



Cisco Nexus 1000V for KVM High Availability Configuration Guide, Release 5.x

First Published: August 01, 2014

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Overview 1

- Cisco Nexus 1000V for KVM and OpenStack 1
- Information About High Availability 2
- System Components 2
- Service-Level High Availability 3
 - Isolation of Processes 3
 - Process Restartability 3
- System-Level High Availability 3
- Network-Level High Availability 3
- VSM-to-VSM Heartbeat 4
 - Control and Management Interface Redundancy 4
 - Partial Communication 5
 - Loss of Communication 5
 - VSM-VEM Communication Loss 5
 - One-Way Communication 6
- Split-Brain Resolution 6
- Checking the Accounting Logs and the Redundancy Traces 7
- VSM Role Collision Detection 7
- Displaying the Role Collision 8
- Recommended Reading 9

CHAPTER 2

Understanding Service-Level High Availability 11

- Information About Cisco NX-OS Service Restarts 11
- Restartability Infrastructure 11
 - System Manager 12
 - Persistent Storage Service 12
 - Message and Transaction Service 12
 - High Availability Policies 12

Process Restartability	13
Stateful Restarts	14
Stateless Restarts	14
Switchovers	14
Restarts on Standby Supervisor Services	15
Restarts on Switching Module Services	15
Troubleshooting Restarts	15
MIBs	15
RFCs	16
Technical Assistance	16

CHAPTER 3

Configuring System-Level High Availability	17
Information About System-Level High Availability	18
Information About Single and Dual Supervisor Roles	18
HA Supervisor Roles	18
Dual Supervisor Active and Standby Redundancy States	19
Dual Supervisor Synchronization	19
Information About VSM Restarts and Switchovers	20
Restarts on Standalone VSMs	20
Restarts on Dual VSMs	20
Switchovers on Dual VSMs	20
Switchover Characteristics	20
Automatic Switchovers	21
Manual Switchovers	21
Guidelines and Limitations	21
Powering On the VSM VM	21
Powering Off the VSM VM	22
Configuring System-Level High Availability	23
Changing the VSM Role	23
Configuring a Switchover	24
Guidelines and Limitations for Configuring a Switchover	24
Verifying that a System is Ready for a Switchover	24
Manually Switching the Active VSM to Standby	25
Configuring the VSM-to-VSM Heartbeat Interval	26
Adding a Second VSM to a Standalone System	27

Adding a Second VSM to a Standalone System	27
Changing the Standalone VSM to a Primary VSM	28
Verifying the Change to a Dual VSM System	28
Replacing the Secondary in a Dual VSM System	29
Replacing the Primary VSM in a Dual VSM System	29
Changing the Domain ID in a Dual VSM System	30
Verifying the HA Status	31
Related Documents	32
Standards	32
MIBs	32
RFCs	32
Technical Assistance	33
Feature History for System-Level High Availability	33



CHAPTER

1

Overview

This chapter contains the following sections:

- [Cisco Nexus 1000V for KVM and OpenStack, page 1](#)
- [Information About High Availability, page 2](#)
- [System Components, page 2](#)
- [Service-Level High Availability, page 3](#)
- [System-Level High Availability, page 3](#)
- [Network-Level High Availability, page 3](#)
- [VSM-to-VSM Heartbeat, page 4](#)
- [Split-Brain Resolution, page 6](#)
- [Checking the Accounting Logs and the Redundancy Traces, page 7](#)
- [VSM Role Collision Detection, page 7](#)
- [Displaying the Role Collision, page 8](#)
- [Recommended Reading, page 9](#)

Cisco Nexus 1000V for KVM and OpenStack

The Cisco Nexus 1000V for KVM consists of two main components:

- **Virtual Ethernet Module (VEM)**—A software component that is deployed on each kernel-based virtual machine (VM) host. Each VM on the host is connected to the VEM through virtual Ethernet (vEth) ports.
- **Virtual Supervisor Module (VSM)**—The Management component that controls multiple VEMs and helps in the definition of VM-focused network policies. It is deployed either as a virtual appliance on any KVM host or on the Cisco Cloud Services Platform appliance.

Each of these components is tightly integrated with the OpenStack environment:

- The VEM is a hypervisor-resident component and is tightly integrated with the KVM architecture.

- The VSM is integrated with OpenStack using the OpenStack Neutron Plug-in.
- The OpenStack Neutron API has been extended to include two additional user-defined resources:
 - Network profiles are logical groupings of network segments.
 - Policy profiles group port policy information, including security.

