



Send document comments to nexus1k-docfeedback@cisco.com.



Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)

April 15, 2011

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

Text Part Number: OL-21654-01

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 NONINFRINGEMENT 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.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at 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. (1005R)

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

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)
©2010 Cisco Systems, Inc. All rights reserved.



New and Changed Information

No new information was added to the *Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)*.

To find information about new features with this release, see the following documentation:

- [Configuration Guides](#)
- [Release Notes](#)

Send document comments to nexus1k-docfeedback@cisco.com.



CONTENTS

New and Changed Information iii

Preface vii

Audience vii

Organization vii

Document Conventions viii

Available Documents ix

x

Obtaining Documentation and Submitting a Service Request x

CHAPTER 1

Overview 1-1

Information About High Availability 1-1

Service-Level High Availability 1-2

 Isolation of Processes 1-3

 Process Restartability 1-3

System-Level High Availability 1-3

Network-Level High Availability 1-3

VSM to VSM Heartbeat 1-4

 Description 1-4

 Control and Management Interface Redundancy 1-4

 Partial Communication 1-5

 Loss of Communication 1-5

 One-way Communication 1-6

Recommended Reading 1-6

CHAPTER 2

Understanding Service-Level High Availability 2-1

Information About Cisco NX-OS Service Restarts 2-1

Restartability Infrastructure 2-1

 System Manager 2-2

 Persistent Storage Service 2-2

 Message and Transaction Service 2-2

 HA Policies 2-2

Process Restartability 2-3

 Stateful Restarts 2-3

Send document comments to nexus1k-docfeedback@cisco.com.

- Stateless Restarts 2-4
- Switchovers 2-4
- Restarts on Standby Supervisor Services 2-4
- Restarts on Switching Module Services 2-5
- Troubleshooting Restarts 2-5
- Additional References 2-5
 - Related Documents 2-6
 - Standards 2-6
 - MIBs 2-6
 - RFCs 2-6
 - Technical Assistance 2-6

CHAPTER 3

- Configuring System-Level High Availability 3-1**
 - Information About System-Level High Availability 3-1
 - Information About Single and Dual Supervisors 3-1
 - Information About VSM Restarts and Switchovers 3-3
 - Guidelines and Limitations 3-4
 - Configuring System-Level High Availability 3-5
 - Changing the VSM Role 3-5
 - Configuring a Switchover 3-7
 - Adding a Second VSM to a Standalone System 3-11
 - Replacing the Standby in a Dual VSM System 3-15
 - Replacing the Active in a Dual VSM System 3-16
 - Changing the Domain ID in a Dual VSM System 3-17
 - Verifying the HA Status 3-18
 - Additional References 3-22
 - Related Documents 3-22
 - Standards 3-22
 - MIBs 3-23
 - RFCs 3-23
 - Feature History for System-Level High Availability 3-23

INDEX



Preface

The *Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)* describes how the built-in high availability (HA) and redundancy features can limit the impact of failures—both hardware and software— within a system.

This preface includes the following sections:

- [Audience, page vii](#)
- [Organization, page vii](#)
- [Document Conventions, page viii](#)
- [Available Documents, page ix](#)
- [Obtaining Documentation and Submitting a Service Request, page x](#)

Audience

This guide is for network administrators with the following experience and knowledge:

- An understanding of virtualization
- Using VMware tools to configure a vswitch



Note

Note: Knowledge of VMware vNetwork Distributed Switch is not a prerequisite.

Organization

This guide is organized as follows:

Chapter and Title	Description
Chapter 1, “Overview”	Provides an overview of high availability features.
Chapter 2, “Understanding Service-Level High Availability”	Describes service-level high availability and restarts.
Chapter 3, “Configuring System-Level High Availability”	Describes system and application high availability and restarts.

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

Document Conventions

This document uses the following conventions:



Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Tip

Means *the following information will help you solve a problem*.

Command descriptions use these conventions:

Convention	Description
boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
[]	Elements in square brackets are optional.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use these conventions:

screen font	Terminal sessions and information that the switch displays are in screen font.
boldface screen font	Information that you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Non-printing characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or number sign (#) at the beginning of a line of code indicates a comment line.

Send document comments to nexus1k-docfeedback@cisco.com.

Available Documents

This section lists the documents used with the Cisco Nexus 1000 and available on [Cisco.com](http://www.cisco.com) at the following url:

http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html

General Information

Cisco Nexus 1000V Documentation Roadmap, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Release Notes, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Compatibility Information, Release 4.0(4)SV1(3)

Cisco Nexus 1010 Management Software Release Notes, Release 4.0(4)SP1(1)

Install and Upgrade

Cisco Nexus 1000V Software Installation Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Software Upgrade Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1010 Virtual Services Appliance Hardware Installation Guide

Configuration Guides

Cisco Nexus 1000V License Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Quality of Service Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Security Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1010 Software Configuration Guide, Release 4.0(4)SP1(1)

Programming Guide

Cisco Nexus 1000V XML API User Guide, Release 4.0(4)SV1(3)

Reference Guides

Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(3)

Cisco Nexus 1000V MIB Quick Reference

Cisco Nexus 1010 Command Reference, Release 4.0(4)SP1(1)

Send document comments to nexus1k-docfeedback@cisco.com.

Troubleshooting and Alerts

Cisco Nexus 1000V Troubleshooting Guide, Release 4.0(4)SV1(3)

Cisco Nexus 1000V Password Recovery Guide

Cisco NX-OS System Messages Reference

Network Analysis Module Documentation

Cisco Network Analysis Module Software Documentation Guide, 4.2

Cisco Nexus 1000V NAM Virtual Service Blade Installation and Configuration Guide

Network Analysis Module Command Reference Guide, 4.2

User Guide for the Cisco Network Analysis Module Virtual Service Blades, 4.2

Cisco Network Analysis Module Software Release Notes, 4.2

Obtaining Documentation and Submitting a Service Request

For information about obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.



CHAPTER 1

Overview

This chapter describes high availability (HA) concepts and features for Cisco NX-OS software and includes the following sections:

- [Information About High Availability, page 1-1](#)
- [Service-Level High Availability, page 1-2](#)
- [System-Level High Availability, page 1-3](#)
- [Network-Level High Availability, page 1-3](#)
- [VSM to VSM Heartbeat, page 1-4](#)
- [Recommended Reading, page 1-6](#)

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.

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.
- Supervisor stateful switchover— Active/standby dual supervisor configuration. The state and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide a seamless and stateful switchover in the event of a VSM failure.

