



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



Cisco Nexus 1000V System Management Configuration Guide, Release 4.2(1) SV1(4a)

January 14, 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

Text Part Number: OL-22824-A1

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 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: 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)

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 System Management Configuration Guide, Release 4.2(1) SV1(4a)
© 2009-2014 Cisco Systems, Inc. All rights reserved.



CONTENTS

New and Changed Information xiii

Preface xv

Audience xv

Document Organization xv

Document Conventions xvi

Recommended Reading xvii

xvii

Available Documents xvii

Obtaining Documentation and Submitting a Service Request xix

xix

CHAPTER 1

System Management Overview 1-1

CDP 1-1

Domains 1-1

Server Connections 1-2

Configuration Management 1-2

File Management 1-2

User Management 1-2

NTP 1-2

Local SPAN and ERSPAN 1-2

SNMP 1-3

NetFlow 1-3

System Messages 1-3

iSCSI Multipath 1-3

Troubleshooting 1-3

CHAPTER 2

Configuring CDP 2-1

Information About CDP 2-1

 High Availability 2-2

Guidelines and Limitations 2-2

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

- Defaults 2-2
- Configuring CDP 2-3
 - CDP Global Configuration 2-3
 - CDP Interface Configuration 2-7
- Monitoring CDP 2-10
 - Clearing CDP Statistics 2-10
- Verifying the CDP Configuration 2-10
- Configuration Example for CDP 2-14
- Additional References 2-14
 - Related Documents 2-14
 - Standards 2-14
- Feature History for CDP 2-14

CHAPTER 3

- Configuring the Domain 3-1**
 - Information About the Domain 3-1
 - About Layer 3 Control 3-1
 - Guidelines and Limitations 3-2
 - Default Settings 3-3
 - Configuring the Domain 3-3
 - Creating a Domain 3-4
 - Changing to Layer 3 Transport 3-6
 - Changing to Layer 2 Transport 3-8
 - Creating a Port Profile for Layer 3 Control 3-9
 - Creating a Control VLAN 3-12
 - Creating a Packet VLAN 3-14
 - Feature History for the VSM Domain 3-16

CHAPTER 4

- Managing Server Connections 4-1**
 - Information About Server Connections 4-1
 - Guidelines and Limitations 4-2
 - Connecting to the vCenter Server 4-2
 - Disconnecting From the vCenter Server 4-4
 - Removing the DVS from the vCenter Server 4-5
 - Removing the DVS from the vCenter Server When the VSM Is Not Connected 4-6
 - Configuring the DVS Admin User or DVS Admin Group 4-6
 - Removing the DVS from the vCenter Server With the DVS Admin Account 4-8
 - Configuring Host Mapping 4-8

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

Information about Host Mapping	4-8
Removing Host Mapping from a Module	4-8
Mapping to a New Host	4-9
Viewing Host Mapping	4-11
Verifying Connections	4-11
Verifying the Domain	4-12
Verifying the Configuration	4-12
Verifying Module Information	4-15
Feature History for Server Connections	4-17

CHAPTER 5
Managing the Configuration 5-1

Information About Configuration Management	5-1
Changing the Switch Name	5-1
Configuring a Message of the Day	5-2
Verifying the Configuration	5-3
Verifying the Software and Hardware Versions	5-3
Verifying the Running Configuration	5-4
Comparing the Startup and Running Configurations	5-6
Verifying the Interface Configuration	5-7
Saving a Configuration	5-10
Erasing a Configuration	5-10
Feature History for Configuration Management	5-11

CHAPTER 6
Working with Files 6-1

Information About Files	6-1
Navigating the File System	6-2
Specifying File Systems	6-2
Identifying the Directory You are Working From	6-2
Changing Your Directory	6-3
Listing the Files in a File System	6-4
Identifying Available File Systems for Copying Files	6-4
Using Tab Completion	6-5
Copying and Backing Up Files	6-6
Creating a Directory	6-7
Removing an Existing Directory	6-8
Moving Files	6-8
Deleting Files or Directories	6-9

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

- Compressing Files 6-10
- Uncompressing Files 6-11
- Directing Command Output to a File 6-12
- Verifying a Configuration File before Loading 6-12
- Rolling Back to a Previous Configuration 6-13
- Displaying Files 6-13
 - Displaying File Contents 6-13
 - Displaying Directory Contents 6-14
 - Displaying File Checksums 6-15
 - Displaying the Last Lines in a File 6-15
- Feature History for File Management 6-15

CHAPTER 7

Managing Users 7-1

- Information About User Management 7-1
- Displaying Current User Access 7-1
- Sending a Message to Users 7-2
- Feature History for User Management 7-2

CHAPTER 8

Configuring NTP 8-1

- Information about NTP 8-1
 - High Availability 8-2
- Prerequisites for NTP 8-2
- Configuration Guidelines and Limitations 8-3
- Default Settings 8-3
- Configuring an NTP Server and Peer 8-3
 - Clearing NTP Statistics or Sessions 8-4
- Verifying the NTP Configuration 8-4
- NTP Example Configuration 8-5
- Additional References 8-5
 - Related Documents 8-5
 - Standards 8-5
- Feature History for NTP 8-5

CHAPTER 9

Configuring Local SPAN and ERSPAN 9-1

- Information About SPAN and ERSPAN 9-1
 - SPAN Sources 9-1
 - SPAN Destinations 9-2

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

SPAN Sessions	9-5
SPAN Guidelines and Limitations	9-5
Default Settings	9-6
Configuring SPAN	9-6
Configuring a Local SPAN Session	9-7
Configuring an ERSPAN Port Profile	9-9
Configuring an ERSPAN Session	9-13
Shutting Down a SPAN Session	9-16
Resuming a SPAN Session	9-17
Configuring the Allowable ERSPAN Flow IDs	9-19
Verifying the SPAN Configuration	9-20
Example Configurations	9-20
Example Configuration for a SPAN Session	9-20
Example Configuration for an ERSPAN Session	9-21
Additional References	9-22
Related Documents	9-22
Standards	9-22
Feature History for SPAN and ERSPAN	9-23

CHAPTER 10

Configuring SNMP	10-1
Information About SNMP	10-1
SNMP Functional Overview	10-1
SNMP Notifications	10-2
SNMPv3	10-2
High Availability	10-5
Guidelines and Limitations	10-5
Default Settings	10-5
Configuring SNMP	10-5
Configuring SNMP Users	10-6
Enforcing SNMP Message Encryption	10-7
Creating SNMP Communities	10-8
Configuring SNMP Notification Receivers	10-8
Configuring the Notification Target User	10-9
Enabling SNMP Notifications	10-9
Disabling LinkUp/LinkDown Notifications on an Interface	10-11
Enabling a One-time Authentication for SNMP over TCP	10-11
Assigning the SNMP Switch Contact and Location Information	10-11
Disabling SNMP	10-12

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

- Modifying the AAA Synchronization Time 10-13
- Verifying the SNMP Configuration 10-13
- SNMP Example Configuration 10-13
- Additional References 10-14
 - Related Documents 10-14
 - Standards 10-14
 - MIBs 10-15
- Feature History for SNMP 10-16

CHAPTER 11

- Configuring NetFlow 11-1**
 - Information About NetFlow 11-1
 - What is a Flow 11-2
 - Flow Record Definition 11-2
 - Accessing NetFlow Data 11-5
 - Exporting Flows to the NetFlow Collector Server 11-7
 - What NetFlow Data Looks Like 11-8
 - Network Analysis Module 11-8
 - High Availability 11-8
 - Prerequisites for NetFlow 11-8
 - Configuration Guidelines and Limitations 11-9
 - Default Settings 11-9
 - Enabling the NetFlow Feature 11-10
 - Configuring NetFlow 11-11
 - Defining a Flow Record 11-11
 - Defining a Flow Exporter 11-14
 - Defining a Flow Monitor 11-16
 - Assigning a Flow Monitor to an Interface 11-19
 - Adding a Flow Monitor to a Port Profile 11-20
 - Verifying the NetFlow Configuration 11-21
 - Configuration Example for NetFlow 11-25
 - Additional References 11-26
 - Related Documents 11-26
 - Standards 11-27
 - Feature History for NetFlow 11-27

CHAPTER 12

- Configuring System Message Logging 12-1**
 - Information About System Message Logging 12-1
 - System Message Logging Facilities 12-2

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

Guidelines and Limitations	12-5
Default Settings	12-5
Configuring System Message Logging	12-5
Configuring System Message Logging to Terminal Sessions	12-6
Restoring System Message Logging Defaults for Terminal Sessions	12-7
Configuring System Message Logging for Modules	12-8
Restoring System Message Logging Defaults for Modules	12-9
Configuring System Message Logging for Facilities	12-9
Restoring System Message Logging Defaults for Facilities	12-11
Configuring syslog Servers	12-11
Restoring System Message Logging Defaults for Servers	12-12
Using a UNIX or Linux System to Configure Logging	12-13
Displaying Log Files	12-13
Verifying the System Message Logging Configuration	12-14
System Message Logging Example Configuration	12-18
Additional References	12-18
Related Documents	12-18
Standards	12-18
Feature History for System Message Logging	12-18

CHAPTER 13

Configuring iSCSI Multipath	13-1
Information About iSCSI Multipath	13-1
Overview	13-1
Supported iSCSI Adapters	13-2
iSCSI Multipath Setup on the VMware Switch	13-3
Guidelines and Limitations	13-4
Prerequisites	13-5
Default Settings	13-5
Configuring iSCSI Multipath	13-5
Uplink Pinning and Storage Binding	13-5
Process for Uplink Pinning and Storage Binding	13-6
Creating a Port Profile for a VMkernel NIC	13-6
Creating VMkernel NICs and Attaching the Port Profile	13-8
Manually Pinning the NICs	13-9
Identifying the iSCSI Adapters for the Physical NICs	13-11
Binding the VMkernel NICs to the iSCSI Adapter	13-13
Converting to a Hardware iSCSI Configuration	13-13
Process for Converting to a Hardware iSCSI Configuration	13-14

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

- Removing the Binding to the Software iSCSI Adapter 13-14
- Adding the Hardware NICs to the DVS 13-15
- Changing the VMkernel NIC Access VLAN 13-15
 - Process for Changing the Access VLAN 13-15
 - Changing the Access VLAN 13-16
- Verifying the iSCSI Multipath Configuration 13-18
- Additional References 13-19
 - Related Documents 13-19
 - Standards 13-19
- Feature History for iSCSI Multipath 13-19

CHAPTER 14

Configuring VSM Backup and Recovery 14-1

- Information About VSM Backup and Recovery 14-1
- Guidelines and Limitations 14-1
- Configuring VSM Backup and Recovery 14-2
 - Backing Up the VSM 14-2
 - Recovering the VSM 14-8
- Additional References 14-22
 - Related Documents 14-22
 - Standards 14-23
- Feature History for VSM Backup and Recovery 14-23

CHAPTER 15

Virtualized Workload Mobility (DC to DC vMotion) 15-1

- Information About Virtualized Workload Mobility (DC to DC vMotion) 15-1
 - Stretched Cluster 15-1
 - Split Cluster 15-2
- Prerequisites for Virtualized Workload Mobility (DC to DC vMotion) 15-2
- Guidelines and Limitations 15-2
 - Physical Site Considerations 15-2
 - Handling Inter-Site Link Failures 15-3
 - Handling Additional Distance/Latency Between the VSM and VEM 15-3
- Migrating a VSM 15-3
 - Migrating a VSM Hosted on an ESX or ESXi Host 15-4
- Verifying the Virtualized Workload Mobility (DC to DC vMotion) Configuration 15-4
- Monitoring Virtualized Workload Mobility (DC to DC vMotion) 15-4
- Configuration Limits 15-4
- Feature History for Virtualized Workload Mobility (DC to DC vMotion) 15-5

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

CHAPTER 16 **Configuration Limits** 16-1

INDEX

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



New and Changed Information

This chapter lists new and changed content by release, and where it is located in this document.

Content	Description	Changed in Release	Where Documented
Virtualized Workload Mobility (DC to DC vMotion)	This feature is addressing Cisco Nexus 1000 across two physical data centers.	4.2(1)SV1(4a)	Chapter 15, “Virtualized Workload Mobility (DC to DC vMotion)”
DVS Deletion	Allows for the deletion of the DVS from the vCenter Server when there is no connectivity to the VSMs.	4.2(1)SV1(4a)	“Managing Server Connections”
VSM Backup	Allows for the restoration of VSMs when both VSMs have been deleted in an HA environment.	4.2(1)SV1(4a)	“Configuring VSM Backup and Recovery”
Enable NetFlow feature	You can enable/disable the NetFlow feature.	4.2(1)SV1(4)	“Configuring NetFlow”
Add port profile as Local SPAN source	You can specify a port profile as a source for Local SPAN monitor traffic.	4.2(1)SV1(4)	“Configuring Local SPAN and ERSPAN”
Add port profile as ERSPAN source	You can specify a port profile as a source for ERSPAN monitor traffic.	4.2(1)SV1(4)	“Configuring Local SPAN and ERSPAN”
Hardware iSCSI Multipath	You can use a hardware iSCSI adapter for multipathing.	4.2(1)SV1(4)	“Configuring iSCSI Multipath”
SNMP MIBs added	List of supported MIBs.	4.2(1)SV1(4)	“MIBs”
Network Analysis Module (NAM)	NAM support for NetFlow data sources	4.0(4)SV1(3)	“Configuring NetFlow”
	NAM support for ERSPAN data sources	4.0(4)SV1(3)	“Configuring Local SPAN and ERSPAN”
ERSPAN Type-III header	The ERSPAN Type-III extended format header frame enhances support for network management, intrusion detection, and lawful intercept.	4.0(4)SV1(3)	“Configuring Local SPAN and ERSPAN”
Layer 3 Control	Allows a VSM to be Layer 3 accessible and control hosts that reside in a separate Layer 2 network.	4.0(4)SV1(2)	“Configuring the Domain”

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

Content	Description	Changed in Release	Where Documented
iSCSI Multipath	Allows multiple routes between a server and its storage devices.	4.0(4)SV1(2)	"Configuring iSCSI Multipath"
Recommended Reading	Lists reading recommended before configuring the Cisco Nexus 1000V.	4.0(4)SV1(2)	"Preface"
Configuration Limits	Lists the configuration limits for the Cisco Nexus 1000V.	4.0(4)SV1(2)	"Configuration Limits"



Preface

The System Management Configuration document provides procedures for managing the system, such as configuring system message logging, managing the configuration file, managing server connections, and so forth.

This preface describes the following aspects of this document:

- [Audience, page xv](#)
- [Document Organization, page xv](#)
- [Document Conventions, page xvi](#)
- [Available Documents, page xvii](#)
- [Obtaining Documentation and Submitting a Service Request, page xix](#)

Audience

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

- An understanding of virtualization
- Using VMware tools to configure a virtual switch



Note

Knowledge of the VMware vNetwork Distributed Switch is not required.

Document Organization

This document is organized into the following chapters:

Chapter and Title	Description
Chapter 1, “System Management Overview”	Describes the available system management features.
Chapter 2, “Configuring CDP”	Provides procedures for configuring Cisco Discovery Protocol (CDP) for sending and receive information to and from other connected devices.

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

Chapter and Title	Description
Chapter 3, “Configuring the Domain”	Describes how to configure the Cisco Nexus 1000V domain, including creating the domain and assigning VLANs.
Chapter 4, “Managing Server Connections”	Describes how to create a connection and connect to a server, how to disconnect from a server, and how to view server connections.
Chapter 5, “Managing the Configuration”	Describes how to manage the configuration file.
Chapter 6, “Working with Files”	Describes how to manage files including copying and moving files.
Chapter 7, “Managing Users”	Describes how to manage users on the system including displaying current users and sending messages to users.
Chapter 8, “Configuring NTP”	Provides procedures for configuring Network Time Protocol (NTP) to synchronize timekeeping among a set of distributed time servers and clients. This synchronization allows you to correlate events when you receive system logs and other time-specific events from multiple network devices.
Chapter 9, “Configuring Local SPAN and ERSPAN”	Describes how to configure the Ethernet switched port analyzer (SPAN).
Chapter 10, “Configuring SNMP”	Describes how to configure the SNMP including users, message encryption, notifications, authentication over TCP, and so forth.
Chapter 11, “Configuring NetFlow”	Describes how to configure NetFlow.
Chapter 12, “Configuring System Message Logging”	Describes how to configure system message logging.
Chapter 13, “Configuring iSCSI Multipath”	Describes how to configure iSCSI Multipath to set up multiple routes between a server and its storage devices.
Chapter 14, “Configuring VSM Backup and Recovery”	Describes how to configure the backup and recovery procedures on the Visual Supervisor Module (VSM).
Chapter 15, “Virtualized Workload Mobility (DC to DC vMotion)”	Describes an environment where Cisco Nexus 1000V exists across two data centers.
Chapter 16, “Configuration Limits”	Lists the configuration limits for system management.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
{ }	Elements in braces are required choices.
[]	Elements in square brackets are optional.

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

x y z	Alternative, mutually exclusive elements are 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 the device displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions for notes and cautions:



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.

Recommended Reading

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

Cisco Nexus 1000V Getting Started Guide, Release 4.2(1)SV1(4a)

Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)

Cisco VN-Link: Virtualization-Aware Networking White Paper

Available Documents

This section lists the documents used with the Cisco Nexus 1000V 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 4.2(1)SV1(4a)

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

Cisco Nexus 1000V Release Notes, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Compatibility Information, Release 4.2(1)SVI(4a)

Cisco Nexus 1010 Management Software Release Notes, Release 4.2(1)SP1(3)

Install and Upgrade

Cisco Nexus 1000V Virtual Supervisor Module Software Installation Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Software Upgrade Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V VEM Software Installation and Upgrade Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1010 Virtual Services Appliance Hardware Installation Guide

Cisco Nexus 1010 Software Installation and Upgrade Guide, Release 4.2(1)SP1(3)

Configuration Guides

Cisco Nexus 1000V License Configuration Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Getting Started Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.2(1)SVI(4)

Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.2(1)SVI(4)

Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Quality of Service Configuration Guide, Release 4.2(1)SVI(4)

Cisco Nexus 1000V Security Configuration Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V System Management Configuration Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1010 Software Configuration Guide, Release 4.2(1)SP1(3)

Programming Guide

Cisco Nexus 1000V XML API User Guide, Release 4.2(1)SVI(4)

Reference Guides

Cisco Nexus 1000V Command Reference, Release 4.2(1)SVI(4)

Cisco Nexus 1000V MIB Quick Reference

Cisco Nexus 1010 Command Reference, Release 4.2(1)SP1(3)

Troubleshooting and Alerts

Cisco Nexus 1000V Troubleshooting Guide, Release 4.2(1)SVI(4a)

Cisco Nexus 1000V Password Recovery Guide

Cisco NX-OS System Messages Reference

Virtual Security Gateway Documentation

Cisco Virtual Security Gateway for Nexus 1000V Series Switch

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

Virtual Network Management Center

Cisco Virtual Network Management Center

Network Analysis Module Documentation

Cisco Prime Network Analysis Module Software Documentation Guide, 5.1

Cisco Prime Network Analysis Module (NAM) for Nexus 1010 Installation and Configuration Guide, 5.1

Cisco Prime Network Analysis Module Command Reference Guide 5.1

Cisco Prime Network Analysis Module Software 5.1 Release Notes

Cisco Prime Network Analysis Module Software 5.1 User Guide

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: <http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

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



CHAPTER 1

System Management Overview

This chapter describes the following system management features:

- [CDP, page 1-1](#)
- [Domains, page 1-1](#)
- [Server Connections, page 1-2](#)
- [Configuration Management, page 1-2](#)
- [File Management, page 1-2](#)
- [User Management, page 1-2](#)
- [NTP, page 1-2](#)
- [Local SPAN and ERSPAN, page 1-2](#)
- [SNMP, page 1-3](#)
- [NetFlow, page 1-3](#)
- [System Messages, page 1-3](#)
- [Troubleshooting, page 1-3](#)

CDP

Cisco Discovery Protocol (CDP) runs over the data link layer and is used to advertise information to all attached Cisco devices, and to discover and view information about attached Cisco devices. CDP runs on all Cisco-manufactured equipment.

For more information about CDP, see [Chapter 2, “Configuring CDP.”](#)

Domains

You must create a domain name for Cisco Nexus 1000V and then add control and packet VLANs for communication and management. This process is part of the initial setup of the a Cisco Nexus 1000V when installing the software. If you need to create a domain later, you can do so using the **setup** command or the procedures in [Chapter 3, “Configuring the Domain.”](#)

You can establish Layer 3 Control in your VSM domain so that your VSM is Layer 3 accessible and able to control hosts that reside in a separate Layer 2 network. For more information, see the [“About Layer 3 Control” section on page 3-1.](#)

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

Server Connections

In order to connect to vCenter Server or an ESX server, you must first define the connection in the Cisco Nexus 1000V. [Chapter 4, “Managing Server Connections”](#) describes how to connect and disconnect with vCenter Server and viewing connections.

Configuration Management

The Cisco Nexus 1000V provides you with the capability to change the switch name, configure messages of the day, and display, save, and erase configuration files. For more information about managing the configuration, see [Chapter 5, “Managing the Configuration.”](#)

File Management

Using a single interface, you can manage the file system including:

- Flash memory file systems
- Network file systems (TFTP and FTP)
- Any other endpoint for reading or writing data (such as the running configuration)

For more information about working with files, see [Chapter 6, “Working with Files.”](#)

User Management

You can identify the users currently connected to the device and send a message to either a single user or all users. For more information, see [Chapter 7, “Managing Users.”](#)

NTP

The Network Time Protocol (NTP) synchronizes timekeeping among a set of distributed time servers and clients. This synchronization allows you to correlate events when you receive system logs and other time-specific events from multiple network devices.

For more information about NTP, see [Chapter 8, “Configuring NTP.”](#)

Local SPAN and ERSPAN

The Ethernet switched port analyzer (SPAN) lets you monitor traffic in and out of your device, and duplicate packets from source ports to destination ports.

For information about configuring SPAN, see [Chapter 9, “Configuring Local SPAN and ERSPAN.”](#)

You can also use the Cisco Network Analysis Module (NAM) to monitor ERSPAN data sources for application performance, traffic analysis, and packet header analysis.

To use NAM for monitoring the Cisco Nexus 1000V ERSPAN data sources see the *Cisco Nexus 1010 Network Analysis Module Installation and Configuration Note, 4.2*.

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

SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

For more information about SNMP, see [Chapter 10, “Configuring SNMP.”](#)

NetFlow

NetFlow gives visibility into traffic transiting the virtual switch by characterizing IP traffic based on its source, destination, timing, and application information. This information is used to assess network availability and performance, assist in meeting regulatory requirements (compliance), and help with troubleshooting.

For more information, see [Chapter 11, “Configuring NetFlow.”](#)

You can also use the Cisco Network Analysis Module (NAM) to monitor NetFlow data sources. For more information see the *Cisco Nexus 1010 Network Analysis Module Installation and Configuration Note*, 4.2.

System Messages

You can use system message logging to control the destination and to filter the severity level of messages that system processes generate. You can configure logging to a terminal session, a log file, and syslog servers on remote systems.

System message logging is based on RFC 3164. For more information about the system message format and the messages that the device generates, see the *Cisco NX-OS System Messages Reference*.

For information about configuring system messages, see [Chapter 12, “Configuring System Message Logging.”](#)

iSCSI Multipath

The iSCSI multipath feature sets up multiple routes between a server and its storage devices for maintaining a constant connection and balancing the traffic load.

For more information, see [Configuring iSCSI Multipath, page 13-1](#).

Troubleshooting

Ping and traceroute are among the available troubleshooting tools.

For more information, see the *Cisco Nexus 1000V Troubleshooting Guide, Release 4.2(1)SV1(4a)*.

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



CHAPTER 2

Configuring CDP

This chapter describes how to configure the Cisco Discovery Protocol (CDP), and includes the following sections:

- [Information About CDP, page 2-1](#)
- [Guidelines and Limitations, page 2-2](#)
- [Defaults, page 2-2](#)
- [Configuring CDP, page 2-3](#)
- [Monitoring CDP, page 2-10](#)
- [Verifying the CDP Configuration, page 2-10](#)
- [Configuration Example for CDP, page 2-14](#)
- [Additional References, page 2-14](#)

Information About CDP

Cisco Discovery Protocol (CDP) runs over the data link layer and is used to advertise information to all attached Cisco devices, and to discover and view information about attached Cisco devices. CDP runs on all Cisco-manufactured equipment.

CDP gathers protocol addresses of neighboring devices and discovers the platform of those devices. CDP runs over the data link layer only. Two systems that support different Layer 3 protocols can learn about each other.

Each device you configure for CDP sends periodic advertisements to a multicast address. Each device advertises at least one address at which it can receive SNMP messages. The advertisements also contain hold-time information, which indicates the length of time that a receiving device should hold CDP information before discarding it. You can configure the advertisement or refresh timer and the hold timer.

CDP Version 2 (CDPv2) allows you to track instances where the native VLAN ID or port duplex states do not match between connecting devices.

CDP advertises the following type-length-value fields (TLVs):

- Device ID
- Address
- Port ID
- Capabilities

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

- Version
- Platform
- Native VLAN
- Full/Half Duplex
- MTU
- SysName
- SysObjectID
- Management Address
- Physical Location

All CDP packets include a VLAN ID. The CDP packet is untagged, so it goes over the native/access VLAN, which is then also added to the packet.

For more information on VLANs, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.2(1)SV1(4)*.

High Availability

Stateless restarts are supported for CDP. After a reboot or a supervisor switchover, the running configuration is applied.

Guidelines and Limitations

CDP has the following configuration guidelines and limitations:

- CDP can discover up to 256 neighbors per port if the port is connected to a hub with 256 connections.
- CDP must be enabled globally before you can configure CDP on an interface. CDP is enabled globally by default, but can be disabled using the [“Enabling or Disabling CDP Globally” procedure on page 2-3](#).
- You can configure CDP on physical interfaces and port channels only.

Defaults

[Table 2-1](#) lists the CDP default settings.

Table 2-1 CDP Defaults

Parameters	Default
CDP	Enabled globally and on all interfaces
CDP version	Version 2
CDP device ID	System name
CDP timer	60 seconds
CDP hold timer	180 seconds

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

Configuring CDP

This section includes the following topics:

- [CDP Global Configuration, page 2-3](#)
- [Enabling CDP on an Interface, page 2-7](#)
- [Disabling CDP on an Interface, page 2-8](#)

CDP Global Configuration

This section includes the following topics:

- [Enabling or Disabling CDP Globally, page 2-3](#)
- [Advertising a CDP Version, page 2-4](#)
- [Configuring CDP Options, page 2-5](#)

Enabling or Disabling CDP Globally

Use this procedure to enable or disable CDP globally. Although CDP is enabled globally by default, should it be disabled, you can use this procedure to enable it again.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- CDP must be enabled globally before you can configure it on an interface.
- When you globally disable the CDP feature, all CDP configurations are removed.

SUMMARY STEPS

1. `config t`
2. `[no] cdp enable`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: <code>n1000v# config t</code> <code>n1000v(config)#</code>	Places you in the CLI Global Configuration mode.

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

	Command	Purpose
Step 2	<pre>[no] cdp enable</pre> <p>Example: n1000v(config)# cdp enable</p> <p>Example: n1000v(config)# no cdp enable</p>	Enables or disables the CDP feature globally.

Advertising a CDP Version

Use this procedure to designate the CDP version to advertise on the device.

BEFORE YOU BEGIN

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

- You know the version of CDP currently supported on the device.
- Only one version of CDP (version 1 or version 2) is advertised at a time for all uplinks and port channels on the switch.
- For more information about CDP, see the [“Information About CDP”](#) section on page 2-1.

SUMMARY STEPS

1. **config t**
2. **cdp advertise {v1 | v2}**
3. (Optional) **show cdp global**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	<pre>config t</pre> <p>Example: n1000v# config t n1000v(config)#</p>	Places you in the CLI Global Configuration mode.
Step 2	<pre>cdp advertise {v1 v2}</pre> <p>Example 1: n1000v(config)# cdp advertise v1 n1000v(config)#</p> <p>Example 2: n1000v(config)# cdp advertise v2 n1000v(config)#</p>	Assigns the CPD version to advertise. <ul style="list-style-type: none"> • CDP Version 1 • CDP Version 2
Step 3	<pre>show cdp global</pre>	(Optional) Displays the CDP configuration, indicating the CDP version that is being advertised or sent to other devices.

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

Command	Purpose
<p>Example 1: n1000v(config)# show cdp global Global CDP information: CDP enabled globally Sending CDP packets every 60 seconds Sending a holdtime value of 180 seconds Sending CDPv2 advertisements is disabled Sending DeviceID TLV in Default Format</p> <p>Example 2: n1000v(config)# show cdp global Global CDP information: CDP enabled globally Sending CDP packets every 60 seconds Sending a holdtime value of 180 seconds Sending CDPv2 advertisements is enabled Sending DeviceID TLV in Default Format</p>	
<p>Step 4 copy running-config startup-config</p> <p>Example: n1000v(config)# copy running-config startup-config</p>	<p>(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.</p>

Configuring CDP Options

Use this procedure to configure the following for CDP:

- the device ID format to use



Note Only the **system-name** device ID format is supported.

- the maximum hold time for neighbor information
- the refresh time for sending advertisements

BEFORE YOU BEGIN

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

- You can view output from upstream cat6k switch using the **show cdp neighbor** command.
- If you are setting the holdtime, you know how long you want CDP to retain neighbor information.
- If you are setting the CDP timer, you know how often you want CDP to advertise.
- For more information about CDP, see the [“Information About CDP” section on page 2-1](#).

SUMMARY STEPS

- config t**
- (Optional) **cdp format device-id system-name**
- show cdp neighbors** from the upstream device
- show cdp neighbors** from your device
- (Optional) **cdp timer** *seconds*

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

6. (Optional) **cdp holdtime** *seconds*
7. (Optional) **show cdp global**
8. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose																																				
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.																																				
Step 2	cdp format device-id system-name Example: n1000v(config)# cdp format device-id system-name n1000v(config)#	(Optional) Specifies that CDP uses the system name for the device ID format.																																				
Step 3	show cdp neighbors Example: swordfish-6k-2#show cdp neighbors Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone <table border="1"> <thead> <tr> <th>Device ID</th> <th>Local Intrfce</th> <th>Holdtme</th> <th>Capability</th> <th>Platform</th> <th>Port ID</th> </tr> </thead> <tbody> <tr> <td>02000c000000</td> <td>Gig 1/16</td> <td>14</td> <td>S</td> <td>Soft Swit</td> <td>Eth 2/4</td> </tr> <tr> <td>02000c000000</td> <td>Gig 1/17</td> <td>14</td> <td>S</td> <td>Soft Swit</td> <td>Eth 2/5</td> </tr> <tr> <td>02000c000000</td> <td>Gig 1/14</td> <td>14</td> <td>S</td> <td>Soft Swit</td> <td>Eth 2/2</td> </tr> <tr> <td>02000c000000</td> <td>Gig 1/15</td> <td>14</td> <td>S</td> <td>Soft Swit</td> <td>Eth 2/3</td> </tr> <tr> <td>02000c000000</td> <td>Gig 1/18</td> <td>13</td> <td>S</td> <td>Soft Swit</td> <td></td> </tr> </tbody> </table>	Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID	02000c000000	Gig 1/16	14	S	Soft Swit	Eth 2/4	02000c000000	Gig 1/17	14	S	Soft Swit	Eth 2/5	02000c000000	Gig 1/14	14	S	Soft Swit	Eth 2/2	02000c000000	Gig 1/15	14	S	Soft Swit	Eth 2/3	02000c000000	Gig 1/18	13	S	Soft Swit		Displays your device from the upstream device.
Device ID	Local Intrfce	Holdtme	Capability	Platform	Port ID																																	
02000c000000	Gig 1/16	14	S	Soft Swit	Eth 2/4																																	
02000c000000	Gig 1/17	14	S	Soft Swit	Eth 2/5																																	
02000c000000	Gig 1/14	14	S	Soft Swit	Eth 2/2																																	
02000c000000	Gig 1/15	14	S	Soft Swit	Eth 2/3																																	
02000c000000	Gig 1/18	13	S	Soft Swit																																		
Step 4	show cdp neighbors Example: n1000v(config)# show cdp neighbors Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge S - Switch, H - Host, I - IGMP, r - Repeater, V - VoIP-Phone, D - Remotely-Managed-Device, s - Supports-STP-Dispute <table border="1"> <thead> <tr> <th>Device ID</th> <th>Local Intrfce</th> <th>Hldtme</th> <th>Capability</th> <th>Platform</th> <th>Port ID</th> </tr> </thead> <tbody> <tr> <td>swordfish-6k-2</td> <td>Eth2/2</td> <td>169</td> <td>R S I</td> <td>WS-C6503-E</td> <td>Gig1/14</td> </tr> <tr> <td>swordfish-6k-2</td> <td>Eth2/3</td> <td>139</td> <td>R S I</td> <td>WS-C6503-E</td> <td>Gig1/15</td> </tr> <tr> <td>swordfish-6k-2</td> <td>Eth2/4</td> <td>135</td> <td>R S I</td> <td>WS-C6503-E</td> <td>Gig1/16</td> </tr> <tr> <td>swordfish-6k-2</td> <td>Eth2/5</td> <td>177</td> <td>R S I</td> <td>WS-C6503-E</td> <td>Gig1/17</td> </tr> <tr> <td>swordfish-6k-2</td> <td>Eth2/6</td> <td>141</td> <td>R S I</td> <td>WS-C6503-E</td> <td>Gig1/18</td> </tr> </tbody> </table>	Device ID	Local Intrfce	Hldtme	Capability	Platform	Port ID	swordfish-6k-2	Eth2/2	169	R S I	WS-C6503-E	Gig1/14	swordfish-6k-2	Eth2/3	139	R S I	WS-C6503-E	Gig1/15	swordfish-6k-2	Eth2/4	135	R S I	WS-C6503-E	Gig1/16	swordfish-6k-2	Eth2/5	177	R S I	WS-C6503-E	Gig1/17	swordfish-6k-2	Eth2/6	141	R S I	WS-C6503-E	Gig1/18	Displays the upstream device from your device,
Device ID	Local Intrfce	Hldtme	Capability	Platform	Port ID																																	
swordfish-6k-2	Eth2/2	169	R S I	WS-C6503-E	Gig1/14																																	
swordfish-6k-2	Eth2/3	139	R S I	WS-C6503-E	Gig1/15																																	
swordfish-6k-2	Eth2/4	135	R S I	WS-C6503-E	Gig1/16																																	
swordfish-6k-2	Eth2/5	177	R S I	WS-C6503-E	Gig1/17																																	
swordfish-6k-2	Eth2/6	141	R S I	WS-C6503-E	Gig1/18																																	
Step 5	cdp holdtime <i>seconds</i> Example: n1000v(config)# cdp holdtime 10	(Optional) Sets the maximum amount of time that CDP holds onto neighbor information before discarding it. <ul style="list-style-type: none"> • The range is from 10 to 255 seconds. • The default is 180 seconds. 																																				

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

	Command	Purpose
Step 6	cdp timer <i>seconds</i> Example: n1000v(config)# cdp timer 5	(Optional) Sets the refresh time for CDP to send advertisements to neighbors. <ul style="list-style-type: none"> • The range is from 5 to 254 seconds. • The default is 60 seconds.
Step 7	show cdp global Example: n1000v(config)# show cdp global Global CDP information: CDP enabled globally Sending CDP packets every 5 seconds Sending a holdtime value of 10 seconds Sending CDPv2 advertisements is disabled Sending DeviceID TLV in Mac Address Format	Displays the global CDP configuration.
Step 8	copy running-config startup-config Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

CDP Interface Configuration

This section includes the following procedures:

- [Enabling CDP on an Interface, page 2-7](#)
- [Disabling CDP on an Interface, page 2-8](#)

Enabling CDP on an Interface

Use this procedure to enable CDP on a specific interface. Although CDP is enabled by default on all interfaces, should it become disabled, you can use this procedure to enable it again.

BEFORE YOU BEGIN

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

- The CDP feature is enabled globally. CDP is enabled globally by default, but can also be re-enabled using the “[Enabling or Disabling CDP Globally](#)” procedure on page 2-3.
- For more information about CDP, see the “[Information About CDP](#)” section on page 2-1.

SUMMARY STEPS

1. **config t**
2. **interface** *interface-type number*
3. **no cdp enable**
4. **cdp enable**
5. **show cdp interface** *interface-type number*
6. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	interface <i>interface-type number</i> Example: n1000v(config)# interface port-channel 2 n1000v(config-if)#	Places you in the CLI Interface Configuration mode for the specific interface.
Step 3	no cdp enable Example: n1000v(config-if)# no cdp enable	Disables CDP on this interface.
Step 4	cdp enable Example: n1000v(config-if)# cdp enable	Enables CDP on this interface.
Step 5	show cdp interface <i>interface-type number</i> Example: n1000v(config-if)# show cdp interface mgmt0 mgmt0 is up CDP disabled on interface Sending CDP packets every 60 seconds Holdtime is 180 seconds	(Optional) Displays CDP information for the specified interface.
Step 6	copy running-config startup-config Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

This example shows how to enable CDP on port channel 2:

```
n1000v# config t
n1000v(config)# interface port-channel 2
n1000v(config-if)# no cdp enable
n1000v(config-if)# cdp enable
n1000v(config-if)# copy running-config startup-config
```

Disabling CDP on an Interface

Use this procedure to disable CDP on a specific interface.

BEFORE YOU BEGIN

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

- CDP is currently enabled on the device.



Note If CDP is disabled on the device, then it is also disabled for all interfaces.

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

- CDP is currently enabled on the specific interface you want to configure.
- For more information about CDP, see the “[Information About CDP](#)” section on page 2-1.

SUMMARY STEPS

1. **config t**
2. **interface** *interface-type number*
3. **no cdp enable**
4. (Optional) **show cdp interface** *interface-type number*
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	interface <i>interface-type number</i> Example: n1000v(config)# interface mgmt0 n1000v(config-if)#	Places you in the CLI Interface Configuration mode for the specified interface.
Step 3	no cdp enable Example: n1000v(config-if)# no cdp enable	Disables CDP on the specified interface.
Step 4	show cdp interface <i>interface-type number</i> Example: n1000v(config-if)# show cdp interface mgmt0	(Optional) Displays CDP information for an interface.
Step 5	copy running-config startup-config Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

This example shows how to disable CDP on mgmt0:

```
n1000v# config t
n1000v(config)# interface mgmt0
n1000v(config-if)# no cdp enable
n1000v(config-if)# show cdp interface mgmt0
mgmt0 is up
    CDP disabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
n1000v(config-if)# copy running-config startup-config
```

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

Monitoring CDP

To monitor CDP traffic, use the following command:

Command	Purpose
<code>show cdp traffic interface <i>interface-type slot/port</i></code>	Displays the CDP traffic statistics on an interface. See Example 2-7 on page 2-13

Clearing CDP Statistics

To clear CDP statistics, use one of the following commands.

Command	Purpose
<code>clear cdp counters</code>	Clears CDP statistics on all interfaces.
<code>clear cdp counters interface <i>number</i></code>	Clears CDP statistics on the specified interface.
<code>clear cdp table</code>	Clears the CDP cache for one or all interfaces.