Using OpenStack, you create VMs, networks, and subnets on the Cisco Nexus 1000V for KVM, by defining components such as the following:

- Tenants
- Network segments, such as VLANs, VLAN trunks, and VXLANs
- IP address pools (subnets)

Using the Cisco Nexus 1000V for KVM VSM, you create port profiles (called policy profiles in OpenStack), which define the port policy information, including security settings.

When a VM is deployed, a port profile is dynamically created on the Cisco Nexus 1000V for KVM for each unique combination of policy port profile and network segment. All other VMs deployed with the same policy to this network reuse this dynamic port profile.

**Note**

You must consistently use OpenStack for all VM network and subnet configuration. If you use *both* OpenStack and the VSM to configure VM networks and subnets, the OpenStack and the VSM configurations can become out-of-sync and result in faulty or inoperable network deployments.

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 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.

System Components

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

- One or two VSMs that run within Virtual Machines (VMs).

- Virtual Ethernet Modules (VEMs) that run within virtualization servers. VEMs are represented as modules within the VSM.
- A remote management component.

Service-Level High Availability

Isolation of Processes

The Cisco NX-OS software has independent processes, known as services, that perform a function or set of functions for a subsystem or feature set. Each service and service instance runs as an independent, protected process. This way of operating provides a highly fault-tolerant software infrastructure and fault isolation between services. A failure in a service instance does not affect any other services that are running at that time. Additionally, each instance of a service can run as an independent process, which means that two instances of a routing protocol can run as separate processes.

Process Restartability

Cisco NX-OS processes run in a protected memory space independently of each other and the kernel. This process isolation provides fault containment and enables rapid restarts. Process restartability ensures that process-level failures do not cause system-level failures. In addition, most services can perform stateful restarts. These stateful restarts allow a service that experiences a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network.

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 VSMs are configured as either primary or secondary as a part of the Cisco Nexus 1000V installation.

The state and configuration remain constantly synchronized between the two VSMs to provide a stateful switchover if the active VSM fails.

Network-Level High Availability

The Cisco Nexus 1000V high availability 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 previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, 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.

VSM-to-VSM Heartbeat

The primary and secondary VSM use a VSM to VSM heartbeat to do the following within their domain:

- Broadcast their presence
- Detect the presence of another VSM
- Negotiate active and standby redundancy states

When a VSM first boots up, it broadcasts discovery frames to the domain to detect the presence of another VSM. If no other VSM is found, the booting VSM becomes active. If another VSM is found to be active, the booting VSM becomes the standby VSM. If another VSM is found to be initializing (for example, during a system reload), the primary VSM has priority over the secondary to become the active VSM.



Note

From SV3(1.1) onwards, VSM validates the source MAC address of High Availability (HA) packets it receives on control and management interfaces. During initial contact VSM learns the peer VSM MAC addresses and stores in a permanent location. So only HA packets from a learnt VSM will be accepted.

After the initial contact and role negotiation, the active and standby VSMs unicast the following in heartbeat messages:

- Redundancy state
- Control flags requesting action by the other VSM

The following intervals apply when sending heartbeat messages.

Interval	Description
1 second	Interval at which heartbeat requests are sent.
6-30 seconds (Default is 15 seconds)	Interval after which missed heartbeats indicate degraded communication on the control interface so that heartbeats are also sent on the management interface. The default is 15 seconds; however, you can configure the interval to be between 6 and 30 seconds.
varies	The standby VSM resets automatically when the active VSM is no longer able to synchronize with it. This means the interval varies depending on how long the active VSM can buffer the data to be synchronized when the communication is interrupted.

Control and Management Interface Redundancy

The VSM communicates with the peer VSM over layer 2 only on the control and management interfaces. If the active VSM does not receive a heartbeat response over the control interface for a period of half of the inter-VSM maximum heartbeat loss interval (eight heartbeats by default), communication is seen as degraded and the VSM begins sending requests over the management interface in addition to the control interface. In

this case, the management interface provides redundancy by preventing both VSMs from becoming active. This process is called an active-active or split-brain situation.



Note The communication is not fully redundant, however, because the management interface only handles heartbeat requests and responses.

AIPC and the synchronization of data between VSMs is done through the control interface only.

Partial Communication

The secondary VSM is not immediately rebooted when communication over the control interface is interrupted because the HA mechanism tolerates brief interruptions in communication. When communication is first interrupted on the control interface, the heartbeat messages are sent over the management interface. If communication over the management interface is successful, the VSMs enter into a degraded mode, as displayed in the **show system internal redundancy trace** command output. If communication is interrupted on both interfaces for too long, the two VSMs get out of synchronization and the standby VSM is forced to reboot.