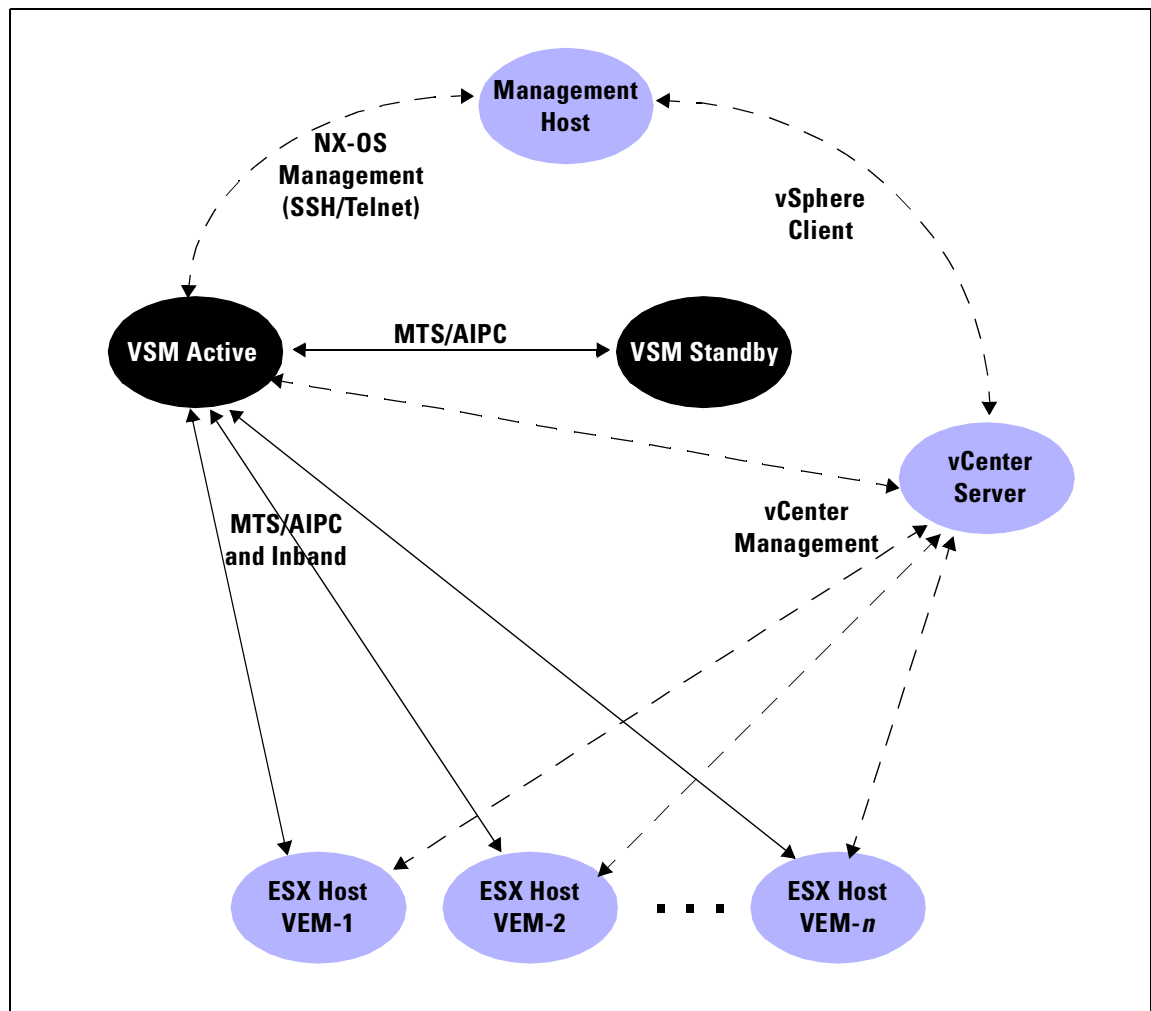
The Cisco Nexus 1000V system is made up of the following:

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

Send document comments to nexus1k-docfeedback@cisco.com.

Figure 1-1 shows the HA components and the communication links between them.

Figure 1-1 HA Components and Communication Links



Service-Level High Availability

The Cisco NX-OS software compartmentalizes processes for fault isolation, redundancy, and efficiency.

This section includes the following topics:

- [Isolation of Processes, page 1-3](#)
- [Process Restartability, page 1-3](#)

For additional details about service-level HA, see [Chapter 2, “Understanding Service-Level High Availability.”](#)

Send document comments to nexus1k-docfeedback@cisco.com.

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 will not affect any other services 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, which allows 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.

For more information about system-level high availability, see the [“Configuring System-Level High Availability” section on page 3-1](#).

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 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(4)SV1(3)*.

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

VSM to VSM Heartbeat

This section includes the following topics:

- [Description, page 1-4](#)
- [Control and Management Interface Redundancy, page 1-4](#)
- [Partial Communication, page 1-5](#)
- [Loss of Communication, page 1-5](#)

Description

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 standby. If another VSM is found to be initializing (for example, during a system reload), the primary VSM has priority over the secondary to become active.

After the initial contact and role negotiation, the active and standby VSMs unicast the following in heartbeat messages at one-second intervals:

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

The remainder of this section describe how the VSM to VSM heartbeat mechanism works including observance of the following intervals.

Interval	Description
1 second	Interval at which heartbeat requests are sent.
3 seconds	Interval after which missed heartbeats indicate degraded communication on the control interface so that heartbeats are also sent on the management interface.
6 seconds	Interval of communication loss after which the active VSM instructs the standby to reload.

Control and Management Interface Redundancy

If the active VSM does not receive a heartbeat response over the control interface for a period of 3 seconds, then communication is seen as degraded and the VSM also begins sending requests over the management interface. In this case, the management interface provides redundancy only in the sense that it acts to prevent both VSMs from becoming active, also called an active-active or split-brain situation. 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.

Send document comments to nexus1k-docfeedback@cisco.com.

Partial Communication

If communication over the control interface degrades sufficiently, the two VSMS lose synchronization which can only be restored by reloading the standby VSM. When this communication between VSMS is interrupted for over six seconds, but the VSMS can still communicate over the management interface, then the active VSM instructs the standby to reload. If communication remains degraded, the standby VSM may reload multiple times before control communication is restored.

Loss of Communication

When there is no communication between redundant VSMS or Cisco Nexus 1010s, neither can detect the presence of the other. The active drops the standby. The standby interprets the lack of response as a sign that the active has failed and it also becomes active. This is what is referred to as active-active or split-brain, as both are trying to control the system by connecting to Virtual Center and communicating with the VEMs.

Since redundant VSMS or Cisco Nexus 1010s use the same IP address for their management interface, remote SSH/Telnet connections may fail, as a result of the path to this IP address changing in the network. For this reason, it is recommended that you use the consoles during a split-brain conflict.

During a split-brain conflict, avoid making changes to the configuration. When control and management communication is restored between the redundant devices, the split-brain conflict is resolved by reloading the primary so that it comes back in standby mode and then synchronizes with the secondary. Changes made to a primary configuration during a split-brain conflict are lost after the reload.

