |

Copying Boot Variable Images to the Standby Supervisor Module

You can copy the boot variable images that are in the active supervisor module (but not in the standby supervisor module) to the standby supervisor module. Only those KICKSTART and SYSTEM boot variables that are set for the standby supervisor module can be copied. For module (line card) images, all boot variables are copied to the corresponding standby locations (bootflash: or slot0:) if not already present.

Enabling Automatic Copying of Boot Variables

To enable or disable automatic copying of boot variables, follow these steps:

-
- Step 1** Enters configuration mode.
- ```
switch# configure terminal
switch(configure)#
```
- Step 2** Enables (default) automatic copying of boot variables from the active supervisor module to the standby supervisor module.
- ```
switch(configure)# boot auto-copy
Auto-copy administratively enabled
```
- Step 3** Disables the automatic copy feature.
- ```
switch(configure)# boot auto-copy
Auto-copy administratively disabled
```
- 

## Verifying the Copied Boot Variables

Use the **show boot auto-copy** command to verify the current state of the copied boot variables. This example output shows that automatic copying is enabled:

```
switch# show boot auto-copy
Auto-copy feature enabled
```

This example output shows that automatic copying is disabled:

```
switch# show boot auto-copy
Auto-copy feature disabled
```

Use the **show boot auto-copy list** command to verify what files are being copied. This example output displays the image being copied to the standby supervisor module's bootflash. Once this is successful, the next file will be image2.bin.




---

**Note** This command only displays files on the active supervisor module.

---

```
switch# show boot auto-copy list
File: /bootflash:/image1.bin
Bootvar: kickstart
File:/bootflash:/image2.bin
Bootvar: system
```

This example output displays a typical message when the **auto-copy** option is disabled or if no files are copied:

```
switch# show boot auto-copy list
No file currently being auto-copied
```

## Displaying HA Status Information

Use the **show system redundancy status** command to view the HA status of the system. Tables Redundancy States to Internal States [Table 3: Redundancy States](#), on page 7 and [Table 5: Internal States](#), on page 8 explain the possible output values for the redundancy, supervisor, and internal states.

```
switch# show system redundancy status
Redundancy mode

 administrative: HA
 operational: HA
This supervisor (sup-1)

 Redundancy state: Active
 Supervisor state: Active
 Internal state: Active with HA standby
Other supervisor (sup-2)

 Redundancy state: Standby
 Supervisor state: HA standby
 Internal state: HA standby
```

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is Active with HA standby and the other supervisor module is HA standby, the switch is operationally HA and can do automatic synchronization.
- If the internal state of one of the supervisor modules is none, the switch cannot do automatic synchronization.

The following table lists the possible values for the redundancy states.

**Table 3: Redundancy States**

| State        | Description                                                                                                                                                                                                                                                                                                                                    |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Not present  | The supervisor module is not present or is not plugged into the chassis.                                                                                                                                                                                                                                                                       |
| Initializing | The diagnostics have passed and the configuration is being downloaded.                                                                                                                                                                                                                                                                         |
| Active       | The active supervisor module and the switch is ready to be configured.                                                                                                                                                                                                                                                                         |
| Standby      | A switchover is possible.                                                                                                                                                                                                                                                                                                                      |
| Failed       | The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state.<br><br><b>Note</b> You should try to initialize the supervisor module until it comes up as HA-standby. This state is a temporary state. |
| Offline      | The supervisor module is intentionally shut down for debugging purposes.                                                                                                                                                                                                                                                                       |

| State   | Description                                                                                                    |
|---------|----------------------------------------------------------------------------------------------------------------|
| At BIOS | The switch has established connection with the supervisor and the supervisor module is performing diagnostics. |
| Unknown | The switch is in an invalid state. If it persists, call TAC.                                                   |

The following table lists the possible values for the supervisor module states.

**Table 4: Supervisor States**

| State      | Description                                                           |
|------------|-----------------------------------------------------------------------|
| Active     | The active supervisor module in the switch is ready to be configured. |
| HA standby | A switchover is possible.                                             |
| Offline    | The switch is intentionally shut down for debugging purposes.         |
| Unknown    | The switch is in an invalid state and requires a support call to TAC. |

The following table lists the possible values for the internal redundancy states.

**Table 5: Internal States**

| State                          | Description                                                                                                                                 |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| HA standby                     | The HA switchover mechanism in the standby supervisor module is enabled (see the <a href="#">Synchronizing Supervisor Modules</a> section). |
| Active with no standby         | A switchover is not possible.                                                                                                               |
| Active with HA standby         | The active supervisor module in the switch is ready to be configured. The standby supervisor module is in the HA-standby state.             |
| Shutting down                  | The switch is being shut down.                                                                                                              |
| HA switchover in progress      | The switch is in the process of changing over to the HA switchover mechanism.                                                               |
| Offline                        | The switch is intentionally shut down for debugging purposes.                                                                               |
| HA synchronization in progress | The standby supervisor module is in the process of synchronizing its state with the active supervisor modules.                              |
| Standby (failed)               | The standby supervisor module is not functioning.                                                                                           |
| Active with failed standby     | The active supervisor module and the second supervisor module is present but is not functioning.                                            |
| Other                          | The switch is in a transient state. If it persists, call TAC.                                                                               |

From Cisco MDS NX-OS Release 8.5(1), use the **show hardware fabric crc status** command to display the status of the internal CRC detection and isolation function.



```
switch# show hardware fabric crc status
Hardware Fabric CRC Action : log-only
Hardware Fabric CRC Feature threshold per module stage : 3
Hardware Fabric CRC Feature sampling time in hours : 24
```

## Displaying the System Uptime

The system uptime refers to the time that the chassis was powered on and has at least one supervisor module controlling the switch. Use the **reset** command to reinitialize the system uptime. On switches that use dual supervisors, nondisruptive upgrades and switchovers do not reinitialize the system uptime, which means that the system uptime is contiguous across such upgrades and switchovers.

The kernel uptime refers to the time since the NX-OS software was loaded on the supervisor module. Use the **reset** and **reload** commands to reinitialize the kernel uptime.

The active supervisor uptime refers to the time since the NX-OS software was loaded on the active supervisor module. The active supervisor uptime can be lower than the kernel uptime after nondisruptive switchovers.

You can use the **show system uptime** command to view the start time of the system, uptime of the kernel, and the active supervisor.

This example shows how to display the supervisor uptime:

```
switch# show system uptime
System start time: Fri Aug 27 09:00:02 2004
System uptime: 1546 days, 2 hours, 59 minutes, 9 seconds
Kernel uptime: 117 days, 1 hours, 22 minutes, 40 seconds
Active supervisor uptime: 117 days, 0 hours, 30 minutes, 32 seconds
```

For more information on high availability, see chapter 1, [High Availability Overview](#).

