



System-Level High Availability

This chapter describes the Cisco NX-OS high availability (HA) system and application restart operations and includes the following sections:

- [About Cisco NX-OS System-Level High Availability, on page 1](#)
- [Physical Redundancy, on page 2](#)
- [Supervisor Restarts and Switchovers, on page 7](#)
- [Displaying HA Status Information, on page 13](#)
- [Additional References for System-Level High Availability, on page 15](#)

About Cisco NX-OS System-Level High Availability

Cisco NX-OS system-level HA mitigates the impact of hardware or software failures and is supported by the following features:

- Redundant hardware components:
 - Power supply
 - Fan tray (Cisco Nexus 9500 platform only) or modules (Cisco Nexus 9200/9300/9300-EX/9300-FX)
 - Switch fabric (Cisco Nexus 9504, 9508, and 9516 chassis only)
 - System controller (Cisco Nexus 9504, 9508, and 9516 chassis only)
 - Supervisor (Cisco Nexus 9504, 9508, and 9516 chassis only)

For details about physical requirements and redundant hardware components, see the [Hardware Installation Guide](#) for your specific Cisco Nexus 9000 Series chassis.

- HA software features:
 - Nonstop forwarding (NSF)—For details about nonstop forwarding, also known as graceful restart, see the [Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide](#).
 - Embedded Event Manager (EEM)—For details about configuring EEM, see the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#).
 - Smart Call Home—For details about configuring Smart Call Home, see the [Cisco Nexus 9000 Series NX-OS System Management Configuration Guide](#).

Physical Redundancy

The Cisco Nexus 9504, 9508, and 9516 chassis include the following physical redundancies:

- Power supply
- Fan tray
- Switch fabric
- System controller
- Supervisor module

The Cisco Nexus 9300 platform switches include the following physical redundancies:

- Power supply
- Fan tray

For additional details about physical redundancies, see the [Hardware Installation Guide](#) for your specific Cisco Nexus 9000 Series chassis.

Power Supply Redundancy

The Cisco Nexus N9K-C9348GC-FXP chassis supports up to two AC power supply modules (each delivering up to 350 W) or two DC power supplies (each delivering up to 350 W).

The Cisco Nexus N9K-C92348GC-X chassis supports up to two AC power supply modules (each delivering up to 400 W) or two DC power supplies (each delivering up to 400 W). It also supports two AC and DC power supply modules (each delivering up to 500 W).

The Cisco Nexus N9K-C93360YC-FX2 chassis supports up to two AC power supply modules (each delivering up to 1200 W) or two DC power supplies (each delivering up to 930 W).

Cisco Nexus N9K-C9316D-GX and Cisco Nexus N9K-C93600CD-GX chassis supports up to two AC power supply modules (each delivering up to 1100 W) or two DC power supplies (each delivering up to 1100 W).

Cisco Nexus N9K-C9364C-GX chassis supports up to two AC power supply modules (each delivering up to 2KW) or two DC power supplies (each delivering up to 2KW).

The Cisco Nexus 9504 chassis supports up to four power supply modules, the Cisco Nexus 9508 chassis supports up to eight power supply modules, and the Cisco Nexus 9516 chassis supports up to ten power supply modules. Each 9500 platform power supply module can deliver up to 3 kW.

The power subsystem allows the power supplies to be configured in one of the available redundancy modes. By installing more modules, you can ensure that the failure of one module does not disrupt system operations. You can replace the failed module while the system is operating. For information on power supply module installation and replacement, see the *Hardware Installation Guide* for your specific Cisco Nexus 9000 Series chassis.

Power Modes

Each of the power redundancy modes imposes different power budgeting and allocation models, which deliver varying usable power yields and capacities. For more information about power budgeting, usable capacity,

planning requirements, and redundancy configuration, see the [Hardware Installation Guide](#) for your specific Cisco Nexus 9000 Series chassis.

The available power supply redundancy modes are described in the following table.