If you need to make changes to the primary during a split-brain conflict, then shut down the secondary. This prevents the two devices from resuming contact which eventually results in a reload. In this case, changes you make to the primary during the split-brain conflict are retained.

**Note**

After a split-brain condition is resolved and the primary and secondary are communicating again, the configuration at the secondary takes precedence. Use care if changes are made when both are active.

**Note**

A transition from active to standby always requires a reload in both the Cisco Nexus 1000V and the Cisco Nexus 1010.

VSM-VEM Communication Loss

Depending on the specific network failure that caused it, each VSM may reach a different, possibly overlapping, subset of VEMs. When the VSM that was in standby state becomes a new active, 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 or not, 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. However, if it also receives a request to switch from the new active, it detaches from the original active and attaches to the new one.

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

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. This occurs approximately 15 seconds after it stops receiving heartbeats from the original active. At that point, the VEM attaches to the new active if it has connectivity to it.

One-way Communication

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

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

The redundancy state is inconsistent (**show system redundancy state**) and the two VSMs lose synchronization.

When two-way communication is resumed, the standby VSM detects the inconsistency and reloads.

**Note**

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

Recommended Reading

Before configuring the Cisco Nexus 1000V, Cisco recommends that you read and become familiar with the following documentation:

- *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*
- *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)*
- *Cisco VN-Link: Virtualization-Aware Networking white paper*



CHAPTER 2

Understanding Service-Level High Availability

This chapter describes the Cisco NX-OS service restartability for service-level high availability.

This chapter includes the following sections:

- [Information About Cisco NX-OS Service Restarts, page 2-1](#)
- [Restartability Infrastructure, page 2-1](#)
- [Process Restartability, page 2-3](#)
- [Restarts on Standby Supervisor Services, page 2-4](#)
- [Restarts on Switching Module Services, page 2-5](#)
- [Troubleshooting Restarts, page 2-5](#)
- [Additional References, page 2-5](#)

Information About Cisco NX-OS Service Restarts

The Cisco NX-OS service restart features allow you to restart a faulty service without restarting the supervisor to prevent process-level failures from causing system-level failures. You can restart a service 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. Back-end management and orchestration of processes, services, and applications within a platform are handled by a set of high-level system-control services described in this section.

This section includes the following topics:

- [System Manager, page 2-2](#)
- [Persistent Storage Service, page 2-2](#)
- [Message and Transaction Service, page 2-2](#)
- [HA Policies, page 2-2](#)

Send document comments to nexus1k-docfeedback@cisco.com.

System Manager

The System Manager directs 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 private (it can be read only by that service) or shared (the information can be read by other services). If the information is shared, the service can specify that it is local (the information can be read only by services on the same supervisor) or global (it can be read by services on either supervisor or on modules).

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.

HA 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—Specifies 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, resulting in a supervisor switchover or restart.

Send document comments to nexus1k-docfeedback@cisco.com.

- **Minimum lifetime**—Specifies the time that a service must run after a restart attempt in order to consider the restart attempt as successful. The minimum lifetime will be no less than four minutes.

Process Restartability

Process restartability ensures that a failed service can recover and resume operations without disrupting the data plane or other services. Depending on the service's HA policies, previous restart failures, and the health of other services running on the same supervisor, the System Manager determines the action to be taken when a service fails.

Table 2-1 describes the action taken by the System Manager for various failure conditions.

Table 2-1 System Manager Action on Failure Cases

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)

A failed service is restarted by one of the methods described in this section, depending on the service's HA implementation and HA policies,

This section includes the following topics:

- [Stateful Restarts, page 2-3](#)
- [Stateless Restarts, page 2-4](#)
- [Switchovers, page 2-4](#)

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 sent by other services to the failed service will be available from the MTS when the service resumes.

Whether or not the new instance survives the stateful initialization depends on the cause of 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.

Send document comments to nexus1k-docfeedback@cisco.com.

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, Cisco NX-OS 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.

For detailed information about supervisor switchovers and restarts, see [Chapter 3, “Configuring System-Level High Availability.”](#)

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.

Send document comments to nexus1k-docfeedback@cisco.com.

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 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 on a VSM, the event is logged. To view the log, use the **show processes log** command in that module. The process logs are persistent across supervisor 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 not persistent across supervisor switchovers and resets, but you can configure the system to export core files to an external server using a file transfer utility such as Trivial File Transfer Protocol (TFTP).

For information on collecting and using the generated information that relates to service failures, see the *Cisco Nexus 1000V Troubleshooting Guide, Release 4.0(4)SV1(3)*.

Additional References

For additional information related to implementing service-level HA features, see the following sections:

- [Related Documents, page 2-6](#)
- [Standards, page 2-6](#)
- [MIBs, page 2-6](#)
- [RFCs, page 2-6](#)
- [Technical Assistance, page 2-6](#)

Send document comments to nexus1k-docfeedback@cisco.com.

Related Documents

Related Topic	Document Title
Supervisor switchovers	Chapter 3, “Configuring System-Level High Availability.”
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(3)</i>
Troubleshooting	<i>Cisco Nexus 1000V Troubleshooting Guide, Release 4.0(4)SV1(3)</i>
Cisco Nexus 1000V system description	<i>Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)</i>

Standards

Standards	Title
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

RFCs	Title
No RFCs are supported by this feature	—

Technical Assistance

Description	Link
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.	http://www.cisco.com/public/support/tac/home.shtml



CHAPTER 3

Configuring System-Level High Availability

This chapter describes the Cisco NX-OS high availability (HA) system and application restart operations.

This chapter includes the following sections:

- [Information About VSM Restarts and Switchovers](#), page 3-3
- [Guidelines and Limitations](#), page 3-4
- [Configuring System-Level High Availability](#), page 3-5
- [Verifying the HA Status](#), page 3-18
- [Additional References](#), page 3-22

Information About System-Level High Availability

This section includes the following topics:

- [Information About Single and Dual Supervisors](#), page 3-1
- [Information About VSM Restarts and Switchovers](#), page 3-3

Information About Single and Dual Supervisors

The Cisco Nexus 1000V can be configured with a single virtual supervisor module (VSM) or dual VSMs. [Table 3-1](#) describes the HA supervisor roles for single and dual VSM operation.

Table 3-1 HA Supervisor Roles

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. They are also 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.

Send document comments to nexus1k-docfeedback@cisco.com.

HA Supervisor Roles

The redundancy role indicates not only whether the VSM interacts with other VSMs, but also the module number it occupies. [Table 3-2](#) shows the available HA roles for VSMs.

Table 3-2 HA Supervisor Roles

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.
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 those described in [Table 3-3](#).

Table 3-3 HA Supervisor Redundancy States

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>

Send document comments to nexus1k-docfeedback@cisco.com.