Verifying the CDP Configuration

To verify the CDP configuration, use one of the following commands:

Command	Purpose
<code>show cdp all</code>	Displays all interfaces that have CDP enabled. See Example 2-1 on page 2-10
<code>show cdp entry {all name <i>entry-name</i>}</code>	Displays the CDP database entries. See Example 2-2 on page 2-11
<code>show cdp global</code>	Displays the CDP global parameters. See Example 2-4 on page 2-13
<code>show cdp interface <i>interface-type slot/port</i></code>	Displays the CDP interface status. See Example 2-5 on page 2-13
<code>show cdp neighbors {detail interface <i>interface-type slot/port</i>}</code>	Displays the CDP neighbor status. See Example 2-6 on page 2-13

Example 2-1 show cdp all

```
n1000v# show cdp all
Ethernet2/2 is up
  CDP enabled on interface
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
Ethernet2/3 is up
  CDP enabled on interface
```

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

```

    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
Ethernet2/4 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
Ethernet2/5 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
Ethernet2/6 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds
mgmt0 is up
    CDP enabled on interface
    Sending CDP packets every 60 seconds
    Holdtime is 180 seconds

```

Example 2-2 show cdp entry name

```

n1000v# show cdp entry name swordfish-6k-2
-----
Device ID:swordfish-6k-2
System Name:
Interface address(es):
    IPv4 Address: 172.28.30.2
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/2, Port ID (outgoing port): GigabitEthernet1/14
Holdtime: 152 sec

Version:
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team

```

Example 2-3 show cdp entry all

```

n1000v# show cdp entry all
-----
Device ID:swordfish-6k-2
System Name:
Interface address(es):
    IPv4 Address: 172.28.30.2
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/2, Port ID (outgoing port): GigabitEthernet1/14
Holdtime: 140 sec

Version:
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team

Advertisement Version: 1
-----

```

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

```

Device ID: swordfish-6k-2
System Name:
Interface address(es):
  IPv4 Address: 172.28.30.2
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/3, Port ID (outgoing port): GigabitEthernet1/15
Holdtime: 129 sec

Version:
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team

Advertisement Version: 1

-----
Device ID: swordfish-6k-2
System Name:
Interface address(es):
  IPv4 Address: 7.7.8.1
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/4, Port ID (outgoing port): GigabitEthernet1/16
Holdtime: 154 sec

Version:
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team

Advertisement Version: 1

-----
Device ID: swordfish-6k-2
System Name:
Interface address(es):
  IPv4 Address: 7.7.8.1
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/5, Port ID (outgoing port): GigabitEthernet1/17
Holdtime: 156 sec

Version:
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team

Advertisement Version: 1

-----
Device ID: swordfish-6k-2
System Name:
Interface address(es):
  IPv4 Address: 172.28.15.229
Platform: cisco WS-C6503-E, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet2/6, Port ID (outgoing port): GigabitEthernet1/18
Holdtime: 171 sec

Version:

```

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

```
Cisco IOS Software, s72033_rp Software (s72033_rp-IPBASE-M), Version 12.2(33)SXH2a,
RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2008 by Cisco Systems, Inc.
Compiled Fri 25-Apr-08 09:11 by prod_rel_team
```

Advertisement Version: 1

Example 2-4 show cdp global

```
n1000v(config)# show cdp global
Global CDP information:
  CDP enabled globally
  Sending CDP packets every 60 seconds
  Sending a holdtime value of 180 seconds
  Sending CDPv2 advertisements is disabled
  Sending DeviceID TLV in Default Format
```

Example 2-5 show cdp interface

```
n1000v(config)# show cdp interface ethernet 2/3
Ethernet2/3 is up
  CDP enabled on interface
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
```

Example 2-6 show cdp neighbors interface

```
n1000v(config)# show cdp neighbors interface ethernet 2/3
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
```

Device ID	Local Intrfce	Hldtme	Capability	Platform	Port ID
swordfish-6k-2	Eth2/3	173	R S I	WS-C6503-E	Gig1/15

Example 2-7 show cdp traffic interface

```
n1000v(config)# show cdp traffic interface ethernet 2/3
```

Traffic statistics for Ethernet2/3

Input Statistics:

```
Total Packets: 98
Valid CDP Packets: 49
  CDP v1 Packets: 49
  CDP v2 Packets: 0
Invalid CDP Packets: 49
  Unsupported Version: 49
  Checksum Errors: 0
  Malformed Packets: 0
```

Output Statistics:

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

```
Total Packets: 47
  CDP v1 Packets: 47
  CDP v2 Packets: 0
Send Errors: 0
```

Configuration Example for CDP

This example enables the CDP feature and configures the refresh and hold timers:

```
config t
cdp enable
cdp timer 50
cdp holdtime 100
```

Additional References

This section includes the following additional information related to CDP:

- [Related Documents, page 2-14](#)
- [Standards, page 2-14](#)

Related Documents

Related Topic	Document Title
VLAN	<i>Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.2(1)SV1(4)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</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.	—

Feature History for CDP

This section provides the CDP feature release history.

Feature Name	Releases	Feature Information
CDP	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 3

Configuring the Domain

This chapter describes how to configure the Cisco Nexus 1000V domain, including creating the domain, assigning VLANs, configuring Layer 3 Control, and so forth.

This chapter includes the following topics:

- [Information About the Domain, page 3-1](#)
- [Guidelines and Limitations, page 3-2](#)
- [Default Settings, page 3-3](#)
- [Configuring the Domain, page 3-3](#)
- [Feature History for the VSM Domain, page 3-16](#)

Information About the Domain

You must create a domain name for Cisco Nexus 1000V and then add control and packet VLANs for communication and management. This process is part of the initial setup of the a Cisco Nexus 1000V when installing the software. If you need to create a domain later, you can do so using the **setup** command or the procedures described in this chapter.

About Layer 3 Control

Layer 3 control, or IP connectivity, is supported between the VSM and VEM for control and packet traffic. With Layer 3 control, a VSM can be Layer 3 accessible and control hosts that reside in a separate Layer 2 network. All hosts controlled by a VSM, however, must still reside in the same Layer 2 network. Since a VSM cannot control a host that is outside of the Layer 2 network it controls, the host on which it resides must be controlled by another VSM.

To implement Layer 3 control, you must make the following configurations:


- Configure the VSM domain transport mode as Layer 3.
For more information, see the [“Changing to Layer 3 Transport” procedure on page 3-6](#)
- Configure a port profile using the [“Creating a Port Profile for Layer 3 Control” procedure on page 3-9](#).
- Create an VMware kernel NIC interface on each host and apply the Layer 3 control port profile to it. For more information, see your VMware documentation.

[Figure 3-1](#) illustrates the following example of Layer 3 control.

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

- VSM0 controls VEM_0_1.
- VEM_0_1, in turn, hosts VSM1 and VSM2.
- VSM1 and VSM2 control VEMs in other Layer 2 networks.

Figure 3-1 Example of Layer 3 Control IP Connectivity



196104

Guidelines and Limitations

The VSM domain has the following configuration guidelines and limitations:

- UDP port 4785 is required for Layer 3 communication between the VSM and VEM. If you have a firewall in your network, and are configuring Layer 3 control, then make sure UDP port 4785 is open on your upstream switch or firewall device. For more information, see the documentation for your upstream switch or firewall device.
- In a Layer 2 network, you can switch between the Layer 2 and Layer 3 transport modes, but when you do so, the modules may be out of service briefly.
- The capability attribute (Layer 3 control) cannot be inherited from the port profile.
- Different hosts can use different VLANs for Layer 3 control.

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

- A port profile used for Layer 3 control must be an access port profile. It cannot be a trunk port profile.
- We recommend that if you are using the VMware kernel NIC for Layer 3 Control, you do not use it for any other purpose. For example, do not also use the Layer 3 Control VMware kernel NIC for VMotion or NFS mount.
- Control VLANs, packet VLANs, and management VLANs must be configured as regular VLANs and not as private VLANs.
- If you have a firewall in your network, ensure that TCP ports 80 and 443 are open for traffic destined to the vCenter Server and TCP port 80 is open for traffic destined to the Cisco Nexus 1000V Virtual Supervisor Module (VSM).

Default Settings

Table 3-1 lists the default settings in the domain configuration.

Table 3-1 Domain Defaults

Parameter	Default
Control VLAN (svs-domain)	VLAN 1
Packet VLAN (svs-domain)	VLAN 1
VMware port group name (port-profile)	The name of the port profile
SVS mode (svs-domain)	Layer 2
Switchport mode (port-profile)	Access
State (port-profile)	Disabled
State (VLAN)	Active
Shut state (VLAN)	No shutdown

Configuring the Domain

This section includes the following procedures:

- [Creating a Domain, page 3-4](#)
- [Changing to Layer 3 Transport, page 3-6](#)
- [Changing to Layer 2 Transport, page 3-8](#)
- [Creating a Port Profile for Layer 3 Control, page 3-9](#)
- [Creating a Control VLAN, page 3-12](#)
- [Creating a Packet VLAN, page 3-14](#)

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

Creating a Domain

Use this procedure to create a domain name for the Cisco Nexus 1000V that identifies the VSM and VEMs; and then add control and packet VLANs for communication and management. This process is part of the initial setup of the Cisco Nexus 1000V when installing the software. If you need to create a domain after initial setup, you can do so using this procedure.

BEFORE YOU BEGIN

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

- If two or more VSMS share the same control and/or packet VLAN, the domain helps identify the VEMs managed by each VSM.
- You are logged in to the CLI in EXEC mode.
- You must have a unique domain ID for this Cisco Nexus 1000V instance.
- You must identify the VLANs to be used for control and packet traffic.
- We recommend using one VLAN for control traffic and a different VLAN for packet traffic.
- We recommend using a distinct VLAN for each instances of Cisco Nexus 1000V (different domains)
- The **svs mode** command in the SVS Domain Configuration mode is not used and has no effect on a configuration.
- For information about changing a domain ID after adding a second VSM see the *Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.2(1)SV1(4)*.

SUMMARY STEPS

1. **config t**
2. **svs-domain**
3. **domain id** *domain-id*
4. **control vlan** *vlan-id*
5. **packet vlan** *vlan-id*
6. **exit**
7. **show svs domain**
8. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	svs-domain Example: n1000v(config)# svs-domain n1000v(config-svs-domain)#	Places you into the SVS Domain Configuration mode.

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

	Command	Purpose
Step 3	domain id <i>number</i> Example: n1000v(config-svs-domain)# domain id 100 n1000v(config-svs-domain)#	Creates the domain ID for this Cisco Nexus 1000V instance.
Step 4	control vlan <i>number</i> Example: n1000v(config-svs-domain)# control vlan 190 n1000v(config-vlan)#	Assigns the control VLAN for this domain.
Step 5	packet vlan <i>number</i> Example: n1000v(config-vlan)# packet vlan 191 n1000v(config-vlan)#	Assigns the packet VLAN for this domain.
Step 6	show svcs domain Example: n1000v(config-vlan)# show svcs domain	Displays the domain configuration.
Step 7	exit Example: n1000v(config-vlan)# exit n1000v(config)#	Returns you to CLI Global Configuration mode.
Step 8	copy running-config startup-config Example: n1000v(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

Example:
n1000v# config t
n1000v(config)# svcs-domain
n1000v(config-svs-domain)# domain id 100
n1000v(config-svs-domain)# control vlan 190
n1000v(config-svs-domain)# packet vlan 191
n1000v(config-vlan)# exit

n1000v (config)# show svcs domain
SVS domain config:
  Domain id: 100
  Control vlan: 190
  Packet vlan: 191
  L2/L3 Aipc mode: L2
  L2/L3 Aipc interface: mgmt0
  Status: Config push to VC successful.

n1000v(config)#
n1000v(config)# copy run start
[#####] 100%
n1000v(config)#

```

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

Changing to Layer 3 Transport

Use this procedure to change the transport mode from Layer 2 to Layer 3 for the VSM domain control and packet traffic.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- This procedure requires you to disable the control and packet VLANs. You cannot change to Layer 3 Control before disabling the control and packet VLANs.
- You have already configured the Layer 3 interface (mgmt 0 or control 0) and assigned an IP address.
- When control 0 is used for Layer 3 transport, proxy-arp must be enabled on the control 0 VLAN gateway router.

For information about configuring an interface, see the *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)*.

SUMMARY STEPS

1. **show svcs domain**
2. **config t**
3. **svcs-domain**
4. **no control vlan**
5. **no packet vlan**
6. **show svcs domain**
7. **svcs mode L2 | svcs mode L3 interface { mgmt0 | control0 }**
8. **show svcs domain**
9. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	show svcs domain Example: n1000v(config)# show svcs domain SVS domain config: Domain id: 100 Control vlan: 100 Packet vlan: 101 L2/L3 Control mode: L2 L3 control interface: NA Status: Config push to VC successful.	Displays the existing domain configuration, including control and packet VLAN IDs.
Step 2	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.

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

	Command	Purpose
Step 3	svs-domain Example: n1000v(config)# svs-domain n1000v(config-svs-domain)#	Places you in the CLI SVS Domain Configuration mode.
Step 4	no packet vlan Example: n1000v(config-svs-domain)# no packet vlan n1000v(config-svs-domain)#	Removes the packet VLAN configuration.
Step 5	no control vlan Example: n1000v(config-svs-domain)# no control vlan n1000v(config-svs-domain)#	Removes the control VLAN configuration.
Step 6	show svs domain Example: n1000v(config)# show svs domain SVS domain config: Domain id: 100 Control vlan: 1 Packet vlan: 1 L2/L3 Control mode: L2 L2/L3 Control interface: NA Status: Config push to VC successful. switch(config-svs-domain)#	Displays the existing domain configuration, with the default control and packet VLAN IDs.
Step 7	svs mode L3 interface { mgmt0 control0 } Example: n1000v(config-svs-domain)# svs mode l3 interface mgmt0 n000v(config-svs-domain)#	<p>Configures Layer 3 transport mode for the VSM domain.</p> <p>If configuring Layer 3 transport, then you must designate which interface to use; and the interface must already have an IP address configured.</p> <p>This example shows how to configure Layer 3 transport over the management 0 interface.</p>
Step 8	show svs domain Example: SVS domain config: Domain id: 100 Control vlan: 1 Packet vlan: 1 L2/L3 Control mode: L3 L3 control interface: mgmt0 Status: Config push to VC successful. n1000v(config-svs-domain)#	(Optional) Displays the new Layer 3 control mode configuration for this VSM domain.
Step 9	copy running-config startup-config Example: n1000v(config-svs-domain)# copy running-config startup-config [#####] 100% n1000v(config-svs-domain)#	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

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

Changing to Layer 2 Transport

Use this procedure to change the transport mode to Layer 2 for the VSM domain control and packet traffic. The transport mode is Layer 2 by default, but if it is changed, you can use this procedure to configure it again as Layer 2.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- This procedure requires you to configure a control VLAN and a packet VLAN. You cannot configure these VLANs if the VSM domain capability is Layer 3 Control. You will first change the capability to Layer 3 Control, and then configure the control VLAN and packet VLAN.

SUMMARY STEPS

1. **show svcs domain**
2. **config t**
3. **svcs-domain**
4. **svcs mode L2 | svcs mode L3 interface { mgmt0 | control0 }**
5. **show svcs domain**
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	show svcs domain Example: SVS domain config: Domain id: 100 Control vlan: 1 Packet vlan: 1 L2/L3 Control mode: L3 L3 control interface: mgmt0 Status: Config push to VC successful. n1000v(config-svs-domain)#	Displays the existing domain configuration, including control and packet VLAN IDs and the Layer 3 interface configuration.
Step 2	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 3	svcs-domain Example: n1000v(config)# svcs-domain n1000v(config-svs-domain)#	Places you in the CLI SVS Domain Configuration mode.

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

	Command	Purpose
Step 4	svs mode L2 Example: n1000v(config-svs-domain)# svs mode l2 n1000v(config-svs-domain)#	Configures Layer 2 transport mode for the VSM domain.
Step 5	control vlan vlanID Example: n1000v(config-svs-domain)# control vlan 100	Configures the specified VLAN ID as the control VLAN for the VSM domain.
Step 6	packet vlan vlanID Example: n1000v(config-svs-domain)# packet vlan 101	Configures the specified VLAN ID as the packet VLAN for the VSM domain.
Step 7	show svs domain Example: SVS domain config: Domain id: 100 Control vlan: 100 Packet vlan: 101 L2/L3 Control mode: L2 L3 control interface: NA Status: Config push to VC successful. n1000v(config-svs-domain)#	(Optional) Displays the new Layer 2 control mode configuration for this VSM domain.
Step 8	copy running-config startup-config Example: n1000v(config-svs-domain)# copy running-config startup-config [#####] 100% n1000v(config-svs-domain)#	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Creating a Port Profile for Layer 3 Control

Use this procedure to allow the VSM and VEM to communicate over IP for control and packet traffic.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- The transport mode for the VSM domain has already been configured as Layer 3. For more information, see the [“Changing to Layer 2 Transport” procedure on page 3-8](#).
- All VEMs must belong to the same Layer 2 domain.
- The VEM VM kernel NIC must connect to this Layer 3 control port profile when adding the host to the Cisco Nexus 1000V DVS.
- Only one VM kernel NIC can be assigned to this Layer 3 control port profile per host.
- You know the VLAN ID for the VLAN you are adding to this Layer 3 control port profile.
 - The VLAN must already be created on the Cisco Nexus 1000V.
 - The VLAN assigned to this Layer 3 control port profile must be a system VLAN.

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

- One of the uplink ports must already have this VLAN in its system VLAN range.
- The port profile must be an access port profile. It cannot be a trunk port profile. This procedure includes steps to configure the port profile as an access port profile.
- More than one port profile can be configured as **capability L3 control**.
- Different hosts can use different VLANs for Layer 3 control.

SUMMARY STEPS

1. **config t**
2. **port-profile** *name*
3. **capability l3control**
4. **vmware port-group** [*name*]
5. **switchport mode access**
6. **switchport access vlan** *vlanID*
7. **no shutdown**
8. **system vlan** *vlanID*
9. **state enabled**
10. (Optional) **show port-profile** *name*
11. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	port-profile <i>name</i> Example: n1000v(config)# port-profile l3control-150 n1000v(config-port-prof)#	Creates a port profile and places you into Port Profile Configuration mode for the named port profile. The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V.
Step 3	capability l3control Example: n1000v(config-port-prof)# capability l3control n1000v(config-port-prof)#	Allows the port to be used for IP connectivity. In vCenter Server, the Layer 3 control port profile must be selected and assigned to the VM kernel NIC physical port.

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

	Command	Purpose
Step 4	vmware port-group [<i>name</i>] Example: n1000v(config-port-prof)# vmware port-group n1000v(config-port-prof)#	Designates the port-profile as a VMware port group. The port profile is mapped to a VMware port group of the same name. When a vCenter Server connection is established, the port group created in Cisco Nexus 1000V is then distributed to the virtual switch on the vCenter Server. name: Port group name. If you do not specify a name, then the port group name will be the same as the port profile name. If you want to map the port profile to a different port group name, use the alternate name.
Step 5	switchport mode access] Example: n1000v(config-port-prof)# switchport mode access n1000v(config-port-prof)#	Designates that the interfaces are switch access ports (the default).
Step 6	switchport access vlan <i>vlanID</i> Example: n1000v(config-port-prof)# switchport access vlan 150 n1000v(config-port-prof)#	Assigns the system VLAN ID to the access port for this Layer 3 control port profile.
Step 7	no shutdown Example: n1000v(config-port-prof)# no shutdown n1000v(config-port-prof)#	Administratively enables all ports in the profile.
Step 8	system vlan <i>vlanID</i> Example: n1000v(config-port-prof)# system vlan 150 n1000v(config-port-prof)#	Adds the system VLAN to this Layer 3 control port profile. This ensures that, when the host is added for the first time or rebooted later, the VEM will be able to reach the VSM. One of the uplink ports must have this VLAN in its system VLAN range.
Step 9	state enabled Example: n1000v(config-port-prof)# state enabled n1000v(config-port-prof)#	Enables the Layer 3 control port profile. The configuration for this port profile is applied to the assigned ports, and the port group is created in the VMware vSwitch on the vCenter Server.

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

	Command	Purpose
Step 10	<p>show port-profile name <i>name</i></p> <p>Example:</p> <pre>n1000v(config-port-prof)# show port-profile name l3control-150 port-profile l3control-150 description: type: vethernet status: enabled capability l3control: yes pinning control-vlan: 8 pinning packet-vlan: 8 system vlans: 150 port-group: l3control-150 max ports: 32 inherit: config attributes: switchport mode access switchport access vlan 150 no shutdown evaluated config attributes: switchport mode access switchport access vlan 150 no shutdown assigned interfaces: n1000v(config-port-prof)#</pre>	(Optional) Displays the current configuration for the port profile.
Step 11	<p>copy running-config startup-config</p> <p>Example:</p> <pre>n1000v(config-port-prof)# copy running-config startup-config</pre>	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Creating a Control VLAN

Use this procedure to add a control VLAN to the domain.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- If Layer 3 Control is configured on your VSM, you can not create a control VLAN. You must first disable Layer 3 Control.
- You have already configured and enabled the required switched virtual interface (SVI) using the document, *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SVI(4a)* The SVI is also called the VLAN interface and provides communication between VLANs.
- You are familiar with how VLANs are numbered. For more information, see the document, *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.2(1)SVI(4)*.
- Newly-created VLANs remain unused until Layer 2 ports are assigned to them.

SUMMARY STEPS

1. **config t**

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

2. `vlan vlan-id`
3. `name vlan-name`
4. `state vlan-state`
5. `exit`
6. `show vlan id vlan-id`
7. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: n1000v# <code>config t</code> n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	<code>vlan 30</code> Example: n1000v(config)# <code>vlan 30</code> n1000v(config-vlan)#	Creates VLAN ID 30 for control traffic and places you into CLI VLAN Configuration mode. Note If you enter a VLAN ID that is assigned to an internally allocated VLAN, the CLI returns an error message.
Step 3	<code>name cp_control</code> Example: n1000v(config-vlan)# <code>name cp_control</code> n1000v(config-vlan)#	Adds the descriptive name, <code>cp_control</code> , to this VLAN.
Step 4	<code>state active</code> Example: n1000v(config-vlan)# <code>state active</code> n1000v(config-vlan)#	Changes the operational state of the VLAN to active.
Step 5	<code>show vlan id 30</code> Example: n1000v(config-vlan)# <code>show vlan id 30</code>	Displays the configuration for VLAN ID 30.
Step 6	<code>copy running-config startup-config</code> Example: n1000v(config-vlan)# <code>copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

```

Example:
n1000v# config t
n1000v(config)# vlan 30
n1000v(config-vlan)# name cp_control
n1000v(config-vlan)# state active
n1000v(config)# show vlan id 30

```

```

VLAN Name                Status    Ports
-----
30    cp_control            active

```

```

VLAN Type MTU
-----

```

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

```

5      enet 1500

Remote SPAN VLAN
-----
Disabled

Primary  Secondary  Type           Ports
-----  -
n1000v(config)# copy run start
[#####] 100%
n1000v(config)#

```

Creating a Packet VLAN

Use this procedure to add the packet VLAN to the domain.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You have already configured and enabled the required switched virtual interface (SVI) using the document, *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SVI(4a)*. The SVI is also called the VLAN interface and provides communication between VLANs.
- You are familiar with how VLANs are numbered. For more information, see the document, *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.2(1)SVI(4)*.
- Newly-created VLANs remain unused until Layer 2 ports are assigned to them.

SUMMARY STEPS

1. **config t**
2. **vlan *vlan-id***
3. **name *vlan-name***
4. **state *vlan-state***
5. **exit**
6. **show vlan id *vlan-id***
7. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: n1000v# <code>config t</code> n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	<code>vlan 31</code> Example: n1000v(config)# <code>vlan 31</code> n1000v(config-vlan)#	Creates VLAN ID 31 for packet traffic and places you into CLI VLAN Configuration mode. Note If you enter a VLAN ID that is assigned to an internally allocated VLAN, the CLI returns an error message.
Step 3	<code>name cp_packet</code> Example: n1000v(config-vlan)# <code>name cp_packet</code> n1000v(config-vlan)#	Adds the descriptive name, <code>cp_packet</code> , to this VLAN.
Step 4	<code>state active</code> Example: n1000v(config-vlan)# <code>state active</code> n1000v(config-vlan)#	Changes the operational state of the VLAN to active.
Step 5	<code>show vlan id 31</code> Example: n1000v(config-vlan)# <code>show vlan id 30</code>	Displays the configuration for VLAN ID 31.
Step 6	<code>exit</code> Example: n1000v(config-vlan)# <code>exit</code> n1000v(config)#	Returns you to CLI Global Configuration mode.
Step 7	<code>copy running-config startup-config</code> Example: n1000v(config)# <code>copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

Example:
n1000v# `config t`
n1000v(config)# `vlan 31`
n1000v(config-vlan)# `name cp_packet`
n1000v(config-vlan)# `state active`
n1000v(config-vlan)# `exit`
n1000v(config)# `show vlan id 31`

```

VLAN Name                Status    Ports
-----
31    cp_packet              active
VLAN Type MTU
----
5    enet 1500
Remote SPAN VLAN
-----

```

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

Disabled

Primary Secondary Type Ports

```
-----
n1000v(config)# copy run start
[#####] 100%
n1000v(config)#
```

Feature History for the VSM Domain

This section provides the VSM domain feature release history.

Feature Name	Releases	Feature Information
Layer 3 Control	4.0(4)SV1(2)	Added the following information: <ul style="list-style-type: none"> • About Layer 3 Control, page 3-1 • Guidelines and Limitations, page 3-2 • Changing to Layer 2 Transport, page 3-8 • Changing to Layer 3 Transport, page 3-6 • Creating a Port Profile for Layer 3 Control, page 3-9
VSM Domain	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 4

Managing Server Connections

This chapter describes how to create a connection and connect to a server, how to disconnect from a server, and how to view server connections.

This chapter includes the following topics:

- [Information About Server Connections, page 4-1](#)
- [Guidelines and Limitations, page 4-2](#)
- [Connecting to the vCenter Server, page 4-2](#)
- [Disconnecting From the vCenter Server, page 4-4](#)
- [Removing the DVS from the vCenter Server, page 4-5](#)
- [Removing the DVS from the vCenter Server When the VSM Is Not Connected, page 4-6](#)
- [Configuring Host Mapping, page 4-8](#)
- [Verifying Connections, page 4-11](#)
- [Verifying the Domain, page 4-12](#)
- [Verifying the Configuration, page 4-12](#)
- [Verifying Module Information, page 4-15](#)
- [Feature History for Server Connections, page 4-17](#)

Information About Server Connections

In order to connect to vCenter Server or an ESX server, you must first define the connection in the Cisco Nexus 1000V including the following:

- A connection name
- The protocol used
- The server IP address
- The server DNS name
- The datacenter name

All communication with vCenter Server is secured by the TLS protocol.

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

Guidelines and Limitations

Server connections have the following configuration guidelines and limitations:

- A single VSM can only connect to one vCenter server at a time. A single VSM cannot connect to multiple vCenter servers at once.

Connecting to the vCenter Server

Use this procedure to configure a connection and then connect to vCenter server or an ESX server.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You know the datacenter name.
- The vCenter Server management station is installed and running.
- You know the vCenter Server IP address or hostname.
- The ESX servers are installed and running.
- The management port is configured.
- The vCenter Server is reachable from the Cisco Nexus 1000V.
- The Cisco Nexus 1000V appliance is installed.
- If you are configuring a connection using a hostname, DNS is already configured.
- You have already registered an extension with the vCenter Server. The extension includes the extension key and public certificate for the VSM. vCenter Server uses these to verify the authenticity of the request it receives from the VSM. For instructions about adding and registering an extension, see the *Cisco Nexus 1000V Software Installation Guide, Release 4.2(1)SV1(4a)*.

SUMMARY STEPS

1. **config t**
2. **svs connection *name***
3. **protocol vmware-vim**
4. **remote {ip address *address A.B.C.D* | hostname *name*}**
5. **vmware dvs datacenter-name *name***
6. **connect**

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

DETAILED STEPS

	Command	Description
Step 1	<p>config t</p> <p>Example: n1000v# config t n1000v(config)#</p>	Places you into global configuration mode.
Step 2	<p>svs connection name</p> <p>Example: n1000v (config#) svs connection VC n1000v(config-svs-conn#)</p>	Places you into connection configuration mode for adding this connection between Cisco Nexus 1000V and either a particular ESX server or the vCenter Server. By using a name, information for multiple connections can be stored in the configuration.
Step 3	<p>protocol vmware-vim [http]</p> <p>Example: n1000v(config-svs-conn#) protocol vmware-vim n1000v(config-svs-conn#)</p>	<p>Specifies that this connection uses the VIM protocol. This command is stored locally.</p> <ul style="list-style-type: none"> • http: Specifies that the VIM protocol runs over HTTP. The default is to use HTTP over SSL (HTTPS).
Step 4	<p>Do one of the following:</p> <ul style="list-style-type: none"> • If you are configuring an IP address, go to Step 5. • If you are configuring a hostname, go to Step 6. 	
Step 5	<p>remote ip address ipaddress</p> <p>Example: n1000v(config-svs-conn#) remote ip address 192.168.0.1 n1000v(config-svs-conn#)</p> <p>Go to Step 7.</p>	Specifies the IP address of the ESX server or vCenter Server for this connection. This command is stored locally.
Step 6	<p>remote hostname hostname</p> <p>Example: n1000v(config-svs-conn#) remote hostname vcMain n1000v(config-svs-conn#)</p>	<p>Specifies the DNS name of the ESX server or vCenter Server for this connection. This command is stored locally.</p> <p>Note DNS is already configured.</p>

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

	Command	Description
Step 7	<pre>vmware dvs datacenter-name [folder/] name</pre> <p>Example:</p> <pre>n1000v(config-svs-conn#) vmware dvs datacenter-name Hamilton-DC n1000v(config-svs-conn#)</pre>	<p>Identifies the datacenter name in the vCenter Server where Cisco Nexus 1000V is to be created as a distributed virtual switch (DVS). You can use this command before or after connecting. The datacenter name is stored locally.</p> <p>Note The Nexus 1000V folder name should be same in the vCenter Server and in the VSM. If the Nexus 1000V folder is renamed in the vCenter Server, it must also be renamed in the VSM.</p>
Step 8	<pre>connect</pre> <p>Example:</p> <pre>n1000v(config-svs-conn#) connect</pre>	<p>Initiates the connection. If the username and password have not been configured for this connection, the user is prompted for a username and password.</p> <p>The default is no connect. There can be only one active connection at a time. If a previously-defined connection is up, an error message displays and the command is rejected until the user closes the previous connection by entering no connect.</p>

Examples

```
n1000v# config t
n1000v (config)# svcs connection VC
n1000v (config-svs-conn#) protocol vmware-vim
n1000v (config-svs-conn#) remote ip address 192.168.0.1
n1000v (config-svs-conn#) vmware dvs datacenter-name Hamilton-DC
n1000v (config-svs-conn#) connect

n1000v# show svcs connections
connection VC:
  ip address: 192.168.0.1
  protocol: vmware-vim https
  certificate: default
  datacenter name: Hamilton-DC
  DVS uuid: ac 36 07 50 42 88 e9 ab-03 fe 4f dd d1 30 cc 5c
  config status: Enabled
  operational status: Connected
n1000v#
```

Disconnecting From the vCenter Server

Use this procedure to disconnect from the vCenter Server, for example, after correcting a vCenter Server configuration.

BEFORE YOU BEGIN

- You are logged in to the Cisco Nexus 1000V in EXEC mode.
- You have configured an Cisco Nexus 1000V connection using the [“Connecting to the vCenter Server” procedure on page 4-2](#).

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

- The Cisco Nexus 1000V is connected to vCenter Server/ESX.

DETAILED STEPS

	Command	Description
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into global configuration mode.
Step 2	svs connection name Example: n1000v (config#) svs connection vcWest n1000v(config-svs-conn) #	Places you into a global configuration submode for the connection to vCenter Server.
Step 3	no connect Example: n1000v(config-svs-conn) # no connect n1000v(config-svs-conn) #	Closes the connection.

Removing the DVS from the vCenter Server

Use this procedure to remove the DVS from the vCenter Server.



Note

If you do not have connectivity to the VSM, see the [“Removing the DVS from the vCenter Server When the VSM Is Not Connected”](#) section on page 4-6.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You have configured a connection using the [“Connecting to the vCenter Server”](#) procedure on page 4-2.
- The Cisco Nexus 1000V is connected to vCenter Server/ESX.
- The Server Administrator has already removed from the VI client all of the hosts connected to Cisco Nexus 1000V. For more information, see the VMware documentation.

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

DETAILED STEPS

	Command	Description
Step 1	<pre>config t</pre> <p>Example: <pre>n1000v# config t n1000v(config)#</pre></p>	Places you into global configuration mode.
Step 2	<pre>svs connection name</pre> <p>Example: <pre>n1000v(config#) svs connection vcWest n1000v(config-svs-conn)#</pre></p>	Places you into a global configuration submode for the connection to vCenter Server.
Step 3	<pre>no vmware dvs</pre> <p>Example: <pre>n1000v(config-svs-conn)# no vmware dvs n1000v(config-svs-conn)#</pre></p>	Removes the DVS associated with the specified connection from the vCenter Server.

Removing the DVS from the vCenter Server When the VSM Is Not Connected

Use this procedure to remove the DVS from the vCenter Server when the VSM does not have connectivity to the vCenter Server.

Configuring the ability to delete the DVS when the VSM is not connected to the vCenter Server is a two-step process:

1. Configure the admin user or group. See the [“Configuring the DVS Admin User or DVS Admin Group” procedure on page 4-6](#).
2. Remove the DVS from the vCenter Server. See the [“Removing the DVS from the vCenter Server With the DVS Admin Account” procedure on page 4-8](#).

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You have logged in the vCenter Server.
- The admin user or group account has been configured on the vCenter Server.

Configuring the DVS Admin User or DVS Admin Group

Use this procedure to configure a DVS admin user or DVS admin group.

BEFORE YOU BEGIN

- Ensure that the System Administrator has created an admin user or admin group on the vCenter Server to manage and delete the DVS. This user should not be given any other permissions like deploying VMs or hosts, etc. The admin user name configured on the VSM should be the same as the user name on the vCenter Server.

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

Summary Steps

1. `config t`
2. `show svcs connections`
3. `svcs connection name`
4. `admin {user username | group groupname}`
5. `show svcs connections`

Detailed Steps

Step 1 Determine the name of the DVS.

```
switch# show svcs connections

connection VC:
  ipaddress: 10.104.63.16
  remote port: 80
  protocol: VMware-vim https
  certificate: default
  datacenter name: N1K-DC
  admin:
    DVS uuid: a2 ...
    config status: Enabled
    operational status: Connected
    sync status: Complete
    version: VMware vCenter Server 4.1.0 build 258902
```

Step 2 Configure the admin user in the vCenter Server.

```
switch# config t
switch(config)# svcs connection VC
switch(config-svcs-conn) # admin user NAuser
switch(config-svcs-conn) #
```



Note You can also configure an admin group by entering the `admin group groupname` command.

Step 3 Verify that the admin user has been created.

```
switch# show svcs connections

connection VC:
  ipaddress: 10.104.63.16
  remote port: 80
  protocol: VMware-vim https
  certificate: default
  datacenter name: N1K-DC
  admin: NAuser(user)
  DVS uuid: a2 ...
  config status: Enabled
  operational status: Connected
  sync status: Complete
  version: VMware vCenter Server 4.1.0 build 258902
```

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

Removing the DVS from the vCenter Server With the DVS Admin Account

Use this procedure to remove the DVS from the vCenter Server.

-
- Step 1** Log in to the vCenter Server through the VMware vSphere Client with the DVS admin account that was configured in “[Configuring the DVS Admin User or DVS Admin Group](#)” procedure on page 4-6.
- Step 2** In the vSphere Client left pane, choose the data center.
- Step 3** Click **Hosts and Clusters > Networking**.
- Step 4** Right-click the DVS and choose **Remove**.
-

Configuring Host Mapping

This section includes the following topics:

- [Information about Host Mapping, page 4-8](#)
- [Removing Host Mapping from a Module, page 4-8](#)
- [Mapping to a New Host, page 4-9](#)
- [Viewing Host Mapping, page 4-11](#)

Information about Host Mapping

When a VSM detects a new VEM, it automatically assigns a free module number to the VEM and then maintains the mapping between the module number and UUID of a host server. This mapping is used to assign the same module number to a given host server.

Removing Host Mapping from a Module

Use this procedure to remove the mapping of a module to a host server.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You have already removed the host from the Cisco Nexus 1000V DVS on vCenter.

SUMMARY STEPS

1. **config t**
2. **no vem *module-number***
3. **show module vem mapping**
4. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Description
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	no vem <i>module-number</i> Example: n1000v(config)# no vem 4 n1000v(config)# no vem 3 cannot modify slot 3: host module is inserted n1000v((config)#	Removes the specified module from software. Note If the module is still present in the slot, the command is rejected, as shown in this example.
Step 3	show module vem mapping Example: n1000v(config)# show module vem mapping	(Optional) Displays the mapping of modules to host servers.
Step 4	copy running-config startup-config Example: n1000v(config-vem-slot)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Example

This example shows the VEM mapping.

```
n1000v(config)# show module vem mapping
Mod      Status          UUID                                     License Status
---      -
      3      powered-up      93312881-309e-11db-afa1-0015170f51a8      licensed
n1000v(config)#
```

Mapping to a New Host

Use this procedure to map a module number to a different host server UUID.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You have already removed the host from the Cisco Nexus 1000V DVS on vCenter using the [“Removing Host Mapping from a Module” procedure on page 4-8](#).



Note If you do not first remove the existing host server mapping, the new host server is assigned a different module number.

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

SUMMARY STEPS

1. **config t**
2. **vem module *number***
3. **host vmware id *server-bios-uuid***
4. **show module vem mapping**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Description
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	vem module <i>number</i> Example: n1000v(config)# vem 3 n1000v((config-vem-slot)#	Places you into CLI VEM Slot Configuration mode.
Step 3	host vmware id <i>server-bios-uuid</i> Example: n1000v(config-vem-slot)# host vmware id 6dd6c3e3-7379-11db-abcd-000bab086eb6 n1000v(config-vem-slot)#	Assigns a different host server UUID to the specified module.
Step 4	show module vem mapping Example: n1000v(config-vem-slot)# show module vem mapping	(Optional) Displays the mapping of modules to host servers.
Step 5	copy running-config startup-config Example: n1000v(config-vem-slot)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Example

This example shows the VEM mapping.

