



## CHAPTER 17

# High Availability

---

This chapter describes how to identify and resolve problems related to High Availability.

This chapter includes the following sections:

- [Information About High Availability, page 17-1](#)
- [Problems with High Availability, page 17-3](#)
- [High Availability Troubleshooting Commands, page 17-6](#)

## Information About High Availability

The purpose of High Availability (HA) is to limit the impact of failures—both hardware and software—within a system. The Cisco NX-OS operating system is designed for high availability at the network, system, and service levels.

For detailed information about High Availability, see the *Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)*.

The following HA features minimize or prevent traffic disruption in the event of a failure:

- Redundancy— redundancy at every aspect of the software architecture.
- Isolation of processes— isolation between software components to prevent a failure within one process disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- Supervisor stateful switchover— Active/standby dual supervisor configuration. State and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide seamless and stateful switchover in the event of a VSM failure.

The Cisco Nexus 1000V system is made up of the following:

- Virtual Ethernet Modules (VEMs) running within virtualization servers. These are represented as modules within the VSM.
- A remote management component, for example, VMware vCenter Server.
- One or two VSMs running within Virtual Machines (VMs)

**[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)**

## System-Level High Availability

The Cisco Nexus 1000V supports redundant VSM virtual machines — a primary and a secondary — running as an HA pair. Dual VSMs operate in an active/standby capacity in which only one of the VSMs is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two VSMs to provide a stateful switchover if the active VSM fails.

### Single or Dual Supervisors

The Cisco Nexus 1000V system is made up of the following:

- Virtual Ethernet Modules (VEMs) running within virtualization servers (these are represented as modules within the VSM)
- A remote management component, for example, VMware vCenter Server.
- One or two Virtual Supervisor Modules (VSMs) running within Virtual Machines (VMs)

Single VSM Operation	Dual VSM Operation
<ul style="list-style-type: none"> <li>• Stateless—Service restarts from the startup configuration</li> <li>• Stateful—Service resumes from previous state.</li> </ul>	<ul style="list-style-type: none"> <li>• 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. They are also synchronized and kept up to date with the active VSM in order to maintain the runtime context of “ready to run.”</li> <li>• On a switchover, the standby VSM takes over for the active VSM.</li> </ul>

## Network-Level High Availability

The Cisco Nexus 1000V HA at the network level includes port channels and Link Aggregation Control Protocol (LACP). A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to eight physical links. If a member port within a port channel fails, the traffic previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, LACP lets you configure up to 16 interfaces into a port channel. A maximum of eight interfaces can be active, and a maximum of eight interfaces can be placed in a standby state.

For additional information about port channels and LACP, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0*.

**Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).**

## Problems with High Availability

Table 17-1 provides symptoms related to high availability, their possible causes, and recommended solutions.

**Table 17-1** Problems with High Availability

Symptom	Possible Cause	Solution
The active VSM does not see the standby VSM.	Roles are not configured properly. <ul style="list-style-type: none"> <li>primary</li> <li>secondary</li> </ul>	Verify roles, update an incorrect role, and save the configuration. <ol style="list-style-type: none"> <li>Verify the role of each VSM. <b>show system redundancy status</b></li> <li>Update an incorrect role. <b>system redundancy role</b></li> <li>Save the configuration. <b>copy run start</b></li> </ol>
	Network connectivity problems between the VSM and the upstream and virtual switches. Problem could be in the control or management VLAN.	Restore connectivity. <ol style="list-style-type: none"> <li>From the vSphere client, shut down the VSM, which should be in standby mode.</li> <li>From the vSphere client, bring up the standby VSM after network connectivity is restored.</li> </ol>
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs.	Verify that the VSMs are using the same software version. If not, then reinstall the image. <ol style="list-style-type: none"> <li>Verify software version on both VSMs. <b>show version</b></li> <li>Reinstall the secondary VSM with the same version used in the primary.</li> </ol>
	Fatal errors during gsync process. <ul style="list-style-type: none"> <li>Check the gsyncctrl log using the <b>show system internal log sysmgr gsyncctrl</b> command and look for fatal errors.</li> </ul>	Reload the standby VSM. <b>reload module standby_module_number</b> See the “Reloading a Module” procedure on page 17-9.

**Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).**

**Table 17-1** Problems with High Availability (continued)

Symptom	Possible Cause	Solution
The standby VSM reboots periodically.	The VSM has connectivity only through the management interface.  When a VSM is able to communicate through the management interface, but not through the control interface, the active VSM resets the standby to prevent the two VSMs from being in HA mode and out of sync.	Check control VLAN connectivity between the primary and secondary VSM.  <b>show system internal redundancy info</b>  In the output, <b>degraded_mode flag = true</b> .  If there is no connectivity, restore it through the control interface.
	VSMs have different versions.  Enter the <b>debug system internal sysmgr all</b> command and look for the active_verctrl entry that indicates a version mismatch, as the following output shows:  <pre>2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.</pre>	Isolate the standby VSM and boot it.  Use the <b>show version</b> command to check the software version in both VSMs.  Install the image matching the Active VSM on the standby. For more information, see the <i>Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)</i> .
Both VSMs are in active mode.	Network connectivity problems. <ul style="list-style-type: none"> <li>Check for control and management VLAN connectivity between the VSM at the upstream and virtual switches.</li> <li>When the VSM cannot communicate through any of these two interfaces, they will both try to become active.</li> </ul>	If network problems exist: <ol style="list-style-type: none"> <li>From the vSphere client, shut down the VSM, which should be in standby mode.</li> <li>From the vSphere client, bring up the standby VSM after network connectivity is restored.</li> </ol>
	The VSMs have different domain IDs	Verify the domain IDs in each VSM and, if needed, change the incorrect domain ID. <ol style="list-style-type: none"> <li>Verify domain IDs in each VSM. <b>show system internal redundancy info</b></li> <li>Isolate the VSM with the incorrect domain ID so that it cannot communicate with the other VSM.</li> <li>Change the domain ID in the isolated VSM, save configuration, and power off the VSM.</li> <li>Reconnect and power on the isolated VSM.</li> </ol>

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

**Table 17-1**      **Problems with High Availability (continued)**

Symptom	Possible Cause	Solution
Active and standby VSMs are not synchronized	<p><b>Incompatible versions</b></p> <p>The boot variables for active and standby VSMs are set to different image names, or if image names are the same, the files are not the correct files.</p> <p>When active and standby VSMs are running different versions that are not HA compatible, they are unable to synchronize.</p>	<p>Update the software version or the boot variables.</p> <ol style="list-style-type: none"> <li>From each VSM (active and standby), verify the software version.</li> </ol> <p><b>show version</b></p> <ol style="list-style-type: none"> <li>Reload the standby VSM with the version that is running in the active by doing one of the following: <ul style="list-style-type: none"> <li>correcting the boot variable names</li> <li>replacing the incorrect software files</li> </ul> </li> </ol> <p>See the “Reloading a Module” procedure on page 17-9.</p>
	<p><b>Broadcast traffic problem:</b></p> <p>Broadcast traffic from standby to active VSM may prevent the VSMs from synchronizing. Standby VSM tries to contact the active VSM periodically, but if broadcast traffic problems persist for over a minute when the standby is booting up, the system cannot synchronize.</p>	<p>Fix the traffic problem and reload the standby VSM.</p> <ol style="list-style-type: none"> <li>From the standby VSM, verify the broadcast traffic problem.</li> </ol> <p><b>show system internal log sysmgr verctrl</b></p> <p>If so, the following message will appear in the output:</p> <pre>standby_verctrl: no response from the active System Manager</pre> <ol style="list-style-type: none"> <li>Fix network connectivity.</li> <li>Reload standby VSM.</li> </ol> <p><b>reload module standby_module_number</b></p> <p>See the “Reloading a Module” procedure on page 17-9.</p>
	<p><b>False standby removal</b></p> <p>The active VSM falsely detects a disconnect with the standby. The standby is removed and reinserted and synchronization does not occur.</p>	<p>Verify redundancy states and reload the standby VSM.</p> <ol style="list-style-type: none"> <li>Verify active VSM redundancy.</li> </ol> <p><b>show system internal redundancy status</b></p> <p>Output = RDN_DRV_ST_AC_NP</p> <ol style="list-style-type: none"> <li>Verify standby VSM redundancy.</li> </ol> <p><b>show system internal redundancy status</b></p> <p>Output = RDN_DRV_ST_SB_AC</p> <ol style="list-style-type: none"> <li>Reload the standby VSM.</li> </ol> <p><b>reload module standby_module_number</b></p> <p>See the “Reloading a Module” procedure on page 17-9.</p>

**Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).**

## High Availability Troubleshooting Commands

This section lists commands that can be used to troubleshoot problems related to High Availability and includes the following topics:

- [Displaying Cores, page 17-6](#)
- [Checking Redundancy, page 17-7](#)
- [Checking the System Manager State, page 17-8](#)
- [Reloading a Module, page 17-9](#)
- [Attaching to the Standby VSM Console, page 17-10](#)

### Displaying Cores

To list cores, use the following command:

**show cores**

**Example:**

```
n1000V# show cores
VDC No Module-num      Process-name      PID      Core-create-time
-----
1      1      private-vlan      3207      Apr 28 13:29
```

### Displaying Logs of Processes

To display logs of processes, use the following command and the pid from the output of the **show cores** command:

**show processes log [pid pid]**

```
n1000V# show processes log
VDC Process      PID      Normal-exit  Stack  Core  Log-create-time
-----
1 private-vlan    3207      N      Y      N      Tue Apr 28 13:29:48 2009
```

**Example:**

```
n1000V# show processes log pid 3207
```

```
=====
```

```
Service: private-vlan
Description: Private VLAN
```

```
Started at Wed Apr 22 18:41:25 2009 (235489 us)
Stopped at Tue Apr 28 13:29:48 2009 (309243 us)
Uptime: 5 days 18 hours 48 minutes 23 seconds
```

```
Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2) <-- Reason for the process abort
Last heartbeat 46.88 secs ago
System image name: nexus-1000v-mzg.4.0.4.SV1.1.bin
System image version: 4.0(4)SV1(1) S25
```

```
PID: 3207
Exit code: signal 6 (core dumped) <-- Indicates that a cores for the process was generated.
```

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

```
CWD: /var/sysmgr/work
...
```

## Checking Redundancy

This section includes the following topics:

- [Checking System Redundancy Status, page 17-7](#)
- [Checking System Internal Redundancy Status, page 17-7](#)

## Checking System Redundancy Status

To check system redundancy status, use the following command:

### show system redundancy status

```
N1000V# show system redundancy status
Redundancy role
-----
      administrative:  primary <-- Configured redundancy role
      operational:    primary <-- Current operational redundancy role

Redundancy mode
-----
      administrative:  HA
      operational:    HA

This supervisor (sup-1)
-----
      Redundancy state:  Active <-- Redundancy state of this VSM
      Supervisor state:  Active
      Internal state:    Active with HA standby

Other supervisor (sup-2)
-----
      Redundancy state:  Standby <-- Redundancy state of the other VSM
      Supervisor state:  HA standby
      Internal state:    HA standby <-- The standby VSM is in HA mode and in sync
```

## Checking System Internal Redundancy Status

To check the system internal redundancy status, use the following command:

### show system internal redundancy info

```
n1000V# show system internal redundancy info
My CP:
  slot: 0
  domain: 184 <-- Domain id used by this VSM
  role: primary <-- Redundancy role of this VSM
  status: RDN_ST_AC <-- Indicates redundancy state (RDN_ST) of the this VSM is Active (AC)
  state: RDN_DRV_ST_AC_SB
  intr: enabled
  power_off_reqs: 0
  reset_reqs: 0
Other CP:
  slot: 1
```