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, then the active and standby VSMS are not yet synchronized. The following 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

This section includes the following topics:

- [Restarts on Standalone VSMS, page 3-3](#)
- [Restarts on Dual VSMS, page 3-3](#)
- [Switchovers on Dual VSMS, page 3-4](#)

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, however, a switchover may 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.

Send document comments to nexus1k-docfeedback@cisco.com.

Switchovers on Dual VSMs

A dual VSM configuration allows uninterrupted traffic forwarding with 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 in order to provide a seamless and stateful switchover of most services if the active VSM fails.

This section includes the following topics:

- [Switchover Characteristics, page 3-4](#)
- [Automatic Switchovers, page 3-4](#)
- [Manual Switchovers, page 3-4](#)

Switchover Characteristics

A switchover occurs when the active supervisor fails (for example, if repeated failures occur in an essential service or if the system 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 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. To find out if it is, see the [“Verifying that a System is Ready for a Switchover”](#) section on page 3-7.

Once you have verified that the standby VSM is stable, you can manually initiate a switchover (see the [“Manually Switching the Active VSM to Standby”](#) section on page 3-9).

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

Guidelines and Limitations

Follow these guidelines and limitations when configuring system-level high availability:

- Although primary and secondary VSMs can reside in the same host, to improve redundancy, install them in separate hosts and, if possible, connected to different upstream switches.

Send document comments to nexus1k-docfeedback@cisco.com.

- The console for the standby VSM is available through the vSphere client or using the command, **module attach** <x>, but configuration is not allowed and many commands are restricted. The **module attach** <x> command would be run 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.

Configuring System-Level High Availability

This section includes the following topics:

- [Changing the VSM Role, page 3-5](#)
- [Changing the VSM Role, page 3-5](#)
- [Configuring a Switchover, page 3-7](#)
- [Adding a Second VSM to a Standalone System, page 3-11](#)
- [Replacing the Standby in a Dual VSM System, page 3-15](#)
- [Replacing the Active in a Dual VSM System, page 3-16](#)
- [Changing the Domain ID in a Dual VSM System, page 3-17](#)

Changing the VSM Role

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 know or do the following:



Caution

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

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

For an example of changing the port groups and port profiles assigned to the VSM interfaces in the vSphere Client, see the following document:

- *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*

To change a standalone VSM to a secondary VSM, see the [“Adding a Second VSM to a Standalone System”](#) section on page 3-11.

- You are logged into the CLI in EXEC mode.

Send document comments to nexus1k-docfeedback@cisco.com.



Note 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.

- The possible HA roles are standalone, primary, and secondary.
For more information, see the “[HA Supervisor Roles](#)” section on page 3-2.
- The possible HA redundancy states are active and standby.
For more information, see the “[Dual Supervisor Active and Standby Redundancy States](#)” section on page 3-2.
- To activate a change from primary to secondary VSM, you must reload the VSM by doing one of the following:
 - Issue the **reload** command.
 - Power the VM off and then on from the vSphere Client.
- A change from a standalone to a primary VSM takes effect immediately.

SUMMARY STEPS

1. **system redundancy role {standalone | primary | secondary}**
2. **show system redundancy status**
3. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	system redundancy role {standalone primary secondary} Example: n1000v# system redundancy role standalone n1000v#	Designates the HA role of the VSM.
Step 2	show system redundancy status Example: switch# show system redundancy status	(Optional) Displays the current redundancy status for the VSM(s).
Step 3	copy running-config startup-config Example: n1000v(config)# copy running-config startup-config	Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Send document comments to nexus1k-docfeedback@cisco.com.

EXAMPLE

This example shows how to display the system redundancy status of a standalone VSM:

```
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

This section includes the following procedures for configuring a switchover in a dual VSM system:

- [Guidelines and Limitations, page 3-7](#)
- [Verifying that a System is Ready for a Switchover, page 3-7](#)
- [Manually Switching the Active VSM to Standby, page 3-9](#)

Guidelines and Limitations

Follow these guidelines when performing 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 this procedure to verify that both an active and standby VSM are in place and operational before proceeding with a switchover.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged into the CLI in EXEC mode.
- If the standby VSM is not in a stable state (the state must be **ha-standby**), then a manually initiated switchover cannot be performed.

Send document comments to nexus1k-docfeedback@cisco.com.

SUMMARY STEPS

1. `show system redundancy status`
2. `show module`

DETAILED STEPS

	Command	Purpose
Step 1	<p><code>show system redundancy status</code></p> <p>Example: <pre>n1000v# show system redundancy status Redundancy role ----- administrative: primary operational: primary Redundancy mode ----- administrative: HA operational: HA This supervisor (sup-1) ----- Redundancy state: Active Supervisor state: Active Internal state: Active with HA standby Other supervisor (sup-2) ----- Redundancy state: Standby Supervisor state: HA standby Internal state: HA standby</pre></p>	<p>Displays the current redundancy status for the VSM(s). If the output indicates the following, then you can proceed with a system switchover, if needed.</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
Step 2	<p><code>show module</code></p> <pre>n1000v# show module</pre>	<p>Displays information about all available VEMs and VSMs in the system.</p> <p>In the command output, the Status column should display OK for switching modules and an active or ha-standby status for supervisor modules.</p> <p>If the output indicates the following, then you can proceed with a system switchover, if needed:</p> <ul style="list-style-type: none"> • The presence of an active VSM • The presence of a standby VSM in the HA standby redundancy state

Send document comments to nexus1k-docfeedback@cisco.com.

EXAMPLE

This example shows how to display information about all available VEMs and VSMs in the system:

```
n1000v# show module
Mod  Ports  Module-Type                Model                Status
----  -
1    0      Virtual Supervisor Module  Nexus1000V          active *
2    0      Virtual Supervisor Module  Nexus1000V          ha-standby
3    248    Virtual Ethernet Module    NA                   ok

Mod  Sw                Hw
----  -
1    4.0(4)SV1(0.37)  0.0
2    4.0(4)SV1(0.37)  0.0
3    4.0(4)SV1(0.37)  0.4

Mod  MAC-Address(es)                Serial-Num
----  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3    02-00-0c-00-21-00 to 02-00-0c-00-21-80  NA

Mod  Server-IP          Server-UUID                Server-Name
----  -
1    192.168.48.66      NA                          NA
2    192.168.48.66      NA                          NA
3    192.168.48.45      b497bc96-1583-32f1-9062-de3b5d37709c  strider.cisco.com

* this terminal session
```

Manually Switching the Active VSM to Standby

Use this procedure to manually switch an active VSM to standby in a dual supervisor system.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the active VSM CLI in EXEC mode.
- You have completed the steps in the “[Verifying that a System is Ready for a Switchover](#)” section on [page 3-7](#), and have found the system to be ready for a switchover.
- A switchover can only be performed when two VSMs are functioning in the switch.
- If the standby VSM is not in a stable state (ha-standby), then you cannot initiate a manual switchover. You will see the following error message:

```
Failed to switchover (standby not ready to takeover in vdc 1)
```
- Once you enter the **system switchover** command, you cannot start another switchover process on the same system until a stable standby VSM is available.
- If a switchover does not complete successfully within 28 seconds, the supervisors will reset.
- Any unsaved running configuration that was available at active VSM is still unsaved in the new active VSM. You can verify this unsaved running configuration using the **show running-config diff** command. Save that configuration, if needed, as you would do in the other VSM (by entering the **copy running-config startup-config** command).

Send document comments to nexus1k-docfeedback@cisco.com.

SUMMARY STEPS

1. `system switchover`
2. `show running-config diff`
3. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	system switchover Example: n1000v# system switchover	On the active, VSM, initiates a manual switchover to the standby VSM. Note Once you enter this command, you cannot start another switchover process on the same system until a stable standby VSM is available. Note Before proceeding, wait until the switchover completes and the standby supervisor becomes active.
Step 2	show running-config diff Example: n1000v# show running-config diff	(Optional) Verify 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	copy running-config startup-config Example: n1000v(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

EXAMPLES

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.

```
n1000v# 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 )
```


Send document comments to nexus1k-docfeedback@cisco.com.

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

```
n1000v# show running-config diff
*** Startup-config
--- Running-config
*****
*** 1,38 ****
    version 4.0(4)SV1(1)
    role feature-group name new
    role name testrole
    username admin password 5 $1$S7HvKc5G$aguYqHl0dPttBJAhEPwsy1 role network-admin
    telnet server enable
    ip domain-lookup
```

Adding a Second VSM to a Standalone System

Use this section to change a standalone system into a dual supervisor system by adding a second VSM.

This section includes the following topics:

- [Adding a Second VSM to a Standalone System, page 3-11](#)
- [Changing the Standalone VSM to a Primary VSM, page 3-13](#)
- [Verifying the Change to a Dual VSM System, page 3-14](#)

BEFORE YOU BEGIN

Before adding a second VSM to a standalone system, you must know or do the following:

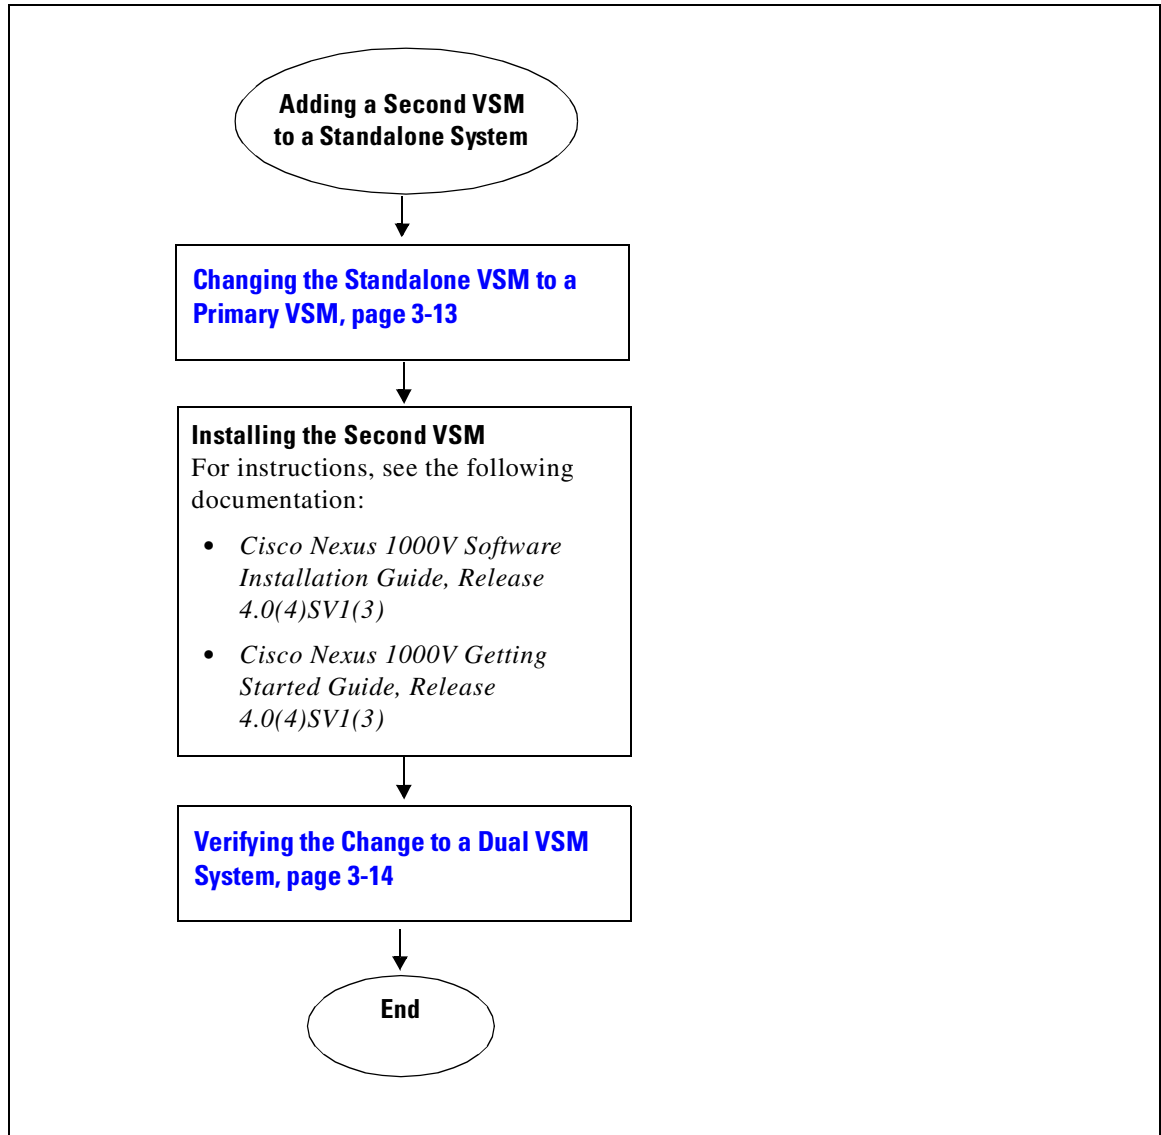
- You are logged into the CLI in EXEC mode.
- You have the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)* document available.
- Although primary and secondary VSMs can reside in the same host, to improve redundancy, install them in separate hosts and, if possible, connected to different upstream switches.
- When installing the second VSM, assign it with the secondary role.
- Set up the port groups for the dual VSM VMs with the same parameters in both hosts.
- 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.

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

Flow Chart: Adding a Second VSM to a Standalone System

The following flow chart (see [Figure 3-1](#)) is designed to guide you through the process of adding a second VSM to a standalone system. After completing each procedure, return to the flow chart to make sure that you complete all required procedures in the correct sequence.

Figure 3-1 Adding a Second VSM to a Standalone System



Send document comments to nexus1k-docfeedback@cisco.com.

Changing the Standalone VSM to a Primary VSM

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

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged into the CLI in EXEC mode.
- A change from a standalone to a primary VSM takes effect immediately.

SUMMARY STEPS

1. `system redundancy role primary`
2. `show system redundancy status`
3. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>system redundancy role primary</code> Example: n1000v# <code>system redundancy role primary</code> n1000v#	Changes the standalone VSM to a primary VSM. The role change occurs immediately.
Step 2	<code>show system redundancy status</code> Example: n1000v# <code>show system redundancy status</code>	Displays the current redundancy state for the VSM.
Step 3	<code>copy running-config startup-config</code> Example: n1000v(config)# <code>copy running-config startup-config</code>	Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

EXAMPLE

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