Note A transition from active to standby always requires a reload in both the Cisco Nexus 1000V and the Cisco Nexus Cloud Services Platform.

Loss of Communication

When there is no communication between redundant VSMs or Cisco Nexus Cloud Services Platforms, they cannot detect the presence of the other. The standby VSM will be removed from the list of inserted modules at the active VSM. The standby interprets the lack of heartbeats as a sign that the active has failed and it also becomes active. This process is what is referred to as active-active or split-brain, as both are trying to control the system by connecting to and communicating with the VEMs.

Because redundant VSMs or Cisco Nexus Cloud Services Platforms use the same IP address for their management interface, remote Secure Shell (SSH)/Telnet connections might fail, as a result of the path to this IP address changing in the network. For this reason, we recommend that you use the consoles during a split-brain conflict.

The following parameters are used to select the VSM to be rebooted during the split-brain resolution: the module count, the last configuration time, and the last active time.

VSM-VEM Communication Loss

Depending on the specific network failure that caused it, each VSM might reach a different, possibly overlapping, subset of Virtual Ethernet Modules (VEMs). When the VSM that was in the standby state becomes a new active VSM, it broadcasts a request to all VEMs to switch to it as the current active device. Whether a VEM switches to the new active VSM, depends on the following:

- The connectivity between each VEM and the two VSMs.
- Whether the VEM receives the request to switch.

A VEM remains attached to the original active VSM even if it receives heartbeats from the new active VSM. However, if the VEM also receives a request to switch from the new active VSM, it detaches from the original active VSM and attaches to the new VSM.

If a VEM loses connectivity to the original active device and only receives heartbeats from the new one, it ignores those heartbeats until it goes into headless mode, which occurs approximately 15 seconds after it stops receiving heartbeats from the original, active VSM. At that point, the VEM attaches to the new active VSM if it has connectivity to it.

One-Way Communication

If a network communication failure occurs where the standby VSM receives heartbeat requests but the active VSM does not receive a response, the following occurs:

- The active VSM declares that the standby VSM is not present.
- The standby VSM remains in a standby state and continues receiving heartbeats from the active VSM.

In this scenario, the redundancy state is inconsistent (**show system redundancy state**) and the two VSMs lose synchronization. When two-way communication is resumed, the standby VSM replies to the active VSM and asks to be reset.

**Note**

If a one-way communication failure occurs in the active to standby direction, it is equivalent to a total loss of communication because a standby VSM sends heartbeats only in response to active VSM requests.

Split-Brain Resolution

When the connectivity between two Virtual Supervisor Modules (VSMs) is broken, this loss of communication can cause both VSMs to take the active role. This condition is called active-active or split-brain condition. When the communication is restored between the VSMs, both VSMs exchange information to decide which one would have a lesser impact on the system, if rebooted.

Both primary and secondary VSMs process the same data to select the VSM (primary or secondary) that needs to be rebooted. When the selected VSM is rebooted and attaches itself back to the system, high availability is back to normal. The VSM uses the following parameters in order of precedence to select the VSM to be rebooted during the split-brain resolution:

- 1 Module count—The number of modules that are attached to the VSM.
- 2 Last configuration time—The time when the last configuration is done on the VSM.
- 3 Last standby-active switch—The time when the VSM last switched from the standby state to the active state. (The VSM with a longer active time gets higher priority.)

Checking the Accounting Logs and the Redundancy Traces

During the split-brain resolution, when a VSM reboots, the accounting logs that are stored on the VSM are lost. You can display the accounting logs that were backed up during the split-brain resolution. You can also check the redundancy traces that are stored on the local and remote VSMs.

Command	Purpose
show system internal active-active accounting logs	Displays the accounting logs that are stored on a local VSM during the last split-brain resolution.
show system internal active-active redundancy traces	Displays the redundancy traces that are stored on a local VSM during the last split-brain resolution.
show system internal active-active remote accounting logs	Displays the remote accounting logs that are stored on a remote VSM during the last split-brain resolution.
show system internal active-active remote redundancy traces	Displays the remote redundancy traces that are stored on a remote VSM during the last split-brain resolution.
clear active-active accounting logs	Clears the accounting logs that are stored on a local VSM during the split-brain resolution.
clear active-active redundancy traces	Clears the redundancy traces that are stored on a local VSM during the split-brain resolution.
clear active-active remote accounting logs	Clears the remote accounting logs that are stored on a remote VSM during the split-brain resolution.
clear active-active remote redundancy traces	Clears the remote redundancy traces that are stored on a remote VSM during the split-brain resolution.