```
n1000v(config-vem-slot)# show module vem mapping
Mod      Status      UUID                                     License Status
---      -
3        powered-up  93312881-309e-11db-afa1-0015170f51a8  licensed
4        absent     6dd6c3e3-7379-11db-abcd-000bab086eb6  licensed

n1000v(config-vem-slot)#
```


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

Viewing Host Mapping

Use this procedure in EXEC mode to view the mapping of modules to host servers.

Summary Steps

1. `show module vem mapping`

Detailed Steps

Step 1 Display the mapping on modules to host servers by entering the following command:

```
n1000v(config)# show module vem mapping
Mod      Status          UUID                               License Status
-----  -
      3      powered-up      93312881-309e-11db-afa1-0015170f51a8  licensed
n1000v(config)#
```

Verifying Connections

Use this procedure to view and verify connections.

BEFORE YOU BEGIN

- You are logged in to the CLI in any command mode.
- You have configured the connection using the [“Connecting to the vCenter Server” procedure on page 4-2](#).
- The Cisco Nexus 1000V is connected to vCenter Server/ESX.

Summary Steps

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

Detailed Steps

	Command	Description
Step 1	<pre>show svcs connections [name]</pre> <p>Example:</p> <pre>n1000v# show svcs connections vc Connection vc: IP address: 172.28.15.206 Protocol: vmware-vim https vmware dvs datacenter-name: HamiltonDC ConfigStatus: Enabled OperStatus: Connected n1000v#</pre>	<p>Displays the current connections to the Cisco Nexus 1000V.</p> <p>Note Network connectivity issues may shut down your connection to the vCenter Server. When network connectivity is restored, the Cisco Nexus 1000V will not automatically restore the connection. In this case, you must restore the connection manually using the following command sequence,</p> <pre>no connect connect</pre>

Verifying the Domain

Use this procedure to view and verify the configured domain.

BEFORE YOU BEGIN

- You are logged in to the CLI in any command mode.
- You have configured a domain using the [“Creating a Domain” procedure on page 3-4](#).

DETAILED STEPS

	Command	Description
Step 1	<pre>show svcs domain</pre> <p>Example:</p> <pre>n1000v# show svcs domain SVS domain config: Domain id: 98 Control vlan: 70 Packet vlan: 71 Sync state: - n1000v#</pre>	<p>Display the domain configured on the Cisco Nexus 1000V.</p>

Verifying the Configuration

Use this procedure to display and verify the running configuration.

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

BEFORE YOU BEGIN

- You are logged in to the CLI in any command mode.
- You have configured Cisco Nexus 1000V connections using the “Connecting to the vCenter Server” procedure on page 4-2.
- The Cisco Nexus 1000V is connected to vCenter Server/ESX.

DETAILED STEPS

	Command	Description
Step 1	<code>show running-config</code>	Display the current configuration. If the Cisco Nexus 1000V is not connected to a vCenter Server or ESX server, the output is limited to connection-related information.

Example:

```
n1000v(config-acl)# show running-config
version 4.0(4)SV1(1)
feature port-security
username adminbackup password 5 $1$0ip/C5Ci$0Odx7oJS1BCFpNRmQK4na. role network-operator
username admin password 5 $1$NlmX5tLD$daXpuxlAPcIHoz53PBhy6/ role network-admin
telnet server enable
ssh key rsa 1024 force
kernel core target 0.0.0.0
kernel core limit 1
system default switchport
ip access-list my66
  10 permit ip 1.1.1.1/32 1.1.1.2/32
snmp-server user admin network-admin auth md5 0x90f3798f3e894496a11ec42ce2efec9c priv
0x90f3798f3e894496a11ec42ce2efec9c localizedkey
snmp-server enable traps entity fru
snmp-server enable traps license
vrf context management
  ip route 0.0.0.0/0 172.28.15.1
switchname srini-cp
vlan 40-43,45-48
vdc srini-cp id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource monitor-session minimum 0 maximum 32
  limit-resource vrf minimum 16 maximum 8192
  limit-resource port-channel minimum 0 maximum 192
  limit-resource u4route-mem minimum 32 maximum 256
  limit-resource u6route-mem minimum 16 maximum 256

interface Ethernet6/2
  inherit port-profile uplinkportprofile1

interface Ethernet6/3
  inherit port-profile uplinkportprofile2

interface Ethernet6/4
  inherit port-profile uplinportprofile3

interface Ethernet7/2
  inherit port-profile uplinkportprofile1

interface mgmt0
```

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

```

ip address 172.28.15.163/24

interface Vethernet1

    inherit port-profile vm100

interface Vethernet2

    inherit port-profile vm100

interface Vethernet3

    inherit port-profile vm100

interface Vethernet4

    inherit port-profile vm100

interface Vethernet5

interface Vethernet6
boot kickstart bootflash:/svs-kickstart-mzg.4.0.1a.S1.0.82.bin sup-1
boot system bootflash:/svs-mzg.4.0.1a.S1.0.82.bin sup-1
boot system bootflash:/isan.bin sup-1
boot kickstart bootflash:/svs-kickstart-mzg.4.0.1a.S1.0.82.bin sup-2
boot system bootflash:/svs-mzg.4.0.1a.S1.0.82.bin sup-2
boot system bootflash:/isan.bin sup-2
ip route 0.0.0.0/0 172.28.15.1
port-profile uplinkportprofile1
    capability uplink
    vmware port-group
    switchport mode trunk
    switchport trunk allowed vlan 1,40-43
    no shutdown
    system vlan 1,40-43
    state enabled
port-profile vm100
    vmware port-group
    switchport mode access
    switchport access vlan 43
    ip port access-group my100 out
    ip port access-group my66 in
    no shutdown
    state enabled
port-profile uplinkportprofile2
    capability uplink
    vmware port-group
    switchport mode trunk
    switchport trunk allowed vlan 45-46
    no shutdown
    state enabled
port-profile uplinportprofile3
    capability uplink
    vmware port-group
    switchport trunk allowed vlan 47-48
    state enabled
port-profile uplinkportprofile3
    no shutdown
svs-domain
    domain id 163

```

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

```
control vlan 41
packet vlan 42
svs connection VCR5
protocol vmware-vim
remote ip address 172.28.30.83
vmware dvs datacenter-name cisco-DC
connect
n1000v(config-acl)#
```

Verifying Module Information

Use this procedure to display and verify module information, including a view of the DVS from Cisco Nexus 1000V.

BEFORE YOU BEGIN

- You are logged in to the CLI in any command mode.
- You have configured the Cisco Nexus 1000V connection using the [“Connecting to the vCenter Server” procedure on page 4-2](#).
- The Cisco Nexus 1000V is connected to vCenter Server/ESX.
- The Server Administrator has already added the host running Cisco Nexus 1000V to the DVS in vCenter Server.

SUMMARY STEPS

1. **show module**
2. **show server-info**
3. **show interface brief**
4. **show interface virtual**

DETAILED STEPS

	Command	Description
Step 1	show module Example: n1000v# show module	Displays module information.
Step 2	show server_info Example: n1000v# show server_info	Displays server information.
Step 3	show interface brief Example: n1000v# show interface brief	Displays interface information, including the uplinks to vCenter Server.
Step 4	show interface virtual Example: n1000v# show interface virtual	Displays virtual interface information.

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

Example

```
n1000v# show module
Mod  Ports  Module-Type                Model                Status
---  ---  -
1    1      Virtual Supervisor Module  Nexus1000V          active *
2    48     Virtual Ethernet Module    --                  ok
3    48     Virtual Ethernet Module    --                  ok

Mod  Sw                Hw      World-Wide-Name(s) (WWN)
---  ---  ---  -
1    4.0(0)S1(0.82)  0.0    --
2    NA              0.0    --
3    NA              0.0    --

Mod  MAC-Address(es)                Serial-Num
---  ---  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
2    02-00-0c-00-02-00 to 02-00-0c-00-02-80  NA
3    02-00-0c-00-03-00 to 02-00-0c-00-03-80  NA

Mod  Server-IP          Server-UUID                Server-Name
---  ---  ---  -
1    172.18.217.180    487701ee-6e87-c9e8-fb62-001a64d20a20  esx-1
2    172.18.117.44     487701ee-6e87-c9e8-fb62-001a64d20a20  esx-2
3    172.18.217.3      4876efdd-b563-9873-8b39-001a64644a24  esx-3
```

* this terminal session

Example

```
n1000v# show server_info

Mod  Status          UUID
---  ---  -
2    powered-up     34303734-3239-5347-4838-323130344654
3    absent         371e5916-8505-3833-a02b-74a4122fc476
4    powered-up     4880a7a7-7b51-dd96-5561-001e4f3a22f9
5    absent         48840e85-e6f9-e298-85fc-001e4f3a2326
6    powered-up     eb084ba6-3b35-3031-a6fe-255506d10cd0
n1000v#
```

Example

```
n1000v# show interface brief

-----
Port  VRF          Status IP Address                Speed  MTU
-----
mgmt0  --          up    172.28.15.211             1000  1500

-----
Ethernet  VLAN  Type Mode  Status Reason                Speed  Port
Interface                                     Ch #
-----
Eth2/2    1     eth trunk up    none                a-1000(D)  --

-----
Interface  VLAN  Type Mode  Status Reason                MTU
-----
```

Example

```
n1000v# show interface virtual
```

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

Port	Adapter	Owner	Mod	Host
Veth49		R-VM-1	2	mcs-srvr35

Feature History for Server Connections

This section provides the server connections feature release history.

Feature Name	Releases	Feature Information
DVS Deletion	4.2(1)SV1(4a)	This feature was added.
Server Connections	4.0(4)SV1(1)	This feature was introduced.

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



CHAPTER 5

Managing the Configuration

This chapter includes the following topics:

- [Information About Configuration Management, page 5-1](#)
- [Changing the Switch Name, page 5-1](#)
- [Configuring a Message of the Day, page 5-2](#)
- [Verifying the Configuration, page 5-3](#)
- [Saving a Configuration, page 5-10](#)
- [Erasing a Configuration, page 5-10](#)
- [Feature History for Configuration Management, page 5-11](#)

Information About Configuration Management

The Cisco Nexus 1000V provides you with the capability to change the switch name, configure messages of the day, and display, save, and erase configuration files.

Changing the Switch Name

Use this procedure to change the switch name or prompt from the default (switch#) to another character string.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in configuration mode.
- If the VSM is connected to vCenter Server then this procedure also changes the DVS the VSM is managing. In case of error in renaming the DVS, a syslog is generated and the DVS on vCenter Server will continue using the old DVS name.

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

DETAILED STEPS

Command	Purpose
Step 1 switchname Example: n1000v(config)# switchname metro metro(config)# exit metro#	Changes the switch prompt.

Configuring a Message of the Day

Use this procedure to configure a message of the day (MOTD) to display before the login prompt on the terminal when a user logs in.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in configuration mode.
- The banner message can be up to 40 lines with up to 80 characters per line.
- Use the following guidelines when choosing your delimiting character:
 - Do not use the *delimiting-character* in the *message* string.
 - Do not use " and % as delimiters.
- The following tokens can be used in the the message of the day:
 - \$(hostname) displays the host name for the switch.
 - \$(line) displays the vty or tty line or name.

DETAILED STEPS

Command	Purpose
Step 1 banner motd [<i>delimiting-character message delimiting-character</i>] Example: n1000v(config)# banner motd #April 16, 2008 Welcome to the svcs# n1000v(config)#	Configures a banner message of the day. <ul style="list-style-type: none"> • up to 40 lines • up to 80 characters per line • enclosed in delimiting character, such as # • can span multiple lines • can use tokens
Step 2 show banner motd Example: n1000v(config)# show banner motd April 16, 2008 Welcome to the Switch	Displays the configured banner message.

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

Verifying the Configuration

Use this section to view the switch configuration. This section includes the following topics:

- [Verifying the Software and Hardware Versions, page 5-3](#)
- [Verifying the Running Configuration, page 5-4](#)
- [Comparing the Startup and Running Configurations, page 5-6](#)
- [Verifying the Interface Configuration, page 5-7](#)

Verifying the Software and Hardware Versions

Use this command to view the versions of software and hardware on your system, for example, to verify the version before and after an upgrade.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 show version Example: n1000v# show version	Displays the versions of system software and hardware that are currently running on the switch.

```

Example:
n1000v# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2009, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

Software
  loader:      version 1.2(2)
  kickstart:   version 4.0(4)SV1(1)
  system:      version 4.0(4)SV1(1)
  kickstart image file is:
  kickstart compile time:  4/2/2009 23:00:00
  system image file is:    bootflash:/svs.bin
  system compile time:     4/2/2009 23:00:00 [04/23/2009 09:55:29]

Hardware
  Cisco Nexus 1000V Chassis ("Virtual Supervisor Module")

```

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

```
Intel(R) Xeon(R) CPU          with 1034780 kB of memory.
Processor Board ID T5056893321
```

```
Device name: n1000v
bootflash:   3897832 kB
```

```
Kernel uptime is 0 day(s), 0 hour(s), 2 minute(s), 55 second(s)
```

```
plugin
Core Plugin, Ethernet Plugin
```

Verifying the Running Configuration

Use this section to view the configuration currently running on the system.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 <code>show running-config</code> Example: n1000v# show running-config	Displays the versions of system software and hardware that are currently running on the switch,

```
Example:
n1000v# show running-config
version 4.0(4)SV1(1)
username admin password 5 $1$ouYE/pRM$/j4/2lg3RMD4PhE.1Z1S.0 role network-admin
telnet server enable
ip domain-lookup
ip host n1000v 172.23.232.141
kernel core target 0.0.0.0
kernel core limit 1
system default switchport
vem 3
  host vmware id 89130a67-e66b-3e57-ad25-547750bcfc7e
snmp-server user admin network-admin auth md5 0xb64ad6879970f0e57600c443287a79f0 priv
0xb64ad6879970f0e57600c443287a79f0 localizedkey
snmp-server enable traps license
vrf context management
  ip route 0.0.0.0/0 172.23.232.1
switchname n1000v
vlan 1,260-269
vdc n1000v id 1
  limit-resource vlan minimum 16 maximum 513
  limit-resource monitor-session minimum 0 maximum 64
  limit-resource vrf minimum 16 maximum 8192
  limit-resource port-channel minimum 0 maximum 256
  limit-resource u4route-mem minimum 32 maximum 80
  limit-resource u6route-mem minimum 16 maximum 48
port-profile Unused_Or_Quarantine_Uplink
```

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

```
description "Port-group created for Nexus1000V internal usage. Do not use."
capability uplink
vmware port-group
shutdown
state enabled
port-profile Unused_Or_Quarantine_Veth
description "Port-group created for Nexus1000V internal usage. Do not use."
vmware port-group
shutdown
state enabled
port-profile system-uplink
capability uplink
vmware port-group
switchport mode trunk
switchport trunk allowed vlan 260-261
no shutdown
system vlan 260-261
state enabled
port-profile vm-uplink
capability uplink
vmware port-group
switchport mode access
switchport access vlan 262
no shutdown
state enabled
port-profile data262
vmware port-group
switchport access vlan 262
no shutdown
state enabled

interface Ethernet3/2
inherit port-profile system-uplink

interface Ethernet3/3
inherit port-profile vm-uplink

interface mgmt0
ip address 172.23.232.141/24

interface control0
line vty
session-limit 32
boot kickstart bootflash:/kick.bin sup-1
boot system bootflash:/svs.bin sup-1
boot kickstart bootflash:/kick.bin sup-2
boot system bootflash:/svs.bin sup-2
svs-domain
domain id 141
control vlan 260
packet vlan 261
svs mode L2
svs connection vc
protocol vmware-vim
remote hostname 172.23.231.201
vmware dvs uuid "2c 6f 3d 50 62 f3 7f 4d-dc 00 70 e2 52 77 ca 15" datacenter-name
HamiltonDC
connect

n1000v#
```

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

Comparing the Startup and Running Configurations

Use this procedure to view the difference between the startup and running configurations.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 show running-config diff Example: n1000v# show running-config diff	Displays the difference between the startup configuration and the running configuration currently on the switch.

Example 5-1 Command output, show running-config diff

```
n1000v# show running-config diff
*** Startup-config
--- Running-config
*****
*** 1,7 ***
  version 4.0(1)
- system mem-thresholds minor 0 severe 0 critical 0
  vrf context management
    ip route 0.0.0.0/0 10.78.1.1
  switchname DCOS-112-S10
  vlan 80,110-111,150,160,170
  vdc DCOS-112-S10 id 1
--- 1,6 ---
*****
*** 116,131 ***
  ip address 10.78.1.112/24
  interface Vethernet49
    inherit port-profile vlan160
- interface Vethernet65
-   inherit port-profile vlan170
  interface Vethernet50
    inherit port-profile vlan160
  interface Vethernet66
    inherit port-profile vlan170
  ip route 0.0.0.0/0 10.78.1.1
  vlan 80-80, 110-110, 111-111, 150-150, 160-160, 170-170

--- 115,130 ---
  ip address 10.78.1.112/24

  interface Vethernet49
    inherit port-profile vlan160

  interface Vethernet50
    inherit port-profile vlan160
```

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

```
+ interface Vethernet65
+   inherit port-profile vlan170
+
+   interface Vethernet66
+     inherit port-profile vlan170
+     ip route 0.0.0.0/0 10.78.1.1
+     vlan 80-80, 110-110, 111-111, 150-150, 160-160, 170-170

n1000v#
```

Verifying the Interface Configuration

This section includes the following procedures:

- [Verifying a Brief Version of an Interface Configuration, page 5-7](#)
- [Verifying a Detailed Version of an Interface Configuration, page 5-8](#)
- [Verifying a Brief Version of all Interfaces, page 5-8](#)
- [Verifying the Running Configuration for all Interfaces, page 5-9](#)

For more information about displaying interfaces, see the document, *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)*

Verifying a Brief Version of an Interface Configuration

Use this procedure to view a brief version of an interface configuration.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 <code>show interface {type} {name} brief</code>	Displays a brief version of information about the specified interface configuration,

Example:

```
n1000v# show interface mgmt 0 brief
```

```
-----
Port    VRF      Status IP Address      Speed    MTU
-----
mgmt0   --      up      10.78.1.63      1000    1500
n1000v#
```

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

Verifying a Detailed Version of an Interface Configuration

Use this procedure to view a detailed version of an interface configuration.

BEFORE YOU BEGIN

Before using the commands in this section, you must know or do the following:

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 <code>show interface {type} {name}</code>	Displays details about the specified interface configuration,

Example:

```
n1000v# show interface mgmt 0
mgmt0 is up
  Hardware: Ethernet, address: 0050.5689.3321 (bia 0050.5689.3321)
  Internet Address is 172.23.232.141/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  full-duplex, 1000 Mb/s
  Auto-Negotiation is turned on
    4961 packets input, 511995 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun, 0 fifo
    245 packets output, 35853 bytes
    0 underrun, 0 output errors, 0 collisions
    0 fifo, 0 carrier errors

n1000v#
```

Verifying a Brief Version of all Interfaces

Use this procedure to view a brief version of all interfaces configured on your system.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 <code>show interface brief</code>	Displays a brief version of all interface configurations on your system,

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

Example:

```
n1000v# show interface brief
```

```
-----
Port      VRF      Status IP Address      Speed      MTU
-----
mgmt0     --       up      172.23.232.141  1000      1500
-----

Ethernet  VLAN   Type Mode   Status Reason      Speed      Port
Interface                               Status Reason      (D)        Ch #
-----
Eth3/2    1      eth trunk up      none      1000(D)   --
Eth3/3    262    eth access up    none      1000(D)   --
-----

Interface  VLAN   Type Mode   Status Reason      MTU
-----
Veth81     630    virt access up    none      1500
Veth82     630    virt access up    none      1500
Veth224    631    virt access up    none      1500
Veth225    1      virt access nonPcpt nonParticipating 1500
n1000v#
```

Verifying the Running Configuration for all Interfaces

Use this procedure to view the running configuration for all interfaces on your system.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in any command mode.
- The output for the command, **show running-config interface** differs from that of the command, **show interface**.

DETAILED STEPS

Command	Description
Step 1 show running-config interface	Displays the running configuration for all interfaces on your system,

Example:

```
n1000v# show running-config interface
version 4.0(1)

interface Ethernet3/2
  switchport
  inherit port-profile sftrunk

interface Ethernet3/6
  switchport
  inherit port-profile vmuplink

interface Ethernet6/2
  switchport
  inherit port-profile alluplink
```

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

```
interface mgmt0
  ip address 10.78.1.63/24

interface Vethernet81
  inherit port-profile vm630

interface Vethernet82
  inherit port-profile vm630

interface Vethernet224
  inherit port-profile vm631

interface Vethernet225

n1000v#
```

Saving a Configuration

Use this procedure to save the running configuration to the startup configuration so that your changes are retained in the configuration file the next time you start the system.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI in any command mode.

DETAILED STEPS

Command	Description
Step 1 <code>copy running-config startup-config</code>	Saves the new configuration into nonvolatile storage, after which the running and the startup copies of the configuration are identical.

```
Example:
n1000v(config)# copy run start
[#####] 100%
n1000v(config)#
```

Erasing a Configuration

Use this procedure to erase a startup configuration.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:



Caution

The **write erase** command erases the entire startup configuration with the exception of loader functions, the license configuration, and the certificate extension configuration.

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

- You are logged in to the CLI.
- The following parameters are used with this command:
 - boot: Erases the boot variables and the mgmt0 IP configuration.
 - debug: Erases the debug configuration.

DETAILED STEPS

Command	Description
Step 1 <code>write erase</code> [boot debug]	The existing startup configuration is completely erased and all settings revert to their factory defaults. The running configuration is not affected.

Feature History for Configuration Management

This section provides the configuration management feature release history.

Feature Name	Releases	Feature Information
Configuration Management	4.0(4)SV1(1)	This feature was introduced.

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



CHAPTER 6

Working with Files

This section includes the following topics:

- [Information About Files, page 6-1](#)
- [Navigating the File System, page 6-2](#)
- [Copying and Backing Up Files, page 6-6](#)
- [Creating a Directory, page 6-7](#)
- [Removing an Existing Directory, page 6-8](#)
- [Moving Files, page 6-8](#)
- [Deleting Files or Directories, page 6-9](#)
- [Compressing Files, page 6-10](#)
- [Uncompressing Files, page 6-11](#)
- [Directing Command Output to a File, page 6-12](#)
- [Verifying a Configuration File before Loading, page 6-12](#)
- [Rolling Back to a Previous Configuration, page 6-13](#)
- [Displaying Files, page 6-13](#)
- [Feature History for File Management, page 6-15](#)

Information About Files

The Cisco Nexus 1000V file system provides a single interface to all the file systems the switch uses, including:

- Flash memory file systems
- Network file systems (TFTP and FTP)
- Any other endpoint for reading or writing data (such as the running configuration)

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

Navigating the File System

This section describes how to navigate the file system and includes the following topics:

- [Specifying File Systems, page 6-2](#)
- [Identifying the Directory You are Working From, page 6-2](#)
- [Changing Your Directory, page 6-3](#)
- [Listing the Files in a File System, page 6-4](#)
- [Identifying Available File Systems for Copying Files, page 6-4](#)
- [Using Tab Completion, page 6-5](#)

Specifying File Systems

The syntax for specifying a file system is `<file system name>:[//server/]`. [Table 6-1](#) describes file system syntax.

Table 6-1 File System Syntax Components

File System Name	Server	Description
bootflash	sup-active sup-local sup-1 module-1	Internal memory located on the active supervisor used for storing system images, configuration files, and other miscellaneous files. Cisco Nexus 1000V CLI defaults to the bootflash: file system.
	sup-standby sup-remote sup-2 module-2	Internal memory located on the standby supervisor used for storing system images, configuration files, and other miscellaneous files.
volatile	—	Volatile random-access memory (VRAM) located on a supervisor module used for temporary or pending changes.

Identifying the Directory You are Working From

Use this procedure to display the directory name of your current CLI location.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.

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

DETAILED STEPS

Step	Command	Purpose
Step 1	<p><code>pwd</code></p> <p>Example: n1000v# pwd bootflash:</p>	Displays the present working directory.

Changing Your Directory

Use this procedure to change your location in the CLI, from one directory or file system to another.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI in any command mode.
- Cisco Nexus 1000V CLI defaults to the bootflash: file system.



Tip Any file saved in the volatile: file system is erased when the switch reboots.

DETAILED STEPS

Step	Command	Purpose
Step 1	<p><code>pwd</code></p> <p>Example: n1000v# pwd volatile: n1000v#</p>	Displays the directory name of your current CLI location.
Step 2	<p><code>cd <i>directory name</i></code></p> <p>Example: n1000v# cd bootflash:</p>	Changes your CLI location to the specified directory.
	<p>Example: n1000v# cd bootflash:mydir</p>	Changes your CLI location to the root directory on the bootflash: file system.
	<p>Example: n1000v# cd mystorage</p>	Changes your CLI location to the mydir directory that resides in the bootflash: file system.
		Changes your CLI location to the mystorage directory that resides within the current directory.
		If the current directory were bootflash: mydir, this command changes the current directory to bootflash: mydir/mystorage.

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

Listing the Files in a File System

Use this procedure to display the contents of a directory or file.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>dir [directory filename]</code>	Displays the contents of a directory or file.

Example:

```
DCOS-112-R5# dir lost+found/
 49241      Jul 01 09:30:00 2008  diagclient_log.2613
 12861      Jul 01 09:29:34 2008  diagmgr_log.2580
    31       Jul 01 09:28:47 2008  dmesg
 1811      Jul 01 09:28:58 2008  example_test.2633
    89       Jul 01 09:28:58 2008  libdiag.2633
 42136     Jul 01 16:34:34 2008  messages
    65       Jul 01 09:29:00 2008  otm.log
    741      Jul 01 09:29:07 2008  sal.log
    87       Jul 01 09:28:50 2008  startupdebug
```

```
Usage for log://sup-local
 51408896 bytes used
 158306304 bytes free
 209715200 bytes total
DCOS-112-R5#
```

Identifying Available File Systems for Copying Files

Use this procedure to identify the file systems you can copy to or from.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.

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

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>copy ?</code>	Displays the source file systems available to the copy command.
Step 2	<p><code>copy filename ?</code></p> <p>Example: n1000v# copy ? bootflash: Select source filesystem core: Select source filesystem debug: Select source filesystem ftp: Select source filesystem licenses Backup license files log: Select source filesystem nvram: Select source filesystem running-config Copy running configuration to destination scp: Select source filesystem sftp: Select source filesystem startup-config Copy startup configuration to destination system: Select source filesystem tftp: Select source filesystem volatile: Select source filesystem</p>	Displays the destination file systems available to the copy command for a specific file.

Using Tab Completion

Use this procedure to have the CLI complete a partial file name in a command.

Command	Purpose
<p>Step 1 <code>show file filesystem name: partial filename <Tab></code></p> <p>Example: n1000v# show file bootflash:nexus-1000v- bootflash:nexus-1000v-dplug-mzg.4.0.4.SV1. 0.42.bin bootflash:nexus-1000v-mzg.4.0.4.SV1.0.42.b in bootflash:nexus-1000v-kickstart-mzg.4.0.4. SV1.0.42.bin</p>	<p>When you type a partial filename and then press Tab, the CLI completes the file name if the characters you typed are unique to a single file.</p> <p>If not, the CLI lists a selection of file names that match the characters you typed.</p> <p>You can then retype enough characters to make the file name unique; and CLI completes the file name for you.</p>
<p>Step 2 <code>show file bootflash:c <Tab></code></p> <p>Example: n1000v# show file bootflash:c<Tab> -----BEGIN RSA PRIVATE KEY----- MIICXgIBAAKBgQDSq93BrlHcg3bX1jXDMY5c9+yZSS T3VhuQBqogvCPDGeLecA+j n1000v#</p>	The CLI completes the file name for you.

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

Copying and Backing Up Files

Use this procedure to copy a file, such as a configuration file, to save it or reuse it at another location. If your internal file systems are corrupted, you could potentially lose your configuration. Save and back up your configuration files periodically. Also, before installing or migrating to a new software configuration, back up the existing configuration files.

BEFORE YOU BEGIN

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

- You are logged in to the CLI through a Telnet, or SSH connection.
- If copying to a remote location, make sure that your device has a route to the destination. Your device and the remote destination must be in the same subnet if you do not have a router or default gateway to route traffic between subnets.
- Using the ping command, make sure that your device has connectivity to the destination.
- Make sure that the source configuration file is in the correct directory on the remote server.
- Make sure that the permissions on the source file are set correctly. Permissions on the file should be set to world-read.



Note

Use the **dir** command to ensure that enough space is available in the destination file system. If enough space is not available, use the **delete** command to remove unneeded files.

File System	Server	File Name
bootflash	sup-active sup-standby sup-1 or module-1 sup-2 or module-2 sup-local sup-remote	User-specified
volatile	—	User-specified
system	—	running-config
tftp ¹	IPv4 address, IPv6 address, or DNS name	User-specified
ftp		
scp (secure copy)		
sftp		
core	<i>slot-number</i>	Process identifier number

1. When downloading and uploading files, a limitation of TFTP restricts file size to 32 MB on the TFTP client and 16 MB on some TFTP servers.

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

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>copy [source filesystem:] filename [destination filesystem:] filename</code>	Copies a file from the specified source location to the specified destination location.
	Example: n1000v# copy system:running-config tftp://10.10.1.1/home/configs/switch3-run.cfg	Saves a copy of the running configuration to a remote switch.
	Example: n1000v# copy bootflash:system_image bootflash://sup-2/system_image	Copies a file from bootflash in the active supervisor module to bootflash in the standby supervisor module.
	Example: n1000v# copy system:running-config bootflash:my-config	Copies a running configuration to the bootflash: file system.
	Example: n1000v# copy scp://user@10.1.7.2/system-image bootflash:system-image	Copies a system image file from the SCP server identified by an IPv4 address to bootflash.
	Example: n1000v# copy sftp://172.16.10.100/myscript.txt volatile:myscript.txt	Copies a script file from the SFTP server identified by an IPv4 address to the volatile: file system.
	Example: n1000v# copy system:running-config bootflash:my-config	Places a back up copy of the running configuration on the bootflash: file system (ASCII file).
	Example: n1000v# copy bootflash:samplefile bootflash:mystorage/samplefile	Copies the file called samplefile from the root directory of the bootflash: file system to the mystorage directory.
	Example: n1000v# copy samplefile mystorage/samplefile	Copies a file within the current file system.
	Example: n1000v# copy tftp://10.10.1.1/home/configs/switch3-run.cfg system:running-config	Copies the source file to the running configuration on the switch, and configures the switch as the file is parsed line by line.

Creating a Directory

Use this procedure to create a directory at the current directory level or at a specified directory level.

Step	Command	Purpose
Step 1	<code>mkdir <i>directory name</i> dir <i>filename</i></code>	Creates a directory at the current directory level
	example: n1000v# mkdir bootflash:test n1000v#	Creates a directory called test in the bootflash: directory.
	example: n1000v# mkdir test n1000v#	Creates a directory called test at the current directory level. If the current directory is bootflash:mydir, this command creates a directory called bootflash:mydir/test.

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

Removing an Existing Directory

Use this section to remove an existing directory from the Flash file system.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.
- This command is only valid on Flash file systems.
- Before you can remove it, the directory must be empty.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>rmkdir {bootflash: debug: volatile:} directory</code>	Removes a directory.
	example: n1000v# <code>rmkdir bootflash:test</code> n1000v#	Removes the directory called test in the bootflash directory.
	example: n1000v# <code>rmkdir test</code> n1000v#	Removes the directory called test at the current directory level. If the current directory is bootflash:mydir, this command deletes the bootflash:mydir/test directory.

Moving Files

Use this procedure to move a file from one location to another location.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.
- The copy will not complete if there is not enough space in the destination directory.



Caution

If a file with the same name already exists in the destination directory, that file is overwritten by the moved file.

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

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>move {source path and filename} {destination path and filename}</code>	Deletes a directory.
	Example: n1000v# move bootflash:samplefile bootflash:mystorage/samplefile	Moves the file from one directory to another in the same file system (bootflash:).
	Example: n1000v# move samplefile mystorage/samplefile	Moves the file from one directory to another in the current file system.

Deleting Files or Directories

Use this procedure to delete files or directories on a Flash Memory device.

BEFORE YOU BEGIN



Caution

When deleting, if you specify a directory name instead of a file name, the entire directory and its contents are deleted.

- When you delete a file, the software erases the file.
- If you attempt to delete the configuration file or image specified by the CONFIG_FILE or BOOTLDR environment variable, the system prompts you to confirm the deletion.
- If you attempt to delete the last valid system image specified in the BOOT environment variable, the system prompts you to confirm the deletion.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>delete [bootflash: debug: log: volatile:] filename or directory name</code>	Deletes a specified file or directory.
	Example: n1000v# delete bootflash:dns_config.cfg	
	Example: n1000v# delete dns_config.cfg	Deletes the named file from the current working directory.
	Example: n1000v# delete bootflash:my-dir	Deletes the named directory and its contents.

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

Compressing Files

Use this procedure to compress (zip) a specified file using LZ77 coding.

BEFORE YOU BEGIN

- You are logged in to the CLI.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>show command > [path] filename</code> Example: n1000v# show system internal l2fm event-history errors n1000v#	Directs show command output to a file.
Step 2	<code>dir</code> Example: n1000v# dir	Displays the contents of the current directory, including the new file created in the first step.
Step 3	<code>gzip [path] filename</code> Example: n1000v# gzip bootflash:errorsfile n1000v#	Compresses the specified file
Step 4	<code>dir</code> Example: n1000v# dir	Displays the contents of the specified directory, including the newly-compressed file. Shows the difference in the file size of the newly-compressed file.

```

Example:
n1000v# show system internal l2fm event-history errors >errorsfile
n1000v# dir
      2687      Jul 01 18:17:20 2008  errorsfile
    16384      Jun 30 05:17:51 2008  lost+found/
      4096      Jun 30 05:18:29 2008  routing-sw/
         49      Jul 01 17:09:18 2008  sample_test.txt
    1322843     Jun 30 05:17:56 2008  nexus-1000v-dplug-mzg.4.0.4.SV1.0.42.bin
    21629952     Jun 30 05:18:02 2008  nexus-1000v-kickstart-mzg.4.0.4.SV1.0.42.bin
    39289400     Jun 30 05:18:14 2008  nexus-1000v-mzg.4.0.4.SV1.0.42.bin

```

```

Usage for bootflash://
 258408448 bytes used
 2939531264 bytes free
 3197939712 bytes total
n1000v# gzip bootflash:errorsfile
n1000v# dir
      1681      Jun 30 05:21:08 2008  cisco_svs_certificate.pem
         703      Jul 01 18:17:20 2008  errorsfile.gz
    16384      Jun 30 05:17:51 2008  lost+found/
      4096      Jun 30 05:18:29 2008  routing-sw/
         49      Jul 01 17:09:18 2008  sample_test.txt
    1322843     Jun 30 05:17:56 2008  nexus-1000v-dplug-mzg.4.0.4.SV1.0.42.bin
    21629952     Jun 30 05:18:02 2008  nexus-1000v-kickstart-mzg.4.0.4.SV1.0.42.bin
    39289400     Jun 30 05:18:14 2008  nexus-1000v-mzg.4.0.0.S1.0.34.bin

```

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

```
Usage for bootflash://
 258408448 bytes used
 2939531264 bytes free
 3197939712 bytes total
n1000v#
```

Uncompressing Files

Use this procedure to uncompress (unzip) a specified file that is compressed using LZ77 coding.

BEFORE YOU BEGIN

- You are logged in to the CLI.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>gunzip [path] filename</code>	Uncompresses the specified file.
Step 2	<code>dir</code>	Displays the contents of a directory, including the newly uncompresses file.

Example:

```
n1000v# gunzip bootflash:errorsfile.gz
n1000v# dir bootflash:
 2687      Jul 01 18:17:20 2008  errorsfile
16384      Jun 30 05:17:51 2008  lost+found/
 4096      Jun 30 05:18:29 2008  routing-sw/
  49       Jul 01 17:09:18 2008  sample_test.txt
1322843    Jun 30 05:17:56 2008  nexus-1000v-dplug-mzg.4.0.0.SV1.0.42.bin
21629952   Jun 30 05:18:02 2008  nexus-1000v-kickstart-mzg.4.0.4.SV1.0.42.bin
39289400   Jun 30 05:18:14 2008  nexus-1000v-mzg.4.0.0.SV1.0424.bin
```

```
Usage for bootflash://sup-local
 258408448 bytes used
 2939531264 bytes free
 3197939712 bytes total
DCOS-112-R5#
```

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

Directing Command Output to a File

Use this procedure to direct command output to a file.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>show running-config > [path filename]</code>	Directs the output of the command, show running-config , to a path and filename.
	Example: n1000v# show running-config > volatile:switch1-run.cfg	Directs the output of the command, show running-config , to the file, switch1-run.cfg, on the volatile file system.
	Example: n1000v# show running-config > bootflash:switch2-run.cfg	Directs the output of the command, show running-config , to the file, switch2-run.cfg, in bootflash.
	Example: n1000v# show running-config > tftp://10.10.1.1/home/configs/switch3-run.cfg	Directs the output of the command, show running-config , to the file, switch3-run.cfg, on a TFTP server.
	Example: n1000v# show interface > samplefile	Directs the output of the command, show interface, to the file, samplefile, at the same directory level, for example, in bootflash.

Verifying a Configuration File before Loading

Use this procedure to verify the integrity of an image before loading it. This command can be used for both the system and kickstart images.

DETAILED STEPS

Step	Command	Purpose
Step 1	<code>copy source path and file system:running-config</code>	Copies the source file to the running configuration on the switch, and configures the switch as the file is parsed line by line.
	Example: n1000v# copy tftp://10.10.1.1/home/configs/switch3-run.cfg system:running-config	
Step 2	<code>show version image [bootflash: modflash: volatile:]</code>	Validates the specified image.
	Example: n1000v# show version image bootflash:isan.bin image name: nexus-1000v-mz.4.0.4.SV1.1.bin bios: version unavailable system: version 4.0(4)SV1(1) compiled: 4/2/2009 23:00:00 [04/23/2009 09:55:29] n1000v#	

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

Rolling Back to a Previous Configuration

Use this procedure to recover your configuration from a previously saved version.

BEFORE YOU BEGIN



Note

Each time a **copy running-config startup-config** command is used, a binary file is created and the ASCII file is updated. A valid binary configuration file reduces the overall boot time significantly. A binary file cannot be uploaded, but its contents can be used to overwrite the existing startup configuration. The **write erase** command clears the binary file.

DETAILED STEPS

Step	Command	Purpose
Step 1	<pre>copy running-config bootflash: {filename}</pre> <p>Example:</p> <pre>n1000v# copy running-config bootflash:June03-Running</pre>	Reverts to a snapshot copy of a previously saved running configuration (binary file).
	<pre>copy bootflash: {filename} startup-config</pre> <p>Example:</p> <pre>n1000v# copy bootflash:my-config startup-config</pre>	Reverts to a configuration copy that was previously saved in the bootflash: file system (ASCII file).

Displaying Files

This section describes how to display information about files and includes the following procedures:

- [Displaying File Contents, page 6-13](#)
- [Displaying Directory Contents, page 6-14](#)
- [Displaying File Checksums, page 6-15](#)
- [Displaying the Last Lines in a File, page 6-15](#)

Displaying File Contents

Use this procedure to display the contents of a specified file.

BEFORE YOU BEGIN

- You are logged in to the CLI.

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

DETAILED STEPS

Step	Command	Purpose
Step 1	<pre>show file [bootflash: debug: volatile:] filename Example: n1000v# show file bootflash:sample_test.txt config t Int veth1/1 no shut end show int veth1/1 n1000v#</pre>	Displays the contents of the specified file.

Displaying Directory Contents

Use this procedure to display the contents of a directory or file system.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.

Step	Command	Purpose
Step 1	<pre>pwd Example: n1000v# pwd bootflash:</pre>	Displays the present working directory.
Step 2	<pre>dir</pre>	Displays the contents of the directory.