```
n1000v# 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
```

Send document comments to nexus1k-docfeedback@cisco.com.

```

Internal state:   Active with no standby

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

```

Verifying the Change to a Dual VSM System

Use this procedure to verify a change from a single VSM to a dual VSM system.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged into the CLI in EXEC mode.
- You have already changed the single VSM role from standalone to primary (see the [“Changing the Standalone VSM to a Primary VSM”](#) section on page 3-13).
- You have already installed the second VSM using the following documents:
 - *Cisco Nexus 1000V Software Installation Guide, Release 4.0(4)SV1(3)*
 - *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*

SUMMARY STEPS

1. **show system redundancy status**
2. **show module**

DETAILED STEPS

	Command	Purpose
Step 1	show system redundancy status Example: n1000v# show system redundancy status	Displays the current redundancy status for VSMs in the system.
Step 2	show module Example: n1000v# show module	Displays information about all available VSMs and VEMs in the system.

EXAMPLES

This example shows how to display the current redundancy status for VSMs in the system. In this example, the primary and secondary VSMs are shown following a change from a single VSM system to a dual VSM system.

```

n1000v# show system redundancy status
Redundancy role
-----
administrative: primary
operational: primary

```

Send document comments to nexus1k-docfeedback@cisco.com.

```

Redundancy mode
-----
administrative: HA
operational: HA
This supervisor (sup-1)
-----
Redundancy state: Active
Supervisor state: Active
Internal state: Active with HA standby
Other supervisor (sup-2)
-----
Redundancy state: Standby
Supervisor state: HA standby
Internal state: HA standby

```

This example shows how to display information about all available VSMs and VEMs in the system. In this example, the primary and secondary VSMs are shown following a change from a single VSM system to a dual VSM system. In addition, there is one VEM in module 3.

```

n1000v# show module
Mod  Ports  Module-Type                Model                Status
---  -
1    0      Virtual Supervisor Module  Nexus1000V          active *
2    0      Virtual Supervisor Module  Nexus1000V          ha-standby
3    248    Virtual Ethernet Module    NA                   ok

Mod  Sw                Hw
---  -
1    4.0(4)SV1(0.37)  0.0
2    4.0(4)SV1(0.37)  0.0
3    4.0(4)SV1(0.37)  0.4

Mod  MAC-Address(es)                Serial-Num
---  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3    02-00-0c-00-21-00 to 02-00-0c-00-21-80  NA

Mod  Server-IP          Server-UUID                Server-Name
---  -
1    192.168.48.66      NA                          NA
2    192.168.48.66      NA                          NA
3    192.168.48.45      b497bc96-1583-32f1-9062-de3b5d37709c  strider.cisco.com
* this terminal session

```

Replacing the Standby in a Dual VSM System

Use this procedure to replace a standby/secondary VSM in a dual VSM system.



Note

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

Step 1 Power off the standby VSM.

Step 2 Install the new VSM as a standby, with the same domain ID as the existing VSM, using the procedure in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

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

Once the new VSM is added to the system, it will synchronize with the existing VSM.

Replacing the Active in a Dual VSM System

Use this procedure to replace an active/primary VSM in a dual VSM system.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

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



Note

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

DETAILED STEPS

-
- Step 1** Power off the active VSM.
The secondary VSM becomes active.
- Step 2** On the vSphere Client, change the port group configuration for the new primary VSM to prevent communication with the secondary VSM and the VEMs during setup.
For an example of changing the port groups and port profiles assigned to the VSM interfaces in the vSphere Client, see the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*
- Step 3** Install the new VSM as a primary, with the same domain ID as the existing VSM, using the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.
- Step 4** Save the configuration.
- Step 5** Power off the VM.
- Step 6** On the vSphere Client, change the port group configuration for the new primary VSM to permit communication with the secondary VSM and the VEMs.
- Step 7** Power up the new primary VSM.
The new primary VSM starts and automatically synchronizes all configuration data with the secondary, 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.
-

Send document comments to nexus1k-docfeedback@cisco.com.

Changing the Domain ID in a Dual VSM System

Use this procedure to change the domain ID in a dual VSM system.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You have access to the console of both the active and standby VSM.
- 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. In order to avoid these mechanisms while changing the domain ID, you must isolate the standby VSM from the active VSM. This procedure has a step for isolating the VSMS.



Note

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

DETAILED STEPS

- Step 1** On the vSphere Client for the standby VSM, do one of the following to isolate the VSMS and prevent their communication while completing this procedure:
- Change the port group configuration for the interfaces using port groups that prevent the VSMS from communicating with each other.
 - Unmark the “Connected” option for the interfaces.

The standby VSM becomes active but cannot communicate with the other active VSM or the VEM.

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

Example:

```
n1000v# config t
n1000v(config)# svcs-domain
n1000v(config-svs-domain)# domain id 100
n1000v(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:

```
n1000v# config t
n1000v(config)# svcs-domain
n1000v(config-svs-domain)# domain id 100
n1000v(config-svs-domain)# copy running-config startup-config
```

The domain id is changed on the active VSM and the VEM connected to it.

- Step 5** On the vSphere Client for the standby VSM, do one of the following to permit communication with the active VSM:
- Change the port group configuration for the interfaces.
 - Make sure the "Connect at power on" option is marked for the interfaces.

Once powered up, the standby VSM will be able to communicate with the active VSM.

Send document comments to nexus1k-docfeedback@cisco.com.

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

Verifying the HA Status

Use the following commands to display and verify the HA status of the system.

Command	Purpose
<code>show system redundancy status</code>	Displays the HA status of the system. See Example 3-1 .
<code>show module</code>	Displays information about all available VSMs and VEMs in the system. See Example 3-2 .
<code>show processes</code>	Displays the state of all processes and the process start count. <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 etc) See Example 3-3 .

