



# Configuring System-Level High Availability

---

This chapter contains the following sections:

- [Information About System-Level High Availability, page 1](#)
- [Information About VSM Restarts and Switchovers, page 4](#)
- [Guidelines and Limitations, page 5](#)
- [Configuring System-Level High Availability, page 5](#)
- [Adding a Second VSM to a Standalone System, page 9](#)
- [Replacing the Secondary VSM in a Dual VSM System, page 11](#)
- [Replacing the Primary VSM in a Dual VSM System, page 11](#)
- [Changing the Domain ID in a Dual VSM System, page 12](#)
- [Verifying the HA Status, page 13](#)
- [MIBs, page 14](#)
- [Feature History for System-Level High Availability, page 14](#)

## Information About System-Level High Availability

### Information About Single and Dual Supervisor Roles

The Cisco Nexus 1000V can be configured with a single Virtual Supervisor Module (VSM) or dual VSMs. The following table describes the HA supervisor roles for single and dual VSM operation.

Single VSM Operation	Dual VSM Operation
<ul style="list-style-type: none"> <li>• Stateless—In case of failure, service restarts from the startup configuration.</li> <li>• Stateful—In case of failure, service resumes from previous state.</li> </ul>	<ul style="list-style-type: none"> <li>• Redundancy is provided by one active VSM and one standby VSM.</li> <li>• The active VSM runs all the system applications and controls the system.</li> <li>• On the standby VSM, the applications are started and initialized in standby mode. The applications are synchronized and kept up to date with the active VSM in order to be ready to run.</li> <li>• On a switchover, the standby VSM takes over for the active VSM.</li> <li>• The control interface of the VSMs are used to pass heartbeats between the two VSMs.</li> <li>• The management interface is used to prevent split-brain scenarios.</li> </ul>

## HA Supervisor Roles

The redundancy role indicates not only whether the VSM interacts with other VSMs, but also the module number it occupies. The following table shows the available HA roles for VSMs.

role	Module Number	Description
Standalone	1	<ul style="list-style-type: none"> <li>• This role does not interact with other VSMs.</li> <li>• You assign this role when there is only one VSM in the system.</li> <li>• This role is the default.</li> </ul>
Primary	1	<ul style="list-style-type: none"> <li>• This role coordinates the active/standby state with the secondary VSM.</li> <li>• This role takes precedence during bootup when negotiating active/standby mode. That is, if the secondary VSM does not have the active role at bootup, the primary VSM takes the active role.</li> <li>• You assign this role to the first VSM that you install in a dual VSM system.</li> </ul>

role	Module Number	Description
Secondary	2	<ul style="list-style-type: none"> <li>This role coordinates the active/standby state with the primary VSM.</li> <li>You assign this role to the second VSM that you install in a dual VSM system.</li> </ul>

### Dual Supervisor Active and Standby Redundancy States

Independent of its role, the redundancy state of a VSM can be one of the following described in this table.

Redundancy State	Description
Active	Controls the system and is visible to the outside world.
Standby	<p>Synchronizes its configuration with that of the active VSM so that it is continuously ready to take over in case of a failure or manual switchover.</p> <p>You cannot use Telnet or Secure Shell (SSH) protocols to communicate with the standby VSM. Instead, you can use the <b>attach module</b> command from the active VSM to access the standby VSM console. Only a subset of the CLI commands are available from the standby VSM console.</p>

### Dual Supervisor Synchronization

The active and standby VSMs are in the operationally HA state and can automatically synchronize when the internal state of one supervisor module is Active with HA Standby and the internal state of the other supervisor module is HA Standby.