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

```

status: RDN_ST_SB <-- Indicates redundancy state (RDN_ST) of the other VSM is Standby
(SB)
active: true
ver_rcvd: true
degraded_mode: false <-- When true, it indicates that communication through the control
interface is faulty
Redun Device 0: <-- This device maps to the control interface
name: ha0
pdev: ad7b6c60
alarm: false
mac: 00:50:56:b7:4b:59
tx_set_ver_req_pkts: 11590
tx_set_ver_rsp_pkts: 4
tx_heartbeat_req_pkts: 442571
tx_heartbeat_rsp_pkts: 6
rx_set_ver_req_pkts: 4
rx_set_ver_rsp_pkts: 1
rx_heartbeat_req_pkts: 6
rx_heartbeat_rsp_pkts: 442546 <-- Counter should be increasing, as this indicates that
communication between VSM is working properly.
rx_drops_wrong_domain: 0
rx_drops_wrong_slot: 0
rx_drops_short_pkt: 0
rx_drops_queue_full: 0
rx_drops_inactive_cp: 0
rx_drops_bad_src: 0
rx_drops_not_ready: 0
rx_unknown_pkts: 0
Redun Device 1: <-- This device maps to the mgmt interface
name: hal
pdev: ad7b6860
alarm: true
mac: ff:ff:ff:ff:ff:ff
tx_set_ver_req_pkts: 11589
tx_set_ver_rsp_pkts: 0
tx_heartbeat_req_pkts: 12
tx_heartbeat_rsp_pkts: 0
rx_set_ver_req_pkts: 0
rx_set_ver_rsp_pkts: 0
rx_heartbeat_req_pkts: 0
rx_heartbeat_rsp_pkts: 0 <-- When communication between VSM through the control
interface is interrupted but continues through the mgmt interface, the
rx_heartbeat_rsp_pkts will increase.
rx_drops_wrong_domain: 0
rx_drops_wrong_slot: 0
rx_drops_short_pkt: 0
rx_drops_queue_full: 0
rx_drops_inactive_cp: 0
rx_drops_bad_src: 0
rx_drops_not_ready: 0
rx_unknown_pkts: 0

```

## Checking the System Manager State

To check the system internal sysmgr state, use the following command:

```
show system internal sysmgr state
```

```
N1000V# show system internal sysmgr state
```

```
The master System Manager has PID 1988 and UUID 0x1.
```



***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

```
Last time System Manager was gracefully shutdown.
The state is SRV_STATE_MASTER_ACTIVE_HOTSTDBY entered at time Tue Apr 28 13:09:13 2009.

The '-b' option (disable heartbeat) is currently disabled.

The '-n' (don't use rlimit) option is currently disabled.

Hap-reset is currently enabled.

Watchdog checking is currently disabled.

Watchdog kgdb setting is currently enabled.
```

Debugging info:

```
The trace mask is 0x00000000, the syslog priority enabled is 3.
The '-d' option is currently disabled.
The statistics generation is currently enabled.
```

HA info:

```
slotid = 1    supid = 0
cardstate = SYSMGR_CARDSTATE_ACTIVE .
cardstate = SYSMGR_CARDSTATE_ACTIVE (hot switchover is configured enabled).
Configured to use the real platform manager.
Configured to use the real redundancy driver.
Redundancy register: this_sup = RDN_ST_AC, other_sup = RDN_ST_SB.
EOBC device name: eth0.
Remote addresses:  MTS - 0x00000201/3      IP - 127.1.1.2
MSYNC done.
Remote MSYNC not done.
Module online notification received.
Local super-state is: SYSMGR_SUPERSTATE_STABLE
Standby super-state is: SYSMGR_SUPERSTATE_STABLE
Swover Reason : SYSMGR_SUP_REMOVED_SWOVER <-- Reason for the last switchover
Total number of Switchovers: 0 <-- Number of switchovers
                        >> Duration of the switchover would be listed, if any.
```

Statistics:

```
Message count:          0
Total latency:          0          Max latency:          0
Total exec:             0          Max exec:             0
```

## Reloading a Module

To reload a module, use the following command:

```
reload module
```



**Note** Using the **reload** command without specifying a module reloads the whole system.

**Example:**

```
n1000V# reload module 2
```

In this example, the secondary VSM is reloaded.

***Send document comments to [nexus1k-docfeedback@cisco.com](mailto:nexus1k-docfeedback@cisco.com).***

## Attaching to the Standby VSM Console

The standby VSM console is not accessible externally, but can be accessed from the active VSM through the **attach module** *module-number* command.

To attach to the standby VSM console, use the following command.

**attach module**

**Example:**

```
n1000V# attach module 2
```

This example shows how to attach to the console of the secondary VSM.