Example 3-1 `show system redundancy status`

```
n1000v# show system redundancy status
Redundancy role
-----
administrative: primary
operational: primary
Redundancy mode
-----
administrative: HA
operational: HA
This supervisor (sup-1)
-----
Redundancy state: Active
Supervisor state: Active
Internal state: Active with HA standby
Other supervisor (sup-2)
-----
Redundancy state: Standby
Supervisor state: HA standby
Internal state: HA standby
```

Example 3-2 `show module`

```
n1000v# show module
Mod  Ports  Module-Type                Model          Status
---  ---
1    0      Virtual Supervisor Module  Nexus1000V    active *
2    0      Virtual Supervisor Module  Nexus1000V    ha-standby
3    248    Virtual Ethernet Module    NA            ok
```


Send document comments to nexus1k-docfeedback@cisco.com.

```

Mod  Sw                Hw
---  -
1    4.0(4)SV1(0.37)  0.0
2    4.0(4)SV1(0.37)  0.0
3    4.0(4)SV1(0.37)  0.4

Mod  MAC-Address(es)                Serial-Num
---  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
3    02-00-0c-00-21-00 to 02-00-0c-00-21-80  NA

Mod  Server-IP                Server-UUID                Server-Name
---  -
1    192.168.48.66            NA                          NA
2    192.168.48.66            NA                          NA
3    192.168.48.45            b497bc96-1583-32f1-9062-de3b5d37709c  strider.cisco.com
* this terminal session

```

Example 3-3 show processes

```
n1000v# show processes
```

```

PID    State  PC                Start_cnt  TTY  Type  Process
-----
1      S      77f8a468          1          -    O    init
2      S      0                  1          -    O    ksoftirqd/0
3      S      0                  1          -    O    desched/0
4      S      0                  1          -    O    events/0
5      S      0                  1          -    O    khelper
10     S      0                  1          -    O    kthread
18     S      0                  1          -    O    kblockd/0
35     S      0                  1          -    O    khubd
119    S      0                  1          -    O    pdflush
120    S      0                  1          -    O    pdflush
122    S      0                  1          -    O    aio/0
121    S      0                  1          -    O    kswapd0
707    S      0                  1          -    O    kseriod
754    S      0                  1          -    O    kide/0
762    S      0                  1          -    O    scsi_eh_0
1083   S      0                  1          -    O    kjournald
1088   S      0                  1          -    O    kjournald
1603   S      0                  1          -    O    kjournald
1610   S      0                  1          -    O    kjournald
1920   S      77f6c18e          1          -    O    portmap
1933   S      0                  1          -    O    nfsd
1934   S      0                  1          -    O    nfsd
1935   S      0                  1          -    O    nfsd
1936   S      0                  1          -    O    nfsd
1937   S      0                  1          -    O    nfsd
1938   S      0                  1          -    O    nfsd
1939   S      0                  1          -    O    nfsd
1940   S      0                  1          -    O    nfsd
1941   S      0                  1          -    O    lockd
1942   S      0                  1          -    O    rpciod
1947   S      77f6e468          1          -    O    rpc.mountd
1957   S      77f6e468          1          -    O    rpc.statd
1984   S      77dfe468          1          -    VG   sysmgr
2265   S      0                  1          -    O    mping-thread
2266   S      0                  1          -    O    mping-thread
2280   S      0                  1          -    O    redun_kthread
2281   S      0                  1          -    O    redun_timer_kth

```

Send document comments to nexus1k-docfeedback@cisco.com.

```

2341      S      0      1      -      O      stun_kthread
2817      S      0      1      -      O      sf_rdn_kthread
2818      S 77f37468      1      -      VU      xinetd
2819      S 77f6e468      1      -      VU      tftpd
2820      S 7784f1b6      1      -      VL      syslogd
2821      S 77ec2468      1      -      VU      sdwrapd
2822      S 77dbf468      1      -      VU      platform
2830      S      0      1      -      O      ls-notify-mts-t
2842      S 77ea5be4      1      -      VU      pfm_dummy
3270      S 77f836be      1      -      O      klogd
3274      S 77d84be4      1      -      VL      vshd
3275      S 77a41f43      1      -      VL      smm
3276      S 77e41468      1      -      VL      session-mgr
3277      S 77c26468      1      -      VL      psshelper
3278      S 77f75468      1      -      VU      lmgrd
3279      S 77e5cbe4      1      -      VG      licmgr
3280      S 77eb2468      1      -      VG      fs-daemon
3281      S 77eb8468      1      -      VL      feature-mgr
3282      S 77e72468      1      -      VU      confcheck
3283      S 77e9e468      1      -      VU      capability
3284      S 77c26468      1      -      VU      psshelper_gsvc
3294      S 77f75468      1      -      O      cisco
3311      S 77856f43      1      -      VL      clis
3360      S 77cbd468      1      -      VL      xmlma
3361      S 77e5b468      1      -      VL      vmm
3362      S 77b44468      1      -      VG      vdc_mgr
3363      S 77e71468      1      -      VU      ttyd
3364      R 77e9e5f5      1      -      VL      sysinfo
3365      S 77b5a468      1      -      VL      sksd
3366      S 77e9b468      1      -      VG      res_mgr
3367      S 77e44468      1      -      VG      plugin
3368      S 77ccc468      1      -      VL      mvsh
3369      S 77dfc468      1      -      VU      module
3370      S 77ccb468      1      -      VL      evms
3371      S 77ccc468      1      -      VL      evmc
3373      S 77ec1468      1      -      VU      core-dmon
3374      S 7761c40d      1      -      VL      ascii-cfg
3375      S 77cd9be4      1      -      VL      securityd
3376      S 77ca3468      1      -      VU      cert_enroll
3377      S 77b11be4      1      -      VL      aaa
3380      S 77a38f43      1      -      VL      l3vm
3381      S 77a2ef43      1      -      VL      u6rib
3383      S 77a2ef43      1      -      VL      urib
3384      S 77e13468      1      -      VU      ExceptionLog
3385      S 77df0468      1      -      VU      bootvar
3386      S 77dbc468      1      -      VG      ifmgr
3387      S 77ea0468      1      -      VU      tcap
3390      S 77f2abe4      1      -      VU      core-client
3418      S 77a3ff43      1      -      VL      adjmgr
3431      S 77f836be      1      1      O      getty
3432      S 77a7deee      1      S0      O      vsh
3434      S 77f1deee      1      -      O      gettylogin1
3454      S 77a41f43      1      -      VL      arp
3455      S 7786d896      1      -      VL      icmpv6
3456      S 778e1f43      1      -      VL      netstack
3510      S 776c340d      1      -      VL      radius
3511      S 77f58be4      1      -      VL      ip_dummy
3512      S 77f58be4      1      -      VL      ipv6_dummy
3513      S 7780640d      1      -      VU      ntp
3514      S 77f58be4      1      -      VL      pktmgr_dummy
3515      S 7786540d      1      -      VL      snmpd
3517      S 777f540d      1      -      VL      cdp
3706      S 77f836be      1      S1      O      getty
3711      S 77b66468      1      -      VL      aclmgr