Table 1: Power Redundancy Modes

Redundancy Mode	Description
Combined (nonredundant)	This mode does not provide power redundancy. The available power is the total power capacity of all power supplies.
insrc-redundant (grid redundancy)	<p>This mode provides grid redundancy when you connect half of the power supplies to one grid and the other half of the power supplies to the second grid. The available power is the amount of power available through a grid.</p> <p>To enable grid redundancy, you must connect the power supplies to the correct power grid slots. For example, on the Cisco Nexus 9508 switch, slots 1, 2, 3, and 4 are in grid A, and slots 5, 6, 7, and 8 are in grid B. To configure and operate in grid redundancy mode, you must connect half of your power supplies to the slots in grid A and the rest of your power supplies to the slots in grid B. For more information on power grid slot assignments for your power supplies, see the Hardware Installation Guide for your specific Cisco Nexus 9000 Series platform.</p>
ps-redundant (N+1 redundancy)	This mode provides an extra power supply if an active power supply goes down. One power supply of all the available power supplies is considered an extra power supply, and the total available power is the amount provided by the active power supply units.

Use the **power redundancy-mode {combined | insrc_redundant | ps-redundant}** command to specify one of these power modes.

Fan Tray Redundancy

The Cisco Nexus 9000 Series switches contain redundant system fan trays for cooling the system. For the number of supported fan trays per chassis, see [Physical Redundancy](#).

The fan speeds are variable and depend on the temperature of the ASICs in the system. If fans are removed or go bad, the other fan modules can start running at a higher speed to compensate for the missing or failed fans. If the system temperature increases above the thresholds, the system shuts down.

- If a single fan fails within a fan tray, the fan speed of the other fans in the tray does not increase.
- If multiple fans fail within a fan tray, the fan speed increases to 100% on all the fan trays.
- If an entire fan tray is removed, the fan speed for the other two fan trays increases to 100% as soon as the tray is removed.

- If multiple fan trays are removed and not replaced within 2 minutes, the device will shut down. The switch can be recovered by power cycle. When the device comes back, if it still detects the multiple fan tray failure, it will shut down again after 2 minutes. If required, you can use EEM to overwrite this policy.
- If a fan tray fails, leave the failed unit in place to ensure proper airflow until you can replace it. The fan trays are hot swappable, but you must replace one fan tray at a time. Otherwise, the device reboots after 2 minutes if multiple fan trays are missing.



Note There is no time limit for replacing a single fan tray, but to ensure proper airflow, replace the fan tray as soon as possible.

Switch Fabric Redundancy

Cisco NX-OS provides switching fabric availability through redundant switch fabric module implementation. You can configure a single Cisco Nexus 9504, 9508, or 9516 chassis with one to six switch fabric modules for capacity and redundancy. Each line card installed in the system automatically connects to and uses all functionality of the installed switch fabric modules. A failure of a switch fabric module triggers an automatic reallocation and balancing of traffic across the remaining active switch fabric modules. Replacing the failed fabric module reverses this process. After you insert the replacement fabric module and bring it online, traffic is again redistributed across all installed fabric modules and redundancy is restored.

Fabric modules are hot swappable. Hot swapping can temporarily disrupt traffic. To prevent the disruption of traffic when you hot-swap fabric modules, use the **poweroff module slot-number** command before you remove a fabric module and the **no poweroff module slot-number** command after you reinsert the fabric module.

X9400 line cards: To achieve the maximum bandwidth allowed per card requires four fabric modules (N9K-C95xx-FM-S for the N9K-X9432C-S or N9K-C95xx-FM for the other X9400 line cards) in four fabric module slots (FM2, FM3, FM4, and FM6). Additional fabric modules will not provide additional redundancy for these line cards.

X9500 line cards: To achieve the maximum bandwidth allowed per card requires three fabric modules (N9K-C95xx-FM) in the even fabric module slots (FM2, FM4, and FM6). Additional fabric modules will provide additional redundancy for these line cards. Each even fabric module provides redundancy for each odd fabric module failure (FM2 provides redundancy for FM1, FM4 provides redundancy for FM3, and FM6 provides redundancy for FM5).