VSM Role Collision Detection

In the Cisco Nexus 1000V, if a secondary VSM is configured or installed with the same role as the primary VSM and with the same domain ID, the secondary VSM and the primary VSM exchange heartbeats to discover each other. Both VSMs detect and report a role collision when they exchange heartbeats. When a collision is detected, the VSMs report the MAC address of the VSM with which the local VSM is colliding.

Due to this issue, the HA-paired VSM cannot communicate with the correct VSM. This problem can occur on a primary VSM or a secondary VSM depending on whether the newly configured or the installed VSM has the primary or the secondary role assigned to it.

The collisions are detected on the control and the management interfaces. The maximum number of colliding VSMs reported is eight.



Note After the eighth role collision, the problem is still logged and the MAC address entry is overwritten. The **show system redundancy status** command displays the overwrite details.



Note The colliding VSMs might also report a collision detection from the original VSM. If the colliding VSMs use the same IP address for their management interfaces, the remote SSH/Telnet connections might fail. Therefore, we recommend that you use the consoles during a role collision detection.

Enter the **show system redundancy status** command on both the primary and secondary VSM consoles to display the MAC addresses of the detected VSMs with the same role and domain ID, if any. When the VSM stops communicating in the domain, the collision time is not updated anymore. After an hour elapses since the last collision, the collision MAC entries are removed.

Displaying the Role Collision

Use the **show system redundancy status** CLI command to display the VSM role collision:

Command	Purpose
show system redundancy status	Displays a detected role collision A warning is highlighted in the CLI output. Along with the MAC addresses, the latest collision time is also displayed in the output. If no collisions are detected, the highlighted output does not appear.

This example shows how to display the detected traffic collision:

```
switch# show system redundancy status

Redundancy role
-----
      administrative:  secondary
      operational:    secondary

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

This supervisor (sup-2)
-----
      Redundancy state:  Active
      Supervisor state:  Active
      Internal state:   Active with HA standby

Other supervisor (sup-1)
-----
      Redundancy state:  Standby
      Supervisor state:  HA standby
      Internal state:   HA standby

WARNING! Conflicting sup-2(s) detected in same domain
-----
      MAC                Latest Collision Time
00:50:56:97:02:3b      2012-Sep-11 18:59:17
```

```
00:50:56:97:02:3c      2012-Sep-11 18:59:17
00:50:56:97:02:2f      2012-Sep-11 18:57:42
00:50:56:97:02:35      2012-Sep-11 18:57:46
00:50:56:97:02:29      2012-Sep-11 18:57:36
00:50:56:97:02:30      2012-Sep-11 18:57:42
00:50:56:97:02:36      2012-Sep-11 18:57:46
00:50:56:97:02:2a      2012-Sep-11 18:57:36
```

NOTE: Please run the same command on sup-1 to check for conflicting(if any) sup-1(s) in the same domain.

Recommended Reading

- *Cisco Nexus 1000V for KVM Software Installation Guide*
- *Cisco Nexus 1000V for KVM Port Profile Configuration Guide*



Understanding Service-Level High Availability

This chapter contains the following sections:

- [Information About Cisco NX-OS Service Restarts, page 11](#)
- [Restartability Infrastructure, page 11](#)
- [Process Restartability, page 13](#)
- [Restarts on Standby Supervisor Services , page 15](#)
- [Restarts on Switching Module Services, page 15](#)
- [Troubleshooting Restarts, page 15](#)
- [MIBs, page 15](#)
- [RFCs, page 16](#)
- [Technical Assistance, page 16](#)

Information About Cisco NX-OS Service Restarts

The Cisco NX-OS service restart feature restarts a faulty service without restarting the supervisor to prevent process-level failures from causing system-level failures. A service is restarted depending on current errors, failure circumstances, and the high-availability policy for the service. A service can undergo either a stateful or stateless restart. Cisco NX-OS allows services to store run-time state information and messages for a stateful restart. In a stateful restart, the service can retrieve this stored state information and resume operations from the last checkpoint service state. In a stateless restart, the service can initialize and run as if it had just been started with no prior state.

Restartability Infrastructure

Cisco NX-OS allows stateful restarts of most processes and services. The back-end management and orchestration of processes, services, and applications within a platform are handled by a set of high-level system-control services.

System Manager

The System Manager directs the overall system function, service management, and system health monitoring, and enforces high-availability policies. The System Manager is responsible for launching, stopping, monitoring, and restarting services and for initiating and managing the synchronization of service states and supervisor states for stateful switchovers.