```

Send document comments to nexus1k-docfeedback@cisco.com.

```

3718      S 77d18468      1      -      VU  aclcomp
3871      S 778b440d      1      -      VL  ufdm
3872      S 77d08468      1      -      VU  sf_nf_srv
3873      S 779dff43      1      -      VL  rpm
3874      S 7789340d      1      -      VG  pltfm_config
3875      S 77ef4468      1      -      VU  pixmc
3876      S 77dd5468      1      -      VG  pixm
3877      S 7786640d      1      -      VL  nfm
3878      S 77dc9468      1      -      VU  msp
3879      S 77d82468      1      -      VL  monitor
3880      S 7786240d      1      -      VL  mfdm
3881      S 7784140d      1      -      VL  l2fm
3882      S 77d90468      1      -      VL  ipqosmgr
3883      S 77bf8468      1      -      VU  copp
3885      S 75f39497      1      -      VU  vms
3891      S 779ca27b      1      -      VL  igmp
3929      S 77b3d468      1      -      VL  eth_port_channel
3930      S 77cd5468      1      -      VL  vlan_mgr
3934      S 7777e40d      1      -      VL  ethpm
3960      S 77b58468      1      -      VL  eth-port-sec
3961      S 77a93468      1      -      VL  stp
3998      S 77d7f468      1      -      VL  private-vlan
3999      S 77d4e468      1      -      VU  vim
4009      S 77da9468      1      -      VL  lacp
4016      S 77d5d468      1      -      VU  portprofile
4221      S 77f58be4      1      -      VL  tcpudp_dummy
4226      S 77c12468      1      -      VU  pdl_srv_tst
4242      S 77e55468      1      -      VU  ethanalyzer
4243      S 77afb40d      1      -      VL  dcos-thttpd
4244      S 77ad740d      1      -      VL  dcos-xinetd
4261      S 77b0240d      1      -      O   ntpd
4542      S          0      1      -      O   mts-sync-thr
7372      S 77f426be      1      S0    O   more
7373      S 77aa4be4      1      S0    O   vsh
7374      R 77f716be      1      -      O   ps
-         NR      -          0      -      VL  tacacs+
-         NR      -          0      -      VL  eigrp
-         NR      -          0      -      VL  isis
-         NR      -          0      -      VL  ospf
-         NR      -          0      -      VL  ospfv3
-         NR      -          0      -      VL  rip
-         NR      -          0      -      VL  eigrp
-         NR      -          0      -      VL  isis
-         NR      -          0      -      VL  ospf
-         NR      -          0      -      VL  ospfv3
-         NR      -          0      -      VL  rip
-         NR      -          0      -      VL  eigrp
-         NR      -          0      -      VL  isis
-         NR      -          0      -      VL  ospf
-         NR      -          0      -      VL  ospfv3
-         NR      -          0      -      VL  rip
-         NR      -          0      -      VL  eigrp
-         NR      -          0      -      VL  isis
-         NR      -          0      -      VL  ospf
-         NR      -          0      -      VL  ospfv3
-         NR      -          0      -      VL  rip
-         NR      -          0      -      VL  amt
-         NR      -          0      -      VL  bgp
-         NR      -          0      -      VL  eou
-         NR      -          0      -      VL  glbp
-         NR      -          0      -      VL  hsrp_engine
-         NR      -          0      -      VU  installer
-         NR      -          0      -      VL  interface-vlan
-         NR      -          0      -      VU  lisp

```

Send document comments to nexus1k-docfeedback@cisco.com.

```

- NR - 0 - VL msdp
- NR - 0 - VL pim
- NR - 0 - VL pim6
- NR - 0 - VL scheduler
- NR - 0 - VU vbuilder

```

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 etc)

Additional References

For additional information related to implementing system-level HA features, see the following sections:

- [Related Documents, page 3-22](#)
- [Standards, page 3-22](#)
- [MIBs, page 3-23](#)
- [RFCs, page 3-23](#)

Related Documents

Related Topic	Document Title
Software upgrades	<i>Cisco Nexus 1000V Software Upgrade Guide, Release 4.0(4)SV1(3)</i>
Getting Started with Cisco Nexus 1000V and the CLI	<i>Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(3)</i>

Standards

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

Send document comments to nexus1k-docfeedback@cisco.com.

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> 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

RFCs	Title
No RFCs are supported by this feature	—

Feature History for System-Level High Availability

This section provides the System-Level High Availability release history.

Feature Name	Releases	Feature Information
System-Level High Availability	4.0(4)SV1(1)	This feature was introduced.

Send document comments to nexus1k-docfeedback@cisco.com.



INDEX

A

audience [iv-vii](#)
automatic synchronization
 about [3-3](#)

C

conventions [iv-viii](#)

D

documentation
 additional publications [iv-ix](#)
 related documents [iv-viii](#)

F

failure, switchover [3-9](#)

H

HA policy
 description [2-2](#)
 maximum retries [2-2](#)
 minimum lifetime [2-3](#)
high availability
 description [1-1](#)
 displaying status [3-18](#)
 network level [1-3](#)
 supervisor module switchover mechanism [3-4](#)
 switchover characteristics [3-4](#)
 system level [1-3](#)

M

maximum retries. See HA policy
message and transaction service. See MTS
minimum lifetime. See HA policy
MTS
 description [2-2](#)

P

persistent storage service. See PSS
policy. See HA policy
primary role, VSM [3-5](#)
processes
 restartability [2-3](#)
PSS
 description [2-2](#)
 global and local synchronization [2-2](#)
 private and shared [2-2](#)

R

related documents [iv-ix](#)
restart
 stateful, description [2-4](#)
 stateless, description [2-4](#)
restartability
 infrastructure [2-1](#)
 process [2-3](#)
restarts
 troubleshooting [2-5](#)
role, VSM
 primary [3-5](#)

Send document comments to nexus1k-docfeedback@cisco.com.

- secondary 3-5
- standalone 3-5

S

- secondary role, VSM 3-5
- services
 - restartability 2-3
- standalone role, VSM 3-5
- stateful restart
 - description 2-4
- stateless restart
 - description 2-4
- supervisor module
 - role
 - secondary 3-5
- supervisor modules
 - changing the domain ID 3-17
 - replacing standby supervisor 3-15, 3-16
 - role
 - primary 3-5
 - standalone 3-5
 - switchover mechanisms 3-4
- switchovers 3-9
 - characteristics 3-4
 - guidelines 3-7

V

- VSMs
 - manual switchover 3-9