X9600 line cards: The maximum bandwidth allowed per card requires six fabric modules (N9K-C95xx-FM).

X9600-R line cards: The maximum bandwidth allowed per N9K-X9636C-R requires five fabric modules (N9K-C95xx-FM-R), and the maximum bandwidth allowed per N9K-X9636Q-R requires four fabric modules (N9K-C95xx-FM-R). Additional fabric modules will provide additional redundancy for these line cards. The maximum bandwidth allowed per N9K-X96136YC-R line card requires six N9K-C9504-FM-R fabric modules for redundancy. The maximum bandwidth allowed per N9K-X9636C-R (P-100) requires 5 fabric modules. The maximum bandwidth allowed per N9K-X9636-RX requires 6 fabric modules for redundancy.

X9700-EX line cards: The maximum bandwidth allowed per card requires four fabric modules (N9K-C95xx-FM-E) in four fabric module slots (FM2, FM3, FM4, and FM6). Additional fabric modules will not provide additional redundancy for these line cards.

X9700-FX line cards: The maximum bandwidth allowed per N9K-X9788TC-FX requires 2 fabric modules. N9K-X9732C-FX requires 4 fabric modules (N9K-C9508-FM-E2 and N9K-C9516-FM-E2) for maximum

bandwidth. N9K-X9732C-FX is redundant with 5 fabric modules (the fabric modules can be either 95xx-FM-E or 95xx-FM-E2). N9K-X9736C-FX requires 5 fabric modules for maximum bandwidth. N9K-X9736C-FX is not redundant with additional fabric modules (the fabric modules can be either 95xx-FM-E or 95xx-FM-E2). Fabric Module 25 is the fifth fabric module for both N9K-X9732C-FX and N9K-X9736C-FX. FM25 will be powered down if there are any EX linecards in the chassis.



Note To achieve Fabric module redundancy, N9K-9732C-FX should not be mixed with any older modules. If any other module is detected, Fabric Module 25 is powered down by the system.

Line card and Fabric module failures

To keep line cards and Fabric module powered down whenever they fail or crash, use the **system module failure-action shutdown** command to prevent the cards from rebooting. This command is useful if your topology is configured for network-level redundancy and you want to prevent a second disruption from occurring in the network because a line card or fabric module is trying to come up.

You can use the **show module module** command to verify that the line card has been powered down. If desired, use the **no poweroff module module** command to manually bring the module (fabric or line card) back up.

```
switch(config)# system module failure-action shutdown
2014 Sep 8 23:31:51 switch %$ VDC-1 %$ %SYSMGR-SLOT1-2-SERVICE_CRASHED:
Service "ipfib" (PID 2558) hasn't caught signal 11 (core will be saved).

2014 Sep 8 23:32:25 switch %$ VDC-1 %$ %PLATFORM-2-MOD_PWRDN:
Module 1 powered down (Serial number SAL1815Q1DP)

switch(config)# show module 1
Mod  Ports  Module-Type                               Model                Status
---  ---
1     52      48x1/10G-T 4x40G Ethernet Module         N9K-X9564TX          powered-dn

switch(config)# no poweroff module 1
2014 Sep 8 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-PFM_MODULE_POWER_ON:
Manual power-on of Module 1 from Command Line Interface

2014 Sep 8 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-MOD_DETECT:
Module 1 detected (Serial number SAL1815Q1DP) Module-Type 48x1/10G-T
4x40G Ethernet Module Model N9K-X9564TX

2014 Sep 8 23:34:31 switch %$ VDC-1 %$ %PLATFORM-2-MOD_PWRUP:
Module 1 powered up (Serial number SAL1815Q1DP)
```

System Controller Redundancy

Two redundant system controllers in the Cisco Nexus 9504, 9508, and 9516 chassis offload chassis management functions from the supervisor modules. The controllers are responsible for managing power supplies and fan trays and act as a central point for the Gigabit Ethernet Out-of-Band Channel (EOBC) between the supervisors, fabric modules, and line cards.