```
Example:
n1000v# pwd
bootflash:
n1000v# dir

Usage for volatile://
    0 bytes used
 20971520 bytes free
 20971520 bytes total
n1000v#
```

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

Displaying File Checksums

Use this procedure to display checksums for checking file integrity.

Step	Command	Purpose
Step 1	show file <i>filename</i> [cksum md5sum] Example: n1000v# show file bootflash:cisco_svs_certificate.pem cksum 266988670 Example: n1000v# show file bootflash:cisco_svs_certificate.pem md5sum d3013f73aea3fda329f7ea5851ae81ff n1000v#	Provides the checksum or MD5 checksum of the file for comparison with the original file. Provides the Message-Digest Algorithm 5 (MD5) checksum of the file. MD5 is an electronic fingerprint for the file.

Displaying the Last Lines in a File

Use this command to display the last lines (tail end) of a specified file.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.

DETAILED STEPS

Step	Command	Purpose
Step 1	tail {path}[filename] {Number of lines}	Displays the requested number of lines from the end of the specified file. Allowable range for number of lines: 0 - 80

```

Example:
n1000v# tail bootflash:errorsfile 5

20) Event:E_DEBUG, length:34, at 171590 usecs after Tue Jul  1 09:29:05 2008
    [102] main(326): stateless restart

n1000v#

```

Feature History for File Management

This section provides the file management feature release history.

Feature Name	Releases	Feature Information
File Management	4.0(4)SV1(1)	This feature was introduced.

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



CHAPTER 7

Managing Users

This section includes the following procedures:

- [Information About User Management, page 7-1](#)
- [Displaying Current User Access, page 7-1](#)
- [Sending a Message to Users, page 7-2](#)
- [Feature History for User Management, page 7-2](#)

Information About User Management

You can identify the users currently connected to the device and send a message to either a single user or all users.

For information about assigning user roles, see the *Cisco Nexus 1000V Security Configuration Guide, Release 4.2(1)SV1(4)*.

Displaying Current User Access

Use this procedure to display all users currently accessing the switch.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.

DETAILED STEPS

Command	Description
Step 1 <code>show users</code>	Displays a list of users who are currently accessing the system.

Example:

```
n1000v# show users
NAME      LINE      TIME      IDLE      PID COMMENT
admin     pts/0     Jul  1 04:40 03:29     2915 (::ffff:64.103.145.136)
```

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

```

admin pts/2 Jul 1 10:06 03:37 6413 (::ffff:64.103.145.136)
admin pts/3 Jul 1 13:49 . 8835 (171.71.55.196)*
n1000v#

```

Sending a Message to Users

Use this command to send a message to all active CLI users currently using the system.

BEFORE YOU BEGIN

Before using this command, you must know or do the following:

- You are logged in to the CLI.

DETAILED STEPS

Command	Description
Step 1 <code>send {session device} line</code>	<p>Sends a message to users currently logged in to the system.</p> <ul style="list-style-type: none"> session: sends the message to a specified pts/tty device type. line: a message of up to 80 alphanumeric characters in length.

Example:

```
n1000v# send Hello. Shutting down the system in 10 minutes.
```

```
Broadcast Message from admin@switch
(/dev/pts/34) at 8:58 ...
```

```
Hello. Shutting down the system in 10 minutes.
```

```
n1000v#
```

Feature History for User Management

This section provides the user management feature release history.

Feature Name	Releases	Feature Information
User Management	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 8

Configuring NTP

This chapter describes how to configure the Network Time Protocol (NTP) and includes the following topics:

- [Information about NTP, page 8-1](#)
- [Prerequisites for NTP, page 8-2](#)
- [Configuration Guidelines and Limitations, page 8-3](#)
- [Default Settings, page 8-3](#)
- [Configuring an NTP Server and Peer, page 8-3](#)
- [Verifying the NTP Configuration, page 8-4](#)
- [NTP Example Configuration, page 8-5](#)
- [Additional References, page 8-5](#)
- [Feature History for NTP, page 8-5](#)

Information about NTP

The Network Time Protocol (NTP) synchronizes timekeeping among a set of distributed time servers and clients. This synchronization allows you to correlate events when you receive system logs and other time-specific events from multiple network devices.

NTP uses the User Datagram Protocol (UDP) as its transport protocol. All NTP communication uses the Universal Time Coordinated (UTC) standard. An NTP server usually receives its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of each other.

NTP uses a stratum to describe how many NTP hops away that a network device is from an authoritative time source. A stratum 1 time server has an authoritative time source (such as an atomic clock) directly attached to the server. A stratum 2 NTP server receives its time through NTP from a stratum 1 NTP server, which in turn connects to the authoritative time source.

NTP avoids synchronizing to a network device that may keep accurate time. NTP never synchronizes to a system that is not in turn synchronized itself. NTP compares the time reported by several network devices and does not synchronize to a network device that has a time that is significantly different than the others, even if its stratum is lower.

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

Cisco NX-OS cannot act as a stratum 1 server. You cannot connect to a radio or atomic clock. We recommend that the time service that you use for your network is derived from the public NTP servers available on the Internet.


If the network is isolated from the Internet, Cisco NX-OS allows you to configure a network device so that the device acts as though it is synchronized through NTP, when in fact it has determined the time using other means. Other network devices can then synchronize to that network device through NTP.

NTP Peers

NTP allows you to create a peer relationship between two networking devices. A peer can provide time on its own or connect to an NTP server. If both the local device and the remote peer point to different NTP servers, your NTP service is more reliable. The local device maintains the right time even if its NTP server fails by using the time from the peer.

Figure 8-1 displays a network with two NTP stratum 2 servers and two switches.

Figure 8-1 NTP Peer and Server Association



In this configuration, switch 1 and switch 2 are NTP peers. switch 1 uses stratum-2 server 1, while switch 2 uses stratum-2 server 2. If stratum-2 server-1 fails, switch 1 maintains the correct time through its peer association with switch 2.

High Availability

Stateless restarts are supported for NTP. After a reboot or a supervisor switchover, the running configuration is applied.

You can configure NTP peers to provide redundancy in case an NTP server fails.

Prerequisites for NTP

If you configure NTP, you must have connectivity to at least one server that is running NTP.

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

Configuration Guidelines and Limitations

NTP has the following configuration guidelines and limitations:

- You should have a peer association with another device only when you are sure that your clock is reliable (which means that you are a client of a reliable NTP server).
- A peer configured alone takes on the role of a server and should be used as backup. If you have two servers, you can configure several devices to point to one server and the remaining devices point to the other server. You can then configure peer association between these two servers to create a more reliable NTP configuration.
- If you only have one server, you should configure all the devices as clients to that server.
- You can configure up to 64 NTP entities (servers and peers).

Default Settings

The following table lists the default settings for CDP and NTP parameters.

Parameter	Default
NTP	Enabled

Configuring an NTP Server and Peer

Use this procedure to configure an NTP server and peer.

BEFORE YOU BEGIN

- You can configure NTP using IPv4 addresses or domain name server (DNS) names.

SUMMARY STEPS

1. **config t**
2. **ntp server** {*ip-address* | *dns-name*}
3. **ntp peer** {*ip-address* | *dns-name*}
4. **show ntp peers**
5. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	<code>ntp server {ip-address dns-name}</code> Example: n1000v(config)# ntp server 192.0.2.10	Forms an association with a server.
Step 3	<code>ntp peer {ip-address dns-name}</code> n1000v(config)# ntp peer 2001:0db8::4101	Forms an association with a peer. You can specify multiple peer associations.
Step 4	<code>show ntp peers</code> Example: n1000v(config)# show ntp peers	(Optional) Displays the configured server and peers. Note A domain name is resolved only when you have a DNS server configured.
Step 5	<code>copy running-config startup-config</code> Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

The following is an example configures an NTP server and peer:

```
n1000v# config t
n1000v(config)# ntp server 192.0.2.10
n1000v(config)# ntp peer 2001:0db8::4101
```

Clearing NTP Statistics or Sessions

Use the following commands to clear NTP statistics or sessions.

Command	Purpose
<code>clear ntp statistics</code>	Clears the NTP statistics.
<code>clear ntp session</code>	Clears the NTP sessions.

Verifying the NTP Configuration

To display NTP configuration information, use one of the following commands:

Command	Purpose
<code>show ntp peer-status</code>	Displays the status for all NTP servers and peers.
<code>show ntp peers</code>	Displays all the NTP peers.
<code>show ntp statistics {io local memory peer {ip-address dns-name}}</code>	Displays the NTP statistics

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

NTP Example Configuration

This example configures an NTP server:

```
config t
ntp server 192.0.2.10
```

Additional References

For additional information related to NTP, see the following sections:

- [Related Documents, page 8-5](#)
- [Standards, page 8-5](#)

Related Documents

Related Topic	Document Title
Interface	<i>Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</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.	—

Feature History for NTP

This section shows the NTP feature release history.

Feature Name	Releases	Feature Information
NTP	4.0(4)SV1(1)	This feature was introduced.

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



CHAPTER 9

Configuring Local SPAN and ERSPAN

This chapter describes how to configure the local and encapsulated remote (ER) switched port analyzer (SPAN) feature to monitor traffic and includes the following topics:

- [Information About SPAN and ERSPAN, page 9-1](#)
- [SPAN Guidelines and Limitations, page 9-5](#)
- [Default Settings, page 9-6](#)
- [Configuring SPAN, page 9-6](#)
- [Verifying the SPAN Configuration, page 9-20](#)
- [Example Configurations, page 9-20](#)
- [Additional References, page 9-22](#)
- [Feature History for SPAN and ERSPAN, page 9-23](#)

Information About SPAN and ERSPAN

The Switched Port Analyzer (SPAN) feature (sometimes called port mirroring or port monitoring) allows network traffic to be analyzed by a network analyzer such as a Cisco SwitchProbe or other Remote Monitoring (RMON) probe.

SPAN lets you monitor traffic on one or more ports, or one or more VLANs, and send the monitored traffic to one or more destination ports where the network analyzer is attached.

This section includes the following topics:

- [SPAN Sources, page 9-1](#)
- [SPAN Destinations, page 9-2](#)
- [SPAN Sessions, page 9-5](#)

SPAN Sources

The interfaces from which traffic can be monitored are called SPAN sources. These include Ethernet, virtual Ethernet, port-channel, port profile, and VLAN. When a VLAN is specified as a SPAN source, all supported interfaces in the VLAN are SPAN sources. When a port profile is specified as a SPAN source, all ports which inherit the port profile are SPAN sources. Traffic can be monitored in the receive direction, the transmit direction, or both directions for Ethernet and virtual Ethernet source interfaces.

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

- Receive source (Rx)—Traffic that enters the switch through this source port is copied to the SPAN destination port.
- Transmit source (Tx)—Traffic that exits the switch through this source port is copied to the SPAN destination port.

Characteristics of SPAN Sources

A Local SPAN source has these characteristics:

- Can be port type Ethernet, virtual Ethernet, port channel, port profile, or VLAN.
- Cannot be a destination port or port profile.
- Can be configured to monitor the direction of traffic —receive, transmit, or both.
- Source ports can be in the same or different VLANs.
- For VLAN SPAN sources, all active ports in the source VLAN are included as source ports.
- Local SPAN sources must be on the same host (VEM) as the destination port.
- For port profile sources, all active interfaces attached to the port profile are included as source ports.

SPAN Destinations

SPAN destinations refer to the interfaces that monitor source ports. This section includes the following topics:

- [Characteristics of Local SPAN Destinations, page 9-2](#)
- [Characteristics of ERSPAN Destinations, page 9-3](#)

Characteristics of Local SPAN Destinations

Each local SPAN session must have at least one destination port (also called a monitoring port) that receives a copy of traffic from the source ports or VLANs. A destination port has these characteristics:

- Can be any physical or virtual Ethernet port, a port channel, or a port profile.
- Cannot be a source port or port profile.
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session or a source port profile.
- Receives copies of transmitted and received traffic for all monitored source ports in the same VEM module. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.
- Must not be private VLAN mode.
- A destination port can only monitor sources on the same host (VEM). See [Figure 9-1, Local SPAN](#).
- Destination ports in access mode receive monitored traffic on all the VLANs.
- Destination ports in trunk mode receive monitored traffic only on the allowed VLANs in the trunk configuration.

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

Characteristics of ERSPAN Destinations

- An ERSPAN destination is specified by an IP address.
- In ERSPAN, the source SPAN interface and destination SPAN interface may be on different devices interconnected by an IP network. ERSPAN traffic is GRE-encapsulated. See [Figure 9-2, ERSPAN](#).

Local SPAN

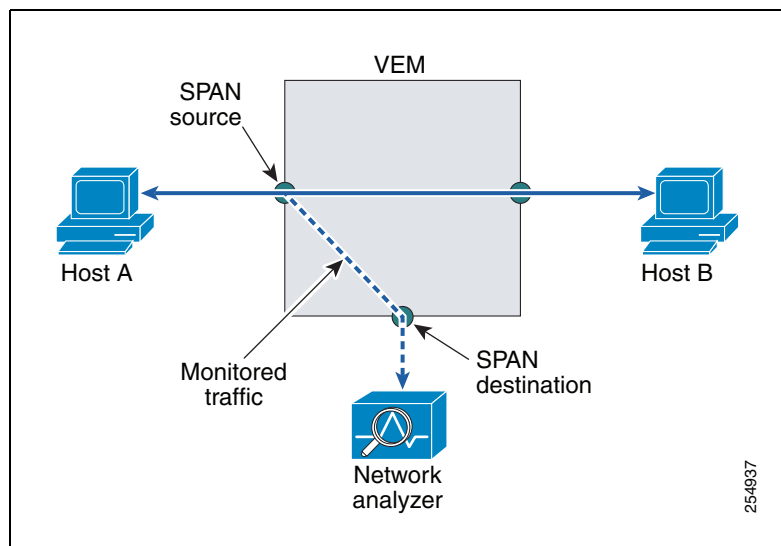
In Local SPAN, the source interface and destination interface are on the same VEM. The network analyzer is attached directly to the SPAN destination port. The SPAN source can be a port, a VLAN interface or port profile. The destination can be a port or port profile.

[Figure 9-1](#) shows that traffic transmitted by host A is received on the SPAN source interface. Traffic (ACLs, QoS, and so forth) is processed as usual. Traffic is then replicated. The original packet is forwarded on toward host B. The replicated packet is then sent to the destination SPAN interface where the monitor is attached.

Local SPAN can replicate to one or more destination ports. Traffic can be filtered so that only traffic of interest is sent out the destination SPAN interface.

Local SPAN can monitor all traffic received on the source interface including BPDUs.

Figure 9-1 Local SPAN



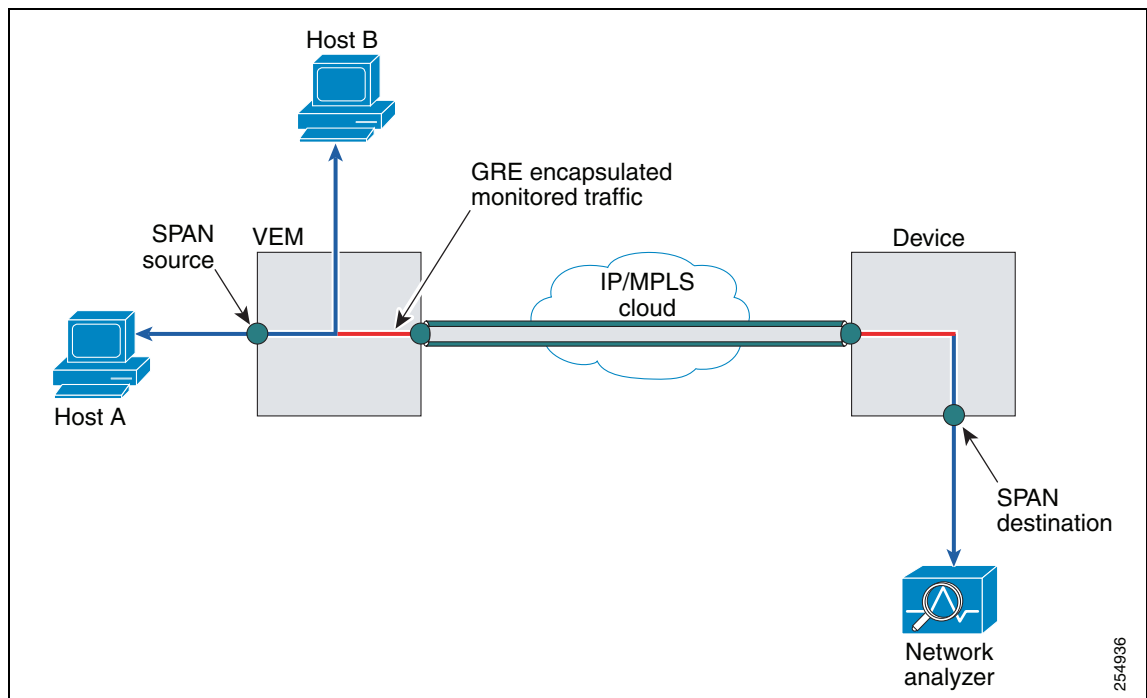
Send document comments to nexus1k-docfeedback@cisco.com.

Encapsulated Remote SPAN

Encapsulated remote (ER) SPAN monitors traffic in multiple network devices across an IP network and sends that traffic in an encapsulated envelope to destination analyzers. In contrast, Local SPAN cannot forward traffic through the IP network. ERSPAN can be used to monitor traffic remotely. ERSPAN sources can be ports, VLANs, or port profiles.

In [Figure 9-2](#), the ingress and egress traffic for host A are monitored using ERSPAN. Encapsulated ERSPAN packets are routed from host A through the routed network to the destination device where they are de-capsulated and forwarded to the attached network analyzer. The destination may also be on the same L2 network as the source.

Figure 9-2 ERSPAN



Network Analysis Module

You can also use the Cisco Network Analysis Module (NAM) to monitor ERSPAN data sources for application performance, traffic analysis, and packet header analysis.

To use NAM for monitoring the Cisco Nexus 1000V ERSPAN data sources see the *Cisco Nexus 1010 Network Analysis Module Installation and Configuration Note*, 4.2.

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

SPAN Sessions

You can create up to 64 total SPAN sessions (Local SPAN plus ERSPAN) on the VEM.

You must configure an ERSPAN session ID that is added to the ERSPAN header of the encapsulated frame to differentiate between ERSPAN streams of traffic at the termination box. You can also configure the range of flow ID numbers. For more information, see [Configuring the Allowable ERSPAN Flow IDs](#), page 9-19.

When trunk ports are configured as SPAN sources and destinations, you can filter VLANs to send to the destination ports from among those allowed. Both sources and destinations must be configured to allow the VLANs.

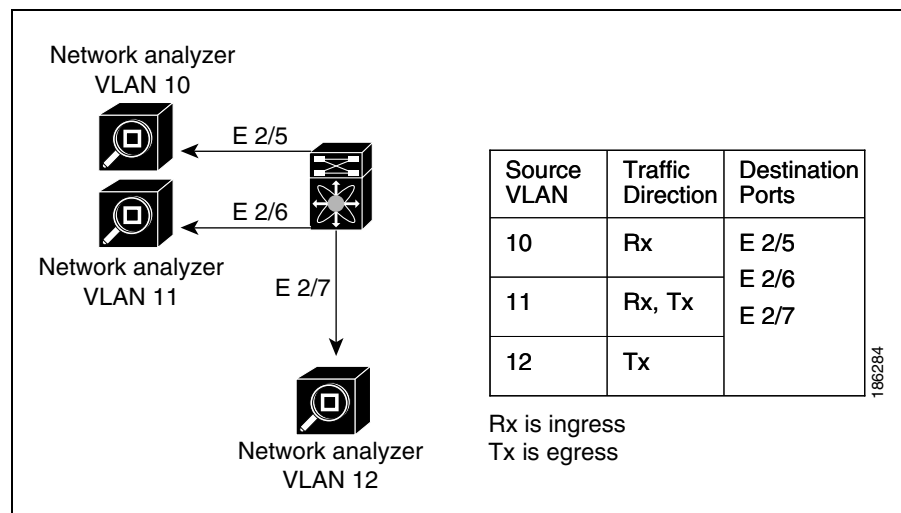
Figure 9-3 shows one example of a VLAN-based SPAN configuration in which traffic is copied from three VLANs to three specified destination ports. You can choose which VLANs to allow on each destination port to limit the traffic transmitted. In Figure 9-3, the device transmits packets from one VLAN at each destination port. The destinations in this example are trunks on which allowed VLANs are configured.



Note

VLAN-based SPAN sessions cause all source packets to be copied to all destinations, whether the packets are required at the destination or not. VLAN traffic filtering occurs at transmit destination ports.

Figure 9-3 VLAN-based SPAN Configuration Example



SPAN Guidelines and Limitations

SPAN has the following configuration guidelines and limitations:

- A maximum of 64 SPAN sessions (Local SPAN plus ERSPAN) can be configured on the VSM.
- A maximum of 32 source VLANs are allowed in a session.
- A maximum of 32 destinations are allowed for a Local SPAN session.
- A maximum of 128 source interfaces are allowed in a session.

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



Caution

Overload Potential

To avoid an overload on uplink ports, use caution when configuring ERSPAN, especially when sourcing VLANs.

- A port can be configured in a maximum of 4 SPAN sessions.
- The destination port used in one SPAN session cannot also be used as the destination port for another SPAN session.
- You cannot configure a port as both a source and destination port.
- In a SPAN session, packets that source ports receive may be replicated even though they are not transmitted on the ports. The following are examples of this behavior:
 - Traffic that results from flooding
 - Broadcast and multicast traffic
- For VLAN SPAN sessions switched on the same VLAN with both receive and transmit configured, two packets (one from receive and one from transmit) are forwarded from the destination port.

Default Settings

Table 9-1 lists the SPAN default settings.

Table 9-1 SPAN Defaults

Parameters	Default
State	SPAN sessions are created in the shut state.
Description	blank
Traffic direction for source interface or port profile	both
Traffic direction for source VLAN	receive (ingress or RX)

Configuring SPAN

This section describes how to configure SPAN and includes the following procedures.

- [Configuring a Local SPAN Session, page 9-7](#)
- [Configuring an ERSPAN Port Profile, page 9-9](#)
- [Configuring an ERSPAN Session, page 9-13](#)
- [Shutting Down a SPAN Session, page 9-16](#)
- [Resuming a SPAN Session, page 9-17](#)
- [Verifying the SPAN Configuration, page 9-20](#)

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

Configuring a Local SPAN Session

Use this procedure to configure a SPAN session.

**Note**

If you are configuring ERSPAN, see the [“Configuring an ERSPAN Session”](#) procedure on page 9-13.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You know the number of the SPAN session you are going to configure.
- The source and destination ports are already configured in either access or trunk mode. For more information, see the *Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)*.
- SPAN sessions are created in the shut state by default.
- When you create a SPAN session that already exists, any additional configuration is added to that session. To make sure the session is cleared of any previous configuration, you can delete the session first (see [Step 2, no monitor session](#)).
- This procedure involves creating the SPAN session in Monitor Configuration mode; and then, optionally, configuring allowed VLANs in Interface Configuration mode.

SUMMARY STEPS

1. **config t**
2. **no monitor session** *session-number*
3. **monitor session** *session-number*
4. **description** *description*
5. **source** {**interface** {*type*} {*id* | *range*} | **vlan** {*id* | *range*} | **port-profile** {*name*}} [**rx** | **tx** | **both**]
6. (Optional) Repeat [Step 5](#) to configure additional SPAN sources.
7. (Optional) **filter vlan** {*number* | *range*}
8. (Optional) Repeat [Step 7](#) to configure all source VLANs to filter.
9. **destination** {**interface** {*type*} {*id*} | **port-profile** {*name*}}
10. (Optional) Repeat [Step 9](#) to configure all SPAN destination ports.
11. **no shut**
12. (Optional) **exit**
13. (Optional) **show monitor session** *session-number*
14. (Optional) **show interface** {*type*} {*id*} **switchport**
15. (Optional) **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	no monitor session session-number Example: n1000v(config)# no monitor session 3	Clears the specified session.
Step 3	monitor session session-number Example: n1000v(config)# monitor session 3 n1000v(config-monitor)#	Creates a session with the given session number and places you in the CLI Monitor Configuration mode to further configure the session.
Step 4	description description Example: n1000v(config-monitor)# description my_span_session_3	For the specified SPAN session, adds a description. <ul style="list-style-type: none"> description: up to 32 alphanumeric characters default = blank (no description)
Step 5	source {interface {type} {id} vlan {id range} port-profile {name}} [rx tx both] Example 1: n1000v(config-monitor)# source interface ethernet 2/1-3, ethernet 3/1 rx Example 2: n1000v(config-monitor)# source interface port-channel 2 Example 3: n1000v(config-monitor)# source interface vethernet 12 both Example 4: n1000v(config-monitor)# source vlan 3, 6-8 tx Example 5: n1000v(config-monitor)# source port-profile my_port_profile	For the specified session, configures the sources and the direction of traffic to monitor. <ul style="list-style-type: none"> type: Specify the interface type—Ethernet or vEthernet. ID: Specify the vEthernet number, the Ethernet slot/port, or the VLAN ID to monitor. range: Specify the VLAN range to monitor name: Specify the name of the existing port profile. This port profile is different from the port profile created to carry ERSPAN packets through the IP network as defined in the “Configuring an ERSPAN Port Profile” section on page 9-9. traffic direction: Specify direction of traffic monitoring: <ul style="list-style-type: none"> receive (rx) (the VLAN default) transmit (tx) both (the default)
Step 6	(Optional) Repeat Step 5 to configure additional SPAN sources.	
Step 7	filter vlan {id range} Example: n1000v(config-monitor)# filter vlan 3-5, 7	(Optional) For the specified SPAN session, configures the filter from among the source VLANs.
Step 8	(Optional) Repeat Step 7 to configure all source VLANs to filter.	

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

	Command	Purpose
Step 9	destination {interface {type} {id range} port-profile {name}} Example: n1000v(config-monitor)# destination interface ethernet 2/5, ethernet 3/7	For the specified SPAN session, configures the destination(s) for copied source packets. <ul style="list-style-type: none"> • type: Specify the interface type—Ethernet or vEthernet. • ID: Specify the vEthernet number or the Ethernet slot/port to monitor. • name: Specify the name of the port profile to monitor. Note SPAN destination ports must already be configured as either access or trunk ports.
Step 10	(Optional) Repeat Step 9 to configure all SPAN destination ports.	
Step 11	no shut Example: n1000v(config-monitor)# no shut	Enables the SPAN session. By default, the session is created in the shut state.
Step 12	exit Example: n1000v(config-monitor)# exit n1000v(config)#	(Optional) Exits Monitor Configuration mode and places you in CLI Configuration mode.
Step 13	show monitor session session-number Example: n1000v(config-if)# show monitor session 3	(Optional) Displays the configured monitor session.
Step 14	show interface {type} {id} switchport Example: n1000v(config-if)# show interface ethernet 2/5 switchport	(Optional) Displays the configured port including allowed VLANs.
Step 15	copy running-config startup-config Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Configuring an ERSPAN Port Profile

Use this procedure to configure a port profile on the VSM to carry ERSPAN packets through the IP network to a remote destination analyzer.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- This configuration must be completed for all hosts in the vCenter Server.
- You know the name to be used for this port profile.

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



Note The port profile name is used to configure the VMKNIC. A VMKNIC is required on each ESX host to send ERSPAN encapsulated IP packets, and must have IP connectivity to the ERSPAN destination IP address.

- You know the name of the VMware port group to which this profile maps.
- You have the VMware documentation for adding a new virtual adapter.
- You have already created the system VLAN that sends IP traffic to the ERSPAN destination; and you know its VLAN ID which will be used in this configuration.

For more information about system port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)*.

- The port profile used for ERSPAN must be configured for Layer 3 control. This procedure includes a step for making this configuration.
- Only one VM kernel NIC can be assigned to this Layer 3 control port profile per host.
- The port profile must be an access port profile. It cannot be a trunk port profile. This procedure includes steps to configure the port profile as an access port profile.
- For more information about creating a Layer 3 control port profile, see the [“Creating a Port Profile for Layer 3 Control” procedure on page 3-9](#).

SUMMARY STEPS

1. **config t**
2. **port-profile** *name*
3. **capability l3control**
4. **vmware port-group** *name*
5. **switchport mode access**
6. **switchport access vlan** *id*
7. **no shutdown**
8. **system vlan** *id*
9. **state enabled**
10. (Optional) **show port-profile name** *name*
11. (Optional) **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	port-profile port_profile_name Example: n1000v(config)# port-profile erspan_profile n1000v(config-port-prof)#	<p>Creates the port profile and places you into CLI Global Configuration mode for the specified port profile. Saves the port profile in the running configuration.</p> <p>The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V.</p>
Step 3	capability l3control Example: n1000v(config-port-prof)# capability l3control n1000v(config-port-prof)#	Configures the port profile to carry ERSPAN traffic and saves this in the running configuration.
Step 4	vmware port-group name Example: n1000v(config-port-prof)#vmware port-group erspan n1000v(config-port-prof)#	<p>Designates the port profile as a VMware port group and adds the name of the VMware port group to which this profile maps. Saves the settings in the running configuration.</p> <p>The port profile is mapped to a VMware port group of the same name. When a vCenter Server connection is established, the port group created in Cisco Nexus 1000V is then distributed to the virtual switch on the vCenter Server.</p> <ul style="list-style-type: none"> • name: Port group name. If you do not specify a name, then the port group name will be the same as the port profile name. If you want to map the port profile to a different port group name, use the name option followed by the alternate name.
Step 5	switchport mode access Example: n1000v(config-port-prof)# switchport mode access n1000v(config-port-prof)#	Designates the interfaces as switch access ports (the default).
Step 6	switchport access vlan id Example 1: n1000v(config-port-prof)# switchport access vlan 2 n1000v(config-port-prof)#	<p>Assigns a VLAN ID to the access port for this port profile and saves the setting in the running configuration.</p> <p>This VLAN is used to send IP traffic to the ERSPAN destination.</p>
Step 7	no shutdown Example: n1000v(config-port-prof)# no shutdown n1000v(config-port-prof)#	Enables the interface in the running configuration.

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

	Command	Purpose
Step 8	<p>system vlan id</p> <p>Example: n1000v(config-port-prof)# system vlan 2 n1000v(config-port-prof)#</p>	<p>Associates the system VLAN ID with the port profile and saves it in the running configuration.</p> <p>Must match the VLAN ID assigned to the access port. If it does not match, then the following error message is generated:</p> <p>ERROR: System vlan being set does not match the switchport access vlan 2</p>
Step 9	<p>state enabled</p> <p>Example: n1000v(config-port-prof)# state enabled n1000v(config-port-prof)#</p>	<p>Enables the port profile in the running configuration.</p> <p>This port profile is now ready to send out ERSPAN packets on all ESX Hosts with ERSPAN sources</p>
Step 10	<p>show port-profile name port_profile_name</p> <p>Example: n1000v(config-port-prof)# show port-profile name erspan port-profile erspan description: status: enabled capability uplink: no capability l3control: yes system vlans: 2 port-group: access max-ports: 32 inherit: config attributes: switchport access vlan 2 no shutdown evaluated config attributes: switchport access vlan 2 no shutdown assigned interfaces:</p> <p>n1000v(config-port-prof)#</p>	<p>(Optional) Displays the configuration for the specified port profile as it exists in the running configuration.</p>
Step 11	<p>copy running-config startup-config</p> <p>Example: n1000v(config-port-prof)# copy running-config startup-config [#####] 100% n1000v(config-port-prof)#</p>	<p>(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.</p>
Step 12	<p>Using the VMware documentation, go to vSphere Client and configure a VMKNIC on each ESX Host for sending ERSPAN encapsulated packets. Make sure the VMKNIC points to this port profile as a new virtual adapter. This VMKNIC must have IP connectivity to the ERSPAN destination IP address.</p>	

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

Configuring an ERSPAN Session

Use this procedure to configure an ERSPAN session.



Note

If you are configuring Local SPAN, see the “[Configuring a Local SPAN Session](#)” procedure on page 9-7.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You know the number of the SPAN session you are going to configure.
- You have already configured an ERSPAN-capable port profile on the VSM using the “[Configuring an ERSPAN Port Profile](#)” procedure on page 9-9.
- Using the VMware documentation for adding a new virtual adapter, you have already configured the required VMKNIC on each of the ESX hosts. The VMKNIC must have IP connectivity to the ERSPAN destination IP address for sending ERSPAN encapsulated packets.
- SPAN sessions are created in the shut state by default.
- When you create a SPAN session that already exists, any additional configuration is added to that session. To make sure the session is cleared of any previous configuration, you can delete the session first (see [Step 2, no monitor session](#)).
- This procedure involves creating the SPAN session in ERSPAN source configuration mode (config-erspan-source).

SUMMARY STEPS

1. **config t**
2. **no monitor session** *session-number*
3. **monitor session** *session-number* **type erspan-source**
4. **description** *description*
5. **source** {**interface** *type* {*number* | *range*} | **vlan** {*number* | *range*} | **port-profile** *name*} [**rx** | **tx** | **both**]
6. (Optional) Repeat [Step 5](#) to configure additional ERSPAN sources.
7. (Optional) **filter vlan** {*number* | *range*}
8. (Optional) Repeat [Step 7](#) to configure all source VLANs to filter.
9. **destination ip** *ip_address*
10. (Optional) **ip ttl** *ttl_value*
11. (Optional) **ip prec** *ipp_value*
12. (Optional) **ip dscp** *dscp_value*
13. (Optional) **mtu** *mtu_value*
14. (Optional) **header-type** *value*
15. **erspan-id** *flow_id*
16. **no shut**
17. (Optional) **show monitor session** *session_id*

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

18. (Optional) **exit**
19. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	<pre>config t</pre> <p>Example: <pre>n1000v# config t n1000v(config)#</pre></p>	Places you in the CLI Global Configuration mode.
Step 2	<pre>no monitor session session-number</pre> <p>Example: <pre>n1000v(config)# no monitor session 3</pre></p>	Clears the specified session.
Step 3	<pre>monitor session session-number type erspan-source</pre> <p>Example: <pre>n1000v(config)# monitor session 3 type erspan n1000v(config-erspan-source)#</pre></p>	Creates a session with the given session number and places you in the CLI ERSPAN Source Configuration mode. This configuration is saved in the running configuration.
Step 4	<pre>description description</pre> <p>Example: <pre>n1000v(config-erspan-src)# description my_erspan_session_3 n1000v(config-erspan-src)#</pre></p>	For the specified ERSPAN session, adds a description and saves it in the running configuration. <ul style="list-style-type: none"> • description: up to 32 alphanumeric characters default = blank (no description)
Step 5	<pre>source {interface type {number range} vlan {number range} port-profile name}} [rx tx both]</pre> <p>Example 1: <pre>n1000v(config-erspan-src)# source interface ethernet 2/1-3, ethernet 3/1 rx</pre></p> <p>Example 2: <pre>n1000v(config-erspan-src)# source interface port-channel 2</pre></p> <p>Example 3: <pre>n1000v(config-erspan-src)# source interface vethernet 12 both</pre></p> <p>Example 4: <pre>n1000v(config-erspan-src)# source vlan 3, 6-8 tx</pre></p> <p>Example 5: <pre>n1000v(config-erspan-src)# source port-profile my_port_profile</pre></p>	For the specified session, configures the source(s) and the direction of traffic to monitor, and saves them in the running configuration. <ul style="list-style-type: none"> • type: Specify the interface type—ethernet, port-channel, vethernet. • number: Specify the interface slot/port or range; or the VLAN number or range to monitor. • name: name of an existing port profile. • traffic direction: Specify traffic monitoring to be in one of the following directions: <ul style="list-style-type: none"> – receive (rx) (the VLAN default) – transmit (tx) – both (the interface and port profile default value)
Step 6	(Optional) Repeat Step 5 to configure additional ERSPAN sources.	

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

	Command	Purpose
Step 7	<p>filter vlan {number range}</p> <p>Example: n1000v(config-erspan-src)# filter vlan 3-5, 7</p>	<p>(Optional) For the specified ERSPAN session, configures the VLANs, VLAN lists, or VLAN ranges to be monitored; and saves this in the running configuration.</p> <p>On the monitor port, only the traffic from the VLANs which match the VLAN filter list are replicated to the destination.</p>
Step 8	(Optional) Repeat Step 7 to configure all source VLANs to filter.	
Step 9	<p>destination ip ip_address</p> <p>Example: n1000v(config-erspan-src)# destination ip 10.54.54.1 n1000v(config-erspan-src)#</p>	Configures the IP address of the host to which the encapsulated traffic is sent in this monitor session and saves it in the running configuration.
Step 10	<p>ip ttl ttl_value</p> <p>Example: n1000v(config-erspan-src)# ip ttl 64 n1000v(config-erspan-src)#</p>	<p>(Optional) Specifies the IP time-to-live value, from 1-255, for ERSPAN packets in this monitor session, and saves it in the running configuration.</p> <p>The default is 64.</p>
Step 11	<p>ip prec precedence_value</p> <p>Example: n1000v(config-erspan-src)# ip prec 1 n1000v(config-erspan-src)#</p>	<p>(Optional) Specifies the IP precedence value, from 0-7, for the ERSPAN packets in this monitor session, and saves it in the running configuration.</p> <p>The default value is 0.</p>
Step 12	<p>ip dscp dscp_value</p> <p>Example: n1000v(config-erspan-src)# ip dscp 24 n1000v(config-erspan-src)#</p>	<p>(Optional) Specifies the IP DSCP value, from 0-63, for the ERSPAN packets in this monitor session, and saves it in the running configuration.</p> <p>The default is 0.</p>
Step 13	<p>mtu mtu_value</p> <p>Example: n1000v(config-erspan-src)# mtu 1000 n1000v(config-erspan-src)#</p>	<p>(Optional) Specifies an MTU size (50 - 1500) for ERSPANed packets in this monitor session, and saves it in the running configuration. The 1500 MTU size limit includes a 50-byte overhead added to monitored packets by ERSPAN. Packets larger than this size are truncated.</p> <p>The default is 1500.</p> <p>Note If the ERSPAN destination is a Cisco 6500 switch, truncated ERSPAN packets are dropped unless the no mls verify ip length consistent command is configured on the Cisco 6500.</p>
Step 14	<p>header-type value</p> <p>Example: n1000v(config-erspan-src)# header-type 2 n1000v(config-erspan-src)#</p>	<p>(Optional) Specifies the ERSPAN header type (2 or 3) used for ERSPAN encapsulation for this monitor session.</p> <ul style="list-style-type: none"> 2 = ERSPANv2 header type (the default) 3 = ERSPANv3 header type (Used with NAM setups. Any other type of destination works only with the default v2 headers.)

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

	Command	Purpose
Step 15	erspan-id <i>flow_id</i> Example: n1000v(config-erspan-src)# erspan-id 51	Adds an ERSPAN ID (1-1023) to the session configuration and saves it in the running configuration. The session ERSPAN ID is added to the ERSPAN header of the encapsulated frame and can be used at the termination box to differentiate between various ERSPAN streams of traffic.
Step 16	no shut Example: n1000v(config-erspan-src)# no shut	Enables the ERSPAN session and saves it in the running configuration. By default, the session is created in the shut state.
Step 17	show monitor session <i>session_id</i> Example: n1000v(config-erspan-src)# show monitor session 3	(Optional) Displays the ERSPAN session configuration as it exists in the running configuration.
Step 18	copy running-config startup-config Example: n1000v(config-erspan-src)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Shutting Down a SPAN Session

Use this procedure to discontinue the copying of packets for a SPAN session. You can discontinue copying packets from one source and destination; and then resume for another source and destination.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You know which SPAN session that you want to shut down.
- You can shut down a SPAN session from either Global Configuration mode or Monitor Configuration mode.

SUMMARY STEPS

From Global Configuration mode:

1. **config t**
2. **monitor session** {*session-number* | *session-range* | **all**} **shut**
3. **show monitor**
4. **copy running-config startup-config**

From Monitor Configuration mode:

1. **config t**
2. **monitor session** {*session-number* | *session-range* | **all**} [**type erspan-source**]

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

3. **shut**
4. **show monitor**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	monitor session {session-number session-range all} shut Example: n1000v(config)# monitor session 3 shut n1000v(config)# Example: n1000v(config)# monitor session 3 n1000v(config-monitor)# shut	Shuts down the specified SPAN monitor session(s) from either Global Configuration mode or Monitor-Configuration mode. <ul style="list-style-type: none"> • session-number: Specifies a particular SPAN session number. • session range: Specifies a range of SPAN sessions (allowable = from 1 to 64). • all: Specifies all SPAN monitor sessions.
Step 3	show monitor Example: n1000v(config-monitor)# show monitor	(Optional) Displays the status of the SPAN sessions.
Step 4	copy running-config startup-config Example: n1000v(config-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Resuming a SPAN Session

Use this procedure to resume the copying of packets for a SPAN session. You can discontinue copying packets from one source and destination; and then resume for another source and destination.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You know which SPAN session that you want to configure.
- You can resume the SPAN session from either Global Configuration mode or Monitor Configuration mode.

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

SUMMARY STEPS

From Global Configuration mode:

1. `config t`
2. `no monitor session {session-number | session-range | all} shut`
3. `show monitor`
4. `copy running-config startup-config`

From Monitor Configuration mode:

1. `config t`
2. `monitor session {session-number | session-range | all} [type erspan-source]`
3. `no shut`
4. `show monitor`
5. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	<code>[no] monitor session {<i>session-number</i> <i>session-range</i> all} shut</code> Example: n1000v(config)# no monitor session 3 shut n1000v(config)# Example: n1000v(config)# monitor session 3 n1000v(config-monitor)# no shut	Starts the specified SPAN monitor session(s) from either Global Configuration mode or Monitor-Configuration mode. <ul style="list-style-type: none"> • session-number: Specifies a particular SPAN session number. • session range: Specifies a range of SPAN sessions (allowable = from 1 to 64). • all: Specifies all SPAN monitor sessions.
Step 3	<code>show monitor</code> Example: n1000v(config-monitor)# show monitor	(Optional) Displays the status of all configured SPAN sessions for verification.
Step 4	<code>show monitor session <i>session-id</i></code> Example: n1000v(config-monitor)# show monitor session 3	(Optional) Displays detailed configuration and status of a specific SPAN session for verification.
Step 5	<code>copy running-config startup-config</code> Example: n1000v(config-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

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

Configuring the Allowable ERSPAN Flow IDs

Use this procedure to restrict the allowable range of flow IDs that can be assigned to ERSPAN sessions.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You know the restricted range of ERSPAN flow IDs that you want to designate.
- The available ERSPAN flow IDs are 1-1023. You can restrict the range of available IDs using this procedure.

SUMMARY STEPS

1. **config t**
2. **[no] limit-resource erspan-flow-id minimum *min_val* maximum *max_val***
3. **show running monitor**
4. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	[no] limit-resource erspan-flow-id minimum <i>min_val</i> maximum <i>max_val</i> Example: n1000v(config)# limit-resource erspan-flow-id minimum 20 maximum 40 n1000v(config)# Example: n1000v(config)# no limit-resource erspan-flow-id n1000v(config)#	Restricts the allowable range of ERSPAN flow IDs that can be assigned. <ul style="list-style-type: none"> • Allowable range = 1 to 1023 • Defaults: <ul style="list-style-type: none"> – <i>min_val</i> = 1 – <i>max_val</i> = 1023 <p>The no version of this command removes any configured values and restores default values.</p>
Step 3	show running monitor Example: n1000v(config-monitor)# show monitor session 3	(Optional) Displays changes to the default limit-resource erspan-flow-id values for verification.
Step 4	copy running-config startup-config Example: n1000v(config-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

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

Verifying the SPAN Configuration

To verify the SPAN configuration, use the following commands:

Command	Purpose
show monitor session {all <i>session-number</i> <i>range session-range</i> } [brief]	Displays the SPAN session configuration.
show monitor	Displays Ethernet SPAN information.
module vem <i>module-number</i> execute vemcmd show span	Displays the configured SPAN sessions on a VEM module.
show port-profile name <i>port_profile_name</i>	Displays a port profile.

Example Configurations

This section includes the following example configurations:

- [Example Configuration for a SPAN Session, page 9-20](#)
- [Example Configuration for an ERSPAN Session, page 9-21](#)

Example Configuration for a SPAN Session

To configure a SPAN session, follow these steps:

Step 1 Configure destination ports in access or trunk mode, and enable SPAN monitoring.

```
n1000v# config t
n1000v(config)# interface ethernet 2/5
n1000v(config-if)# switchport
n1000v(config-if)# switchport mode trunk
n1000v(config-if)# no shut
n1000v(config-if)# exit
n1000v(config)#
```

Step 2 Configure a SPAN session.

```
n1000v(config)# no monitor session 1
n1000v(config)# monitor session 1
n1000v(config-monitor)# source interface ethernet 2/1-3
n1000v(config-monitor)# source interface port-channel 2
n1000v(config-monitor)# source port-profile my_profile_src
n1000v(config-monitor)# source vlan 3, 6-8 tx
n1000v(config-monitor)# filter vlan 3-5, 7
n1000v(config-monitor)# destination interface ethernet 2/5
n1000v(config-monitor)# destination port-profile my_profile_dst
n1000v(config-monitor)# no shut
n1000v(config-monitor)# exit
n1000v(config)# show monitor session 1
n1000v(config)# copy running-config startup-config

n1000v(config)# show monitor session 1
session 1
-----
```


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

```

type           : local
state          : up
source intf    :
  rx           : Eth2/1 Eth2/2 Eth2/3
  tx           : Eth2/1 Eth2/2 Eth2/3
  both        : Eth2/1 Eth2/2 Eth2/3
source VLANs   :
  rx           :
  tx           : 3,6,7,8
  both        :
source port-profile :
  rx           : my_profile_src
  tx           : my_profile_src
  both        : my_profile_src
filter VLANs   : 3,4,5,7
destination ports : Eth2/5
destination port-profile : my_profile_dst

n1000v# module vem 3 execute vemcmd show span

VEM SOURCE IP NOT CONFIGURED.

HW SSN ID   ERSPAN ID   HDR VER   DST LTL/IP
      1             local    49,51,52,55,56

```

Example Configuration for an ERSPAN Session

The following example shows how to create an ERSPAN session for a source Ethernet interface and destination IP address on the Cisco Nexus 1000V. Packets arriving at the destination IP are identified by the ID 999 in their header.

