CHAPTER 8

Configuring High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework.

This chapter includes the following sections:

About High Availability

The high availability (HA) software framework provides the following:

- Ensures nondisruptive software upgrade capability. See Chapter 6, “Software Images.”
- Provides redundancy for supervisor module failure by using dual supervisor modules.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in switches in the Cisco MDS 9100 Series and the Cisco MDS 9200 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series. See Chapter 13, “Configuring PortChannels.”

- Provides management redundancy using the Virtual Router Redundancy Protocol (VRRP). This feature is also available in switches in the Cisco MDS 9100 Series and in the Cisco MDS 9200 Series.

  See the "The Virtual Router Redundancy Protocol" section on page 36-16.

- Provides switchovers if the active supervisor fails, the standby supervisor, if present, takes over without disrupting storage or host traffic.

  Directors in the Cisco MDS 9500 Series have two supervisor modules in the two center slots (sup-1 and sup-2). When the switch powers up and both supervisor modules are present, the supervisor module that comes up first enters the active mode and the supervisor module that comes up second enters the standby mode. If both supervisor modules come up at the same time, sup-1 becomes active. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Switchover Mechanisms

If the active supervisor module fails, the standby module automatically takes over. You can 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 supervisor modules are not in a stable state (online or powered down), a switchover will not be performed.

HA Switchover Characteristics

An HA switchover has the following characteristics:

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

Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, issue the **system switchover** command. Once issued, another switchover process cannot be started on the same switch until a stable standby module is available.

To ensure that an HA switchover is possible, issue 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.

Switchover Guidelines

Be aware of the following guidelines when performing a 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.

Verifying Switchover Possibilities

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

- Use the s**how 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  Sw           Hw      World-Wide-Name(s) (WWN)
---  -----------  ------  --------------------------------------------------
2    1.3(0.106a)  0.206   20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a)  0.602   --
6    1.3(0.106a)) 0.602   --
8    1.3(0.106a)  0.702   --
9    1.3(0.106a)  0.3     22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00

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.

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This vital process functions on infrastructure that is internal to the switch.

See the "Displaying System Processes" section on page 50-1.

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.

See the "Replacing Modules" section on page 6-27.

Copying Boot Variable Images to the Standby Supervisor

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.

Automatic Copying of Boot Variables

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

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# boot auto-copy`	Enables automatic copying of boot variables from the active supervisor module to the standby supervisor module.
	`switch(config)# no boot auto-copy`	Disables the automatic copy feature (default).

Verifying the Copied Boot Variables

Use the **show boot auto-copy** command to verify the current state of the copied boot variables (see Example 8-1 and Example 8-2).

Example 8-1 Displays the auto-copy Option in an Enabled State

```
switch# show boot auto-copy
Boot variables Auto-Copy ON
```

Example 8-2 Displays the auto-copy Option in a Disabled State

```
switch# show boot auto-copy
Boot variables Auto-Copy OFF
```

Use the **show boot auto-copy list** command to verify what files are being copied. Example 8-3 displays the image being copied to the standby supervisor module's bootflash. Once this is successful, the next file will be image2.bin. This command only displays files on the active supervisor module.

Example 8-3 Displays the Files Being Copied

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

File:/bootflash:/image2.bin
Bootvar: system
```

Example 8-4 displays a typical message when the **auto-copy** option is disabled or if no files are copied.

Example 8-4 Displays the Current auto-copy State

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

Displaying HA Information

Use the **show system redundancy status** command to view the high availability status of the system (see Example 8-5). Tables 8-1 to 8-3 explain the possible output values for the redundancy, supervisor, and internal states.

Example 8-5 Displays Redundancy Status

```
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.

Table 8-1 lists the possible values for the redundancy states.

Table 8-1 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.
Offline	The supervisor module is intentionally shut down for debugging purposes.
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.

Table 8-2 lists the possible values for the supervisor module state.

Table 8-2 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.

Table 8-3 lists the possible values for the internal redundancy states.

Table 8-3 Internal States

State	Description
HA standby	The HA switchover mechanism in the standby supervisor module is enabled (see the "HA Switchover Characteristics" section on page 8-2).
Active with no standby	A switchover is possible.
Active with HA standby	The active supervisor module in the switch is ready to be configured. The standby 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.

Table 8-3 Internal States (continued)

State	Description
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.