If the output of the **show system redundancy** command indicates that the operational redundancy mode of the active VSM is None, the active and standby VSMs are not yet synchronized. This example shows the VSM internal state of dual supervisors as observed in the output of the **show system redundancy status** command:

```

switch# show system redundancy status
Redundancy role
-----
      administrative:  standalone
      operational:    standalone

Redundancy mode
-----
      administrative:  HA
      operational:    None

This supervisor (sup-1)
-----
Redundancy state:    Active
Supervisor state:   Active
Internal state:     Active with no standby

Other supervisor (sup-2)

```

```
-----  
Redundancy state: Not present  
switch#
```

# Information About VSM Restarts and Switchovers

## Restarts on Standalone VSMs

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 restart with no prior state information.

## Restarts on Dual VSMs

When a VSM fails in a system with dual supervisors, the system performs a switchover rather than a system restart in order to maintain a stateful operation. In some cases, a switchover might not be possible at the time of the failure. For example, if the standby VSM is not in a stable standby state, a restart rather than a switchover is performed.

## Switchovers on Dual VSMs

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

### Switchover Characteristics

A switchover occurs when the active supervisor fails (for example, if repeated failures occur in an essential service or if the system that is hosting the VSM fails).

A user-triggered switchover could occur (for example, if you need to perform maintenance tasks on the system hosting the active VSM).

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because the control traffic is not affected.
- It does not disrupt data traffic because the VEMs are not affected.

### Automatic Switchovers

When a stable standby VSM detects that the active VSM has failed, it initiates a switchover and transitions to active. When a switchover begins, another switchover cannot be started until a stable standby VSM is available.

If a standby VSM that is not stable detects that the active VSM has failed, then, instead of initiating a switchover, it tries to restart the system.

## Manual Switchovers

Before you can initiate a manual switchover from the active to the standby VSM, the standby VSM must be stable.

Once you have verified that the standby VSM is stable, you can manually initiate a switchover.

Once a switchover process begins, another switchover process cannot be started until a stable standby VSM is available.

## Guidelines and Limitations

- Although primary and secondary VSMs can reside in the same host, to improve redundancy, install them in separate hosts and, if possible, connect the VSMs to different upstream switches.
- The console for the standby VSM is available by entering the **module attach x** command, but configuration is not allowed and many commands are restricted. Enter this command at the console of the active VSM.
- You cannot use Telnet or Secure Shell (SSH) protocols to communicate with the standby VSM because the management interface IP is unconfigured until the VSM becomes active.
- The active and standby VSMs must be on the same management subnet.

## Configuring System-Level High Availability

### Changing the VSM Role

The Cisco Nexus 1000V VSM software installation provides an opportunity for you to designate the role for each VSM. You can use this procedure to change that initial configuration.

**Caution**

---

Changing the role of a VSM can result in a conflict between the VSM pair. If a primary and secondary VSM see each other as active at the same time, the system resolves this problem by resetting the primary VSM.

---

Use this procedure to change the role of a VSM to one of the following after it is already in service:

- Standalone
- Primary
- Secondary

#### Before You Begin

Before beginning this procedure, you must be logged in to the CLI in EXEC mode.

If you are changing a standalone VSM to a secondary VSM, be sure to first isolate it from the other VSM in the pair to prevent any interaction with the primary VSM during the change. Power the VM off from the SCVMM before reconnecting it as standby.

You must understand the following information:

- The possible HA roles are standalone, primary, and secondary.
- The possible HA redundancy states are active and standby.
- To activate a change from primary to secondary VSM, you must reload the VSM by doing one of the following:
  - Issue the **reload** command.
  - Power the VM off and then on from SCVMM.
- A change from a standalone to a primary VSM takes effect immediately.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	switch# <b>system redundancy role {standalone   primary   secondary}</b>	Designates the HA role of the VSM.
<b>Step 2</b>	switch# <b>show system redundancy status</b>	(Optional) Displays the current redundancy status for the VSMs.
<b>Step 3</b>	switch# <b>copy running-config startup-config</b>	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```

switch# system redundancy role standalone
switch# show system redundancy status
Redundancy role
-----
      administrative:  standalone
      operational:    standalone