```

monitor session 2 type erspan-source
source interface ethernet 3/3
source port-profile my_profile_src
destination ip 10.54.54.1
erspan-id 999
mtu 1000
no shut

show monitor session 2
  session 2
  -----
type           : erspan-source
state          : up
source intf    :
  rx           : Eth3/3
  tx           : Eth3/3
  both        : Eth3/3
source VLANs   :
  rx           :
  tx           :
  both        :
source port-profile :
  rx           : my_profile_src
  tx           : my_profile_src
  both        : my_profile_src
filter VLANs   : filter not specified
destination IP : 10.54.54.1
ERSPAN ID     : 999

```

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

```
ERSPAN TTL           : 64
ERSPAN IP Prec.     : 0
ERSPAN DSCP         : 0
ERSPAN MTU          : 1000
ERSPAN Header Type: 2
```

```
module vem 3 execute vemcmd show span
```

```
VEM SOURCE IP: 10.54.54.10
```

```
HW SSN ID   ERSPAN ID   HDR VER   DST LTL/IP
      1             local    49,51,52,55,56
      2             999      2       10.54.54.1
```

Additional References

For additional information related to implementing SPAN, see the following sections:

- [Related Documents, page 9-22](#)
- [Standards, page 9-22](#)

Related Documents

Related Topic	Document Title
Port profile configuration	<i>Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)</i>
Interface configuration	<i>Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</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.

Feature History for SPAN and ERSPAN

This section provides the SPAN and ERSPAN feature release history.

Feature Name	Releases	Feature Information
Port profile as Local SPAN and ERSPAN source	4.2(1)SV1(4)	You can specify a port profile as a source for local SPAN and ERSPAN monitor traffic.
NAM support for ERSPAN data sources	4.0(4)SV1(3)	NAM support was introduced.
ERSPAN Type III header	4.0(4)SV1(3)	ERSPAN Type III header format was introduced.
SPAN and ERSPAN	4.0(4)SV1(1)	SPAN and ERSPAN were introduced.

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



CHAPTER 10

Configuring SNMP

This chapter describes how to configure the SNMP including users, message encryption, notifications, authentication over TCP, and so forth.

This chapter includes the following sections:

- [Information About SNMP, page 10-1](#)
- [Guidelines and Limitations, page 10-5](#)
- [Default Settings, page 10-5](#)
- [Configuring SNMP, page 10-5](#)
- [Verifying the SNMP Configuration, page 10-13](#)
- [SNMP Example Configuration, page 10-13](#)
- [Additional References, page 10-14](#)
- [Feature History for SNMP, page 10-16](#)

Information About SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

This section includes the following topics:

- [SNMP Functional Overview, page 10-1](#)
- [SNMP Notifications, page 10-2](#)
- [SNMPv3, page 10-2](#)
- [High Availability, page 10-5](#)

SNMP Functional Overview

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.

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

- An SNMP agent—The software component within the managed device that maintains the data for the device and reports these data, as needed, to managing systems. Cisco Nexus 1000V supports the agent and MIB. To enable the SNMP agent, you must define the relationship between the manager and the agent.
- A managed information base (MIB)—The collection of managed objects on the SNMP agent.

SNMP is defined in RFCs 3411 to 3418.



Note

SNMP Role Based Access Control (RBAC) is not supported.

SNMPv1, SNMPv2c, and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security are supported.

SNMP Notifications

A key feature of SNMP is the ability to generate notifications from an SNMP agent. These notifications do not require that requests be sent from the SNMP manager. Notifications can indicate improper user authentication, restarts, the closing of a connection, loss of a connection to a neighbor router, or other significant events.

SNMP notifications are generated as either traps or informs. A trap is an asynchronous, unacknowledged message sent from the agent to the SNMP managers listed in the host receiver table. Informs are asynchronous messages sent from the SNMP agent to the SNMP manager which the manager must acknowledge receipt of.

Traps are less reliable than informs because the SNMP manager does not send any acknowledgment when it receives a trap. Cisco Nexus 1000V cannot determine if the trap was received. An SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If Cisco Nexus 1000V never receives a response, it can send the inform request again.

You can configure Cisco Nexus 1000V to send notifications to multiple host receivers. See the [“Configuring SNMP Notification Receivers” section on page 10-8](#) for more information about host receivers.

SNMPv3

SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are as follows:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

This section includes the following topics:

- [Security Models and Levels for SNMPv1, v2, v3, page 10-3](#)
- [User-Based Security Model, page 10-3](#)

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

- [CLI and SNMP User Synchronization, page 10-4](#)
- [Group-Based SNMP Access, page 10-5](#)

Security Models and Levels for SNMPv1, v2, v3

The security level determines if an SNMP message needs to be protected from disclosure and if the message needs to be authenticated. The various security levels that exist within a security model are as follows:

- noAuthNoPriv—Security level that does not provide authentication or encryption.
- authNoPriv—Security level that provides authentication but does not provide encryption.
- authPriv—Security level that provides both authentication and encryption.

Three security models are available: SNMPv1, SNMPv2c, and SNMPv3. The security model combined with the security level determine the security mechanism applied when the SNMP message is processed.

[Table 10-1](#) identifies what the combinations of security models and levels mean.

Table 10-1 *SNMP Security Models and Levels*

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication.
v3	authNoPriv	HMAC-MD5 or HMAC-SHA	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC) Message Digest 5 (MD5) algorithm or the HMAC Secure Hash Algorithm (SHA).
v3	authPriv	HMAC-MD5 or HMAC-SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard.

User-Based Security Model

SNMPv3 User-Based Security Model (USM) refers to SNMP message-level security and offers the following services:

- Message integrity—Ensures that messages have not been altered or destroyed in an unauthorized manner and that data sequences have not been altered to an extent greater than can occur nonmaliciously.
- Message origin authentication—Ensures that the claimed identity of the user on whose behalf received data was originated is confirmed.

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

- Message confidentiality—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMPv3 authorizes management operations only by configured users and encrypts SNMP messages.

Cisco Nexus 1000V uses two authentication protocols for SNMPv3:

- HMAC-MD5-96 authentication protocol
- HMAC-SHA-96 authentication protocol

Cisco Nexus 1000V uses Advanced Encryption Standard (AES) as one of the privacy protocols for SNMPv3 message encryption and conforms with RFC 3826.

The **priv** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The **priv** option along with the **aes-128** token indicates that this privacy password is for generating a 128-bit AES key. The AES **priv** password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 case-sensitive alphanumeric characters. If you use the localized key, you can specify a maximum of 130 characters.



Note

For an SNMPv3 operation that uses the external AAA server, you must use AES for the privacy protocol in the user configuration on the external AAA server.

CLI and SNMP User Synchronization

SNMPv3 user management can be centralized at the Access Authentication and Accounting (AAA) server level. This centralized user management allows the SNMP agent in Cisco Nexus 1000V to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

Any configuration changes made to the user group, role, or password results in database synchronization for both SNMP and AAA.

Cisco Nexus 1000V synchronizes user configuration in the following ways:

- The authentication passphrase specified in the **snmp-server user** command becomes the password for the CLI user.
- The password specified in the **username** command becomes as the authentication and privacy passphrases for the SNMP user.
- If you delete a user using either SNMP or the CLI, the user is deleted for both SNMP and the CLI.
- User-role mapping changes are synchronized in SNMP and the CLI.
- Role changes (deletions or modifications) from the CLI are synchronized to SNMP.



Note

When you configure passphrase/password in localized key/encrypted format, Cisco Nexus 1000V does not synchronize the password.

Cisco NX-OS holds the synchronized user configuration for 60 minutes by default. See the [“Modifying the AAA Synchronization Time” section on page 10-13](#) for information on how to modify this default value.

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

Group-Based SNMP Access



Note

Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with read access or read-write access.

You can begin communicating with the agent once your username is created, your roles are set up by your administrator, and you are added to the roles.

High Availability

Stateless restarts for SNMP are supported. After a reboot or supervisor switchover, the running configuration is applied.

Guidelines and Limitations

SNMP has the following configuration guidelines and limitations:

- Read-only access to some SNMP MIBs is supported. See the Cisco NX-OS MIB support list at the following URL for more information:
<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>
- SNMP Role Based Access Control (RBAC) is not supported.
- The SNMP set command is supported by the following Cisco MIBs:
 - CISCO-IMAGE-UPGRADE-MIB
 - CISCO-CONFIG-COPY-MIB
- The recommended SNMP polling interval time is 5 minutes.

Default Settings

Table 10-2 lists the default settings for SNMP parameters.

Table 10-2 **Default SNMP Parameters**

Parameters	Default
license notifications	enabled

Configuring SNMP

This section includes the following topics:

- [Configuring SNMP Users, page 10-6](#)

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

- [Enforcing SNMP Message Encryption, page 10-7](#)
- [Creating SNMP Communities, page 10-8](#)
- [Configuring SNMP Notification Receivers, page 10-8](#)
- [Configuring the Notification Target User, page 10-9](#)
- [Enabling SNMP Notifications, page 10-9](#)
- [Disabling LinkUp/LinkDown Notifications on an Interface, page 10-11](#)
- [Enabling a One-time Authentication for SNMP over TCP, page 10-11](#)
- [Assigning the SNMP Switch Contact and Location Information, page 10-11](#)
- [Disabling SNMP, page 10-12](#)
- [Modifying the AAA Synchronization Time, page 10-13](#)



Note

Be aware that the Cisco NX-OS commands for this feature may differ from those used in Cisco IOS.

Configuring SNMP Users

Use this procedure to configure a user for SNMP.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.

SUMMARY STEPS

1. **config t**
2. **snmp-server user *name* [auth {md5 | sha} *passphrase* [auto] [priv [aes-128] *passphrase*] [engineID *id*] [localizedkey]]**
3. **show snmp user**
4. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t Enter configuration commands, one per line. End with CNTL/Z. switch(config)#	Enters global configuration mode.
Step 2	snmp-server user name [auth {md5 sha} passphrase [auto] [priv [aes-128] passphrase] [engineID id] [localizedkey]] Example: switch(config)# snmp-server user Admin auth sha abcd1234 priv abcdefgh	Configures an SNMP user with authentication and privacy parameters. The passphrase can be any case-sensitive alphanumeric string up to 64 characters. If you use the localizedkey keyword, the passphrase can be any case-sensitive alphanumeric string up to 130 characters. The engineID format is a 12-digit colon-separated decimal number.
Step 3	show snmp user Example: switch(config-callhome)# show snmp user	(Optional) Displays information about one or more SNMP users.
Step 4	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

This example shows how to configure the SNMP contact and location information:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# snmp-server user Admin auth sha abcd1234 priv abcdefgh
```

Enforcing SNMP Message Encryption

You can configure SNMP to require authentication or encryption for incoming requests. By default, the SNMP agent accepts SNMPv3 messages without authentication and encryption. When you enforce privacy, Cisco Nexus 1000V responds with an authorizationError for any SNMPv3 PDU request using securityLevel parameter of either noAuthNoPriv or authNoPriv.

Use the following command in global configuration mode to enforce SNMP message encryption for a user:

Command	Purpose
snmp-server user name enforcePriv Example: switch(config)# snmp-server user Admin enforcePriv	Enforces SNMP message encryption for this user.

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

Use the following command in global configuration mode to enforce SNMP message encryption for all users:

Command	Purpose
<pre>snmp-server globalEnforcePriv</pre> <p>Example: switch(config)# snmp-server globalEnforcePriv</p>	Enforces SNMP message encryption for all users.

Creating SNMP Communities

You can create SNMP communities for SNMPv1 or SNMPv2c.

Use the following command in global configuration mode to create an SNMP community string:

Command	Purpose
<pre>snmp-server community name {ro rw}</pre> <p>Example: switch(config)# snmp-server community public ro</p>	Creates an SNMP community string.

Configuring SNMP Notification Receivers

You can configure Cisco Nexus 1000V to generate SNMP notifications to multiple host receivers.

Use the following command in global configuration mode to configure a host receiver for SNMPv1 traps:

Command	Purpose
<pre>snmp-server host ip-address traps version 1 community [udp_port number]</pre> <p>Example: switch(config)# snmp-server host 192.0.2.1 traps version 1 public</p>	Configures a host receiver for SNMPv1 traps. The community can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535.

Use the following command in global configuration mode to configure a host receiver for SNMPv2c traps or informs:

Command	Purpose
<pre>snmp-server host ip-address {traps informs} version 2c community [udp_port number]</pre> <p>Example: switch(config)# snmp-server host 192.0.2.1 informs version 2c public</p>	Configures a host receiver for SNMPv2c traps or informs. The community can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535.

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

Use the following command in global configuration mode to configure a host receiver for SNMPv3 traps or informs:

Command	Purpose
<pre>snmp-server host ip-address {traps informs} version 3 {auth noauth priv} username [udp_port number]</pre> <p>Example: switch(config)# snmp-server host 192.0.2.1 informs version 3 auth NMS</p>	Configures a host receiver for SNMPv2c traps or informs. The username can be any alphanumeric string up to 255 characters. The UDP port number range is from 0 to 65535.



Note

The SNMP manager must know the user credentials (authKey/PrivKey) based on the SNMP engine ID of the Cisco Nexus 1000V device to authenticate and decrypt the SNMPv3 messages.

Configuring the Notification Target User

You must configure a notification target user on the device to send SNMPv3 inform notifications to a notification host receiver.

Cisco Nexus 1000V uses the credentials of the notification target user to encrypt the SNMPv3 inform notification messages to the configured notification host receiver.



Note

For authenticating and decrypting the received INFORM PDU, the notification host receiver should have the same user credentials as configured in Cisco Nexus 1000V to authenticate and decrypt the inform s.

Use the following command in global configuration mode to configure the notification target user:

Command	Purpose
<pre>snmp-server user name [auth {md5 sha} passphrase [auto] [priv [aes-128] passphrase] [engineID id]</pre> <p>Example: switch(config)# snmp-server user NMS auth sha abcd1234 priv abcdefgh engineID 00:00:00:63:00:01:00:10:20:15:10:03</p>	Configures the notification target user with the specified engine ID for notification host receiver. The engineID format is a 12-digit colon-separated decimal number.

Enabling SNMP Notifications

You can enable or disable notifications. If you do not specify a notification name, Cisco Nexus 1000V enables all notifications.

[Table 10-3](#) lists the commands that enable the notifications for Cisco Nexus 1000V MIBs.



Note

The `snmp-server enable traps` command enables both traps and informs, depending on the configured notification host receivers.

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

Table 10-3 Enabling SNMP Notifications

MIB	Related Commands
All notifications	<code>snmp-server enable traps</code>
CISCO-AAA-SERVER-MIB	<code>snmp-server enable traps aaa</code>
ENTITY-MIB	<code>snmp-server enable traps entity</code>
CISCO-ENTITY-FRU-CONTROL-MIB	<code>snmp-server enable traps entity fru</code>
CISCO-LICENSE-MGR-MIB	<code>snmp-server enable traps license</code>
IF-MIB	<code>snmp-server enable traps link</code>
CISCO-PSM-MIB	<code>snmp-server enable traps port-security</code>
SNMPv2-MIB	<code>snmp-server enable traps snmp</code> <code>snmp-server enable traps snmp authentication</code>

The license notifications are enabled by default. All other notifications are disabled by default. Use the following commands in global configuration mode to enable the specified notification:

Command	Purpose
snmp-server enable traps Example: <pre>switch(config)# snmp-server enable traps</pre>	Enables all SNMP notifications.
snmp-server enable traps aaa [server-state-change] Example: <pre>switch(config)# snmp-server enable traps aaa</pre>	Enables the AAA SNMP notifications.
snmp-server enable traps entity [fru] Example: <pre>switch(config)# snmp-server enable traps entity</pre>	Enables the ENTITY-MIB SNMP notifications.
snmp-server enable traps license Example: <pre>switch(config)# snmp-server enable traps license</pre>	Enables the license SNMP notification.
snmp-server enable traps link Example: <pre>switch(config)# snmp-server enable traps link</pre>	Enables the link SNMP notifications.

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

Command	Purpose
snmp-server enable traps port-security Example: switch(config)# snmp-server enable traps port-security	Enables the port security SNMP notifications.
snmp-server enable traps snmp [authentication] Example: switch(config)# snmp-server enable traps snmp	Enables the SNMP agent notifications.

Disabling LinkUp/LinkDown Notifications on an Interface

You can disable linkUp and linkDown notifications on an individual interface. You can use this limit notifications on flapping interface (an interface that transitions between up and down repeatedly).

Use the following command in interface configuration mode to disable linkUp/linkDown notifications for the interface:

Command	Purpose
no snmp trap link-status Example: switch(config-if)# no snmp trap link-status	Disables SNMP link-state traps for the interface. This command is enabled by default.

Enabling a One-time Authentication for SNMP over TCP

You can enable a one-time authentication for SNMP over a TCP session.

Use the following command in global configuration mode to enable one-time authentication for SNMP over TCP:

Command	Purpose
snmp-server tcp-session [auth] Example: switch(config)# snmp-server tcp-session	Enables a one-time authentication for SNMP over a TCP session. The default is disabled.

Assigning the SNMP Switch Contact and Location Information

You can assign the switch contact information, which is limited to 32 characters (without spaces) and the switch location.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.

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

SUMMARY STEPS

1. **config t**
2. **snmp-server contact *name***
3. **snmp-server location *name***
4. **show snmp**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t Enter configuration commands, one per line. End with CNTL/Z. switch(config)#	Enters global configuration mode.
Step 2	snmp-server contact <i>name</i> Example: switch(config)# snmp-server contact Admin	Configures sysContact, which is the SNMP contact name.
Step 3	snmp-server location <i>name</i> Example: switch(config)# snmp-server location Lab-7	Configures sysLocation, which is the SNMP location.
Step 4	show snmp Example: switch(config)# show snmp	(Optional) Displays information about one or more destination profiles.
Step 5	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

This example shows how to configure the SNMP contact and location information:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# snmp contact Admin
switch(config)# snmp location Lab-7
```

Disabling SNMP

You can disable the SNMP protocol on a device.

Use the following command in global configuration mode to disable the SNMP protocol

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

Command	Purpose
<pre>no snmp-server protocol enable</pre> <p>Example: switch(config)# no snmp-server protocol enable</p>	Disables the SNMP protocol. This command is enabled by default.

Modifying the AAA Synchronization Time

You can modify how long Cisco NX-OS holds the synchronized user configuration.

Use the following command in global configuration mode to modify the AAA synchronization time:

Command	Purpose
<pre>snmp-server aaa-user cache-timeout seconds</pre> <p>Example: switch(config)# snmp-server aaa-user cache-timeout 1200.</p>	Configures how long the AAA synchronized user configuration stays in the local cache. The range is from 1 to 86400 seconds. The default is 3600.

Verifying the SNMP Configuration

To display the SNMP configuration, use the following commands:

Command	Purpose
<pre>show running-config snmp [all]</pre>	Displays the SNMP running configuration.
<pre>show snmp</pre>	Displays the SNMP status.
<pre>show snmp community</pre>	Displays the SNMP community strings.
<pre>show snmp context</pre>	Displays the SNMP context mapping.
<pre>show snmp engineID</pre>	Displays the SNMP engineID.
<pre>show snmp group</pre>	Displays SNMP roles.
<pre>show snmp session</pre>	Displays SNMP sessions.
<pre>show snmp trap</pre>	Displays the SNMP notifications enabled or disabled.
<pre>show snmp user</pre>	Displays SNMPv3 users.

SNMP Example Configuration

This example configures sending the Cisco linkUp/Down notifications to one notification host receiver using the Blue VRF and defines two SNMP users, Admin and NMS:

```
config t
snmp-server contact Admin@company.com
snmp-server user Admin auth sha abcd1234 priv abcdefgh
```

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

```
snmp-server user NMS auth sha abcd1234 priv abcdefgh engineID
00:00:00:63:00:01:00:22:32:15:10:03
snmp-server host 192.0.2.1 informs version 3 auth NMS
snmp-server host 192.0.2.1 use-vrf Blue
snmp-server enable traps link cisco
```

Additional References

For additional information related to implementing SNMP, see the following sections:

- [Related Documents, page 10-14](#)
- [Standards, page 10-14](#)
- [MIBs, page 10-15](#)

Related Documents

Related Topic	Document Title
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</i>
MIBs	http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

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

Table 10-4 Supported MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-TC • SNMPv2-MIB • SNMP-COMMUNITY-MIB • SNMP-FRAMEWORK-MIB • SNMP-NOTIFICATION-MIB • SNMP-TARGET-MIB • ENTITY-MIB • IF-MIB • CISCO-ENTITY-EXT-MIB • CISCO-ENTITY-FRU-CONTROL-MIB • CISCO-FLASH-MIB • CISCO-IMAGE-MIB • CISCO-VIRTUAL-NIC-MIB • CISCO-ENTITY-VENDORTYPE-OID-MIB • NOTIFICATION-LOG-MIB • IANA-ADDRESS-FAMILY-NUMBERS-MIB • IANAifType-MIB • IANAiprouteprotocol-MIB • HCNUM-TC 	<p>To locate and download MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml</p>
<ul style="list-style-type: none"> • CISCO-VLAN-MEMBERSHIP-MIB • CISCO-SYSTEM-MIB • CISCO-SYSTEM-EXT-MIB • CISCO-IMAGE-MIB • CISCO-IMAGE-UPGRADE-MIB • CISCO-BRIDGE-MIB • CISCO-CONFIG-COPY-MIB • CISCO-SYSLOG-EXT-MIB • CISCO-PROCESS-MIB • CISCO-AAA-SERVER-MIB • CISCO-AAA-SERVER-EXT-MIB • CISCO-COMMON-ROLES-MIB • CISCO-COMMON-MGMT-MIB 	

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

Feature History for SNMP

This section provides the SNMP feature release history.

Feature Name	Releases	Feature Information
SNMP	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 11

Configuring NetFlow

Use this chapter to configure NetFlow to characterize IP traffic based on its source, destination, timing, and application information, to assess network availability and performance.

This chapter includes the following sections:

- [Information About NetFlow, page 11-1](#)
- [Prerequisites for NetFlow, page 11-8](#)
- [Configuration Guidelines and Limitations, page 11-9](#)
- [Default Settings, page 11-9](#)
- [Enabling the NetFlow Feature, page 11-10](#)
- [Configuring NetFlow, page 11-11](#)
- [Verifying the NetFlow Configuration, page 11-21](#)
- [Configuration Example for NetFlow, page 11-25](#)
- [Additional References, page 11-26](#)
- [Feature History for NetFlow, page 11-27](#)

Information About NetFlow

NetFlow lets you evaluate IP traffic and understand how and where it flows. NetFlow gathers data that can be used in accounting, network monitoring, and network planning.

This section includes the following topics:

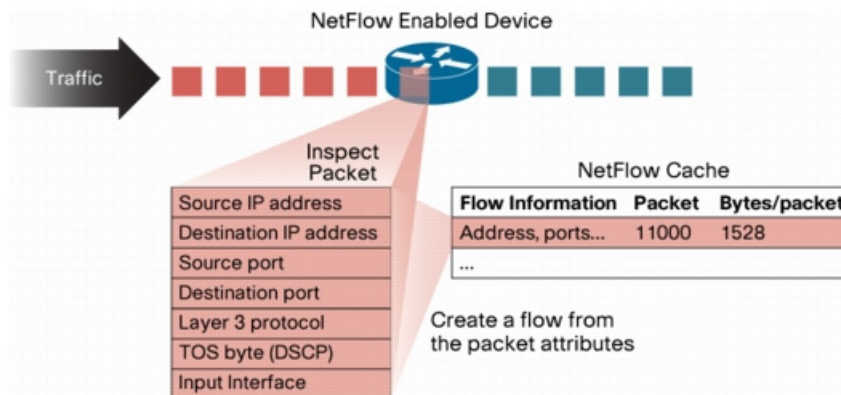
- [What is a Flow, page 11-2](#)
- [Flow Record Definition, page 11-2](#)
- [Accessing NetFlow Data, page 11-5](#)
- [Exporting Flows to the NetFlow Collector Server, page 11-7](#)
- [What NetFlow Data Looks Like, page 11-8](#)
- [High Availability, page 11-8](#)

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

What is a Flow

A flow is a one-directional stream of packets that arrives on a source interface (or subinterface), matching a set of criteria. All packets with the same source/destination IP address, source/destination ports, protocol interface and class of service are grouped into a flow and then packets and bytes are tallied. This condenses a large amount of network information into a database called the NetFlow cache.

Figure 11-1 Creating a Flow in the NetFlow Cache



You create a flow by defining the criteria it gathers. Flows are stored in the NetFlow cache.

Flow information tells you the following:

- Source address tells you who is originating the traffic.
- Destination address tells who is receiving the traffic.
- Ports characterize the application using the traffic.
- Class of service examines the priority of the traffic.
- The device interface tells how traffic is being used by the network device.
- Tallied packets and bytes show the amount of traffic.

Flow Record Definition

A flow record defines the information that NetFlow gathers, such as packets in the flow and the types of counters gathered per flow. You can define new flow records or use the pre-defined Cisco Nexus 1000V flow record.

To create a record, see the [“Defining a Flow Record” procedure on page 11-11](#).

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

The following table describes the criteria defined in a flow record.

Flow record criteria	Description
Match	<p>Defines what information is matched for collection in the flow record.</p> <ul style="list-style-type: none"> • ip: Data collected in the flow record matches one of the following IP options: <ul style="list-style-type: none"> – protocol – tos (type of service) • ipv4: Data collected in the flow record matches one of the following ipv4 address options: <ul style="list-style-type: none"> – source address – destination address • transport: Data collected in the flow record matches one of the following transport options: <ul style="list-style-type: none"> – destination port – source port
Collect	<p>Defines how the flow record collects information.</p> <ul style="list-style-type: none"> • counter: Collects Flow Record information in one of the following formats: <ul style="list-style-type: none"> – bytes: collected in 32-bit counters unless the long 64-bit counter is specified. – packets: collected in 32-bit counters unless the long 64-bit counter is specified. • timestamp sys-uptime: Collects the system up time for the first or last packet in the flow. • transport tcp flags: Collects the TCP transport layer flags for the packets in the flow.

Predefined Flow Records

Cisco Nexus 1000V includes the following pre-defined flow records.

- [Example 11-1 Cisco Nexus 1000V Predefined Flow Record: Netflow-Original, page 11-3](#)
- [Example 11-2 Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Input, page 11-4](#)
- [Example 11-3 Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Output, page 11-4](#)
- [Example 11-4 Cisco Nexus 1000V Predefined Flow Record: Netflow Protocol-Port, page 11-5](#)

Example 11-1 Cisco Nexus 1000V Predefined Flow Record: Netflow-Original

```
n1000v# show flow record netflow-original
Flow record netflow-original:
  Description: Traditional IPv4 input NetFlow with origin ASs
  No. of users: 0
  Template ID: 0
```

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

```

Fields:
  match ipv4 source address
  match ipv4 destination address
  match ip protocol
  match ip tos
  match transport source-port
  match transport destination-port
  match interface input
  match interface output
  match flow direction
  collect routing source as
  collect routing destination as
  collect routing next-hop address ipv4
  collect transport tcp flags
  collect counter bytes
  collect counter packets
  collect timestamp sys-uptime first
  collect timestamp sys-uptime last
n1000v#

```



Note

Although the following lines appear in the output of the **show flow record** command, the commands they are based on are not currently supported in Cisco Nexus 1000V. The use of these commands has no affect on the configuration.

```

collect routing source as
collect routing destination as
collect routing next-hop address ipv4

```

Example 11-2 Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Input

```

n1000v# show flow record netflow ipv4 original-input
Flow record ipv4 original-input:
  Description: Traditional IPv4 input NetFlow
  No. of users: 0
  Template ID: 0
  Fields:
    match ipv4 source address
    match ipv4 destination address
    match ip protocol
    match ip tos
    match transport source-port
    match transport destination-port
    match interface input
    match interface output
    match flow direction
    collect routing source as
    collect routing destination as
    collect routing next-hop address ipv4
    collect transport tcp flags
    collect counter bytes
    collect counter packets
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
n1000v#

```

Example 11-3 Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Output

```

switch# show flow record netflow ipv4 original-output

```


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

```
Flow record ipv4 original-output:
  Description: Traditional IPv4 output NetFlow
  No. of users: 0
  Template ID: 0
  Fields:
    match ipv4 source address
    match ipv4 destination address
    match ip protocol
    match ip tos
    match transport source-port
    match transport destination-port
    match interface input
    match interface output
    match flow direction
    collect routing source as
    collect routing destination as
    collect routing next-hop address ipv4
    collect transport tcp flags
    collect counter bytes
    collect counter packets
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
switch#
```

Example 11-4 Cisco Nexus 1000V Predefined Flow Record: Netflow Protocol-Port

```
switch# show flow record netflow protocol-port
Flow record ipv4 protocol-port:
  Description: Protocol and Ports aggregation scheme
  No. of users: 0
  Template ID: 0
  Fields:
    match ip protocol
    match transport source-port
    match transport destination-port
    match interface input
    match interface output
    match flow direction
    collect counter bytes
    collect counter packets
    collect timestamp sys-uptime first
    collect timestamp sys-uptime last
switch#
```

Accessing NetFlow Data

There are two primary methods used to access NetFlow data:

- [Command Line Interface \(CLI\), page 11-5](#)
- [NetFlow Collector, page 11-6](#)

Command Line Interface (CLI)

To view what is happening in your network now, use the CLI. To see a list of available show commands, see the [“Verifying the NetFlow Configuration” section on page 11-21](#).

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

The CLI uses the following tools to capture and export flow records to the Netflow Collector:

- [Flow Monitor, page 11-6](#)
- [Flow Exporter, page 11-6](#)

Flow Monitor

A flow monitor creates an association between the following NetFlow components:

- a flow record—consisting of matching and collection criteria
- a flow exporter—consisting of the export criteria

This flow monitor association enables a set, consisting of a record and an exporter, to be defined once and re-used many times. Multiple flow monitors can be created for different needs. A flow monitor is applied to a specific interface in a specific direction.

See the “[Defining a Flow Monitor](#)” procedure on page 11-16, and “[Assigning a Flow Monitor to an Interface](#)” procedure on page 11-19.

Flow Exporter

Use the flow exporter to define where and when the flow records are sent from the cache to the reporting server, called the NetFlow Collector.

An exporter definition includes the following.

- Destination IP address
- Source interface
- UDP port number (where the collector is listening)
- Export format



Note

NetFlow export packets use the IP address assigned to the source interface. If the source interface does not have an IP address assigned to it, the exporter will be inactive.

See the “[Defining a Flow Exporter](#)” procedure on page 11-14.

Export Formats



Note

Cisco Nexus 1000V supports the NetFlow Version 9 export format.

Cisco Nexus 1000V supports UDP as the transport protocol for exporting data to up to two exporters per monitor.

NetFlow Collector

You can export NetFlow from the Cisco Nexus 1000V NetFlow cache to a reporting server called the NetFlow Collector. The NetFlow Collector assembles the exported flows and combines them to produce reports used for traffic and security analysis. NetFlow export, unlike SNMP polling, pushes information periodically to the NetFlow reporting collector. The NetFlow cache is constantly filling with flows.

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

Cisco Nexus 1000V searches the cache for flows that have terminated or expired and exports them to the NetFlow collector server. Flows are terminated when the network communication has ended, that is, when a packet contains the TCP FIN flag.

The following steps implement NetFlow data reporting:

- NetFlow records are configured to define the information that NetFlow gathers.
- Netflow monitor is configured to capture flow records to the NetFlow cache.
- NetFlow export is configured to send flows to the collector.
- Cisco Nexus 1000V searches the NetFlow cache for flows that have terminated and exports them to the NetFlow collector server.
- Flows are bundled together based on space availability in the UDP export packet or based on export timer.
- The NetFlow collector software creates real-time or historical reports from the data.

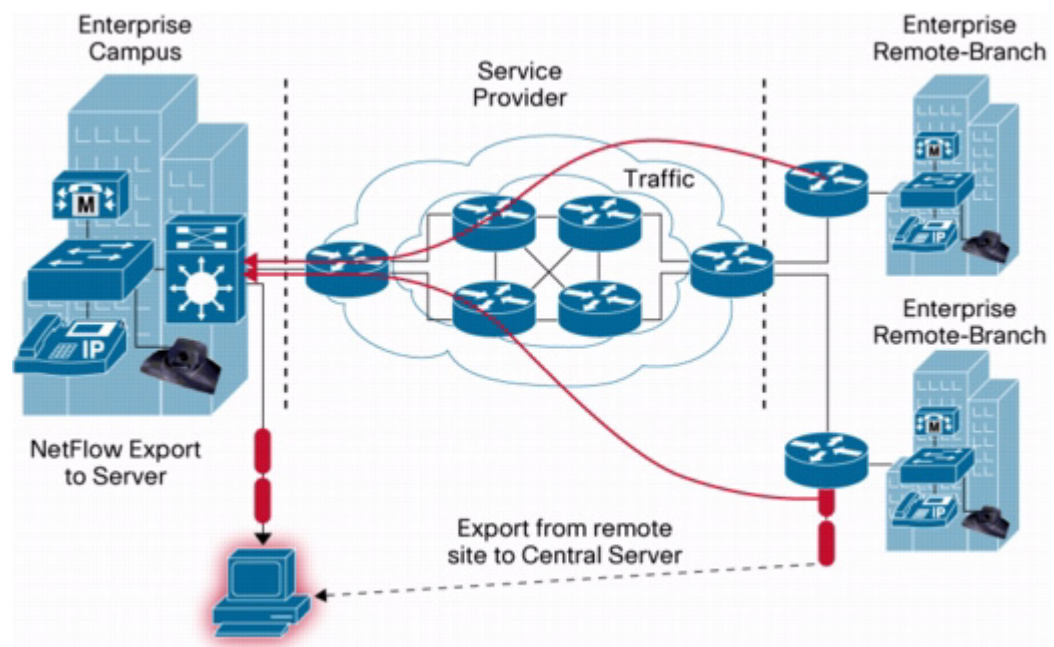
Exporting Flows to the NetFlow Collector Server

Timers determine when a flow is exported to the NetFlow Collector Server.

A flow is ready for export when one of the following occurs:

- The flow is inactive for a certain time during which no new packets are received for the flow.
- The flow has lived longer than the active timer, for example, a long FTP download.
- A TCP flag indicates the flow is terminated. That is, a FIN or RST flag is present.
- The flow cache is full and some flows must be aged out to make room for new flows.

Figure 11-2 *Exporting Flows to the NetFlow Collector Server*

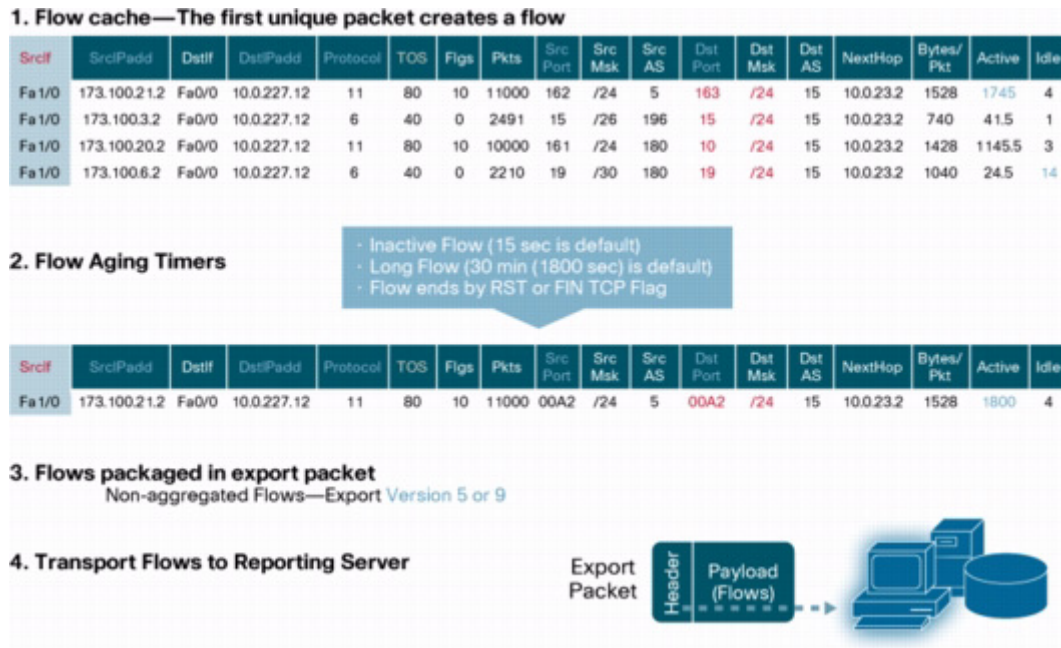


[Send document comments to nexus1k-docfeedback@cisco.com.](mailto:nexus1k-docfeedback@cisco.com)

What NetFlow Data Looks Like

The following figure shows an example of NetFlow data.

Figure 11-3 NetFlow Cache Example



Network Analysis Module

You can also use the Cisco Network Analysis Module (NAM) to monitor NetFlow data sources. NAM enables traffic analysis views and reports such as hosts, applications, conversations, VLAN, and QoS.

To use NAM for monitoring the Cisco Nexus 1000V NetFlow data sources see the *Cisco Nexus 1010 Network Analysis Module Installation and Configuration Note, 4.2*.

High Availability

Cisco Nexus 1000V supports stateful restarts for NetFlow. After a reboot or supervisor switchover, Cisco Nexus 1000V applies the running configuration.

Prerequisites for NetFlow

- You must be aware of resource requirements since NetFlow consumes additional memory and CPU resources.
- Memory and CPU resources are provided by the VEM hosting the flow monitor interface. Resources are limited by the number of CPU cores present on the VEM.

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

Configuration Guidelines and Limitations

NetFlow has the following configuration guidelines and limitations:

- If a source interface is not configured, the NetFlow exporter will remain disabled.
- In Cisco Nexus 1000V, Mgmt0 interface is configured by default as the source interface for an exporter. You can change the source interface if needed.
- Cisco Nexus 1000V includes the following predefined flow records that can be used instead of configuring a new one. For more information, see the “[Flow Record Definition](#)” section on page 11-2:

- netflow-original
Cisco Nexus 1000V predefined traditional IPv4 input NetFlow with origin ASs



Note The routing-related fields in this predefined flow record are ignored.

- netflow ipv4 original-input
Cisco Nexus 1000V predefined traditional IPv4 input NetFlow
- netflow ipv4 original-output
Cisco Nexus 1000V predefined traditional IPv4 output NetFlow
- netflow protocol-port
Cisco Nexus 1000V predefined protocol and ports aggregation scheme
- Up to 256 NetFlow interfaces are allowed per DVS.
- Up to 32 NetFlow interfaces are allowed per host
- A maximum of one flow monitor per interface per direction is allowed.
- Up to 8 flow monitors are allowed per VEM.
- Up to 2 flow exporters are permitted per monitor.
- Up to 32 NetFlow Policies are allowed per DVS.
- Up to 8 NetFlow Policies are allowed per host.
- NetFlow is not supported on port channels.

Default Settings

Table 11-1 lists the default settings for NetFlow parameters.

Table 11-1 *Default NetFlow Parameters*

Parameters	Default
NetFlow version	9
source interface	mgmt0
match	direction and interface (incoming/outgoing)
flow monitor active timeout	1800
flow monitor inactive timeout	15

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

Table 11-1 Default NetFlow Parameters (continued)

Parameters	Default
flow monitor cache size	4096
flow exporter UDP port transport udp command	9995
DSCP	default/best-effort (0)
VRF	default

Enabling the NetFlow Feature

Use this procedure to enable the NetFlow feature.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.

SUMMARY STEPS

- config t**
- feature netflow**
- show feature**
- copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Enters global configuration mode.
Step 2	feature netflow Example: n1000v(config)# feature netflow n1000v(config)#	Enables the NetFlow feature.
Step 3	show feature Example: n1000v(config)# show feature	(Optional) Displays the available features and whether or not they are enabled.
Step 4	copy running-config startup-config Example: n1000v(config-flow-exporter)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

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

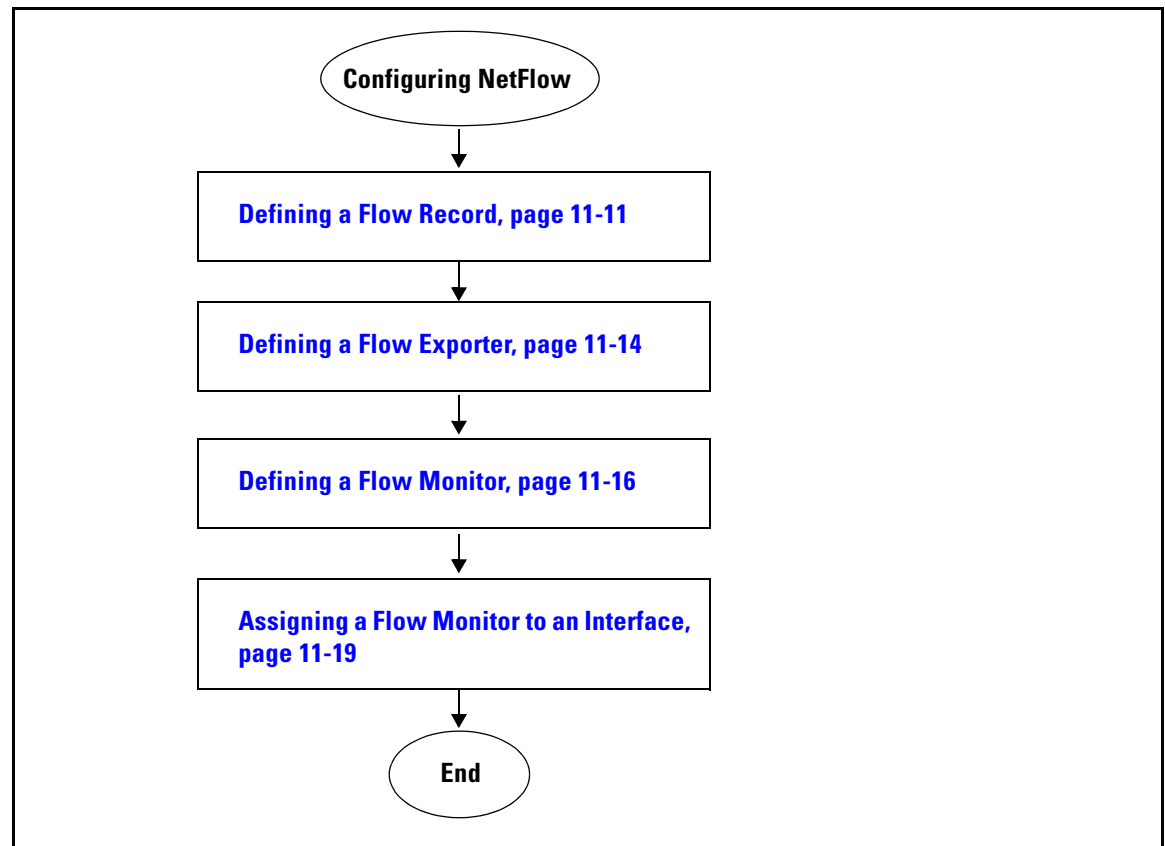
The following is an example for enabling the NetFlow feature:

```
n1000v# config t  
n1000v(config)# feature netflow
```

Configuring NetFlow

The following flow chart is designed to guide you through the netflow configuration process. After completing each procedure, return to the flow chart to make sure you complete all required procedures in the correct sequence.

Flow Chart: Configuring NetFlow



Defining a Flow Record

Use this procedure to create a flow record.



Note

Optionally, you can use the Cisco Nexus 1000V pre-defined record shown in the “[Flow Record Definition](#)” section on page 11-2. See the “[Defining a Flow Monitor](#)” section on page 11-16 to apply a pre-defined record to a flow monitor.

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

BEFORE YOU BEGIN

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

- You know which of the options you want this flow record to match.
- You know which options you want this flow record to collect.

For more information, see the “[Flow Record Definition](#)” section on page 11-2 .



Note

Although the following lines appear in the output of the **show flow record** command, the commands they are based on are not currently supported in Cisco Nexus 1000V. The use of these commands has no affect on the configuration.