Supervisor Module Redundancy

The Cisco Nexus 9504, 9508, and 9516 chassis support dual supervisor modules to provide 1+1 redundancy for the control and management plane. A dual supervisor configuration operates in an active or standby capacity in which only one of the supervisor modules is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two supervisor modules to provide stateful switchover in the event of a supervisor module failure.

A Cisco NX-OS generic online diagnostics (GOLD) subsystem and additional monitoring processes on the supervisor trigger a stateful failover to the redundant supervisor when the processes detect unrecoverable critical failures, service restartability errors, kernel errors, or hardware failures.

If a supervisor-level unrecoverable failure occurs, the currently active failed supervisor triggers a switchover. The standby supervisor becomes the new active supervisor and uses the synchronized state and configuration while the failed supervisor is reloaded. If the failed supervisor is able to reload and pass self-diagnostics, it initializes, becomes the new standby supervisor, and then synchronizes its operating state with the newly active unit.

Supervisor Modules

Two supervisor modules are available for the Cisco Nexus 9500 Series switches: Supervisor A (SUP A) and Supervisor B (SUP B). The following table lists the differences between the two modules.

	Supervisor A	Supervisor B	Supervisor A+	Supervisor B+
CPU	4 core, 1.8 GHz	6 core, 2.1 GHz	4 core, 1.8 GHz	6 core, 1.9 GHz
Memory	16 GB	24 GB	16 GB	32 GB
SSD storage	64 GB	256 GB	256 GB	256 GB
Software release	6.1(2)I1(1) or later release	6.1(2)I3(1) or later release	7.0(3)I7(1)	7.0(3)I7(1)

SUP A and SUP B are not compatible and should not be installed in the same chassis, except for migration purposes. For dual supervisor systems, you should install either two SUP A modules or two SUP B modules (and not a combination of the two) to ensure supervisor module redundancy.

In a dual supervisor system, Cisco NX-OS checks the memory size of both the active and standby supervisors. If the memory size is different for each supervisor (because both SUP A and SUP B are installed), a message appears instructing you to replace SUP A with a second SUP B.

To migrate from SUP A to SUP B, insert SUP B into the device and enter the **system switchover** command. SUP B becomes the active supervisor, and SUP A becomes the standby supervisor, which is not a supported configuration. A warning message appears every hour until you remove SUP A or replace it with a second SUP B.

Supervisor Restarts and Switchovers

Restarts on Single Supervisors

In a system with only one supervisor, when all HA policies have been unsuccessful in restarting a service, the supervisor restarts. The supervisor and all services reset and start with no prior state information.

Restarts on Dual Supervisors

When a supervisor-level failure occurs in a system with dual supervisors, the System Manager performs a switchover rather than a restart to maintain stateful operation. In some cases, however, a switchover might not be possible at the time of the failure. For example, if the standby supervisor module is not in a stable standby state, a restart rather than a switchover is performed.

Switchovers on Dual Supervisors

A dual supervisor configuration allows nonstop forwarding (NSF) with a stateful switchover (SSO) when a supervisor-level failure occurs. The two supervisors operate in an active/standby capacity in which only one of the supervisor modules is active at any given time, while the other acts as a standby backup. The two supervisors constantly synchronize the state and configuration in order to provide a seamless and stateful switchover of most services if the active supervisor module fails.

Switchover Characteristics

An HA switchover has the following characteristics:

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

Switchover Mechanisms

Switchovers occur by one of the following two mechanisms:

- 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.

When a switchover process begins, another switchover process cannot be started on the same switch until a stable standby supervisor module is available.

Switchover Failures

Supervisor switchovers are generally hitless and occur without traffic loss. If for some reason a switchover does not complete successfully, the supervisors reset. A reset prevents loops in the Layer 2 network if the network topology was changed during the switchover. For optimal performance of this recovery function, we recommend that you do not change the Spanning Tree Protocol (STP) default timers.

