



High Availability

This chapter describes how to identify and resolve problems related to high availability.

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.

The following Cisco NX-OS 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 that is 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. The state and configuration remain constantly synchronized between two Virtual Switch Modules (VSMs) to provide a 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) that run within virtualization servers. The VEMs are represented as modules within the VSM.
- One or two VSMs that run within virtual machines (VMs).

Problems with High Availability

Symptom	Possible Causes	Solution
The active VSM does not see the standby VSM.	Roles are not configured properly. Check the role of the two VSMs by entering the show system redundancy status command.	<ol style="list-style-type: none"> 1. Confirm that the roles are the primary and secondary role, respectively. 2. If needed, enter the system redundancy role command to correct the situation. 3. Save the configuration if roles are changed.
	Network connectivity problems. Check the L3 connectivity between the primary and secondary VSMs at the VSM host machines and upstream switches.	<p>If network problems exist, do the following:</p> <ol style="list-style-type: none"> 1. Shut down the VSM, which should be in standby mode. 2. Bring up the standby VSM after network connectivity is restored.
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs. Check that the primary and secondary VSM are using the same image version by entering the show version of the command.	If the active and standby VSM software versions differ, reinstall the secondary VSM with the same version used in the primary.
	Fatal errors during gsync process. Check the gsyncctrl log by entering the show system internal log sysmgr gsyncctrl command and look for fatal errors.	Reload the standby VSM by entering the reload module module-number command, where <i>module-number</i> is the module number for the standby VSM.

Symptom	Possible Causes	Solution
The standby VSM reboots periodically.	<p>The VSM has connectivity only through the management interface.</p> <p>When a VSM is able to communicate through the management interface, but not through the control interface, the active VSM detects the situation and resets the standby VSM to prevent the two VSMs from being in HA mode and out of sync.</p> <p>Check the output of the show system internal redundancy info command and verify if the <i>degraded_mode</i> flag is set to true.</p>	Check the control port connectivity between the primary and secondary VSMs.
	<p>VSMs have different versions.</p> <p>Enter the debug system internal sysmgr all command and look for the <i>active_verctrl</i> entry that indicates a version mismatch, as the following output shows:</p> <pre>2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.</pre>	<p>Isolate the standby VSM and boot it.</p> <p>Enter the show version command to check the software version in both VSMs.</p> <p>Install the image matching the active VSM on the standby.</p>
Both VSMs are in active mode.	<p>Network connectivity problems.</p> <p>Check the L3 connectivity between the primary and secondary VSMs at the VSM host machines and upstream switches.</p> <p>When the VSM cannot communicate through any of these two interfaces, they will both try to become active.</p>	<p>If network problems exist, do the following:</p> <ol style="list-style-type: none"> 1. Shut down the VSM, which should be in standby mode. 2. Bring up the standby VSM after network connectivity is restored.
	<p>Different domain IDs in the two VSMs.</p> <p>Check the <i>domain</i> value by entering the show system internal redundancy info command.</p>	<p>If needed, update the domain ID and save it to the startup configuration.</p> <p>To upgrade the domain ID in a dual VSM system, do the following:</p> <ol style="list-style-type: none"> 1. Isolate the VSM with the incorrect domain ID so that it cannot communicate with the other VSM. 2. Change the domain ID in the isolated VSM, save configuration, and power off the VSM. 3. Reconnect the isolated VSM and power it on.

System-Level High Availability

The Cisco Nexus 1000V supports redundant VSM VMs—a primary and a secondary—that run 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:

- VEMs that run within virtualization servers (these VEMs are represented as modules within the VSM)
- A remote management component, for example, the OpenStack dashboard.
- One or two VSMS that run within VMs.

Single VSM Operation	Dual VSM Operation
<ul style="list-style-type: none"> • Stateless—Service restarts from the startup configuration • Stateful—Service resumes from previous state. 	<ul style="list-style-type: none"> • One active VSM and one standby VSM. • The active VSM runs all the system applications and controls the system. • 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.” • On a switchover, the standby VSM takes over for the active VSM.

Network-Level High Availability

The Cisco Nexus 1000V HA at the network level includes port channels and the 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 that was previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, the LACP allows you to 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 the LACP, see the *Cisco Nexus 1000V for KVM Port Profile Configuration Guide, Release 5.x*.

High Availability Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to high availability.