```
collect routing source as
collect routing destination as
collect routing next-hop address ipv4
```

SUMMARY STEPS

1. **config t**
2. **flow record *name***
3. **description *string***
4. **match {ip {protocol | tos} | ipv4 {destination address | source address} | transport {destination-port | source-port}}**
5. **collect {counter {bytes [long] | packets [long]} | timestamp sys-uptime | transport tcp flags}**
6. **show flow record [*name*]**
7. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	<i>config t</i> Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	flow record <i>name</i> Example: n1000v(config)# flow record RecordTest n1000v(config-flow-record)#	Creates a Flow Record by name, and places you in the CLI Flow Record Configuration mode for that specific record.
Step 3	description <i>string</i> Example: n1000v(config-flow-record)# description Ipv4Flow	(Optional) Adds a description of up to 63 characters to this Flow Record and saves it in the running configuration.

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

	Command	Purpose
Step 4	<p>match {ip{protocol tos} ipv4{destination address source address} transport {destination-port source-port}}</p> <p>Example: n1000v(config-flow-record)# match ipv4 destination address</p>	<p>Defines the Flow Record to match one of the following and saves it in the running configuration.</p> <ul style="list-style-type: none"> ip: Matches one of the following IP options: <ul style="list-style-type: none"> protocol tos (type of service) ipv4: Matches one of the following ipv4 address options: <ul style="list-style-type: none"> source address destination address transport: Matches one of the following transport options: <ul style="list-style-type: none"> destination port source port
Step 5	<p>collect {counter {bytes [long] packets [long]} timestamp sys-uptime transport tcp flags}</p> <p>Example: n1000v(config-flow-record)# collect counter packets</p>	<p>Specifies a collection option to define the information to collect in the Flow Record and saves it in the running configuration.</p> <ul style="list-style-type: none"> counter: Collects Flow Record information in one of the following formats: <ul style="list-style-type: none"> bytes: collected in 32-bit counters unless the long 64-bit counter is specified. packets: collected in 32-bit counters unless the long 64-bit counter is specified. timestamp sys-uptime: Collects the system up time for the first or last packet in the flow. transport tcp flags: Collects the TCP transport layer flags for the packets in the flow.
Step 6	<p>show flow record [name]</p> <p>Example: n1000v(config-flow-exporter)# show flow record RecordTest</p>	<p>(Optional) Displays information about Flow Records.</p>
Step 7	<p>copy running-config startup-config</p> <p>Example: n1000v(config-flow-exporter)# copy running-config startup-config</p>	<p>(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.</p>

The following is an example for creating a flow record:

```
n1000v# config t
n1000v(config)# flow record RecordTest
n1000v(config-flow-record)# description Ipv4flow
n1000v(config-flow-record)# match ipv4 destination address
n1000v(config-flow-record)# collect counter packets
n1000v(config-flow-record)# show flow record RecordTest
Flow record RecordTest:
  Description: Ipv4flow
  No. of users: 0
```

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

```

Template ID: 0
Fields:
  match ipv4 destination address
  match interface input
  match interface output
  match flow direction
  collect counter packets
n1000v(config-flow-record)#

```

Defining a Flow Exporter

Use this procedure to create a Flow Exporter defining where and how Flow Records are exported to the NetFlow Collector Server.

BEFORE YOU BEGIN

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

- A maximum of two flow exporters per monitor are permitted.
- You know destination IP address of the NetFlow Collector Server.
- You know the source interface that Flow Records are sent from.
- You know the transport UDP that the Collector is listening on.
- Export format version 9 is the version supported.

SUMMARY STEPS

1. **config t**
2. **flow exporter** *name*
3. **description** *string*
4. **destination** {*ipv4-address* | *ipv6-address*}
5. **dscp** *value*
6. **source mgmt** *interface_number*
7. **transport udp** *port-number*
8. **version** 9
9. **option** {**exporter-stats** | **interface-table**} **timeout** *seconds*
10. **template data** **timeout** *seconds*
11. **show flow exporter** [*name*]
12. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in CLI Global Configuration mode.
Step 2	flow exporter name Example: n1000v(config)# flow exporter ExportTest n1000v(config-flow-exporter)#	Creates a Flow Exporter, saves it in the running configuration, and then places you in CLI Flow Exporter Configuration mode.
Step 3	description string Example: n1000v(config-flow-exporter)# description ExportV9	Adds a description of up to 63 characters to this Flow Exporter and saves it in the running configuration.
Step 4	destination {ipv4-address ipv6-address} Example: n1000v(config-flow-exporter)# destination 192.0.2.1	Specifies the IP address of the destination interface for this Flow Exporter and saves it in the running configuration.
Step 5	dscp value Example: n1000v(config-flow-exporter)# dscp 0	Specifies the differentiated services codepoint value for this Flow Exporter, between 0 and 63, and saves it in the running configuration.
Step 6	source mgmt interface_number Example: n1000v(config-flow-exporter)# source mgmt 0	Specifies the interface and its number, from which the Flow Records are sent to the NetFlow Collector Server, and saves it in the running configuration.
Step 7	transport udp port-number Example: n1000v(config-flow-exporter)# transport udp 200	Specifies the destination UDP port, between 0 and 65535, used to reach the NetFlow collector, and saves it in the running configuration.
Step 8	version {9} Example: n1000v(config-flow-exporter)# version 9 n1000v(config-flow-exporter-version-9)#	Specifies NetFlow export version 9, saves it in the running configuration, and places you into the export version 9 configuration mode.
Step 9	option {exporter-stats interface-table sampler-table} timeout value Example: n1000v(config-flow-exporter-version-9)# option exporter-stats timeout 1200	Specifies one of the following version 9 exporter resend timers and its value, between 1 and 86400 seconds, and saves it in the running configuration. <ul style="list-style-type: none"> • exporter-stats • interface-table • sampler-table
Step 10	template data timeout seconds Example: n1000v(config-flow-exporter-version-9)# template data timeout 1200	Sets the template data resend timer and its value, between 1 and 86400 seconds, and saves it in the running configuration.

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

	Command	Purpose
Step 11	show flow exporter [name] Example: n1000v(config-flow-exporter)# show flow exporter	(Optional) Displays information about the Flow Exporter.
Step 12	copy running-config startup-config Example: n1000v(config-flow-exporter)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

The following is an example of creating a flow exporter:

```
n1000v(config)# flow exporter ExportTest
n1000v(config-flow-exporter)# description ExportHamilton
n1000v(config-flow-exporter)# destination 192.0.2.1
n1000v(config-flow-exporter)# dscp 2
n1000v(config-flow-exporter)# source mgmt 0
n1000v(config-flow-exporter)# transport udp 200
n1000v(config-flow-exporter)# version 9
n1000v(config-flow-exporter-version-9)# option exporter-stats timeout 1200
n1000v(config-flow-exporter-version-9)# template data timeout 1200
n1000v(config-flow-exporter-version-9)# show flow exporter ExportTest
Flow exporter ExportTest:
  Description: ExportHamilton
  Destination: 192.0.2.1
  VRF: default (1)
  Destination UDP Port 200
  Source Interface Mgmt0
  DSCP 2
  Export Version 9
    Exporter-stats timeout 1200 seconds
    Data template timeout 1200 seconds
  Exporter Statistics
    Number of Flow Records Exported 0
    Number of Templates Exported 0
    Number of Export Packets Sent 0
    Number of Export Bytes Sent 0
    Number of Destination Unreachable Events 0
    Number of No Buffer Events 0
    Number of Packets Dropped (No Route to Host) 0
    Number of Packets Dropped (other) 0
    Number of Packets Dropped (LC to RP Error) 0
    Number of Packets Dropped (Output Drops) 1
    Time statistics were last cleared: Never
n1000v(config-flow-exporter-version-9)#
```

Defining a Flow Monitor

Use this procedure to create a Flow Monitor and associate a Flow Record and a Flow Exporter to it.

BEFORE YOU BEGIN

- A maximum of one flow monitor per interface per direction is permitted.
- You know the name of an existing Flow Exporter to associate with this flow monitor.

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

- You know the name of an existing Flow Record to associate with this flow monitor. You can use either a flow record you previously created, or one of the following Cisco Nexus 1000V predefined flow records:
 - netflow-original
 - netflow ipv4 original-input
 - netflow ipv4 original-output
 - netflow protocol-port

For more information about Flow Records, see the “Flow Record Definition” section on page 11-2

SUMMARY STEPS

1. **config t**
2. **flow monitor** *name*
3. **description** *string*
4. **exporter** *name*
5. **record** *name*
6. **timeout** {**active** *value* | **inactive** *value*}
7. **cache** {**size** *value*}
8. **show flow monitor** [*name*]
9. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	flow monitor <i>name</i> Example: n1000v(config)# flow monitor MonitorTest n1000v(config-flow-monitor)#	Creates a flow monitor, by name, saves it in the running configuration, and then places you in the CLI Flow Monitor Configuration mode,
Step 3	description <i>string</i> Example: n1000v(config-flow-monitor)# description Ipv4Monitor	(Optional) For the specified flow monitor, adds a descriptive string, of up to 63 alphanumeric characters, and saves it in the running configuration.
Step 4	exporter <i>name</i> Example: n1000v(config-flow-monitor)# exporter Exportv9	For the specified flow monitor, adds an existing flow exporter and saves it in the running configuration.

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

	Command	Purpose
Step 5	record { <i>name</i> netflow { <i>ipv4</i> }} Example using Cisco Nexus 1000V pre-defined record: n1000v(config-flow-monitor)# record netflow-original Example using user-defined record: n1000v(config-flow-monitor)# record RecordTest	For the specified flow monitor, adds an existing flow record and saves it in the running configuration. <ul style="list-style-type: none"> name: The name of a flow record you have previously created, or the name of a Cisco provided pre-defined flow record. netflow: Traditional NetFlow collection schemes <ul style="list-style-type: none"> ipv4: Traditional IPv4 NetFlow collection schemes
Step 6	timeout { <i>active value</i> <i>inactive value</i> } Example: n1000v(config-flow-monitor)# timeout inactive 600	(Optional) For the specified flow monitor, specifies an aging timer and its value for aging entries from the cache, and saves them in the running configuration. <ul style="list-style-type: none"> active: Active, or long, timeout. Allowable values are from 60 to 4092 seconds. Default is 1800. inactive: Inactive or normal timeout. Allowable values are from 15 to 4092 seconds. Default is 15.
Step 7	cache { <i>size value</i> } Example: n1000v(config-flow-monitor)# cache size 15000	(Optional) For the specified flow monitor, specifies the cache size, from 256 to 16384, entries, and saves it in the running configuration. Default is 4096. Note This option is used to limit the impact of the monitor cache on memory and performance.
Step 8	show flow monitor [<i>name</i>] Example: n1000v(config-flow-monitor)# show flow monitor Monitor Test	(Optional) Displays information about existing flow monitors.
Step 9	copy running-config startup-config Example: n1000v(config-flow-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

The following is an example of creating a flow exporter:

```
n1000v(config)# flow monitor MonitorTest
n1000v(config-flow-monitor)# description Ipv4Monitor
n1000v(config-flow-monitor)# exporter ExportTest
n1000v(config-flow-monitor)# record RecordTest
n1000v(config-flow-monitor)# cache size 15000
n1000v(config-flow-monitor)# timeout inactive 600
n1000v(config-flow-monitor)# show flow monitor MonitorTest
Flow Monitor monitorTest:
  Use count: 0
  Inactive timeout: 600
  Active timeout: 1800
  Cache Size: 15000
n1000v(config-flow-monitor)#
```

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

Assigning a Flow Monitor to an Interface

Use this procedure to assign a flow monitor to an interface.

BEFORE YOU BEGIN

- You know the name of the flow monitor you want to use for the interface.
- You know the interface type and its number.

SUMMARY STEPS

1. **config t**
2. **interface** *interface-type interface-number*
3. **ip flow monitor** *name* {input | output}
4. **show flow** *interface-type interface-number*
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	interface <i>interface-type interface-number</i> Example: n1000v(config)# interface veth 2 n1000v(config-if)#	Places you in the CLI Interface Configuration mode for the specified interface.
Step 3	ip flow monitor <i>name</i> {input output} Example: n1000v(config-if)# ip flow monitor MonitorTest output	For the specified interface, assigns a flow monitor for input or output packets and saves it in the running configuration.
Step 4	show flow <i>interface-type interface-number</i> Example: n1000v(config-if)# show flow interface veth 2	(Optional) For the specified interface, displays the NetFlow configuration.
Step 5	copy running-config startup-config Example: n1000v(config-if)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

The following is an example showing how to assign a flow monitor to an interface:

```
n1000v(config)# interface veth 2
n1000v(config-if)# ip flow monitor MonitorTest output
n1000v(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
  Direction: Output
```

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

```
n1000v(config-if)#
```

Adding a Flow Monitor to a Port Profile

You can use this procedure to add a flow monitor to a port profile.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You have already created the flow monitor using the “[Defining a Flow Monitor](#)” procedure on [page 11-16](#).
- If using an existing port profile, you have already created the port profile and you know its name.
- If creating a new port profile, you know the type of interface (Ethernet or vEthernet), and you know the name you want to give it.
- For more information about port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)*.

SUMMARY STEPS

1. **config t**
2. **port-profile [type {ethernet | vethernet}] name**
3. **ip flow monitor name {input | output}**
4. **show port-profile [brief | expand-interface | usage] [name profile-name]**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Description
Step 1	config t Example: n1000v# config t n1000v(config)#	Enters global configuration mode.
Step 2	port-profile [type {ethernet vethernet}] name Example: n1000v(config)# port-profile AccessProf n1000v(config-port-prof)#	Enters port profile configuration mode for the named port profile.
Step 3	ip flow monitor name {input output} Example: n1000v(config-port-prof)# ip flow monitor allaccess4 output n1000v(config-port-prof)#	Applies a named flow monitor to the port profile for either incoming (input) or outgoing (output) traffic.

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

	Command	Description
Step 4	<pre>show port-profile [brief expand-interface usage] [name profile-name]</pre> <p>Example: n1000v(config-port-prof)# show port-profile name AccessProf</p>	(Optional) Displays the configuration for verification.
Step 5	<pre>copy running-config startup-config</pre> <p>Example: n1000v(config-port-prof)# copy running-config startup-config</p>	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

EXAMPLES

This example shows how to add a flow monitor to a port profile:

```
n1000v# config t
n1000v(config)# port-profile AccessProf
n1000v(config-port-prof)# ip flow monitor allaccess4 output
n1000v(config-port-prof)# show port-profile name AccessProf
port-profile AccessProf
  type: vethernet
  status: disabled
  capability l3control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: none
  port-group:
  max ports: 32
  inherit:
  config attributes:
    ip flow monitor allaccess4 output
  evaluated config attributes:
    ip flow monitor allaccess4 output
  assigned interfaces:n1000v(config-port-prof)#
```

Verifying the NetFlow Configuration

To verify the NetFlow configuration, use the commands in [Table 11-2](#):

Table 11-2 Verifying the NetFlow Configuration

Command	Purpose
<code>show flow exporter [name]</code>	Displays information about NetFlow flow exporter maps. See Example 11-5 on page 11-22 .
<code>show flow interface [interface-type number]</code>	Displays information about NetFlow interfaces. See Example 11-6 on page 11-23 .

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

Table 11-2 Verifying the NetFlow Configuration (continued)

Command	Purpose
show flow monitor [<i>name</i> [cache module number statistics module number]]	<p>Displays information about NetFlow flow monitors.</p> <p>Note The show flow monitor cache command differs from the show flow monitor statistics command in that the cache command also displays cache entries . Since each processor has its own cache, all output of these commands is based on the number of processors on the server (also called module or host). When more than one processor is involved in processing packets for a single flow, then the same flow appears for each processor.</p> <p>See the following examples:</p> <ul style="list-style-type: none"> • Example 11-7 Show flow monitor, page 11-23 • Example 11-8 Show flow monitor cache module, page 11-23 • Example 11-9 Show flow monitor statistics module, page 11-24
show flow record [<i>name</i>]	Displays information about NetFlow flow records.

Example 11-5 Show flow exporter

```
n1000v(config-flow-exporter-version-9)# show flow exporter ExportTest
Flow exporter ExportTest:
  Description: ExportHamilton
  Destination: 192.0.2.1
  VRF: default (1)
  Destination UDP Port 200
  Source Interface 2
  DSCP 2
  Export Version 9
    Exporter-stats timeout 1200 seconds
    Data template timeout 1200 seconds
  Exporter Statistics
    Number of Flow Records Exported 0
    Number of Templates Exported 0
    Number of Export Packets Sent 0
    Number of Export Bytes Sent 0
    Number of Destination Unreachable Events 0
    Number of No Buffer Events 0
    Number of Packets Dropped (No Route to Host) 0
    Number of Packets Dropped (other) 0
    Number of Packets Dropped (LC to RP Error) 0
    Number of Packets Dropped (Output Drops) 1
    Time statistics were last cleared: Never
n1000v(config-flow-exporter-version-9)#
```

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

Example 11-6 Show flow interface

```
n1000v(config-if)# show flow interface VEth2
Interface veth2:
  Monitor: MonitorTest
  Direction: Output
n1000v(config-if)#
```

Example 11-7 Show flow monitor

```
n1000v(config)# show flow monitor
Flow Monitor MonitorTest:
  Description: Ipv4Monitor
  Use count: 1
  Flow Record: test
  Flow Exporter: ExportTest
  Inactive timeout: 15
  Active timeout: 1800
  Cache Size: 15000
Flow Monitor MonitorIpv4:
  Description: exit
  Use count: 70
  Flow Record: RecordTest
  Flow Exporter: ExportIpv4
  Inactive timeout: 15
  Active timeout: 1800
  Cache Size: 4096
n1000v(config)#
```

Example 11-8 Show flow monitor cache module

```
n1000v# show flow monitor test_mon cache module 5
Cache type: Normal
Cache size (per-processor): 4096
High Watermark: 2
Flows added: 102
Flows aged: 099
- Active timeout 0
- Inactive timeout 099
- Event aged 0
- Watermark aged 0
- Emergency aged 0
- Permanent 0
- Immediate aged 0
- Fast aged 0

Cache entries on Processor0
- Active Flows: 2
- Free Flows: 4094

      IPV4 SRC ADDR   IPV4 DST ADDR  IP PROT  INTF INPUT  INTF OUTPUT  FLOW DIRN
=====
      0.0.0.0 255.255.255.255 17          Veth1          Input
      7.192.192.10 7.192.192.2 1          Veth1          Eth5/2      Input

Cache entries on Processor1
- Active Flows: 0
- Free Flows: 4096

Cache entries on Processor2
- Active Flows: 1
```

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

- Free Flows: 4095

IPV4 SRC ADDR	IPV4 DST ADDR	IP PROT	INTF INPUT	INTF OUTPUT	FLOW DIRN
7.192.192.10	7.192.192.1	1	Veth1	Eth5/2	Input

Cache entries on Processor3

- Active Flows: 0
- Free Flows: 4096

Cache entries on Processor4

- Active Flows: 0
- Free Flows: 4096

Cache entries on Processor5

- Active Flows: 0
- Free Flows: 4096

Cache entries on Processor6

- Active Flows: 0
- Free Flows: 4096

Cache entries on Processor7

- Active Flows: 0
- Free Flows: 4096

Example 11-9 Show flow monitor statistics module

```
NX-1000v# show flow monitor test_mon statistics module 5
Cache type: Normal
Cache size (per-processor): 4096
High Watermark: 2
Flows added: 105
Flows aged: 103
  - Active timeout 0
  - Inactive timeout 103
  - Event aged 0
  - Watermark aged 0
  - Emergency aged 0
  - Permanent 0
  - Immediate aged 0
  - Fast aged 0

Cache entries on Processor0
  - Active Flows: 0
  - Free Flows: 4096

Cache entries on Processor1
  - Active Flows: 1
  - Free Flows: 4095

Cache entries on Processor2
  - Active Flows: 1
  - Free Flows: 4095

Cache entries on Processor3
  - Active Flows: 0
  - Free Flows: 4096

Cache entries on Processor4
  - Active Flows: 0
  - Free Flows: 4096
```

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

```
Cache entries on Processor5
- Active Flows:          0
- Free Flows:           4096

Cache entries on Processor6
- Active Flows:          0
- Free Flows:           4096

Cache entries on Processor7
- Active Flows:          0
- Free Flows:           4096
```

Example 11-10 Show flow record

```
n1000v(config-flow-record)# show flow record RecordTest
Flow record RecordTest:
  Description: Ipv4flow
  No. of users: 0
  Template ID: 0
  Fields:
    match ipv4 destination address
    match interface input
    match interface output
    match flow direction
    collect counter packets
n1000v(config-flow-record)#
```

Configuration Example for NetFlow

The following example shows how to configure a flow monitor using a new flow record and applying it to an interface.

```
n1000v# config t
n1000v(config)# flow record RecordTest
n1000v(config-flow-record)# description Ipv4flow
n1000v(config-flow-record)# match ipv4 destination address
n1000v(config-flow-record)# collect counter packets
n1000v(config-flow-record)# exit
n1000v(config)# flow exporter ExportTest
n1000v(config-flow-exporter)# description ExportHamilton
n1000v(config-flow-exporter)# destination 192.0.2.1
n1000v(config-flow-exporter)# dscp 2
n1000v(config-flow-exporter)# source mgmt 0
n1000v(config-flow-exporter)# transport udp 200
n1000v(config-flow-exporter)# version 9
n1000v(config-flow-exporter-version-9)# option exporter-stats timeout 1200
n1000v(config-flow-exporter-version-9)# template data timeout 1200
n1000v(config-flow-exporter-version-9)# exit
n1000v(config-flow-exporter)# exit
n1000v(config)# flow monitor MonitorTest
n1000v(config-flow-monitor)# description Ipv4Monitor
n1000v(config-flow-monitor)# exporter ExportTest
n1000v(config-flow-monitor)# record RecordTest
n1000v(config-flow-monitor)# exit
n1000v(config)# interface veth 2/1
n1000v(config-if)# ip flow monitor MonitorTest output
n1000v(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
```

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

```
Direction: Output
n1000v(config-if)#
```

The following example shows how to configure flow monitor using a pre-defined record and applying it to an interface.

```
n1000v# config t
n1000v(config)# flow exporter ExportTest
n1000v(config-flow-exporter)# description ExportHamilton
n1000v(config-flow-exporter)# destination 192.0.2.1
n1000v(config-flow-exporter)# dscp 2
n1000v(config-flow-exporter)# source mgmt 0
n1000v(config-flow-exporter)# transport udp 200
n1000v(config-flow-exporter)# version 9
n1000v(config-flow-exporter-version-9)# option exporter-stats timeout 1200
n1000v(config-flow-exporter-version-9)# template data timeout 1200
n1000v(config-flow-exporter-version-9)# exit
n1000v(config-flow-exporter)# exit
n1000v(config)# flow monitor MonitorTest
n1000v(config-flow-monitor)# description Ipv4Monitor
n1000v(config-flow-monitor)# exporter ExportTest
n1000v(config-flow-monitor)# record netflow-original
n1000v(config-flow-monitor)# exit
n1000v(config)# interface veth 2/1
n1000v(config-if)# ip flow monitor MonitorTest output
n1000v(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
  Direction: Output
n1000v(config-if)#
```

Additional References

For additional information related to implementing NetFlow, see the following sections:

- [Related Documents, page 11-26](#)
- [Standards, page 11-27](#)

Related Documents

Related Topic	Document Title
Cisco NetFlow Overview	http://cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html
Port profiles	<i>Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)</i>
Complete command syntax, command mode, command history, defaults, usage guidelines, and examples for Cisco Nexus 1000V commands.	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</i>

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

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.	—

Feature History for NetFlow

This section provides the NetFlow feature release history.

Feature Name	Releases	Feature Information
NAM support for NetFlow data sources	4.0(4)SV1(3)	NAM support for NetFlow data sources was added.
NetFlow	4.0(4)SV1(1)	NetFlow was introduced.

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



CHAPTER 12

Configuring System Message Logging

This chapter describes how to configure system message logging.

This chapter includes the following topics:

- [Information About System Message Logging, page 12-1](#)
- [System Message Logging Facilities, page 12-2](#)
- [Guidelines and Limitations, page 12-5](#)
- [Default Settings, page 12-5](#)
- [Configuring System Message Logging, page 12-5](#)
- [Verifying the System Message Logging Configuration, page 12-14](#)
- [System Message Logging Example Configuration, page 12-18](#)
- [Additional References, page 12-18](#)
- [Feature History for System Message Logging, page 12-18](#)

Information About System Message Logging

You can use system message logging to control the destination and to filter the severity level of messages that system processes generate. You can configure logging to terminal sessions, a log file, and syslog servers on remote systems.

System message logging is based on [RFC 3164](#). For more information about the system message format and the messages that the device generates, see the *Cisco NX-OS System Messages Reference*.

By default, the device outputs messages to terminal sessions. For information about configuring logging to terminal sessions, see the “[Configuring System Message Logging to Terminal Sessions](#)” section on [page 12-6](#).

[Table 12-1](#) describes the severity levels used in system messages. When you configure the severity level, the system outputs messages at that level and lower.

Table 12-1 System Message Severity Levels

Level	Description
0 – emergency	System unusable
1 – alert	Immediate action needed
2 – critical	Critical condition

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

Table 12-1 System Message Severity Levels (continued)

Level	Description
3 – error	Error condition
4 – warning	Warning condition
5 – notification	Normal but significant condition
6 – informational	Informational message only
7 – debugging	Appears during debugging only

The device logs the most recent 100 messages of severity 0, 1, or 2.

You can configure which system messages should be logged based on the facility that generated the message and its severity level. For information about facilities, see the “[System Message Logging Facilities](#)” section on page 12-2. For information about configuring the severity level by module and facility, see the “[Configuring System Message Logging for Modules](#)” section on page 12-8.

syslog servers run on remote systems that are configured to log system messages based on the syslog protocol. You can configure up to three syslog servers. For information about configuring syslog servers, see the “[Configuring syslog Servers](#)” section on page 12-11.



Note

When the device first initializes, messages are sent to syslog servers only after the network is initialized.

System Message Logging Facilities

Table 12-2 lists the facilities that you can use in system message logging configuration.

Table 12-2 System Message Logging Facilities

Facility	Description
aaa	AAA manager
aclmgr	ACL manager
adjmgr	Adjacency Manager
all	Keyword that represents all facilities
arbiter	Arbiter manager
arp	ARP manager
auth	Authorization system
authpriv	Private authorization system
bootvar	Bootvar
callhome	Call home manager
capability	MIG utilities daemon
cdp	CDP manager
cert-enroll	Certificate enroll daemon
cfs	CFS manager

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

Table 12-2 System Message Logging Facilities (continued)

Facility	Description
clis	CLIS manager
cmpproxy	CMP proxy manager
copp	CoPP manager
core	Core daemon
cron	Cron and at scheduling service
daemon	System daemons
dhcp	DHCP manager
diagclient	GOLD diagnostic client manager
diagmgr	GOLD diagnostic manager
eltn	ELTM manager
ethpm	Ethernet PM manager
evmc	EVMC manager
evms	EVMS manager
feature-mgr	Feature manager
fs-daemon	Fs daemon
ftp	File transfer system
glbp	GLBP manager
hsrp	HSRP manager
im	IM manager
ipconf	IP configuration manager
ipfib	IP FIB manager
kernel	OS kernel
l2fm	L2 FM manager
l2nac	L2 NAC manager
l3vm	L3 VM manager
license	Licensing manager
local0	Local use daemon
local1	Local use daemon
local2	Local use daemon
local3	Local use daemon
local4	Local use daemon
local5	Local use daemon
local6	Local use daemon
local7	Local use daemon
lpr	Line printer system
m6rib	M6RIB manager

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

Table 12-2 System Message Logging Facilities (continued)

Facility	Description
mail	Mail system
mfdm	MFDM manager
module	Module manager
monitor	Ethernet SPAN manager
mrrib	MRIB manager
mvsh	MVSH manager
news	USENET news
nf	NF manager
ntp	NTP manag
otm	GLBP manager
pblr	PBLR manager
pfstat	PFSTAT manager
pixm	PIXM manager
pixmc	PIXMC manager
pktmgr	Packet manager
platform	Platform manager
pltfm_config	PLTFM configuration manager
plugin	Plug-in manager
port-channel	Port channel manager
port_client	Port client manager
port_lb	Diagnostic port loopback test manager
qengine	Q engine manager
radius	RADIUS manager
res_mgr	Resource? manager
rpm	RPM manager
security	Security manager
session	Session manager
spanning-tree	Spanning tree manager
syslog	Internal syslog manager
sysmgr	System manager
tcpudp	TCP and UDP manager
u2	U2 manager
u6rib	U6RIB manager
ufdm	UFDM manager
urib	URIB manager
user	User process

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

Table 12-2 System Message Logging Facilities (continued)

Facility	Description
uucp	Unix-to-Unix copy system
vdc_mgr	VDC manager
vlan_mgr	VLAN manager
vmm	VMM manager
vshd	VSHD manager
xbar	XBAR manager
xbar_client	XBAR client manager
xbar_driver	XBAR driver manager
xml	XML agent

Guidelines and Limitations

System messages are logged to the console and the logfile by default.

Default Settings

Table 12-3 lists the default settings for system message logging.

Table 12-3 System Message Logging Defaults

Parameter	Default
Console logging	Enabled at severity level 2
Monitor logging	Enabled at severity level 5
Log file logging	Enabled to log messages at severity level 5
Module logging	Enabled at severity level 5
Facility logging	Enabled; for severity levels, see the “ System Message Logging Facilities ” section on page 12-2.
Time-stamp units	Seconds
syslog server logging	Disabled
syslog server configuration distribution	Disabled

Configuring System Message Logging

This section includes the following topics:

- [Configuring System Message Logging to Terminal Sessions, page 12-6](#)
- [Restoring System Message Logging Defaults for Terminal Sessions, page 12-7](#)

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

- [Configuring System Message Logging for Modules, page 12-8](#)
- [Restoring System Message Logging Defaults for Modules, page 12-9](#)
- [Configuring System Message Logging for Facilities, page 12-9](#)
- [Restoring System Message Logging Defaults for Facilities, page 12-11](#)
- [Configuring syslog Servers, page 12-11](#)
- [Restoring System Message Logging Defaults for Servers, page 12-12](#)
- [Using a UNIX or Linux System to Configure Logging, page 12-13](#)
- [Displaying Log Files, page 12-13](#)



Note

Be aware that NX-OS commands may differ from the Cisco IOS commands.

Configuring System Message Logging to Terminal Sessions

Use this procedure to log messages by severity level to console, telnet, and SSH sessions.

BEFORE YOU BEGIN

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

- By default, logging is enabled for terminal sessions.

SUMMARY STEPS

1. **terminal monitor**
2. **config t**
3. **logging console** [*severity-level*]
4. **show logging console**
5. **logging monitor** [*severity-level*]
6. **show logging monitor**
7. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	terminal monitor Example: n1000v# terminal monitor n1000v#	Enables the device to log messages to the console.
Step 2	config t Example: n1000v# config t n1000v(config)#	Enters global configuration mode.

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

	Command	Purpose
Step 3	logging console [severity-level] Example: n1000v(config)# logging console 2 n1000v(config)#	Configures the device to log messages to the console session based on a specified severity level or higher. Severity levels, which can range from 0 to 7, are listed in Table 12-1. If the severity level is not specified, the default of 2 is used.
Step 4	show logging console	(Optional) Displays the console logging configuration.
Step 5	logging monitor [severity-level] Example: n1000v(config)# logging monitor 3 n1000v(config)#	Enables the device to log messages to the monitor based on a specified severity level or higher. The configuration applies to telnet and SSH sessions. Severity levels, which can range from 0 to 7, are listed in Table 12-1. If the severity level is not specified, the default of 2 is used.
Step 6	show logging monitor	(Optional) Displays the monitor logging configuration.
Step 7	copy running-config startup-config Example: n1000v(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

Example:
n1000v# terminal monitor
n1000v# config t
n1000v(config)# logging console 2
n1000v(config)# show logging console
Logging console:                enabled (Severity: critical)
n1000v(config)# logging monitor 3
n1000v(config)# show logging monitor
Logging monitor:                enabled (Severity: errors)
n1000v(config)#
n1000v(config)# copy running-config startup-config

```

Restoring System Message Logging Defaults for Terminal Sessions

Use the following commands in the CLI Global Configuration mode to restore default settings for system message logging for terminal sessions.

Command	Description
no logging console [severity-level] Example: n1000v(config)# no logging console n1000v(config)#	Disables the device from logging messages to the console.
no logging monitor [severity-level] Example: n1000v(config)# no logging monitor 3 n1000v(config)#	Disables logging messages to telnet and SSH sessions.

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

Configuring System Message Logging for Modules

Use this procedure to configure the severity level and time-stamp units of messages logged by modules.

SUMMARY STEPS

1. **config t**
2. **logging module** [*severity-level*]
3. **show logging module**
4. **logging timestamp** {microseconds | milliseconds | seconds}
5. **show logging timestamp**
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.
Step 2	logging module [<i>severity-level</i>] Example: n1000v(config)# logging module 3	Enables module log messages that have the specified severity level or higher. Severity levels, which range from 0 to 7, are listed in Table 12-1 . If the severity level is not specified, the default of 5 is used.
Step 3	show logging module	(Optional) Displays the module logging configuration.
Step 4	logging timestamp {microseconds milliseconds seconds} Example: n1000v(config)# logging timestamp microseconds	Sets the logging time-stamp units. The default unit is seconds.
Step 5	show logging timestamp	(Optional) Displays the logging time-stamp units configured.
Step 6	copy running-config startup-config Example: n1000v(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure system message logging for modules.

```
n1000v# config t
n1000v(config)# logging module 3
n1000v(config)# show logging module
Logging linecard:                enabled (Severity: errors)
n1000v(config)# logging timestamp microseconds
n1000v(config)# show logging timestamp
Logging timestamp:                Microseconds
```


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

```
n1000v(config)# copy running-config
```

Restoring System Message Logging Defaults for Modules

Use the following commands in the CLI Global Configuration mode to restore default settings for system message logging for modules.

Command	Description
no logging module [<i>severity-level</i>] Example: n1000v(config)# no logging module 3 n1000v(config)#	Restores the default severity level for logging module system messages.
no logging timestamp { <i>microseconds</i> <i>milliseconds</i> <i>seconds</i> } Example: n1000v(config)# no logging timestamp milliseconds	Resets the logging time-stamp unit to the default (seconds).

Configuring System Message Logging for Facilities

Use this procedure to configure the severity level and time-stamp units of messages logged by facilities.

SUMMARY STEPS

1. **config t**
2. **logging level** *facility severity-level*
3. **show logging level** [*facility*]
4. **logging timestamp** {*microseconds* | *milliseconds* | *seconds*}
5. **show logging timestamp**
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you into CLI Global Configuration mode.

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

	Command	Purpose
Step 2	<pre>logging level facility severity-level</pre> <p>Example: <pre>n1000v(config)# logging level aaa 3 n1000v(config)#</pre></p>	Enables logging messages from the specified facility that have the specified severity level or higher. The facilities are listed in the “ System Message Logging Facilities ” section on page 12-2. Severity levels, which range from 0 to 7, are listed in Table 12-1 . To apply the same severity level to all facilities, use the all facility. For defaults, see the show logging level command.
Step 3	<pre>show logging level [facility]</pre> <p>Example: <pre>n1000v(config)# show logging level aaa</pre></p>	(Optional) Displays the logging level configuration and the system default level by facility. If you do not specify a facility, the device displays levels for all facilities.
Step 4	<pre>logging timestamp {microseconds milliseconds seconds}</pre> <p>Example: <pre>n1000v(config)# logging timestamp microseconds</pre></p>	Sets the logging time-stamp units. The default unit is seconds.
Step 5	<pre>show logging timestamp</pre>	(Optional) Displays the logging time-stamp units configured.
Step 6	<pre>copy running-config startup-config</pre> <p>Example: <pre>n1000v(config)# copy running-config startup-config</pre></p>	(Optional) Copies the running configuration to the startup configuration.

The following example shows how to configure system message logging for facilities.

```
n1000v# config t
n1000v(config)# logging level aaa 3
n1000v(config)# show logging level aaa
Facility           Default Severity   Current Session Severity
-----
aaa                 2                   3

0(emergencies)     1(alerts)          2(critical)
3(errors)           4(warnings)        5(notifications)
6(information)     7(debugging)
logging timestamp microseconds
n1000v(config)# show logging timestamp
Logging timestamp:      Microseconds
copy running-config startup-config
```

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

Restoring System Message Logging Defaults for Facilities

Use the following commands to restore system message logging defaults for facilities.

Command	Description
no logging level [<i>facility severity-level</i>] Example: n1000v(config)# no logging level aaa 3 n1000v(config)#	Restores the default logging severity level for the specified facility. If you do not specify a facility and severity level, the device resets all facilities to their default levels.
no logging timestamp { microseconds milliseconds seconds } Example: n1000v(config)# no logging timestamp milliseconds	Resets the logging time-stamp unit to the default (seconds).

Configuring syslog Servers

Use this procedure to configure syslog servers for system message logging.

SUMMARY STEPS

1. **config t**
2. **logging server** *host* [*severity-level* [**use_vrf** *vrf-name*]]
3. **show logging server**
4. **copy running-config startup-config**

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

DETAILED STEPS

	Command	Purpose
Step 1	<pre>config t</pre> <p>Example: n1000v# config t n1000v(config)#</p>	Enters global configuration mode.
Step 2	<pre>logging server host [severity-level [use-vrf vrf-name]]</pre> <p>Example: n1000v(config)# logging server 10.10.2.2 7</p>	<p>Configures a syslog server at the specified host name or IPv4 or IPv6 address. You can limit logging of messages to a particular VRF by using the use_vrf keyword. Severity levels, which range from 0 to 7, are listed in Table 12-1. The default outgoing facility is local7.</p> <p>The example forwards all messages on facility local 7.</p>
Step 3	<pre>show logging server</pre> <p>Example: n1000v(config)# show logging server Logging server: enabled {10.10.2.2} server severity: debugging server facility: local7</p>	(Optional) Displays the syslog server configuration.
Step 4	<pre>copy running-config startup-config</pre> <p>Example: n1000v(config)# copy running-config startup-config</p>	(Optional) Copies the running configuration to the startup configuration.

Restoring System Message Logging Defaults for Servers

Use the following command to restore server system message logging default.

Command	Description
<pre>no logging server host</pre> <p>Example: n1000v(config)# no logging server host</p>	Removes the logging server for the specified host.