Redundancy mode
-----
      administrative:  HA
      operational:    None

This supervisor (sup-1)
-----
      Redundancy state: Active
      Supervisor state: Active
      Internal state:Active with no standby

Other supervisor (sup-2)
-----
      Redundancy state:  Not present
switch#

```

## Configuring a Switchover

### Guidelines and Limitations for Configuring a Switchover

- When you manually initiate a switchover, system messages are generated that indicate the presence of two VSMs and identify which one is becoming active.
- A switchover can only be performed when both VSMs are functioning.

### Verifying that a System is Ready for a Switchover

Use one of the following commands to verify the configuration:

Command	Purpose
<b>show system redundancy status</b>	<p>Displays the current redundancy status for the VSM(s).</p> <p>If the output indicates the following, then you can proceed with a system switchover, if needed:</p> <ul style="list-style-type: none"> <li>• The presence of an active VSM</li> <li>• The presence of a standby VSM in the HA standby redundancy state</li> </ul>
<b>show module</b>	<p>Displays information about all available VEMs and VSMs in the system.</p> <p>If the output indicates the following, then you can proceed with a system switchover, if needed:</p> <ul style="list-style-type: none"> <li>• The presence of an active VSM</li> <li>• The presence of a standby VSM in the HA standby redundancy state</li> </ul>

### Manually Switching the Active VSM to a Standby VSM

Be sure you know the following about manually switching the active VSM to a standby VSM:

- A switchover can only be performed when two VSMs are functioning in the switch.
- If the standby VSM is not in a stable state (ha-standby), then you cannot initiate a manual switchover. You will see the following error message:

```
Failed to switchover (standby not ready to takeover in vdc 1)
```

- If a switchover does not complete successfully within 28 seconds, the supervisors will reset.

## Before You Begin

- Logged in to the active VSM CLI in EXEC mode.
- Completed the steps in the [Verifying that a System is Ready for a Switchover](#) section and have found the system to be ready for a switchover.

## Procedure

	Command or Action	Purpose
<b>Step 1</b>	switch# <b>system switchover</b>	On the active VSM, initiate a manual switchover to the standby VSM.  Note the following: <ul style="list-style-type: none"> <li>• Once you enter this command, you cannot start another switchover process on the same system until a stable standby VSM is available.</li> <li>• Before proceeding, wait until the switchover completes and the standby supervisor becomes active.</li> </ul>
<b>Step 2</b>	switch# <b>show running-config diff</b>	(Optional) Verifies the difference between the running and startup configurations.  Any unsaved running configuration in the active VSM is also unsaved in the VSM that becomes the active after switchover. Save that configuration in the startup if needed.
<b>Step 3</b>	switch# <b>copy running-config startup-config</b>	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to switch an active VSM to a standby VSM and displays the output that appears on the standby VSM as it becomes the active VSM.

```
switch# system switchover
-----
(20 Feb 2013 9:39 PM)
HOGWARTS# sh2013 Feb 21 02:23:56 HOGWARTS %PLATFORM-2-MOD_REMOVE: Module 1 removed (Serial
number )
HOGWARTS# 2013 Feb 21 02:24:33 HOGWARTS %REDUN_MGR-4-CTRL_COMM_STATUS_CHANGE: Control
Connectivity is UP with Primary VSM.
2013 Feb 21 02:24:43 HOGWARTS %BOOTVAR-5-NEIGHBOR_UPDATE_AUTOCOPY: auto-copy supported by
neighbor supervisor, starting...
2013 Feb 21 02:22:20 HOGWARTS %SYSMGR-STANDBY-4-READCONF_STARTED: Configuration update
started (PID 2971).
2013 Feb 21 02:25:19 HOGWARTS %SYSMGR-STANDBY-4-READCONF_STARTED: Configuration update
started (PID 3129).
2013 Feb 21 02:25:22 HOGWARTS %MODULE-5-STANDBY_SUP_OK: Supervisor 1 is standby
2013 Feb 21 02:25:23 HOGWARTS %SYSMGR-STANDBY-5-MODULE_ONLINE: System Manager has received
notification of local module becoming online.
```

This example shows how to display the difference between the running and startup configurations

