



Configuring High Availability

This chapter provides details on the high availability feature that is available on switches with two supervisor modules. It includes the following sections:

- [About High Availability, page 4-2](#)
- [Switchover Mechanisms, page 4-3](#)
- [Configuring System Switchover, page 4-4](#)
- [Switchover Guidelines, page 4-4](#)
- [Process Restartability, page 4-5](#)
- [Synchronizing Supervisor Modules, page 4-5](#)
- [Displaying HA Information, page 4-6](#)
- [Default Settings, page 4-8](#)

Send documentation comments to mdsfeedback-doc@cisco.com

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

- Ensures nondisruptive software upgrade capability. See [Chapter 5, “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 Cisco MDS 9216 switches.
- Protects against link failure using the PortChannel (port aggregation) feature. See [Chapter 11, “Configuring PortChannels.”](#)
- Provides management redundancy using Virtual Routing Redundancy Protocol (VRRP). See the [“Configuring VRRP” section on page 16-12.](#)
- Switchability—When 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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Switchover Mechanisms

When the active supervisor module fails, the standby module automatically takes over. You can also issue a **system switchover** command to manually initiate a switchover from an active supervisor module to a standby supervisor module.

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

To determine version compatibility between switch images, use the **show system switchover impact** command.

```
switch# show system switchover impact
effects of switchover to 1.0(3)
switchover type:      HA
impact on modules:   module(s) will not reset

per module compatibility with system image 1.0(3) running on standby supervisor:
module  compatibility check
1       module running version 1.0(3): is compatible
9       module running version 1.0(3): is compatible
```

Two switchover modes are available in the Cisco MDS 9000 Family: HA switchover (default) or warm switchover.



Note

If the images are not compatible, an HA switchover is not possible.

HA Switchover

When a **show system redundancy status** or a **show module** command displays the `HA-standby` state for the standby supervisor module, an HA switchover (default) is possible. An HA switchover has the following characteristics:

- Is stateful (nondisruptive) since control traffic is not impacted
- Does not impact data traffic since the switching modules are not impacted
- Switching modules are not reset

This is the best possible scenario because there is no system downtime.

Warm Switchover

When a **show system redundancy status** or a **show module** command displays the `warm standby` state for the standby supervisor module, a warm switchover is possible. A warm switchover has the following characteristics:

- Is stateless (disruptive) since control traffic will be impacted.
- Impacts data traffic since switching modules will be impacted.
- Switching modules are reset with an significantly reduced bring up time

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Configuring System Switchover

By default, the system uses a HA switchover. When two supervisor modules are available on the system, you can switch over from the active to the standby supervisor module using a **HA** (nondisruptive) or **warm** (disruptive) switchover. In the HA switchover mode, a HA switchover is performed where possible. If HA switchover is not possible, the warm switchover mode is attempted. If warm switchover mode is configured, then HA switchover is disabled.



Caution

Switching from HA to warm or warm to HA modes cause the standby supervisor module to reset.

To define the switchover mechanism in a switch, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# system switchover warm switch(config)#	Configures the switch to perform a stateless (disruptive) switchover the next time a switchover occurs (EXEC mode command or in response to a failure).
	switch(config)# system switchover HA switch(config)#	Reverts the switch settings to perform the default stateful (nondisruptive) switchover the next time a switchover occurs (EXEC mode command or in response to a failure).
	or issue the following command: switch(config)# no system switchover switch(config)#	Restores the default settings (HA switchover).

Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- Use the **system switchover** command when you need to upgrade the software in a dual supervisor switch (see the “[Performing a System Switchover](#)” section on page 5-19).
- The **system switchover** command returns the following message when the standby supervisor is not present in the switch:

```
switch# system switchover
Failed to switchover: (supervisor has no standby)
```

- You can only perform a switchover when the switch has two supervisor modules functioning in the switch. Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.
- Verify that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    16     1/2 Gbps FC Module        DS-X9016             ok
5     0     Supervisor/Fabric-1       DS-X9530-SF1-K9     active *
6     0     Supervisor/Fabric-1       DS-X9530-SF1-K9     HA-standby
8    32     1/2 Gbps FC Module        DS-X9032             ok
```

Send documentation comments to mdsfeedback-doc@cisco.com

```

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.0(0.253)    1.0         20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de
5    1.0(0.253)    1.0         --
6    1.0(0.253)    1.0         --
8    1.0(0.253)    1.0         20:41:00:05:30:00:38:de to 20:50:00:05:30:00:38:de

```