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

Using a UNIX or Linux System to Configure Logging

Use this procedure on a UNIX or Linux system to configure message logging.

BEFORE YOU BEGIN

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

- The following are the UNIX or Linux fields to configure for syslog:

Field	Description
Facility	Creator of the message, which can be auth, authpriv, cron, daemon, kern, lpr, mail, mark, news, syslog, user, local0 through local7, or an asterisk (*) for all. These facility designators allow you to control the destination of messages based on their origin. Note Check your configuration before using a local facility.
Level	Minimum severity level at which messages are logged, which can be debug, info, notice, warning, err, crit, alert, emerg, or an asterisk (*) for all. You can use none to disable a facility.
Action	Destination for messages, which can be a filename, a host name preceded by the at sign (@), or a comma-separated list of users or an asterisk (*) for all logged-in users.

DETAILED STEPS

-
- Step 1** On the UNIX or Linux system, add the following line to the file, `/var/log/myfile.log`:
- ```
facility.level <five tab characters> action
```
- Example:**
- ```
debug.local7                /var/log/myfile.log
```
- Step 2** Create the log file by entering these commands at the shell prompt:
- ```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```
- Step 3** Make sure the system message logging daemon reads the new changes by checking `myfile.log` after entering this command:
- ```
$ kill -HUP -cat /etc/syslog.pid-
```
-

Displaying Log Files

Use this procedure to display messages in the log file.

SUMMARY STEPS

1. `show logging last number-lines`

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

DETAILED STEPS

	Command	Purpose
Step 1	<code>show logging last <i>number-lines</i></code>	Displays the last number of lines in the logging file. You can specify from 1 to 9999 for the last number of lines.

The following example shows the last five lines in the logging file.

```
n1000v# show logging last 5
2008 Aug 31 09:37:04 CP-beta2 %KERN-3-SYSTEM_MSG: packet_recvms
g: truncated packet (size=1514 left=1500) - kernel
2008 Aug 31 09:37:04 CP-beta2 %KERN-3-SYSTEM_MSG: packet_recvms
g: truncated packet (size=1514 left=1500) - kernel
2008 Aug 31 09:37:05 CP-beta2 %KERN-3-SYSTEM_MSG: packet_recvms
g: truncated packet (size=1514 left=1500) - kernel
2008 Aug 31 09:37:05 CP-beta2 %KERN-3-SYSTEM_MSG: packet_recvms
g: truncated packet (size=1514 left=1500) - kernel
2008 Aug 31 09:37:05 CP-beta2 %KERN-3-SYSTEM_MSG: packet_recvms
g: truncated packet (size=1514 left=1500) - kernel
n1000v#
```

Verifying the System Message Logging Configuration

To verify the system message logging configuration, use one of the following commands:

Command	Purpose
<code>show logging console</code>	Displays the console logging configuration. See Example 12-1 on page 12-15
<code>show logging info</code>	Displays the logging configuration. See Example 12-2 on page 12-15
<code>show logging last <i>number-lines</i></code>	Displays the last number of lines of the log file. See Example 12-3 on page 12-16
<code>show logging level [<i>facility</i>]</code>	Displays the facility logging severity level configuration. See Example 12-4 on page 12-16
<code>show logging module</code>	Displays the module logging configuration. See Example 12-5 on page 12-17
<code>show logging monitor</code>	Displays the monitor logging configuration. See Example 12-6 on page 12-17
<code>show logging server</code>	Displays the syslog server configuration. See Example 12-7 on page 12-17
<code>show logging session</code>	Displays the logging session status. See Example 12-8 on page 12-17

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

Command	Purpose
show logging status	Displays the logging status. See Example 12-9 on page 12-17
show logging timestamp	Displays the logging time-stamp units configuration. See Example 12-10 on page 12-17

Example 12-1 show logging console

```
n1000v# show logging console
Logging console:                disabled
n1000v#
```

Example 12-2 show logging info

```
n1000v# show logging info

Logging console:                enabled (Severity: critical)
Logging monitor:               enabled (Severity: notifications)
Logging linecard:              enabled (Severity: notifications)
Logging timestamp:              Seconds
Logging server:                 disabled
Logging logfile:                enabled
                                Name - g/external/messages: Severity - notifications Size - 4194304
```

Facility	Default Severity	Current Session Severity
-----	-----	-----
aaa	2	2
auth	0	0
authpriv	3	3
bootvar	5	5
callhome	2	2
cdp	2	2
cert_enroll	2	2
cfs	3	3
confcheck	2	2
cron	3	3
daemon	3	3
diagclient	2	2
diagmgr	2	2
eth_port_channel	5	5
ethpm	5	5
evmc	5	5
evms	2	2
feature-mgr	2	2
ftp	3	3
ifmgr	5	5
igmp_1	3	3
ip	2	2
ipv6	2	2
kern	6	6
l2fm	2	2
licmgr	6	6
local0	3	3
local1	3	3
local2	3	3
local3	3	3
local4	3	3

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

```

local5                3                3
local6                3                3
local7                3                3
lpr                   3                3
mail                  3                3
mfdm                  2                2
module                5                5
monitor               7                7
msp                   2                2
mvsh                  2                2
news                  3                3
ntp                   2                2
otm                   3                3
pblr                  2                2
pixm                  2                2
pixmc                 2                2
platform              5                5
portprofile           5                5
private-vlan          3                3
radius                2                2
res_mgr               2                2
rpm                   2                2
sal                   2                2
securityd             2                2
sksd                  3                3
stp                   3                3
syslog                3                3
sysmgr                3                3
ufdm                  2                2
urib                  3                3
user                  3                3
uucp                  3                3
vdc_mgr               6                6
vim                   5                5
vlan_mgr              2                2
vms                   5                5
vshd                  5                5
xmlma                 3                3

0(emergencies)        1(alerts)          2(critical)
3(errors)              4(warnings)        5(notifications)
6(information)        7(debugging)
n1000v$

```

Example 12-3 show logging last

```

n1000v# show logging last 5
2008 Jul 29 17:52:42 S22-DCOS %ETHPORT-5-IF_UP: Interface Ethernet2/5 is up in mode access
2008 Jul 29 17:52:43 S22-DCOS %ETHPORT-5-IF_UP: Interface Ethernet2/2 is up in mode trunk
2008 Jul 29 17:52:43 S22-DCOS %ETHPORT-5-IF_UP: Interface Ethernet2/4 is up in mode access
2008 Jul 29 17:53:04 S22-DCOS %SYSMGR-3-BASIC_TRACE: process_cfg_write: PID 1858 with
message rcvd cfg_action from
sap 0x545 for vdc 1 at time 1217353984 .
2008 Jul 29 17:53:04 S22-DCOS clis[2558]: CLI-3-NVDB: Batched send failed for component:
clic
n1000v#

```

Example 12-4 show logging level aaa

```

n1000v# show logging level aaa
Facility           Default Severity           Current Session Severity
-----           -
aaa                 2                           2

```


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

```
0(emergencies)          1(alerts)           2(critical)
3(errors)               4(warnings)        5(notifications)
6(information)         7(debugging)
n1000v#
```

Example 12-5 show logging module

```
n1000v# show logging module
Logging linecard:          enabled (Severity: notifications)
n1000v#
```

Example 12-6 show logging monitor

```
n1000v# show logging monitor
Logging monitor:          enabled (Severity: errors)
n1000v#
```

Example 12-7 show logging server

```
n1000v# show logging server
Logging server:           enabled
{10.10.2.2}
  server severity:        debugging
  server facility:        local7
n1000v#
```

Example 12-8 show logging session status

```
n1000v# show logging session status
Last Action Time Stamp   : Fri Nov 18 11:28:55 1910
Last Action               : Distribution Enable
Last Action Result       : Success
Last Action Failure Reason : none
n1000v#
```

Example 12-9 show logging status

```
n1000v# show logging status
Fabric Distribute        : Enabled
Session State            : IDLE
n1000v#
```

Example 12-10 show logging timestamp

```
n1000v# show logging timestamp
Logging timestamp:       Seconds
n1000v#
```

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

System Message Logging Example Configuration

The following example shows how to configure system message logging:

```
config t
  logging console 3
  logging monitor 3
  logging logfile my_log 6
  logging module 3
  logging level aaa 2
  logging timestamp milliseconds
  logging distribute
  logging server 172.28.254.253
  logging server 172.28.254.254 5 local3
  logging commit
  copy running-config startup-config
```

Additional References

For additional information related to implementing system message logging, see the following sections:

- [Related Documents, page 12-18](#)
- [Standards, page 12-18](#)

Related Documents

Related Topic	Document Title
System messages	<i>Cisco NX-OS System Messages Reference</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</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.	—

Feature History for System Message Logging

This section provides the system message logging feature release history.

Feature Name	Releases	Feature Information
System Message Logging	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 13

Configuring iSCSI Multipath

Revised: May 7, 2014, OL-20458-01

This chapter describes how to configure iSCSI multipath for multiple routes between a server and its storage devices and includes the following topics:

- [Information About iSCSI Multipath, page 13-1](#)
- [Guidelines and Limitations, page 13-4](#)
- [Prerequisites, page 13-5](#)
- [Default Settings, page 13-5](#)
- [Configuring iSCSI Multipath, page 13-5](#)
- [Verifying the iSCSI Multipath Configuration, page 13-18](#)
- [Additional References, page 13-19](#)
- [Feature History for iSCSI Multipath, page 13-19](#)

Information About iSCSI Multipath

This section includes the following topics:

- [Overview, page 13-1](#)
- [Supported iSCSI Adapters, page 13-2](#)
- [iSCSI Multipath Setup on the VMware Switch, page 13-3](#)

Overview

The iSCSI multipath feature sets up multiple routes between a server and its storage devices for maintaining a constant connection and balancing the traffic load. The multipathing software handles all input and output requests and passes them through on the best possible path. Traffic from host servers is transported to shared storage using the iSCSI protocol that packages SCSI commands into iSCSI packets and transmits them on the Ethernet network.

iSCSI multipath provides path failover. In the event a path or any of its components fails, the server selects another available path. In addition to path failover, multipathing reduces or removes potential bottlenecks by distributing storage loads across multiple physical paths.

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

The Cisco Nexus 1000V DVS performs iSCSI multipathing regardless of the iSCSI target. The iSCSI daemon on an ESX server communicates with the iSCSI target in multiple sessions using two or more VMkernel NICs on the host and pinning them to physical NICs on the Cisco Nexus 1000V. Uplink pinning is the only function of multipathing provided by the Cisco Nexus 1000V. Other multipathing functions such as storage binding, path selection, and path failover are provided by VMware code running in the VMkernel.

Setting up iSCSI Multipath is accomplished in the following steps:

1. Uplink Pinning

Each VMkernel port created for iSCSI access is pinned to one physical NIC.

This overrides any NIC teaming policy or port bundling policy. All traffic from the VMkernel port uses only the pinned uplink to reach the upstream switch.

2. Storage Binding

Each VMkernel port is pinned to the VMware iSCSI host bus adapter (VMHBA) associated with the physical NIC to which the VMkernel port is pinned.

The ESX or ESXi host creates the following VMHBAs for the physical NICs.

- In Software iSCSI, only one VMHBA is created for all physical NICs.
- In Hardware iSCSI, one VMHBA is created for each physical NIC that supports iSCSI offload in hardware.

For detailed information about how to use VMware ESX and VMware ESXi systems with an iSCSI storage area network (SAN), see the [iSCSI SAN Configuration Guide](#).

Supported iSCSI Adapters

This section lists the available VMware iSCSI host bus adapters (VMHBAs) and indicates those supported by the Cisco Nexus 1000V.

VMware iSCSI Host Bus Adapter (VMHBA)	Supported on VSM?	Description	Requires VMkernel networking?
Software	Yes	Allows standard NICs to connect the host to a remote iSCSI target on the IP network.	Yes
Dependent Hardware	Yes	Third-party adapter offloads the iSCSI and network processing from host, but not the iSCSI control processing.	Yes
Independent Hardware	No	Third-party adapter offloads iSCSI control and network processing from host. Configured directly from vSphere client and requires no configuration on VSM.	No

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

iSCSI Multipath Setup on the VMware Switch

Before enabling or configuring multipathing, networking must be configured for the software or hardware iSCSI adapter. This involves creating a VMkernel iSCSI port for the traffic between the iSCSI adapter and the physical NIC.

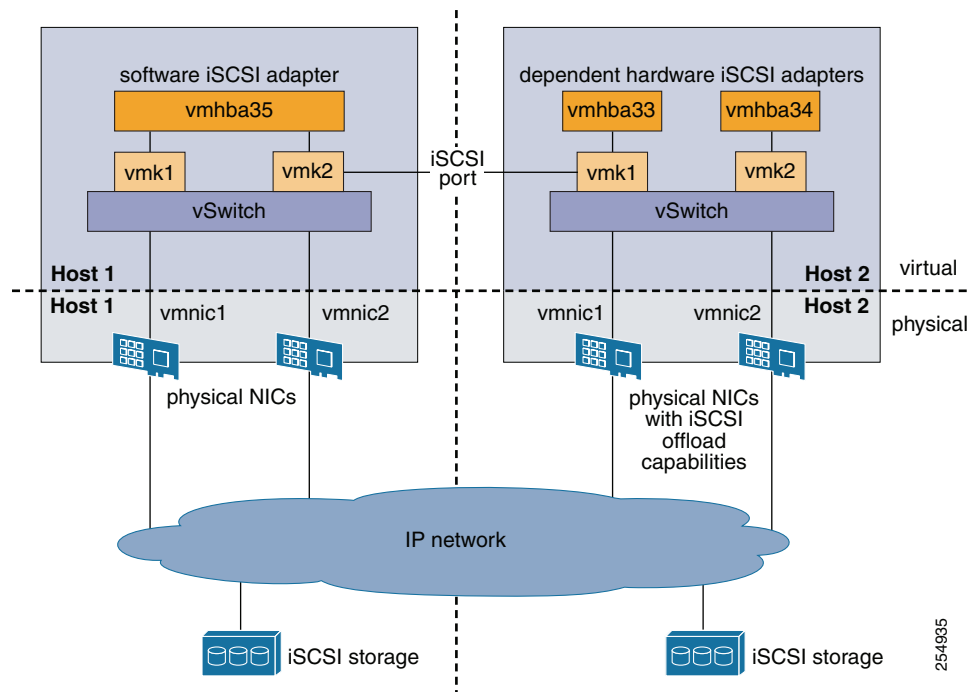
On the vSwitch, uplink pinning is done manually by the admin directly on the vSphere client.

Storage binding is also done manually by the admin directly on the ESX host or using RCLI.

For software iSCSI, only one VMHBA is required for the entire implementation. All VMkernel ports are bound to this adapter. For example, in [Figure 13-1 on page 13-3](#), both vmk1 and vmk2 are bound to VMHBA35.

For hardware iSCSI, a separate adapter is required for each NIC. Each VMkernel port is bound to the adapter of the physical VM NIC to which it is pinned. For example, in [Figure 13-1 on page 13-3](#), vmk1 is bound to VMHBA33, the iSCSI adapter associated with vmnic1 and to which vmk1 is pinned. Similarly vmk2 is bound to VMHBA34.

Figure 13-1 iSCSI Multipath on VMware Virtual Switch



The following are the adapters and NICs used in the hardware and software iSCSI multipathing configuration shown in [Figure 13-1](#).

Software HBA	VMkernel NIC	VM NIC
VMHBA35	1	1
	2	2
Hardware HBA		
VMHBA33	1	1
VMHBA34	2	2

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

Guidelines and Limitations

The following are guidelines and limitations for the iSCSI multipath feature.

- Only port profiles of type vEthernet can be configured with **capability iscsi-multipath**.
- The port profile used for iSCSI multipath must be an access port profile, not a trunk port profile.
- The following are not allowed on a port profile configured with **capability iscsi-multipath**
 - The port profile cannot also be configured with **capability l3 control**.
 - A system VLAN change when the port profile is inherited by VMkernel NIC.
 - An access VLAN change when the port profile is inherited by VMkernel NIC.
 - A port mode change to trunk mode.
- Only VMkernel NIC ports can inherit a port profile configured with **capability iscsi-multipath**.
- The Cisco Nexus 1000V imposes the following limitations if you try to override its automatic uplink pinning.
 - A VMkernel port can only be pinned to one physical NIC.
 - Multiple VMkernel ports can be pinned to a software physical NIC.
 - Only one VMkernel port can be pinned to a hardware physical NIC.
- The iSCSI initiators and storage must already be operational.
- ESX 4.0 Update1 or later supports only software iSCSI multipathing.
- ESX 4.1 or later supports both software and hardware iSCSI multipathing.
- VMkernel ports must be created before enabling or configuring the software or hardware iSCSI for multipathing.
- VMkernel networking must be functioning for the iSCSI traffic.
- Before removing from the DVS an uplink to which an active VMkernel NIC is pinned, you must first remove the binding between the VMkernel NIC and its VMHBA. The following system message displays as a warning:


```
vsm# 2010 Nov 10 02:22:12 sekrishn-bl-vsm %VEM_MGR-SLOT8-1-VEM_SYSLOG_ALERT: sfport
: Removing Uplink Port Eth8/3 (l1 19), when vmknic lveth8/1 (l1 49) is pinned to
this port for iSCSI Multipathing
```
- Hardware iSCSI is new in Cisco Nexus 1000V Release 4.2(1)SV1(4a). If you configured software iSCSI multipathing in a previous release, the following are preserved after upgrade:
 - multipathing
 - software iSCSI uplink pinning
 - VMHBA adapter bindings
 - host access to iSCSI storage

To leverage the hardware offload capable NICs on ESX 4.1, use the [“Converting to a Hardware iSCSI Configuration” procedure on page 13-13](#).
- An iSCSI target and initiator should be in the same subnet.

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

Prerequisites

The iSCSI Multipath feature has the following prerequisites.

- You must understand VMware iSCSI SAN storage virtualization. For detailed information about how to use VMware ESX and VMware ESXi systems with an iSCSI storage area network (SAN), see the [iSCSI SAN Configuration Guide](#).
- You must know how to set up the iSCSI Initiator on your VMware ESX/ESXi host.
- The host is already functioning with one of the following:
 - VMware ESX 4.0.1 Update 01 for software iSCSI
 - VMware ESX 4.1 or later for software and hardware iSCSI
- You must understand iSCSI multipathing and path failover.
- VMware kernel NICs configured to access the SAN external storage are required.

Default Settings

Table 13-1 lists the default settings in the iSCSI Multipath configuration.

Table 13-1 iSCSI Multipath Defaults

Parameter	Default
Type (port-profile)	vEthernet
Description (port-profile)	None
VMware port group name (port-profile)	The name of the port profile
Switchport mode (port-profile)	Access
State (port-profile)	Disabled

Configuring iSCSI Multipath

Use the following procedures to configure iSCSI Multipath.

- [“Uplink Pinning and Storage Binding” procedure on page 13-5](#)
- [“Converting to a Hardware iSCSI Configuration” procedure on page 13-13](#)
- [“Changing the VMkernel NIC Access VLAN” procedure on page 13-15](#)

Uplink Pinning and Storage Binding

Use this section to configure iSCSI multipathing between hosts and targets over iSCSI protocol by assigning the vEthernet interface to an iSCSI multipath port profile configured with a system VLAN.

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

Process for Uplink Pinning and Storage Binding

-
- Step 1** [“Creating a Port Profile for a VMkernel NIC” procedure on page 13-6.](#)
- Step 2** [“Creating VMkernel NICs and Attaching the Port Profile” procedure on page 13-8.](#)
- Step 3** Do one of the following:
- [If you want to override the automatic pinning of NICS, go to “Manually Pinning the NICs” procedure on page 13-9.](#)
 - If not continue with storage binding.
- You have completed uplink pinning. Continue with the next step for storage binding.
- Step 4** [“Identifying the iSCSI Adapters for the Physical NICs” procedure on page 13-11](#)
- Step 5** [“Binding the VMkernel NICs to the iSCSI Adapter” procedure on page 13-13](#)
- Step 6** [“Verifying the iSCSI Multipath Configuration” procedure on page 13-18](#)
-

Creating a Port Profile for a VMkernel NIC

You can use this procedure to create a port profile for a VMkernel NIC.

BEFORE YOU BEGIN

Before starting this procedure, you must know or do the following.

- You have already configured the host with one port channel that includes two or more physical NICs.
- Multipathing must be configured on the interface by using this procedure to create an iSCSI multipath port profile and then assigning the interface to it.
- You are logged in to the CLI in EXEC mode.
- You know the VLAN ID for the VLAN you are adding to this iSCSI multipath port profile.
 - The VLAN must already be created on the Cisco Nexus 1000V.
 - The VLAN that you assign to this iSCSI multipath port profile must be a system VLAN.
 - One of the uplink ports must already have this VLAN in its system VLAN range.
- The port profile must be an access port profile. It cannot be a trunk port profile. This procedure includes steps to configure the port profile as an access port profile.

SUMMARY STEPS

1. `config t`
2. `port-profile type vethernet name`
3. `vmware port-group [name]`
4. `switchport mode access`
5. `switchport access vlan vlanID`
6. `no shutdown`
7. (Optional) `system vlan vlanID`

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

8. `capability iscsi-multipath`
9. `state enabled`
10. (Optional) `show port-profile name`
11. (Optional) `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: n1000v# config t n1000v(config)#	Places you in the CLI Global Configuration mode.
Step 2	port-profile type vethernet name Example: n1000v(config)# port-profile type vethernet VMK-port-profile n1000v(config-port-prof)#	Places you into the CLI Port Profile Configuration mode for the specified port profile. <ul style="list-style-type: none"> • type: Defines the port-profile as Ethernet or vEthernet type. Once configured, this setting cannot be changed. The default is vEthernet type. Note If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports. <ul style="list-style-type: none"> • <i>name:</i> The port profile name can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V.
Step 3	description profile description Example: n1000v(config-port-prof)# description "Port Profile for iSCSI multipath" n1000v(config-port-prof)#	Adds a description to the port profile. This description is automatically pushed to the vCenter Server. <i>profile description:</i> up to 80 ASCII characters Note If the description includes spaces, it must be surrounded by quotations.
Step 4	vmware port-group [name] Example: n1000v(config-port-prof)# vmware port-group VMK-port-profile n1000v(config-port-prof)#	Designates the port-profile as a VMware port group. The port profile is mapped to a VMware port group of the same name. When a vCenter Server connection is established, the port group created in Cisco Nexus 1000V is then distributed to the virtual switch on the vCenter Server. <i>name:</i> The VMware port group name. If you want to map the port profile to a different port group name, use the alternate name.
Step 5	switchport mode access Example: n1000v(config-port-prof)# switchport mode access n1000v(config-port-prof)#	Designates that the interfaces are switch access ports (the default).

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

	Command	Purpose
Step 6	switchport access vlan <i>vlanID</i> Example: n1000v(config-port-prof)# switchport access vlan 254 n1000v(config-port-prof)#	Assigns the system VLAN ID to the access port for this port profile. Note The VLAN assigned to this iSCSI port profile must be a system VLAN.
Step 7	no shutdown Example: n1000v(config-port-prof)# no shutdown n1000v(config-port-prof)#	Administratively enables all ports in the profile.
Step 8	system vlan <i>vlanID</i> Example: n1000v(config-port-prof)# system vlan 254 n1000v(config-port-prof)#	Adds the system VLAN to this port profile. This ensures that, when the host is added for the first time or rebooted later, the VEM will be able to reach the VSM. One of the uplink ports must have this VLAN in its system VLAN range.
Step 9	capability iscsi-multipath Example: n1000v(config-port-prof)# capability iscsi-multipath n1000v(config-port-prof)#	Allows the port to be used for iSCSI multipathing. In vCenter Server, the iSCSI Multipath port profile must be selected and assigned to the VMkernel NIC port.
Step 10	state enabled Example: n1000v(config-port-prof)# state enabled n1000v(config-port-prof)#	Enables the port profile. The configuration for this port profile is applied to the assigned ports, and the port group is created in the VMware vSwitch on the vCenter Server.
Step 11	show port-profile name <i>name</i> Example: n1000v(config-port-prof)# show port-profile name multipath-profile n1000v(config-port-prof)#	(Optional) Displays the current configuration for the port profile.
Step 12	copy running-config startup-config Example: n1000v(config-port-prof)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.
Step 13	“Process for Uplink Pinning and Storage Binding” section on page 13-6	

Creating VMkernel NICs and Attaching the Port Profile

You can use this procedure to create VMkernel NICs and attach a port profile to them which triggers the automatic pinning of the VMkernel NICs to physical NICs.

BEFORE YOU BEGIN

Before starting this procedure, you must know or do the following.

- You have already created a port profile using the procedure, “[Creating a Port Profile for a VMkernel NIC](#)” procedure on page 13-6, and you know the name of this port profile.
- The VMkernel ports are created directly on the vSphere client.

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

- Create one VMkernel NIC for each physical NIC that carries the iSCSI VLAN. The number of paths to the storage device is the same as the number of VMkernel NIC created.
- Step 2 of this procedure triggers automatic pinning of VMkernel NICs to physical NICs, so you must understand the following rules for automatic pinning:
 - A VMkernel NIC is pinned to an uplink only if the VMkernel NIC and the uplink carry the same VLAN.
 - The hardware iSCSI NIC is picked first if there are many physical NICs carrying the iSCSI VLAN.
 - The software iSCSI NIC is picked only if there is no available hardware iSCSI NIC.
 - Two VMkernel NICs are never pinned to the same hardware iSCSI NIC.
 - Two VMkernel NICs can be pinned to the same software iSCSI NIC.

-
- Step 1** Create one VMkernel NIC for each physical NIC that carries the iSCSI VLAN.
- For example, if you want to configure two paths, create two physical NICs on the Cisco Nexus 1000V DVS to carry the iSCSI VLAN. The two physical NICs may carry other vlans. Create two VMkernel NICs for two paths.
- Step 2** Attach the port profile configured with **capability iscsi-multipath** to the VMkernel ports.
- The Cisco Nexus 1000V automatically pins the VMkernel NICs to the physical NICs.
- Step 3** From the ESX host, display the auto pinning configuration for verification.
- ~ # **vemcmd show iscsi pinning**
- Example:**
- ```
~ # vemcmd show iscsi pinning
Vmknict LTL Pinned_Uplink LTL
vmk6 49 vmnic2 19
vmk5 50 vmnic1 18
```
- Step 4** You have completed this procedure. Return to the [“Process for Uplink Pinning and Storage Binding” section on page 13-6](#).
- 

## Manually Pinning the NICs

You can use this procedure to override the automatic pinning of NICs done by the Cisco Nexus 1000V, and manually pin the VMkernel NICs to the physical NICs.



### Note

If the pinning done automatically by Cisco Nexus 1000V is not optimal or if you want to change the pinning, then this procedure describes how to use the vemcmd on the ESX host to override it.

### BEFORE YOU BEGIN

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

- You are logged in to the ESX host.
- You have already created VMkernel NICs and attached a port profile to them, using the [“Creating VMkernel NICs and Attaching the Port Profile” procedure on page 13-8](#).

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

- Before changing the pinning, you must remove the binding between the iSCSI VMkernel NIC and the VMHBA. This procedure includes a step for doing this.
- Manual pinning persists across ESX host reboots. Manual pinning is lost if the VMkernel NIC is moved from the DVS to the vSwitch and back.

**Step 1** List the binding for each VMHBA to identify the binding to remove (iSCSI VMkernel NIC to VMHBA).

**esxcli swiscsi nic list -d vmhbann**

**Example:**

```
esxcli swiscsi nic list -d vmhba33
vmk6
 pNic name: vmnic2
 ipv4 address: 169.254.0.1
 ipv4 net mask: 255.255.0.0
 ipv6 addresses:
 mac address: 00:1a:64:d2:ac:94
 mtu: 1500
 toe: false
 tso: true
 tcp checksum: false
 vlan: true
 link connected: true
 ethernet speed: 1000
 packets received: 3548617
 packets sent: 102313
 NIC driver: bnx2
 driver version: 1.6.9
 firmware version: 3.4.4
vmk5
 pNic name: vmnic3
 ipv4 address: 169.254.0.2
 ipv4 net mask: 255.255.0.0
 ipv6 addresses:
 mac address: 00:1a:64:d2:ac:94
 mtu: 1500
 toe: false
 tso: true
 tcp checksum: false
 vlan: true
 link connected: true
 ethernet speed: 1000
 packets received: 3548617
 packets sent: 102313
 NIC driver: bnx2
 driver version: 1.6.9
 firmware version: 3.4.4
```

**Step 2** Remove the binding between the iSCSI VMkernel NIC and the VMHBA.

**Example:**

```
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk6
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk5
```



**Note** If active iSCSI sessions exist between the host and targets, the iSCSI port cannot be disconnected.

**Step 3** From the EXS host, display the auto pinning configuration.

**~ # vmecmd show iscsi pinning**

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

**Example:**

```
~ # vemcmd show iscsi pinning
Vmknic LTL Pinned_Uplink LTL
vmk6 49 vmnic2 19
vmk5 50 vmnic1 18
```

**Step 4** Manually pin the VMkernel NIC to the physical NIC, overriding the auto pinning configuration.

```
~ # vemcmd set iscsi pinning vmk-ltl vmnic-ltl
```

**Example:**

```
~ # vemcmd set iscsi pinning 50 20
```

**Step 5** Verify the manual pinning.

```
~ # vemcmd show iscsi pinning
```

**Example:**

```
~ # vemcmd show iscsi pinning
Vmknic LTL Pinned_Uplink LTL
vmk6 49 vmnic2 19
vmk5 50 vmnic3 20
```

**Step 6** You have completed this procedure. Return to the [“Process for Uplink Pinning and Storage Binding” section on page 13-6](#).

## Identifying the iSCSI Adapters for the Physical NICs

You can use one of the following procedures in this section to identify the iSCSI adapters associated with the physical NICs.

- [“Identifying iSCSI Adapters on the vSphere Client” procedure on page 13-11](#)
- [“Identifying iSCSI Adapters on the Host Server” procedure on page 13-12](#)

## Identifying iSCSI Adapters on the vSphere Client

You can use this procedure on the vSphere client to identify the iSCSI adapters associated with the physical NICs.

### BEFORE YOU BEGIN

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

- You are logged in to vSphere client.

**Step 1** From the Inventory panel, select a host.

**Step 2** Click the **Configuration** tab.

**Step 3** In the Hardware panel, click **Storage Adapters**.

The dependent hardware iSCSI adapter is displayed in the list of storage adapters.

**Step 4** Select the adapter and click **Properties**.

The iSCSI Initiator Properties dialog box displays information about the adapter, including the iSCSI name and iSCSI alias.

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

- Step 5** Locate the name of the physical NIC associated with the iSCSI adapter.  
The default iSCSI alias has the following format: *driver\_name-vmnic#*, where *vmnic#* is the NIC associated with the iSCSI adapter.
- Step 6** You have completed this procedure. Return to the [“Process for Uplink Pinning and Storage Binding” section on page 13-6](#).
- 

## Identifying iSCSI Adapters on the Host Server

You can use this procedure on the ESX or ESXi host to identify the iSCSI adapters associated with the physical NICs.

### BEFORE YOU BEGIN

Before beginning this procedure, you must do the following:

- You are logged in to the server host.

- Step 1** List the storage adapters on the server.

**esxcfg-scsidevs -a**

**Example:**

```
esxcfg-scsidevs -a
vmhba33 bnx2i unbound iscsi.vmhba33 Broadcom iSCSI Adapter
vmhba34 bnx2i online iscsi.vmhba34 Broadcom iSCSI Adapter
```

- Step 2** For each adapter, list the physical NIC bound to it.

**esxcli swiscsi vmnic list -d adapter-name**

**Example:**

```
esxcli swiscsi vmnic list -d vmhba33 | grep name
vmnic name: vmnic2
esxcli swiscsi vmnic list -d vmhba34 | grep name
vmnic name: vmnic3
```

For the software iSCSI adapter, all physical NICs in the server are listed.

For each hardware iSCSI adaptor, one physical NIC is listed.

- Step 3** You have completed this procedure. Return to the section that pointed you here:
- [“Process for Uplink Pinning and Storage Binding” section on page 13-6](#).
  - [“Process for Converting to a Hardware iSCSI Configuration” section on page 13-14](#).
  - [“Process for Changing the Access VLAN” section on page 13-15](#).
-

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

## Binding the VMkernel NICs to the iSCSI Adapter

You can use this procedure to manually bind the physical VMkernel NICs to the iSCSI adapter corresponding to the pinned physical NICs.

### BEFORE YOU BEGIN

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

- You are logged in to the ESX host.
- You know the iSCSI adapters associated with the physical NICs, found in the [“Identifying the iSCSI Adapters for the Physical NICs” procedure on page 13-11](#).

---

**Step 1** Find the physical NICs to which the VEM has pinned the VMkernel NICs.

**Example:**

```
Vmknic LTL Pinned_Uplink LTL
vmk2 48 vmn2c 18
vmk3 49 vmn3c 19
```

**Step 2** Bind the physical NIC to the iSCSI adapter found when [“Identifying the iSCSI Adapters for the Physical NICs” procedure on page 13-11](#).

**Example:**

```
esxcli swiscsi nic add --adapter vmhba33 --nic vmk2
```

**Example:**

```
esxcli swiscsi nic add --adapter vmhba34 --nic vmk3
```

**Step 3** You have completed this procedure. Return to the section that pointed you here:

- [“Process for Uplink Pinning and Storage Binding” section on page 13-6](#).
  - [“Process for Converting to a Hardware iSCSI Configuration” section on page 13-14](#).
  - [“Process for Changing the Access VLAN” section on page 13-15](#).
- 

## Converting to a Hardware iSCSI Configuration

You can use the procedures in this section on an ESX 4.1 host to convert from a software iSCSI to a hardware iSCSI.

### BEFORE YOU BEGIN

Before starting the procedures in this section, you must know or do the following:

- You have scheduled a maintenance window for this conversion. Converting the setup from software to hardware iSCSI involves a storage update.

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

## Process for Converting to a Hardware iSCSI Configuration

You can use the following steps to convert to a hardware iSCSI configuration:

- 
- Step 1** In the vSphere client, disassociate the storage configuration made on the iSCSI NIC.
- Step 2** Remove the path to the iSCSI targets.
- Step 3** Remove the binding between the VMkernel NIC and the iSCSI adapter using the [“Removing the Binding to the Software iSCSI Adapter” procedure on page 13-14](#)
- Step 4** Move VMkernel NIC from the Cisco Nexus 1000V DVS to the vSwitch.
- Step 5** Install the hardware NICs on the ESX host, if not already installed.
- Step 6** Do one of the following:
- If the hardware NICs are already present on Cisco Nexus 1000V DVS, then continue with the next step.
  - If the hardware NICs are not already present on Cisco Nexus 1000V DVS, then go to the [“Adding the Hardware NICs to the DVS” procedure on page 13-15](#).
- Step 7** Move the VMkernel NIC back from the vSwitch to the Cisco Nexus 1000V DVS.
- Step 8** Find an iSCSI adapter, using the [“Identifying the iSCSI Adapters for the Physical NICs” procedure on page 13-11](#)
- Step 9** Bind the NIC to the adapter, using the [“Binding the VMkernel NICs to the iSCSI Adapter” procedure on page 13-13](#)
- Step 10** Verify the iSCSI multipathing configuration, using the [“Verifying the iSCSI Multipath Configuration” procedure on page 13-18](#)
- 

## Removing the Binding to the Software iSCSI Adapter

You can use this procedure to remove the binding between the iSCSI VMkernel NIC and the software iSCSI adapter.

- 
- Step 1** Remove the iSCSI VMkernel NIC binding to the VMHBA.
- Example:**
- ```
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk6
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk5
```
- Step 2** You have completed this procedure. Return to the [“Process for Converting to a Hardware iSCSI Configuration” section on page 13-14](#).
-

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

Adding the Hardware NICs to the DVS

You can use this procedure, if the hardware NICs are not on Cisco Nexus 1000V DVS, to add the uplinks to the DVS using the vSphere client.

BEFORE YOU BEGIN

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

- You are logged in to vSphere client.
- This procedure requires a server reboot.

-
- Step 1** Select a server from the inventory panel.
 - Step 2** Click the Configuration tab.
 - Step 3** In the Configuration panel, click **Networking**.
 - Step 4** Click the **vNetwork Distributed Switch**.
 - Step 5** Click **Manage Physical Adapters**.
 - Step 6** Select the port profile to use for the hardware NIC.
 - Step 7** Click **Click to Add NIC**.
 - Step 8** In Unclaimed Adapters, select the physical NIC and Click **OK**.
 - Step 9** In the Manage Physical Adapters window, click **OK**.
 - Step 10** Move the iSCSI VMkernel NICs from vSwitch to the Cisco Nexus 1000V DVS.
The VMkernel NICs are automatically pinned to the hardware NICs.
 - Step 11** You have completed this procedure. Return to the [“Process for Converting to a Hardware iSCSI Configuration” section on page 13-14](#).
-

Changing the VMkernel NIC Access VLAN

You can use the procedures in this section to change the access VLAN, or the networking configuration, of the iSCSI VMkernel.

Process for Changing the Access VLAN

You can use the following steps to change the VMkernel NIC access VLAN:

-
- Step 1** In the vSphere client, disassociate the storage configuration made on the iSCSI NIC.
 - Step 2** Remove the path to the iSCSI targets.
 - Step 3** Remove the binding between the VMkernel NIC and the iSCSI adapter using the [“Removing the Binding to the Software iSCSI Adapter” procedure on page 13-14](#)
 - Step 4** Move VMkernel NIC from the Cisco Nexus 1000V DVS to the vSwitch.
 - Step 5** Change the access VLAN, using the [“Changing the Access VLAN” procedure on page 13-16](#).

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

- Step 6** Move the VMkernel NIC back from the vSwitch to the Cisco Nexus 1000V DVS.
 - Step 7** Find an iSCSI adapter, using the “Identifying the iSCSI Adapters for the Physical NICs” procedure on page 13-11
 - Step 8** Bind the NIC to the adapter, using the “Binding the VMkernel NICs to the iSCSI Adapter” procedure on page 13-13
 - Step 9** Verify the iSCSI multipathing configuration, using the “Verifying the iSCSI Multipath Configuration” procedure on page 13-18
-

Changing the Access VLAN

BEFORE YOU BEGIN

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

- You are logged in to the ESX host.
 - You are not allowed to change the access VLAN of an iSCSI multipath port profile if it is inherited by a VMkernel NIC. Use the **show port-profile name profile-name** command to verify inheritance.
-

- Step 1** Remove the path to the iSCSI targets from the vSphere client.
- Step 2** List the binding for each VMHBA to identify the binding to remove (iSCSI VMkernel NIC to VMHBA).

esxcli swiscsi nic list -d vmhbann

Example:

```
esxcli swiscsi nic list -d vmhbba33
vmk6
  pNic name: vmnic2
  ipv4 address: 169.254.0.1
  ipv4 net mask: 255.255.0.0
  ipv6 addresses:
  mac address: 00:1a:64:d2:ac:94
  mtu: 1500
  toe: false
  tso: true
  tcp checksum: false
  vlan: true
  link connected: true
  ethernet speed: 1000
  packets received: 3548617
  packets sent: 102313
  NIC driver: bnx2
  driver version: 1.6.9
  firmware version: 3.4.4
vmk5
  pNic name: vmnic3
  ipv4 address: 169.254.0.2
  ipv4 net mask: 255.255.0.0
  ipv6 addresses:
  mac address: 00:1a:64:d2:ac:94
  mtu: 1500
  toe: false
  tso: true
  tcp checksum: false
  vlan: true
```

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

```
link connected: true
ethernet speed: 1000
packets received: 3548617
packets sent: 102313
NIC driver: bnx2
driver version: 1.6.9
firmware version: 3.4.4
```

Step 3 Remove the iSCSI VMkernel NIC binding to the VMHBA.

Example:

```
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk6 esxcli swiscsi nic remove --adapter
vmhba33 --nic vmk5
```

Step 4 Remove the **capability iscsi-multipath** configuration from the port profile.

no capability iscsi-multipath

Example:

```
n1000v# config t
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# no capability iscsi-multipath
```

Step 5 Remove the system VLAN.

no system vlan *vlanID*

Example:

```
n1000v# config t
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# no system vlan 300
```

Step 6 Change the access VLAN in the port profile.

switchport access vlan *vlanID*

Example:

```
n1000v# config t
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# switchport access vlan 300
```

Step 7 Add the system VLAN.

system vlan *vlanID*

Example:

```
n1000v# config t
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# system vlan 300
```

Step 8 Add the **capability iscsi-multipath** configuration back to the port profile.

capability iscsi-multipath

Example:

```
n1000v# config t
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# capability iscsi-multipath
```

Step 9 You have completed this procedure. Return to the [“Process for Changing the Access VLAN”](#) section on page 13-15.

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

Verifying the iSCSI Multipath Configuration

You can use the commands in this section to verify the iSCSI multipath configuration.

Command	Purpose
<code>~ # vemcmd show iscsi pinning</code>	Displays the auto pinning of VMkernel NICs See Example 13-1 on page 13-18 .
<code>esxcli swiscsi nic list -d vmhba33</code>	Displays the iSCSI adapter binding of VMkernel NICs. See Example 13-2 on page 13-18
<code>show port-profile [brief expand-interface usage] [name profile-name]</code>	Displays the port profile configuration. See Example 13-3 on page 13-18

Example 13-1 ~ # vemcmd show iscsi pinning