If three system-initiated switchovers occur within 20 minutes, all nonsupervisor modules shut down to prevent switchover cycling. The supervisors remain operational to allow you to collect system logs before resetting the switch.

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 run this command, you cannot start another switchover process on the same system until a stable standby supervisor module is available.



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

To ensure that an HA switchover is possible, use 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, you can manually initiate a switchover.

Switchover Guidelines

Follow these guidelines when performing a switchover:

- When you manually initiate a switchover, it takes place immediately.
- A switchover can be performed only when two supervisor modules are functioning in the switch.
- The modules in the chassis must be functioning.

Verifying Switchover Possibilities

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

- 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
---  -
1    32     32p 40G Ethernet Module    N9K-X9432PQ        ok
2    52     48x1/10G SFP+ 4x40G Ethernet Module N9K-X9564PX        ok
5    52     48x1/10G SFP+ 4x40G Ethernet Module N9K-X9464PX        ok
6    36     36p 40G Ethernet Module    N9K-X9536PQ        ok
7    36     36p 40G Ethernet Module    N9K-X9536PQ        ok
10   32     32p 40G Ethernet Module    N9K-X9432PQ        ok
11   52     48x1/10G-T 4x40G Ethernet Module N9K-X9564TX        ok
12   52     48x1/10G-T 4x40G Ethernet Module N9K-X9464TX        ok
15   52     48x1/10G SFP+ 4x40G Ethernet Module N9K-X9464PX        ok
21   0      Fabric Module              N9K-C9516-FM       ok
22   0      Fabric Module              N9K-C9516-FM       ok
23   0      Fabric Module              N9K-C9516-FM       ok
24   0      Fabric Module              N9K-C9516-FM       ok
25   0      Fabric Module              N9K-C9516-FM       ok
26   0      Fabric Module              N9K-C9516-FM       ok
27   0      Supervisor Module          N9K-SUP-A           ha-standby
```

28	0	Supervisor Module	N9K-SUP-A	active *
29	0	System Controller	N9K-SC-A	active
30	0	System Controller	N9K-SC-A	standby

Mod	Sw	Hw	Slot
---	-----	-----	----
1	6.1(2) I3(1)	0.1050	LC1
2	6.1(2) I3(1)	0.2010	LC2
5	6.1(2) I3(1)	0.1010	LC5
6	6.1(2) I3(1)	0.2060	LC6
7	6.1(2) I3(1)	0.2060	LC7
10	6.1(2) I3(1)	0.1010	LC10
11	6.1(2) I3(1)	0.2100	LC11
12	6.1(2) I3(1)	0.1010	LC12
15	6.1(2) I3(1)	0.1050	LC15
21	6.1(2) I3(1)	0.3010	FM1
22	6.1(2) I3(1)	0.3040	FM2
23	6.1(2) I3(1)	0.3040	FM3
24	6.1(2) I3(1)	0.3040	FM4
25	6.1(2) I3(1)	0.3010	FM5
26	6.1(2) I3(1)	0.3040	FM6
27	6.1(2) I3(1)	1.1	SUP1
28	6.1(2) I3(1)	1.1	SUP2
29	6.1(2) I3(1)	1.2	SC1
30	6.1(2) I3(1)	1.2	SC2

Mod	MAC-Address (es)	Serial-Num
---	-----	-----
1	74-26-ac-10-cb-0c to 74-26-ac-10-cb-9f	SAL1817REX2
2	00-22-bd-fd-93-57 to 00-22-bd-fd-93-9a	SAL1733B92R
5	74-26-ac-eb-99-0c to 74-26-ac-eb-99-4f	SAL1814PTNM
6	c0-8c-60-62-60-98 to c0-8c-60-62-61-2b	SAL1812NTG1
7	c0-8c-60-62-5f-70 to c0-8c-60-62-60-03	SAL1812NTFD
10	74-26-ac-e9-32-68 to 74-26-ac-e9-32-fb	SAL1811NH4K
11	78-da-6e-74-15-14 to 78-da-6e-74-15-57	SAL1746G7XE
12	74-26-ac-ec-2b-50 to 74-26-ac-ec-2b-93	SAL1816QUQX
15	c0-8c-60-62-a3-b4 to c0-8c-60-62-a3-f7	SAL1816QGXE
21	NA	SAL1801K507
22	NA	SAL1813P9Y2
23	NA	SAL1813P9YM
24	NA	SAL1813P9Y9
25	NA	SAL1801K50F
26	NA	SAL1813NZN3
27	c0-67-af-a1-0e-d6 to c0-67-af-a1-0e-e7	SAL1803KWXY
28	c0-67-af-a1-0d-a4 to c0-67-af-a1-0d-b5	SAL1804L578
29	NA	SAL1801JU2Z
30	NA	SAL1801JU4V

Mod	Online Diag Status
---	-----
1	Pass
2	Pass
5	Pass
6	Pass
7	Pass
10	Pass
11	Pass
12	Pass
15	Pass
21	Pass
22	Pass
23	Pass
24	Pass
25	Pass