The System Manager directs the overall system function, service management, and system health monitoring. The System Manager is responsible for launching, stopping, monitoring, and restarting services.

The System Manager directs the overall system function, service management, and system health monitoring, and enforces high-availability policies. The System Manager is responsible for launching, stopping, monitoring, and restarting services and for initiating and managing the synchronization of service states and supervisor states for stateful switchovers.

Persistent Storage Service

Cisco NX-OS services use the persistent storage service (PSS) to store and manage the operational run-time information and configuration of platform services. The PSS component works with system services to recover states in the event of a service restart. PSS functions as a database of state and run-time information, which allows services to make a checkpoint of their state information whenever needed. A restarting service can recover the last known operating state that preceded a failure, which allows for a stateful restart.

Each service that uses PSS can define its stored information as one of the following:

- Private—It can be read only by that service.
- Shared—The information can be read by other services.

The service can specify that it is one of the following:

- Local—The information can be read only by services on the same supervisor.
- Global—It can be read by services on either supervisor.

Message and Transaction Service

The message and transaction service (MTS) is a high-performance interprocess communications (IPC) message broker that specializes in high-availability semantics. MTS handles message routing and queuing between services on and across modules and between supervisors. MTS facilitates the exchange of messages such as event notification, synchronization, and message persistency between system services and system components. MTS can maintain persistent messages and logged messages in queues for access even after a service restart.

High Availability Policies

Cisco NX-OS allows each service to have an associated set of internal HA policies that define how a failed service will be restarted. Each service can have four defined policies—a primary and secondary policy when two supervisors are present, and a primary and secondary policy when only one supervisor is present. If no HA policy is defined for a service, the default HA policy to be performed upon a service failure will be a switchover if two supervisors are present or a supervisor reset if only one supervisor is present.

Each HA policy specifies three parameters:

- Action to be performed by the System Manager:
 - Stateful restart
 - Stateless restart
 - Supervisor switchover (or restart)
- Maximum retries—The number of restart attempts to be performed by the System Manager. If the service has not restarted successfully after this number of attempts, the HA policy is considered to have failed, and the next HA policy is used. If no other HA policy exists, the default policy is applied, which results in a supervisor switchover or restart.
- Minimum lifetime—The time that a service must run after a restart attempt in order to consider the restart attempt as successful. The minimum lifetime is no less than four minutes.

Process Restartability

Cisco NX-OS processes run in a protected memory space independently of each other and the kernel. This process isolation provides fault containment and enables rapid restarts. Process restartability ensures that process-level failures do not cause system-level failures. In addition, most services can perform stateful restarts. These stateful restarts allow a service that experiences a failure to be restarted and to resume operations transparently to other services within the platform and to neighboring devices within the network.

A failed service is restarted by different methods depending on the service's HA implementation and HA policies.

The following table describes the action taken by the System Manager for various failure conditions.

Failure	Action
Service/process exception	Service restart
Service/process crash	Service restart
Unresponsive service/process	Service restart
Repeated service failure	Supervisor reset (single) or switchover (dual)
Unresponsive System Manager	Supervisor reset (single) or switchover (dual)
Kernel failure	Supervisor reset (single) or switchover (dual)
Watchdog timeout	Supervisor reset (single) or switchover (dual)

Stateful Restarts

When a restartable service fails, it is restarted on the same supervisor. If the new instance of the service determines that the previous instance was abnormally terminated by the operating system, the service then determines whether a persistent context exists. The initialization of the new instance attempts to read the persistent context to build a run-time context that makes the new instance appear like the previous one. After the initialization is complete, the service resumes the tasks that it was performing when it stopped. During the restart and initialization of the new instance, other services are unaware of the service failure. Any messages that are sent by other services to the failed service are available from the MTS when the service resumes.

Whether or not the new instance survives the stateful initialization depends on the cause of the failure of the previous instance. If the service is unable to survive a few subsequent restart attempts, the restart is considered as failed. In this case, the System Manager executes the action specified by the service's HA policy, forcing either a stateless restart, no restart, or a supervisor switchover or reset.

During a successful stateful restart, there is no delay while the system reaches a consistent state. Stateful restarts reduce the system recovery time after a failure.

The events before, during, and after a stateful restart are as follows:

- 1 The running services make a checkpoint of their run-time state information to the PSS.
- 2 The System Manager monitors the health of the running services that use heartbeats.
- 3 The System Manager restarts a service instantly when it crashes or hangs.
- 4 After restarting, the service recovers its state information from the PSS and resumes all pending transactions.
- 5 If the service does not resume a stable operation after multiple restarts, the System Manager initiates a reset or switchover of the supervisor.
- 6 Cisco NX-OS will collect the process stack and core for debugging purposes with an option to transfer core files to a remote location.