```

Mod  MAC-Address(es)          Serial-Num
---  -
2    00-05-30-00-0f-e4 to 00-05-30-00-0f-e8    jab0636063v
5    00-05-30-00-19-66 to 00-05-30-00-19-6a    jab06370593
6    00-05-30-02-20-02 to 00-05-30-02-20-06    jab06371593
8    00-05-30-00-1a-12 to 00-05-30-00-1a-16    jab06370574

```

* this terminal session

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



Note A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled. If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that the 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 [“Displaying System Processes” section on page 26-2](#).

Synchronizing Supervisor Modules

The **system auto-sync image** option is disabled by default on switches in the Cisco MDS 9000 Series. This command can only be operational if the following cases apply:

- the **system switchover HA** command is configured.
- two supervisor modules are up and running

You can synchronize the standby supervisor module software image with the bootflash image using the **system auto-sync image** command in configuration mode. The current running image and configuration files are synchronized from the active to the standby supervisor module (see the [“Specifying Kickstart and System Boot Variables” section on page 5-18](#)).



Note If both supervisors modules are running the same software image, the **system auto-sync image** command will have no effect.

Send documentation comments to mdsfeedback-doc@cisco.com

To enable or disable automatic synchronization, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# system auto-sync image switch(config)#	Enables automatic synchronization.
	switch(config)# no system auto-sync image Automatic synchronization of BOOT and KICKSTART is now disabled.	Disables automatic synchronization (default).

When you log into the switch after the basic upgrade, the standby supervisor module synchronizes its image automatically with the running image on the active supervisor module. To upgrade the image, you must disable this option. By disabling this option, you are ensuring that the synchronization does not take place with undesired images. Enabling this option synchronizes the running image on both supervisor modules.



Note

During a synchronization, the boot variables are not synchronized. The boot variables are independent of the two supervisor modules (see “Performing a System Switchover” section on page 5-19).

Use the **show auto-sync** command to view the status of the auto-sync configuration. See [Example 4-1](#).

Example 4-1 Displays Auto Synchronization Status

```
switch# show system auto-sync
auto-sync is disabled
auto-sync not started
```

You can view the output of the **show system redundancy** command to verify if HA switchover and automatic synchronization are enabled and operational.

Displaying HA Information

Use the **show system redundancy status** command to view the high availability status of the system. See [Example 4-2](#). Tables 4-1 to 4-3 explain the possible redundancy, supervisor, and internal states output in this command.

Example 4-2 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

The following conditions identify when automatic synchronization is possible:

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

Table 4-1 lists the possible values for the redundancy states.

Table 4-1 Redundancy States

State	Description
Not present	The supervisor module is not present or is not plugged into the switch.
Initializing	The diagnostics have passed and the configuration is being downloaded.
Active	This module is the active supervisor module and the switch is ready to be configured.
Standby	This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3).
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 switch is intentionally shut down for debugging purposes.
At BIOS	The module 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 4-2 lists the possible values for the Supervisor state.

Table 4-2 Supervisor States

State	Description
Active	This module is the active supervisor module and the switch is ready to be configured.
HA standby	This module is the standby supervisor module and the HA switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3).
Warm standby	This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “ HA Switchover ” section on page 4-3).
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.

Send documentation comments to mdsfeedback-doc@cisco.com

Table 4-3 lists the possible values for the internal state.

Table 4-3 Internal States

State	Description
Warm standby	This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “HA Switchover” section on page 4-3).
HA standby	This module is the standby supervisor module and the HA switchover mechanism is enabled (see the “HA Switchover” section on page 4-3).
Active with no standby	This module is the active supervisor module, and the second supervisor module is not present in the switch.
Active with HA standby	This module is the active supervisor module and the switch is ready to be configured. The standby module is in the HA-standby state.
Active with warm standby	This module is the standby supervisor module and the warm switchover mechanism is enabled (see the “HA Switchover” section on page 4-3).
Shutting down	The switch is being shut down.
Warm switchover in progress	The switch is in the process of changing over to the warm switchover mechanism.
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 supervisor modules.
Standby (failed)	The standby supervisor module is not functioning.
Active with failed standby	This module is 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.

Default Settings

Table 4-4 lists the default settings for high availability features.

Table 4-4 Default High Availability Setting

Parameters	Default
Switchover mode	HA