```

26    Pass
27    Pass
28    Pass
29    Pass
30    Pass

```

```
* 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.

- 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 are as follows:

```

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

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

```

Replacing the Active Supervisor Module in a Dual Supervisor System

You can nondisruptively replace the active supervisor module in a dual supervisor system.

SUMMARY STEPS

1. switch # **system switchover**
2. switch# **reload module slot-number force**
3. switch# **copy bootflash:nx-os-image bootflash:nx-os-image**
4. switch# **configure terminal**
5. switch (config)# **boot nxos bootflash:nx-os-image [sup-number]**
6. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	switch # system switchover	Initiates a manual switchover to the standby supervisor. Note Wait until the switchover completes and the standby supervisor becomes active.
Step 2	switch# reload module slot-number force	Boots the supervisor module replacement immediately.

	Command or Action	Purpose
		Note If you do not force the boot, the replacement supervisor module should be booted by the active supervisor module 6 minutes after insertion. For information on replacing a supervisor module, see the Hardware Installation Guide for your specific Cisco Nexus 9000 Series chassis.
Step 3	switch# copy bootflash:nx-os-image bootflash:nx-os-image	Copies the nx-os image from the active supervisor module to the standby supervisor module.
Step 4	switch# configure terminal	Enters global configuration mode.
Step 5	switch (config)# boot nxos bootflash:nx-os-image [sup-number]	Configures the standby supervisor boot variables.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to replace the active supervisor module in a dual supervisor system:

```
switch# system switchover
Raw time read from Hardware Clock: Y=2013 M=2 D=2 07:35:48
writing reset reason 7,

NX9 SUP Ver 3.17.0
Serial Port Parameters from CMOS
PMCON_1: 0x200
PMCON_2: 0x0
PMCON_3: 0x3a
PM1_STS: 0x1
Performing Memory Detection and Testing
Testing 1 DRAM Patterns
Total mem found : 4096 MB
Memory test complete.
NumCpus = 2.
Status 61: PCI DEVICES Enumeration Started
Status 62: PCI DEVICES Enumeration Ended
Status 9F: Dispatching Drivers
Status 9E: IOFPGA Found
Status 9A: Booting From Primary ROM
Status 98: Found Cisco IDE
Status 98: Found Cisco IDE
Status 90: Loading Boot Loader
Reset Reason Registers: 0x1 0x10
Filesystem type is ext2fs, partition type 0x83
Filesystem type is ext2fs, partition type 0x83

GNU GRUB version 0.97

Loader Version 3.17.0
```

```

current standby sup
-----
switch(standby)# 2014 Aug  2 07:35:46 switch %$ VDC-1 %$ %KERN-2-SYSTEM_MSG: Switchover
started by redundancy driver - kernel
2014 Aug  2 07:35:47 switch %$ VDC-1 %$ %SYSMGR-2-HASWITCHOVER_PRE_START: This supervisor
is becoming active (pre-start phase).
2014 Aug  2 07:35:47 switch %$ VDC-1 %$ %SYSMGR-2-HASWITCHOVER_START: This supervisor is
becoming active.
2014 Aug  2 07:35:48 switch %$ VDC-1 %$ %SYSMGR-2-SWITCHOVER_OVER: Switchover completed.

switch# reload module 27 force
switch# copy bootflash:n9000-dk9.6.1.2.I3.1.bin bootflash:n9000-dk9.6.1.2.I3.1.bin
switch# config terminal
switch# boot nxos bootflash:n9000-dk9.6.1.2.I3.1.bin sup-1
switch# copy running-config startup-config