When a stateful restart occurs, Cisco NX-OS sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap.

Stateless Restarts

Cisco NX-OS infrastructure components manage stateless restarts. During a stateless restart, the System Manager identifies the failed process and replaces it with a new process. The service that failed does not maintain its run-time state upon the restart, so the service can either build the run-time state from the running configuration, or if necessary, exchange information with other services to build a run-time state.

When a stateless restart occurs, Cisco NX-OS sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap.

Switchovers

If a standby supervisor is available, performs a supervisor switchover rather than a supervisor restart whenever multiple failures occur at the same time, because these cases are considered unrecoverable on the same supervisor. For example, if more than one HA application fails, that is considered an unrecoverable failure.

In a system with dual VSMs, after a switchover, the active supervisor resets and comes back up as a standby supervisor.

Restarts on Standby Supervisor Services

When a service fails on a supervisor that is in the standby state, the System Manager does not apply the HA policies and restarts the service after a delay of 30 seconds. The delay ensures that the active supervisor is not overwhelmed by repeated standby service failures and synchronizations. If the service being restarted requires synchronization with a service on the active supervisor, the standby supervisor is taken out of hot standby mode until the service is restarted and synchronized. Services that are not restartable cause the standby supervisor to reset.

When a standby service restart occurs, Cisco NX-OS sends a syslog message of level LOG_ERR. If SNMP traps are enabled, the SNMP agent sends a trap.

Restarts on Switching Module Services

Service failures on nonsupervisor module services do not require a supervisor switchover.

On the VEMs, the Data Path Agent (DPA) is restarted if it crashes. This situation causes the module to be removed and readded on the VSM.

Troubleshooting Restarts

When a service fails, the system generates information that can be used to determine the cause of the failure. The following sources of information are available:

- Every service restart generates a syslog message of level LOG_ERR.
- If SNMP traps are enabled, the SNMP agent sends a trap when a service is restarted
- When a service failure occurs, the event is logged. To view the log, use the **show processes log** command in that module. The process logs are persistent across switchovers and resets.
- When a service fails, a system core image file is generated. You can view recent core images by entering the **show cores** command on the active supervisor. Core files are persistent across VSM switchovers and resets and are saved in the `bootflash:core:` directory. Even though the core files are persistent across system reloads, the `show cores` command only displays the core files that were produced since the last system reload. You can configure the system to export core files to an external server using a file transfer utility, such as the Trivial File Transfer Protocol (TFTP).

MIBs

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

RFCs

No RFCs are supported by this feature.

Technical Assistance

Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.

Go to the following URL: <http://www.cisco.com/cisco/web/support/index.html>



Configuring System-Level High Availability

This chapter contains the following sections:

- [Information About System-Level High Availability, page 18](#)
- [Information About VSM Restarts and Switchovers, page 20](#)
- [Guidelines and Limitations, page 21](#)
- [Powering On the VSM VM, page 21](#)
- [Powering Off the VSM VM, page 22](#)
- [Configuring System-Level High Availability, page 23](#)
- [Adding a Second VSM to a Standalone System, page 27](#)
- [Replacing the Secondary in a Dual VSM System, page 29](#)
- [Replacing the Primary VSM in a Dual VSM System, page 29](#)
- [Changing the Domain ID in a Dual VSM System, page 30](#)
- [Verifying the HA Status, page 31](#)
- [Related Documents, page 32](#)
- [Standards, page 32](#)
- [MIBs, page 32](#)
- [RFCs, page 32](#)
- [Technical Assistance, page 33](#)
- [Feature History for System-Level High Availability, page 33](#)

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

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"> • This role does not interact with other VSMs. • You assign this role when there is only one VSM in the system. • This role is the default.

role	Module Number	Description
Primary	1	<ul style="list-style-type: none"> This role coordinates the active/standby state with the secondary VSM. 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. You assign this role to the first VSM that you install in a dual VSM system.
Secondary	2	<ul style="list-style-type: none"> This role coordinates the active/standby state with the primary VSM. You assign this role to the second VSM that you install in a dual VSM system.

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 attach module 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.

Instead, you can use the following methods:

- Virtual Network Computing (VNC) . Use the port 5901 and the host node's management IP to connect to the VSM VM.
 - The **virsh console VSM_VM_domain_name** command from the server console that is hosting the VSM VM.
 - A tool similar to **virt-manager** on the server that is hosting the VSM virtual machines. With this tool, you can use a GUI interface to login to the VSM virtual machines and then execute the commands for setting the redundancy role.
- The active and standby VSMs must be on the same management subnet.