```
switch# show running-config diff
*** Startup-config
--- Running-config
```



```
*****
*** 1,13 ****

! !Command: show startup-config
! !Time: Thu Feb 21 02:26:23 2013
! !Startup config saved at: Wed Feb 20 12:43:56 2013

version 5.2(1)SM1(5.1)
svs switch edition advanced

hostname HOGWARTS

no feature telnet
feature private-vlan
feature netflow
--- 1,13 ----
```

## Adding a Second VSM to a Standalone System

### Adding a Second VSM to a Standalone System

- Although primary and secondary VSMs can reside in the same host, to improve redundancy, install them in separate hosts and, if possible, connect them to different upstream switches.
- When you install the second VSM, assign the secondary role to it.
- The secondary VSM needs to have its control, management, and packet network interfaces in the same VLANs as the primary VSM.
- To add a secondary VSM to the standalone VSM, convert the standalone VSM to a primary VSM.
- After the secondary VSM is installed, the following occurs automatically:
  - The secondary VSM is reloaded and added to the system.
  - The secondary VSM negotiates with the primary VSM and becomes the standby VSM.
  - The standby VSM synchronizes the configuration and state with the primary VSM.

### Changing the Standalone VSM to a Primary VSM

Use this procedure to change the role of a VSM from standalone in a single VSM system to primary in a dual VSM system.

A change from a standalone to a primary VSM takes effect immediately.

#### Before You Begin

Before beginning this procedure, you must be logged in to the CLI in EXEC mode.

**Procedure**

	<b>Command or Action</b>	<b>Purpose</b>
<b>Step 1</b>	switch# <b>system redundancy role primary</b>	Changes the standalone VSM to a primary VSM. The role change occurs immediately.  The role change occurs immediately.
<b>Step 2</b>	switch# <b>show system redundancy status</b>	(Optional) Displays the current redundancy state for the VSM.
<b>Step 3</b>	switch# <b>copy running-config startup-config</b>	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to display the current system redundancy status for the VSM.

```
switch# system redundancy role primary
switch# show system redundancy status
Redundancy role
-----
      administrative:  primary
      operational:    primary

Redundancy mode
-----
      administrative:  HA
      operational:    None

This supervisor (sup-1)
-----
      Redundancy state:  Active
      Supervisor state:  Active
      Internal state:   Active with no standby