```

Replacing the Standby Supervisor Module in a Dual Supervisor System

You can nondisruptively replace the standby supervisor module in a dual supervisor system.

SUMMARY STEPS

1. switch# **reload module** *slot-number* **force**
2. switch# **copy bootflash:***nx-os-image* **bootflash:***nx-os-image*
3. switch# **configure terminal**
4. switch (config)# **boot nxos bootflash:***nx-os-image* [*sup-number*]
5. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	switch# reload module <i>slot-number</i> force	Boots the supervisor module replacement immediately. Note If you do not force the boot, the replacement supervisor module should be booted by the active supervisor module 6 minutes after insertion. For information on replacing a supervisor module, see the Hardware Installation Guide for your specific Cisco Nexus 9000 Series chassis.
Step 2	switch# copy bootflash: <i>nx-os-image</i> bootflash: <i>nx-os-image</i>	Copies the nx-os image from the active supervisor module to the standby supervisor module.
Step 3	switch# configure terminal	Enters global configuration mode.
Step 4	switch (config)# boot nxos bootflash: <i>nx-os-image</i> [<i>sup-number</i>]	Configures the standby supervisor boot variables.

	Command or Action	Purpose
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to replace the standby supervisor module in a dual supervisor system:

```
switch# reload module 27 force
switch# copy bootflash:n9000-dk9.6.1.2.I3.1.bin bootflash:n9000-dk9.6.1.2.I3.1.bin
switch# config terminal
switch# boot nxos bootflash:n9000-dk9.6.1.2.I3.1.bin sup-1
switch# copy running-config startup-config
```

Displaying HA Status Information

Use the **show system redundancy status** command to view the HA status of the system.

```
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 system is operationally HA and can perform automatic synchronization.
- If the internal state of one of the supervisor modules is none, the system cannot perform automatic synchronization.

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

Table 2: 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.

State	Description
Active	The active supervisor module and the switch are ready to be configured.
Standby	A switchover is possible.
Failed	The system detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three 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 system has established a connection with the supervisor, and the supervisor module is performing diagnostics.
Unknown	The system is in an invalid state. If it persists, call TAC.

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

Table 3: 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 system is intentionally shut down for debugging purposes.
Unknown	The system 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 4: Internal States

State	Description
HA standby	The HA switchover mechanism in the standby supervisor module is enabled.
Active with no standby	A switchover is impossible.
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 system is being shut down.
HA switchover in progress	The system is in the process of entering the active state.
Offline	The system 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.

State	Description
Active with failed standby	The active supervisor module and the second supervisor module are present, but the second supervisor module is not functioning.
Other	The system is in a transient state. If it persists, call TAC.

Additional References for System-Level High Availability

This section describes additional information related to system-level high availability.

Related Documents

Related Topic	Document Title
Hardware	Cisco Nexus 9000 Series Switch Hardware Installation Guide
Power mode configuration and Cisco NX-OS fundamentals	Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide
Nonstop forwarding (NSF)	Cisco Nexus 9000 Series NX-OS Unicast Routing Configuration Guide
EEM and Smart Call Home	Cisco Nexus 9000 Series NX-OS System Management Configuration Guide
Licensing	Cisco NX-OS Licensing Guide

MIBs

MIBs	MIBs Link
MIBs related to system-level high availability	For more information about MIBs and to download MIBs, refer to Cisco Nexus 7000 Series and 9000 Series MIB Quick Reference .