Powering On the VSM VM

This procedure powers on the Virtual Supervisor Module (VSM) virtual machine (VM).

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	# <code>ssh server_node_IP_address</code>	Log in to the server node that is hosting the VSM VM.
Step 2	switch# <code>virsh list --all</code>	(Optional) Get the VSM VM domain name.
Step 3	switch# <code>virsh start VSM_VM_domain_name</code>	Power on the VSM VM.

This example shows how to power on the VSM VM:

```
# ssh 192.0.2.0
switch# virsh list --all
switch# virsh start vml
Domain vml started
switch#
```

Powering Off the VSM VM

This procedure powers off the Virtual Supervisor Module (VSM) virtual machine (VM).

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	# <code>ssh server_node_IP_address</code>	Log in to the server node that is hosting the VSM VM.
Step 2	switch# <code>virsh list</code>	(Optional) Get the VSM VM name and ID.
Step 3	switch# <code>virsh destroy VSM_VM_domain_name_or_id</code>	Power off the VSM VM.

This example shows how to power off the VSM VM:

```
# ssh 192.0.2.0
switch# virsh list
switch# virsh destroy vml
Domain vml destroyed
switch#
```

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 as described in [Split-Brain Resolution, on page 6](#).

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 before reconnecting it as standby. See [Powering Off the VSM VM, on page 22](#) for more information.

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:
 - Enter the **reload** command.
 - Power the VM off and then on. See [Powering Off the VSM VM, on page 22](#) and [Powering On the VSM VM, on page 21](#) for more information.
- A change from a standalone to a primary VSM takes effect immediately.

Procedure

	Command or Action	Purpose
Step 1	switch# system redundancy role {standalone primary secondary}	Designates the HA role of the VSM.

	Command or Action	Purpose
Step 2	switch# show system redundancy status	(Optional) Displays the current redundancy status for the VSMs.
Step 3	switch# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to change the VSM role:

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

Manually Switching the Active VSM to Standby

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

- A switchover can be performed only when two VSMs are functioning in the switch.
- If the standby VSM is not in a stable state (ha-standby), you cannot initiate a manual switchover and 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 reset.

Before You Begin

- Log in to the active VSM CLI in EXEC mode.
- Complete the steps in "[Verifying that a System is Ready for a Switchover](#)" and verify that the system is ready for a switchover.

Procedure

	Command or Action	Purpose
Step 1	<code>switch# system switchover</code>	<p>On the active VSM, initiates a manual switchover to the standby VSM.</p> <p>Once you enter this command, you cannot start another switchover process on the same system until a stable standby VSM is available.</p>

	Command or Action	Purpose
		Before proceeding, wait until the switchover completes and the standby supervisor becomes active.
Step 2	<code>switch# show running-config diff</code>	(Optional) Verifies the difference between the running and startup configurations. Any unsaved running configuration in an active VSM is also unsaved in the VSM that becomes active after switchover. Save that configuration in the startup if needed.
Step 3	<code>switch# copy running-config startup-config</code>	(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 the standby VSM and displays the output that appears on the standby VSM as it becomes the active VSM:

```
switch# system switchover
-----
2009 Mar 31 04:21:56 n1000v %% VDC-1 %% %SYSMGR-2-HASWITCHOVER_PRE_START:
This supervisor is becoming active (pre-start phase).
2009 Mar 31 04:21:56 n1000v %% VDC-1 %% %SYSMGR-2-HASWITCHOVER_START:
This supervisor is becoming active.
2009 Mar 31 04:21:57 n1000v %% VDC-1 %% %SYSMGR-2-SWITCHOVER_OVER: Switchover completed.
2009 Mar 31 04:22:03 n1000v %% VDC-1 %% %PLATFORM-2-MOD_REMOVE: Module 1 removed (Serial
number )
```

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

```
switch# show running-config diff
*** Startup-config
--- Running-config
*****
*** 180,189 ****
--- 179,190 ----
inherit port-profile bundle-uplink
no shutdown

interface Ethernet3/2
inherit port-profile bundle-uplink
+
+ interface control0
clock summer-time EDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60
line console
line vty
boot kickstart bootflash:/nexus-1000v-sk-upgrade-kickstart.5.2.1.SK1.2.0.437.gbin.upg sup-1
boot system bootflash:/nexus-1000v-sk-upgrade.5.2.1.SK1.2.0.437.gbin.upg sup-1
switch#
```