To list process logs and cores, enter these commands:

- **show cores**

```
switch# show cores
Module Instance Process-name PID Date(Year-Month-Day Time)
-----
1 1 private-vlan 3207 Apr 28 13:29
```

- **show processes log [pid pid]**

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

switch# 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: switchh-dk9.5.2.1.SM15.0.1.bin
System image version: 5.2(1)SK1(1.1)

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

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

To check the redundancy status, enter this command:

- **show system redundancy status**

```
switch# show 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: N/A

      Supervisor state: N/A
      Internal state:  N/A

System start time:          Thu Sep  4 16:48:55 2014

System uptime:              7 days, 11 hours, 39 minutes, 24 seconds
Kernel uptime:             7 days, 11 hours, 39 minutes, 11 seconds
Active supervisor uptime:  7 days, 11 hours, 38 minutes, 45 seconds
```

To check the system internal redundancy status, enter this command:

- **show system internal redundancy info**

```
switch# show system internal redundancy info
My CP:
  slot: 0
  domain: 36
  role: primary
  status: RDN_ST_AC
  state: RDN_DRV_ST_AC_NP
  intr: enabled
  power_off_reqs: 0
  reset_reqs: 1
  inter_vsm_max_heartbeat_loss: 15
  product_type: 2
Other CP:
  slot: 1
  status: RDN_ST_NP
  active: true
  ver_rcvd: false
  degraded_mode: true
  prod_type_rcvd: false
  peer_mac_rcvd: false
Redun Device 0:
  name: ha0
  pdev: c9949800
  alarm: false
  mac: ff:ff:ff:ff:ff:ff
  tx_set_ver_req_pkts: 646867
  tx_set_ver_rsp_pkts: 0
  tx_peer_mac_req_pkts: 0
  tx_peer_mac_rsp_pkts: 0
  tx_heartbeat_req_pkts: 0
  tx_heartbeat_rsp_pkts: 0
  rx_set_ver_req_pkts: 0
  rx_set_ver_rsp_pkts: 0
  rx_peer_mac_req_pkts: 0
  rx_peer_mac_rsp_pkts: 0
  rx_heartbeat_req_pkts: 0
  rx_heartbeat_rsp_pkts: 0
  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_drops_wrong_ver: 0
  rx_unknown_pkts: 0
  tx_rdn_mgr_params_msg_pkts: 0
  tx_rdn_mgr_params_ack_pkts: 0
  rx_rdn_mgr_params_msg_pkts: 0
  rx_rdn_mgr_params_ack_pkts: 0
Redun Device 1:
  name: ha1
  pdev: c994c800
  alarm: false
  mac: ff:ff:ff:ff:ff:ff
  tx_set_ver_req_pkts: 646867
  tx_set_ver_rsp_pkts: 0
  tx_peer_mac_req_pkts: 0
  tx_peer_mac_rsp_pkts: 0
  tx_heartbeat_req_pkts: 0
```

```

tx_heartbeat_rsp_pkts: 0
rx_set_ver_req_pkts: 0
rx_set_ver_rsp_pkts: 0
rx_peer_mac_req_pkts: 0
rx_peer_mac_rsp_pkts: 0
rx_heartbeat_req_pkts: 0
rx_heartbeat_rsp_pkts: 0
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_drops_wrong_ver: 0
rx_unknown_pkts: 0
tx_rdn_mgr_params_msg_pkts: 0
tx_rdn_mgr_params_ack_pkts: 0
rx_rdn_mgr_params_msg_pkts: 0
rx_rdn_mgr_params_ack_pkts: 0

```

To check the system internal sysmgr state, enter this command:

- **show system internal sysmgr state**

```

switch# show system internal sysmgr state
The master System Manager has PID 1323 and UUID 0x1.
Last time System Manager was gracefully shutdown.
The state is SRV_STATE_MASTER_ACTIVE_ALONE entered at time Thu Sep  4 16:49:07 2014.

```

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

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

Hap-reset is currently enabled.

Process restart capability is currently disabled.

Watchdog checking is currently enabled.

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_NP.
EOBC device name: eth0.
Remote addresses: MTS - [not available]    IP - [not available]
MSYNC not 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_UNKNOWN_SWOVER
Total number of Swovers: 0
Swover threshold settings: 20 swovers within 1200 seconds
Swovers within threshold interval: 0
Last swover time: 0 seconds after system start time
Cumulative time between last 0 swovers: 0
Start done received for 3 plugins, Total number of plugins = 3

```

Statistics:

```

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

```

To reload a module, enter this command:

- **reload module**

```
switch# reload module 2
```

This command reloads the secondary VSM.



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

To attach to the standby VSM console, enter this command:

- **attach module**

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

```
switch# attach module 2
```

This command attaches to the console of the secondary VSM.