Other supervisor (sup-2)
-----
      Redundancy state:  Not present
switch# copy running-config startup-config
```

**Verifying the Change to a Dual VSM System**

Use one of the following commands to verify the configuration:

<b>Command</b>	<b>Purpose</b>
<b>show system redundancy status</b>	Displays the current redundancy status for VSMs in the system.
<b>show module</b>	Displays information about all available VSMs and VEMs in the system.

## Replacing the Secondary VSM in a Dual VSM System



**Note** Equipment Outage—This procedure requires that you power down and reinstall a VSM. During this time, your system will be operating with a single VSM.

---

### Procedure

**Step 1** Power off the secondary VSM.

**Step 2** Install the new VSM as a secondary VSM with the same domain ID as the existing VSM. See the *Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide*.  
Once the new VSM is added to the system, it will synchronize with the existing VSM.

## Replacing the Primary VSM in a Dual VSM System

Use this procedure to replace the primary VSM in a dual VSM system.



**Note** Equipment Outage—This procedure requires powering down and reinstalling a VSM. During this time, your system will be operating with a single VSM.

---

### Before You Begin

- Log in to the CLI in EXEC mode.
- Configure the port groups so that the new primary VSM cannot communicate with the secondary VSM or any of the VEMs during setup. VSMs with a primary or secondary redundancy role have built-in mechanisms for detecting and resolving the conflict between two VSMs in the active state. To avoid these mechanisms during the configuration of the new primary VSM, you must isolate the new primary VSM from the secondary VSM.

---

### Procedure

**Step 1** Power off the active primary VSM.  
The secondary VSM becomes the primary VSM.

**Step 2** On the SCVMM, change the port group configuration for the new primary VSM to prevent communication with the secondary VSM and the VEMs during setup.  
For an example of changing the port groups and port profiles assigned to the VSM interfaces in the vSphere Client, see the *Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide*

- Step 3** Install the new VSM as a primary, with the same domain ID as the existing VSM. See the *Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide*.
- Step 4** Save the configuration.
- Step 5** Power off the VM.
- Step 6** On the SCVMM, change the port group configuration for the new primary VSM to permit communication with the secondary VSM and the VEMs.
- Step 7** Power up the new primary VSM.  
The new primary VSM starts and automatically synchronizes all configuration data with the secondary VSM, which is currently the active VSM. Because the existing VSM is active, the new primary VSM becomes the standby VSM and receives all configuration data from the existing active VSM.
- 

## Changing the Domain ID in a Dual VSM System

### Before You Begin

- Have access to the console of both the primary VSM and the secondary VSM.
- Isolate the secondary VSM from the primary VSM to avoid the built-in mechanisms that detect and resolve conflict between two VSMs in a dual VSM system. This procedure includes the step for isolating the VSMs.



**Note** Equipment Outage—This procedure requires powering down a VSM. During this time, your system will be operating with a single VSM.

---

### Procedure

---

- Step 1** On the SCVMM for the secondary VSM, do one of the following to isolate the VSMs and prevent their communication while completing this procedure:
- Change the port group configuration for the interfaces using port groups that prevent the VSMs from communicating with each other.
  - Change the “Connected” option to “Not Connected” for the VSM VM vNic.

The secondary VSM becomes active but cannot communicate with the primary VSM or the VEM

- Step 2** At the console of the secondary VSM, change the domain id and save the configuration.

#### Example:

```
switch# configure terminal
switch(config)# svcs-domain
switch(config-svs-domain)# domain id 100
switch(config-svs-domain)# copy running-config startup-config
```

The domain ID is changed on the secondary VSM and the VEM connected to it.

**Step 3** Power down the secondary VSM.

**Step 4** At the console of the primary VSM, change the domain id and save the configuration.

**Example:**

```
switch# configure terminal
switch(config)# svcs-domain
switch(config-svs-domain)# domain id 100
switch(config-svs-domain)# copy running-config startup-config
```

The domain ID is changed on the primary VSM and the VEM connected to it.

**Step 5** On the SCVMM for the secondary VSM, change the port group configuration for the interfaces to permit communication with the primary VSM.

When the secondary VSM is powered up, it will be able to communicate with the primary VSM.

**Step 6** Power up the secondary VSM.

Both VSMs are now using the new domain ID and will synchronize with each other.

## Verifying the HA Status

Use one of the following commands to verify the configuration:

Command	Purpose
<b>show system redundancy status</b>	Displays the HA status of the system.
<b>show module</b>	Displays information about all available VSMs and VEMs in the system.
<b>show processes</b>	<p>Displays the state of all processes and the start count of the process.</p> <p>The states and types are described as follows:</p> <ul style="list-style-type: none"> <li>• State: R (runnable), S (sleeping), Z (defunct)</li> <li>• Type: U (unknown), O (non sysmgr), VL (vdc-local), VG (vdc-global), VU (vdc-unaware), NR (not running), ER (terminated etc)</li> </ul>

## MIBs

MIBs	MIBs Link
CISCO-PROCESS-MIB	To locate and download MIBs, go to the following URL: <a href="http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml">http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</a>

## Feature History for System-Level High Availability

Feature Name	Releases	Feature Information
System -Level High Availability	5.2(1)SM1(5.1)	This feature was introduced.