Configuring the VSM-to-VSM Heartbeat Interval

If the communication network between two VSMs in an HA pair experiences interruptions longer than the 15-second default, you can change the VSM-to-VSM heartbeat interval so that it is less sensitive to falsely detecting active VSM failures.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# system inter-sup-heartbeat time <i>time-interval</i>	Configures the VSM-to-VSM heartbeat with a time interval from 6 to 30 seconds. The default is 15 seconds.
Step 3	switch(config)# show running-config grep heartbeat	Displays the heartbeat interval setting in the running configuration.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to set the VSM-to-VSM interval to 10 seconds and verify this setting:

```
switch# configure terminal
switch(config)# system inter-sup-heartbeat time 10
switch(config)# show running-config | grep heartbeat
system inter-sup-heartbeat time 10
switch(config)# copy running-config startup-config
```

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

You can 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

Log in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch# system redundancy role primary	Changes the standalone VSM to a primary VSM. The role change occurs immediately.
Step 2	switch# show system redundancy status	(Optional) Displays the current redundancy state for the VSM.
Step 3	switch# copy running-config startup-config	(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:

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

Replacing the Secondary 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 operate with a single VSM.

Procedure

- Step 1** Power off the secondary VSM.
- Step 2** Install the new VSM as a secondary, with the same domain ID as the existing VSM, using the procedure in the “Installing and Configuring the VSM VM” section in the *Cisco Nexus 1000V for KVM Software Installation 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

You can replace a primary 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 operate with a single VSM.

You can alternately replace a primary VSM in a dual VSM system by using the Cisco OpenStack Installer (COI). See [Replacing the Primary VSM in a Dual VSM System Using the Cisco OpenStack Installer](#).

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 the 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 primary VSM.
The secondary VSM becomes active.
- Step 2** Install the new VSM as a primary, with the same domain ID as the existing VSM, using the “Installing the Primary VSM as a VM” section in the *Cisco Nexus 1000V for KVM Software Installation Guide*.
- Step 3** Power off the new primary VSM and restore communication with the secondary VSM.
Communication must be restored between the two VSMs before the new primary VSM is powered up again.
- Step 4** 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.
-

What to Do Next

If you no longer need the original primary VSM VM, you can remove it.

Changing the Domain ID in a Dual VSM System

Before You Begin

- Have access to the console of both the active and standby VSM.
- Isolate the standby VSM from the active VSM to avoid the built-in mechanisms that detect and resolve conflict between two VSMs with a primary or secondary redundancy role. This procedure has a step for isolating the VSMs.



Note Equipment Outage—This procedure requires that you power down a VSM. During this time, your system will operate with a single VSM.

Procedure

- Step 1** From the active VSM, shut down the interface that is used for the management and control on the standby VSM host node.

Example:

This example shows how to shut down an interface:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# shutdown
switch(config-if)# copy running-config startup-config
```

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

Example:

This example shows how to change the domain ID and save the configuration:

```
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 standby VSM and the VEM connected to it

- Step 3** Power down the standby VSM.
Step 4 At the console of the active VSM, change the domain ID and save the configuration.

Example:

This example shows how to change the domain ID and save the configuration:

```
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 active VSM and the VEM that is connected to it.

- Step 5** Bring up the VSM management and control interface that you shut down in step 1.

Example:

This example shows how to bring up an interface:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# no shutdown
switch(config-if)# copy running-config startup-config
```

- Step 6** Power up the standby VSM.
 Both VSMs are now using the new domain ID and will synchronize.

Verifying the HA Status

Use one of the following commands to verify the configuration:

Command	Purpose
<code>show system redundancy status</code>	Displays the HA status of the system.
<code>show module</code>	Displays information about all available VSMs and VEMs in the system.

Command	Purpose
<code>show processes</code>	<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"> • State: R (runnable), S (sleeping), Z (defunct) • Type: U (unknown), O (non sysmgr), VL (vdc-local), VG (vdc-global), VU (vdc-unaware), NR (not running), ER (terminated)

Related Documents

Related Topic	Document Title
Software upgrades	<i>Cisco Nexus 1000V for KVM Software Installation Guide</i>
Cisco Nexus 1000V commands	<i>Cisco Nexus 1000V for KVM Command Reference</i>

Standards

No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.

MIBs

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

RFCs

No RFCs are supported by this feature.

Technical Assistance

Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.

Go to the following URL: <http://www.cisco.com/cisco/web/support/index.html>

Feature History for System-Level High Availability

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Feature Name	Releases	Feature Information
System -Level High Availability	Release 5.2(1)SK1(2.1)	This feature was introduced.