```
~ # vemcmd show iscsi pinning
Vmknict   LTL       Pinned_Uplink   LTL
vmk6      49        vmnic2          19
vmk5      50        vmnic1          18
```

Example 13-2 esxcli swiscsi nic list -d vmhba33

```
esxcli swiscsi nic list -d vmhba33
vmk6
  pNic name: vmnic2
  ipv4 address: 169.254.0.1
  ipv4 net mask: 255.255.0.0
  ipv6 addresses:
  mac address: 00:1a:64:d2:ac:94
  mtu: 1500
  toe: false
  tso: true
  tcp checksum: false
  vlan: true
  link connected: true
  ethernet speed: 1000
  packets received: 3548617
  packets sent: 102313
  NIC driver: bnx2
  driver version: 1.6.9
  firmware version: 3.4.4
```

Example 13-3 show port-profile name iscsi-profile

```
n1000v# show port-profile name iscsi-profile
port-profile iscsi-profile
  type: Vethernet
  description:
  status: enabled
  max-ports: 32
  inherit:
  config attributes:
  evaluated config attributes:
  assigned interfaces:
  port-group:
```

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

```
system vlans: 254
capability l3control: no
capability iscsi-multipath: yes
port-profile role: none
port-binding: static
n1000v#
```

Additional References

For additional information related to implementing iSCSI Multipath, see the following sections:

- [Related Documents, page 13-19](#)
- [Standards, page 13-19](#)

Related Documents

Related Topic	Document Title
VMware SAN Configuration	VMware SAN Configuration Guide
Port Profile Configuration	<i>Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.2(1)SV1(4)</i>
Interface Configuration	<i>Cisco Nexus 1000V Interface Configuration Guide, Release 4.2(1)SV1(4a)</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples for all Cisco Nexus 1000V commands.	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4)</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.	—

Feature History for iSCSI Multipath

Table 13-2 lists the release history for the iSCSI Multipath feature.

Table 13-2 Feature History for iSCSI Multipath

Feature Name	Releases	Feature Information
Hardware iSCSI Multipath	4.2(1)SV1(4)	Added support for hardware iSCSI Multipath.
iSCSI Multipath	4.0(4)SV1(2)	The iSCSI Multipath feature was added.

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



CHAPTER 14

Configuring VSM Backup and Recovery

This chapter describes how to configure the backup and recovery procedures on the Virtual Supervisor Module (VSM).

This chapter includes the following sections:

- [Information About VSM Backup and Recovery, page 14-1](#)
- [Guidelines and Limitations, page 14-1](#)
- [Configuring VSM Backup and Recovery, page 14-2](#)
- [Additional References, page 14-22](#)
- [Feature History for VSM Backup and Recovery, page 14-23](#)

Information About VSM Backup and Recovery

You can use the VSM backup and recovery procedure to create a template from which the VSMs can be re-created in the event that both VSMs fail in a high availability (HA) environment.



Note

We recommend that you do periodic backups after the initial backup to ensure that you have the most current configuration. See the [“Performing a Periodic Backup” section on page 14-8](#).

Guidelines and Limitations

VSM backup and recovery has the following configuration guidelines and limitations:

- Backing up the VSM VM is a onetime task.
- Backing up the VSM VM requires coordination between the network administrator and the server administrator.
- The following procedures are applicable starting with Release 4.0(4)SV1(3) and later releases.
- These procedures are not for upgrades and downgrades.
- These procedures require that the restoration is done on the VSM with the same release as the one from which the backup was made.
- Configuration files do not have enough information to re-create a VSM.

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

Configuring VSM Backup and Recovery

This section includes the following topics:

- [Performing a Backup of the VSM VM, page 14-2](#)
- [Performing a Periodic Backup, page 14-8](#)
- [Recovering the VSM, page 14-8](#)

**Note**

Be aware that Cisco NX-OS commands might differ from the Cisco IOS commands.

Backing Up the VSM

This section includes the following topics:

- [Performing a Backup of the VSM VM, page 14-2](#)
- [Performing a Periodic Backup, page 14-8](#)

Performing a Backup of the VSM VM

This section describes how to create a backup of the VSM.

BEFORE YOU BEGIN

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

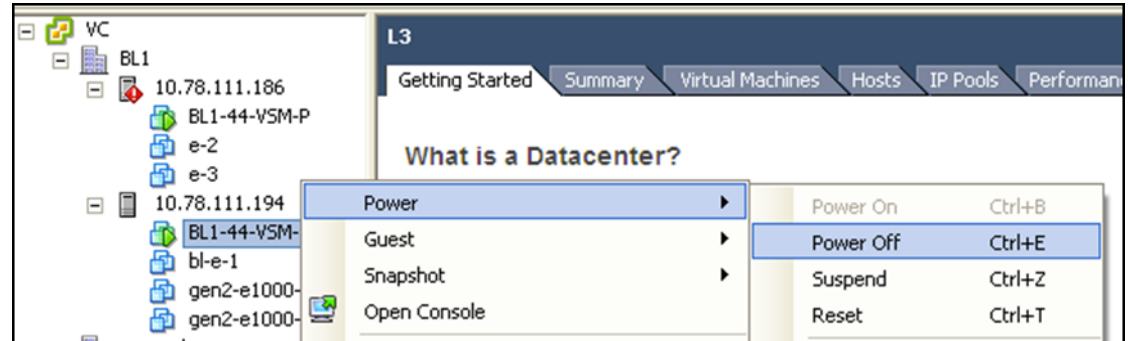
- If the VSM is on a Virtual Ethernet Module (VEM) host, you must configure the management VLAN as a system VLAN.
- Enter the **copy running-config startup-config** command at the VSM before beginning this procedure.
- This procedure is required when there is a Certificate change, Extension key change, after an upgrade to a new release, and installation of the license.

PROCEDURE

-
- Step 1** Open the vSphere Client.
The vSphere Client window opens. See [Figure 14-1](#).

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

Figure 14-1 vSphere Client Window

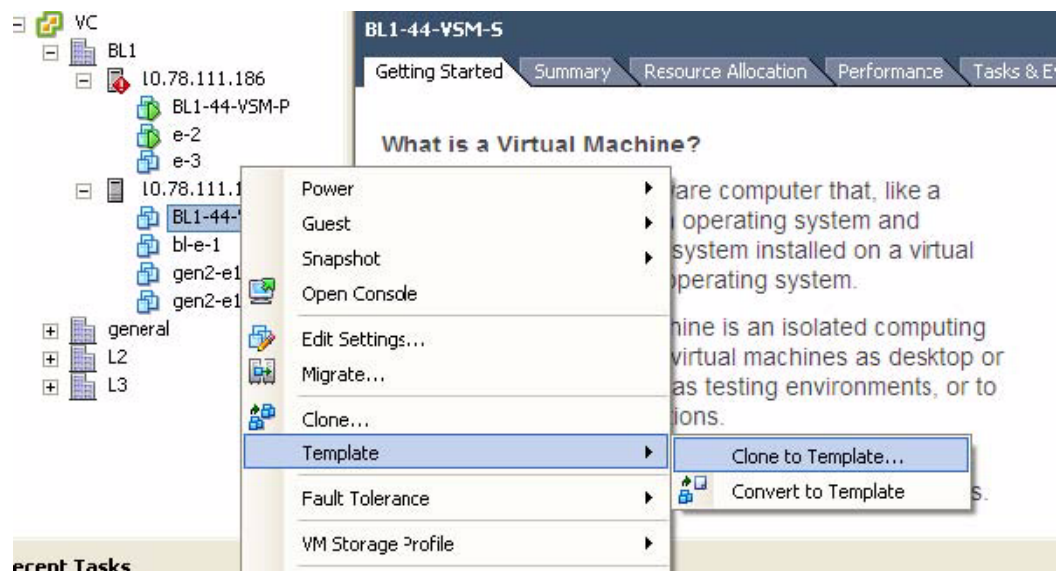


Step 2 In the left navigation pane, right-click the standby VSM.

A drop-down list appears.

Step 3 Choose **Power > Power Off**.

Figure 14-2 Clone to Template Window



Step 4 In the left navigation pane, right-click the standby VSM.


A drop-down list appears.

Step 5 Choose **Template > Clone to Template**.

The Clone Virtual Machine to Template screen opens. See [Figure 14-3](#).

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

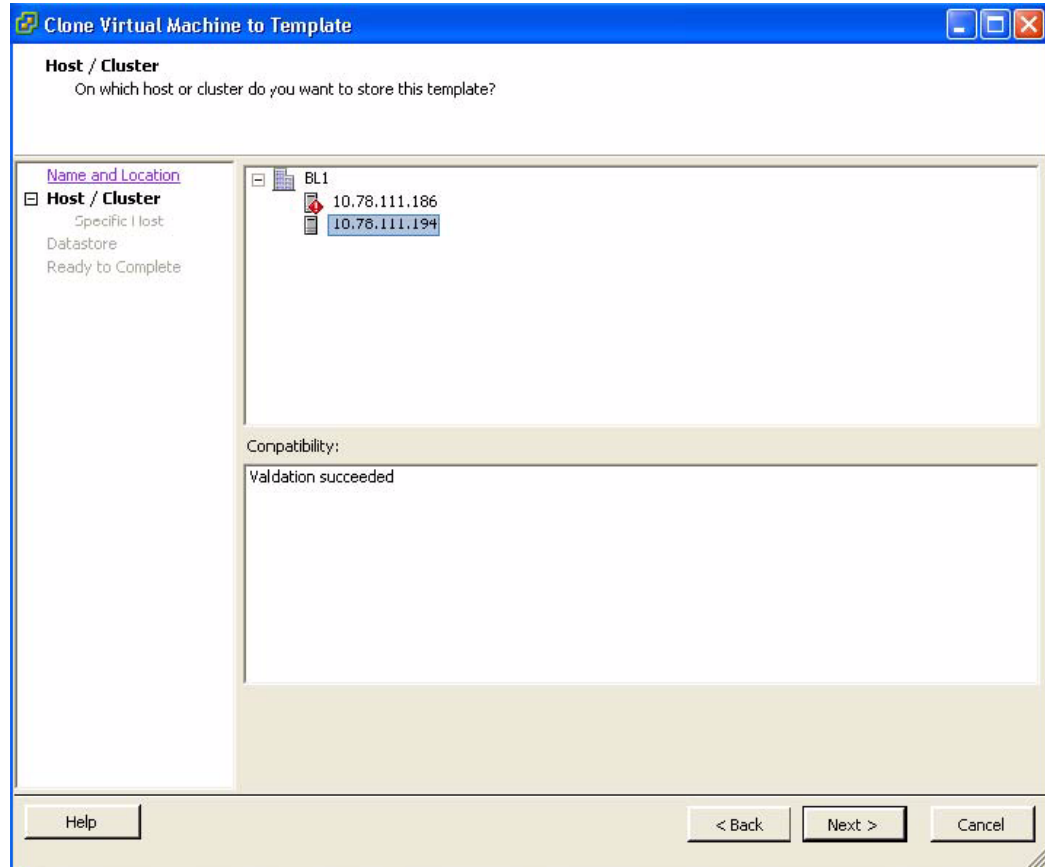
Figure 14-3 Clone Virtual Machine to Template Screen



- Step 6** In the Template Name field, enter a name.
- Step 7** In the Template Inventory Location pane, choose a location for the template.
- Step 8** Click **Next**.
- The Choosing the Host screen opens. See [Figure 14-4](#).

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

Figure 14-4 Choosing the Host Screen



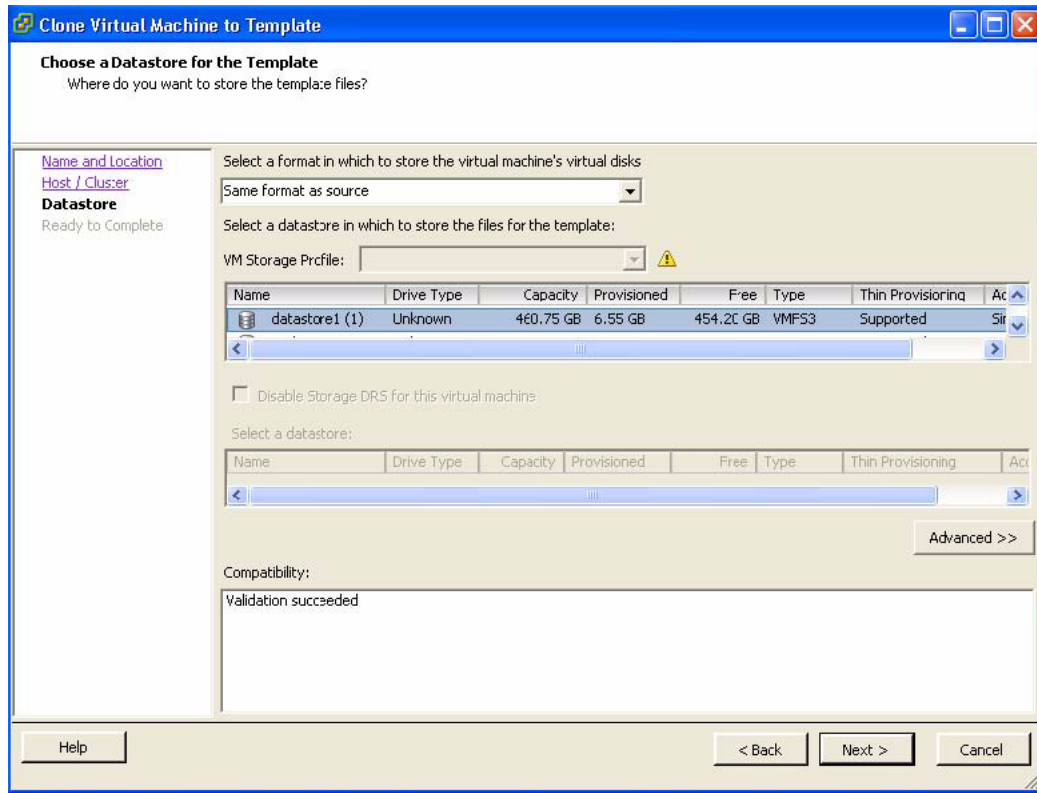
Step 9 Choose the host on which the template will be stored.

Step 10 Click **Next**.

The Choosing a Datastore screen opens. See [Figure 14-5](#).

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

Figure 14-5 Choosing a Datastore Screen



Step 11 In the Select a format in which to store the virtual machine's virtual disks drop-down list, choose **Same format as source**.

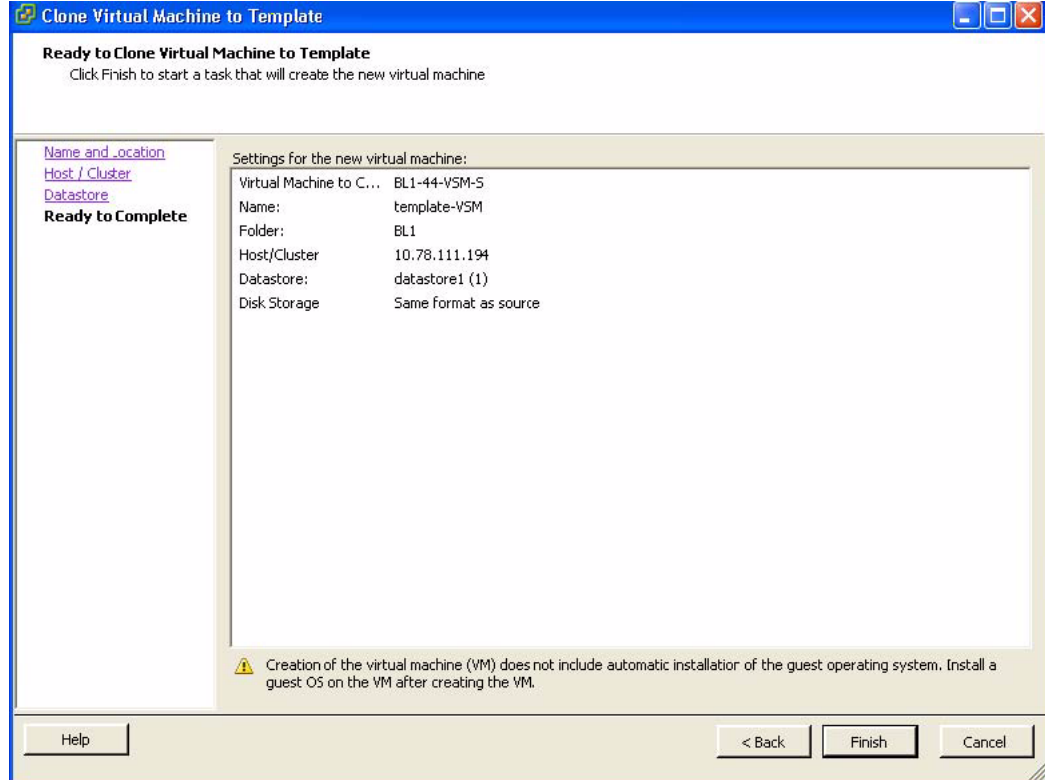
Step 12 Choose a datastore.

Step 13 Click **Next**.

The Confirming Settings screen opens. See [Figure 14-6](#).

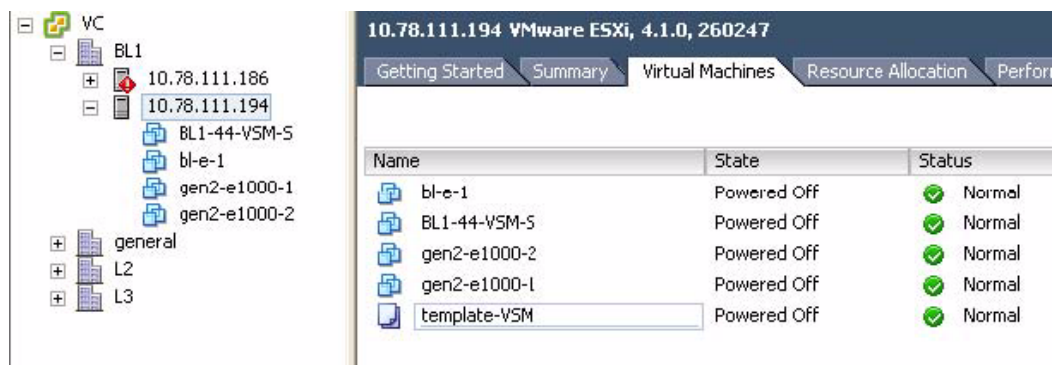
Send document comments to nexus1k-docfeedback@cisco.com.

Figure 14-6 Confirming Settings Screen



- Step 14** Confirm the settings for the new virtual machine and click **Finish**.
 The backup template is created and appears under the Virtual Machines tab.
 The Template Virtual Machine window opens. See [Figure 14-7](#).

Figure 14-7 Template Virtual Machine Window



The template creation is complete.

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

Performing a Periodic Backup

This section describes how to back up the active VSM after the initial backup of the standby VSM has been performed.

The following lists some instances when you should run this procedure:

- You have performed an upgrade.
- You have made a significant change to the configuration.

PROCEDURE

Step 1 Back up the VSM by entering a command similar to the following:

```
switch# copy running-config scp://root@10.78.19.15/tftpboot/config/
Enter destination filename: [switch-running-config]
Enter vrf (If no input, current vrf 'default' is considered):
The authenticity of host '10.78.19.15 (10.78.19.15)' can't be established.
RSA key fingerprint is 29:bc:4c:26:e3:6f:53:91:d4:b9:fe:d8:68:4a:b4:a3.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.78.19.15' (RSA) to the list of known hosts.
root@10.78.19.15's password:
switch-running-config 100% 6090 6.0KB/s 00:00
switch#
```

Recovering the VSM

This section describes how to deploy a VSM by using the backup template.

This section includes the following topics:

- [Deploying the Backup VSM VM, page 14-8](#)
- [Erasing the Old Configuration, page 14-15](#)
- [Restoring the Backup Configuration on the VSM, page 14-16](#)

PROCEDURE

- Step 1** To deploy the backed up VSM VM, see the “[Deploying the Backup VSM VM](#)” section on page 14-8.
- Step 2** To erase the old configuration, see the “[Erasing the Old Configuration](#)” section on page 14-15.
- Step 3** To restore the backup configuration, see the “[Restoring the Backup Configuration on the VSM](#)” section on page 14-16.

Deploying the Backup VSM VM

This section describes how to deploy the backup VSM VM when the primary and secondary VSMs are not present.

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

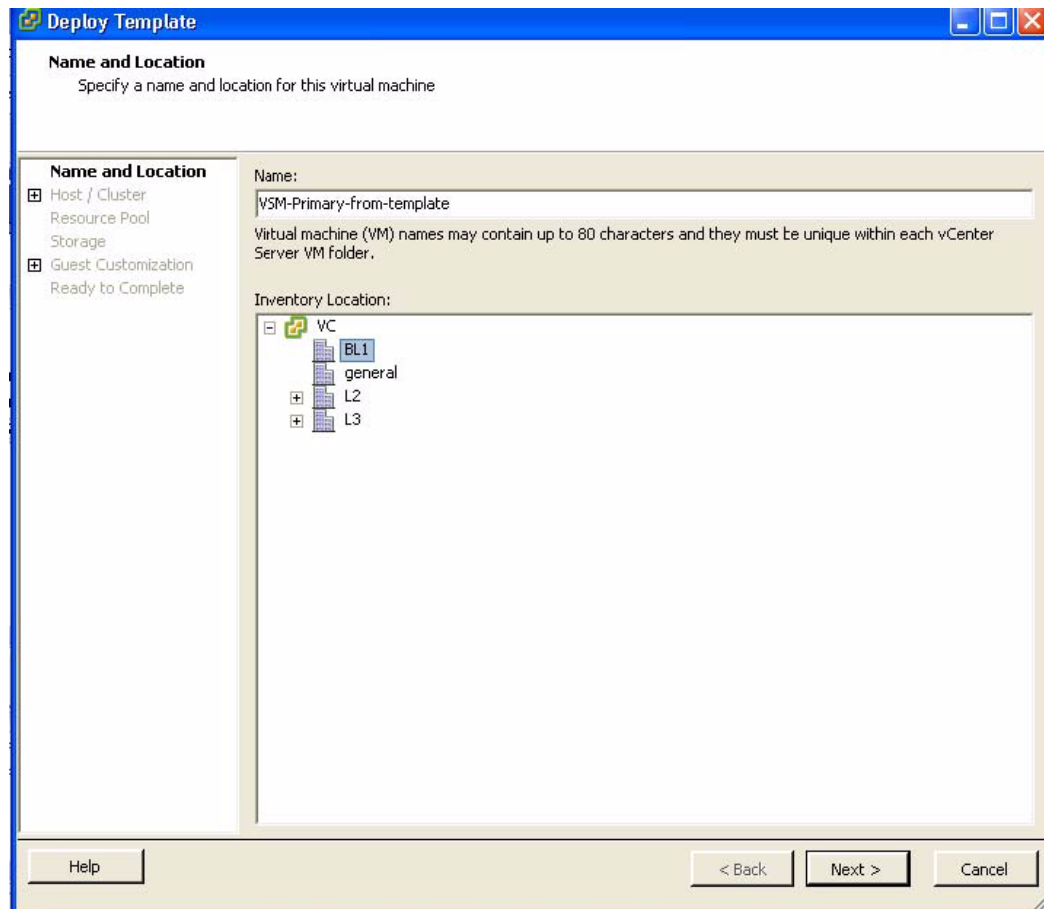


Note While deploying the VSM VM, do not power it on.

PROCEDURE

- Step 1** Open the vSphere Client.
The vSphere Client window opens.
- Step 2** In the left navigation pane, choose the host of the standby VSM.
- Step 3** Click the **Virtual Machines** tab.
- Step 4** Right-click the **template_VSM**.
- Step 5** Choose **Deploy Virtual Machine from this Template**.
The Deploy Template Wizard screen opens. See [Figure 14-8](#).

Figure 14-8 Deploy Template Wizard Screen

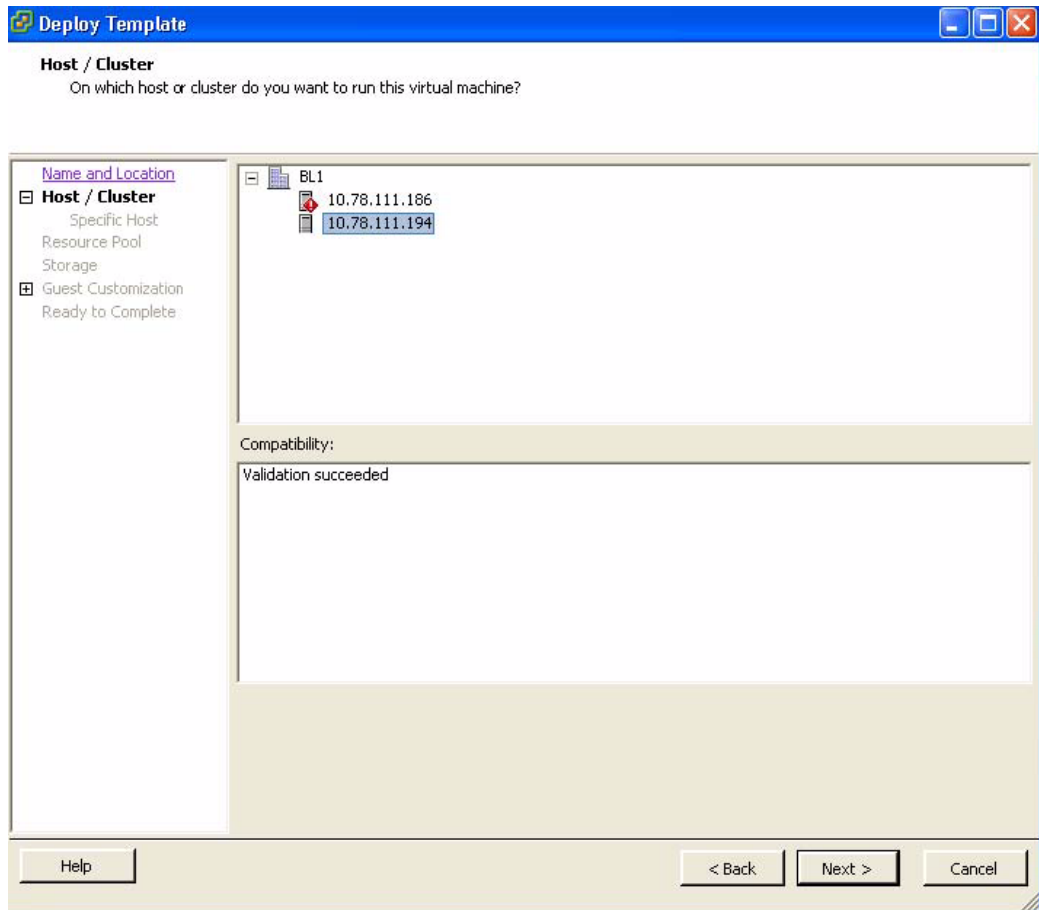


- Step 6** In the Name field, enter a name for the VSM.
- Step 7** In the Inventory Location pane, choose a cluster.
- Step 8** Click **Next**.

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

The Choosing a Host screen opens. See [Figure 14-9](#).

Figure 14-9 Choosing a Host Screen



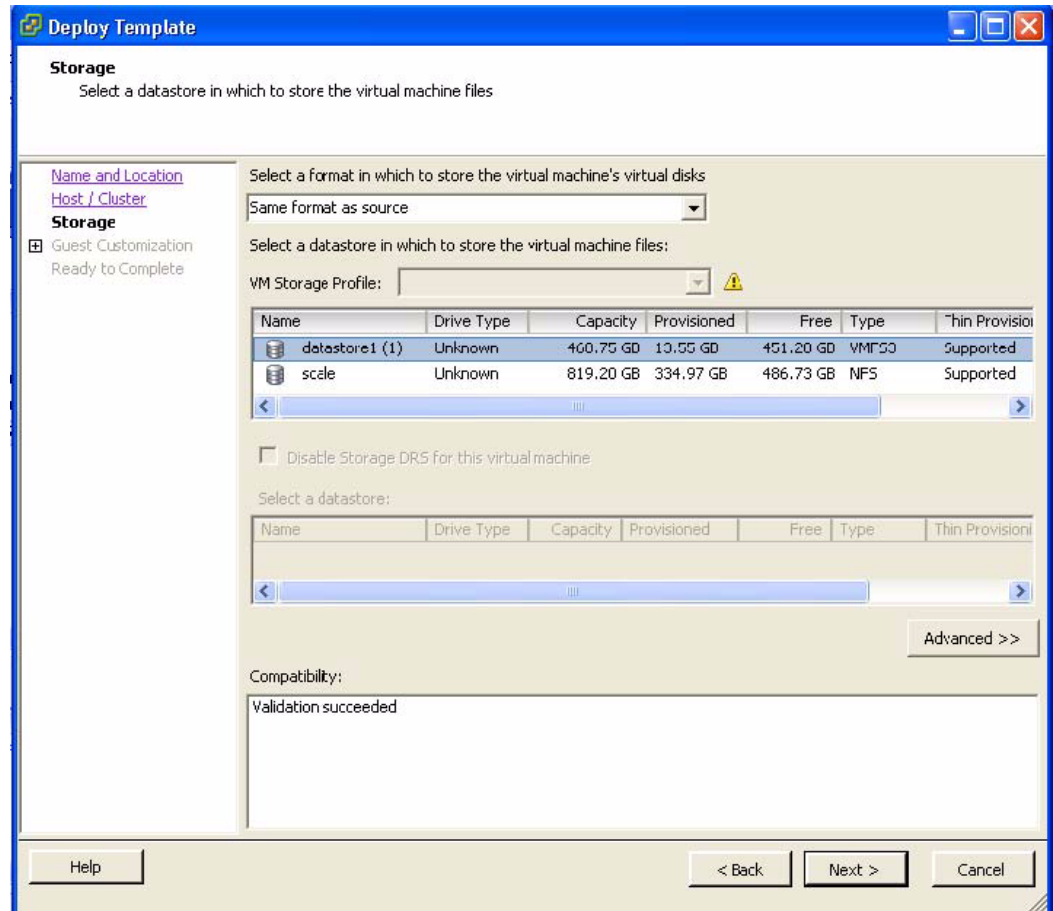
Step 9 Choose a host.

Step 10 Click **Next**.

The Choosing a Datastore screen opens. See [Figure 14-10](#).

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


Figure 14-10 Choosing the Datastore Screen



- Step 11** In the Select a format in which to store the virtual machine's virtual disks drop-down list, choose **Same format as source**.
- Step 12** Choose a datastore.
- Step 13** Click **Next**.
- The Guest Customization screen opens. See [Figure 14-11](#).

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

Figure 14-11 Guest Customization Screen




Note Make sure that the **Power on this virtual machine after creation** check box is not checked.

Step 14 Click **Next**.

The Deploy Template - Ready to Complete screen opens. See [Figure 14-12](#).

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

Figure 14-12 Deploy Template - Ready to Complete Screen



Step 15 Confirm the settings for the new virtual machine and click **Finish**.



Note If the management VLAN is not available on the VEM, you must add the management interface to the vSwitch.

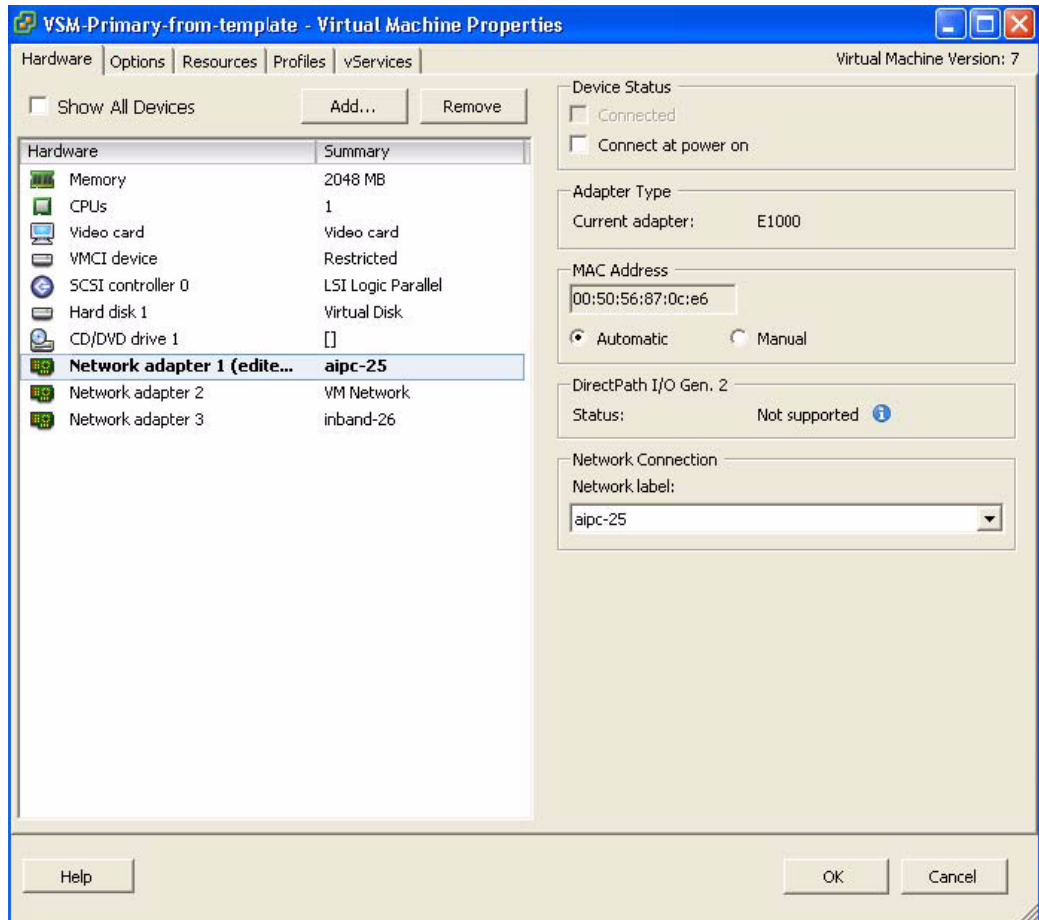
Step 16 Right-click the newly deployed VM.

Step 17 Choose **Edit Settings**.

The Virtual Machine Properties window opens. See [Figure 14-13](#).

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

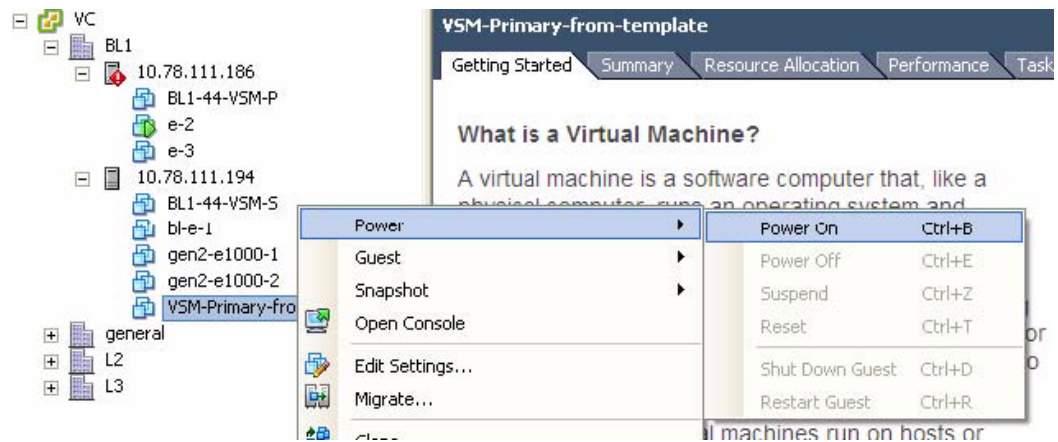
Figure 14-13 Virtual Machine Properties Window



- Step 18** In the Hardware / Summary pane, choose **Network adapter 1**.
- Step 19** In the Hardware / Summary pane, uncheck the **Connect at power on** check box.
- Step 20** In the Hardware / Summary pane, choose **Network adapter 2**.
- Step 21** In the Device Status area, uncheck the **Connect at power on** check box.
- Step 22** Click **OK**.
The Power On window opens. See [Figure 14-14](#).

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

Figure 14-14 Power On Window



- Step 23** Right-click the newly deployed VSM.
A drop-down list appears.
- Step 24** Choose **Power > Power On**.
Deploying the backup VSM VM is complete.

Erasing the Old Configuration

This section describes how to erase the startup configuration of the newly deployed VSM.

PROCEDURE

- Step 1** Launch the virtual machine console of the newly deployed VSM.
- Step 2** Set the redundancy role to primary by entering the following command:
- ```
switch# system redundancy role primary
Setting will be activated on next reload
switch#
```
- Step 3** Copy the running configuration to the startup configuration by entering the following command:
- ```
switch# copy running-config startup-config
scp:          sftp:          startup-config
[#####] 100%
switch#
```
- Step 4** Erase the startup configuration by entering the following command:
- ```
switch# write erase
Warning: The command will erase the startup-configurations.
Do you wish to proceed anyway? (y/n) [n] y
```
- Step 5** Reboot the primary and secondary VSMs by entering the following command:
- ```
switch# reload
This command will reboot the system. (y/n)? [n] y
```

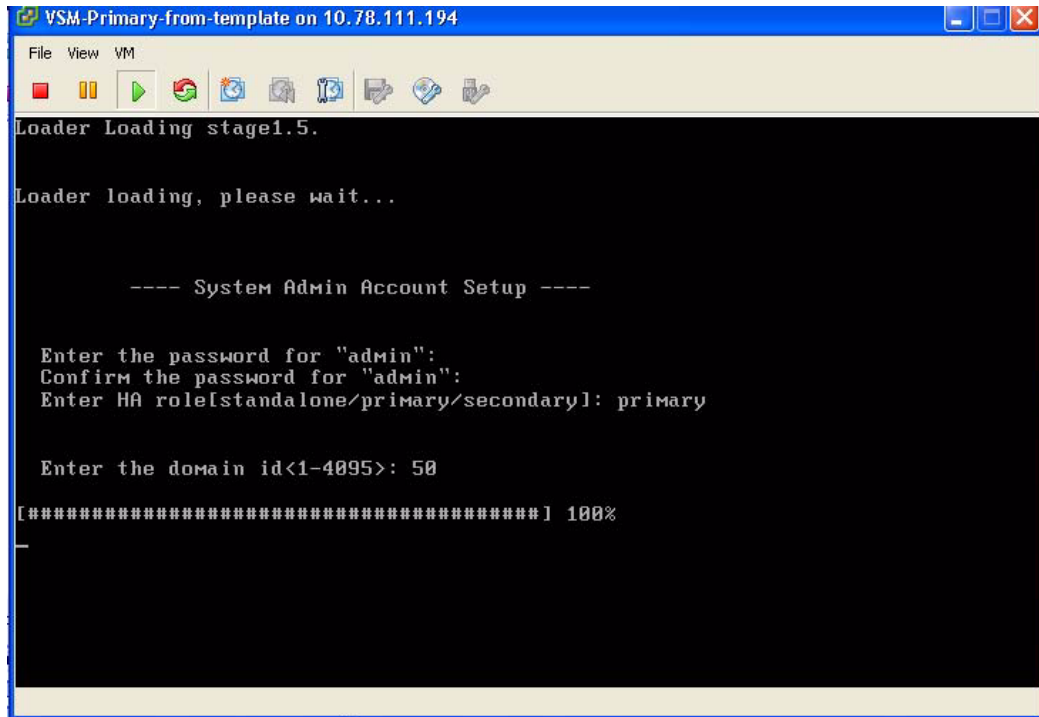
Send document comments to nexus1k-docfeedback@cisisco.com.

Restoring the Backup Configuration on the VSM

This section describes how to restore the backup configuration on the VSM.

- Step 1** When the VSM reboots, the System Admin Account Setup window opens. See [Figure 14-15](#).

Figure 14-15 System Admin Account Setup Window



- Step 2** Enter and confirm the Administrator password.

```

---- System Admin Account Setup ----
Enter the password for "admin":
Confirm the password for "admin":

```

- Step 3** Enter the domain ID.

```

Enter the domain id<1-4095>: 50

```

- Step 4** Enter the HA role.

If you do not specify a role, standalone is assigned by default.

```

Enter HA role[standalone/primary/secondary]: primary

```

```

[#####] 100%

```

```

---- Basic System Configuration Dialog ----

```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime

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

to skip the remaining dialogs.

- Step 5** Enter yes when you are prompted to enter the basic configuration dialog.
Would you like to enter the basic configuration dialog (yes/no): **yes**
- Step 6** Enter no when asked to create another login account.
Create another login account (yes/no) [n]: **no**
- Step 7** Enter no when asked to configure a read-only SNMP community string.
Configure read-only SNMP community string (yes/no) [n]: **no**
- Step 8** Enter no when asked to configure a read-write SNMP community string.
Configure read-write SNMP community string (yes/no) [n]: **no**
- Step 9** Enter a name for the switch.
Enter the switch name:
- Step 10** Enter yes when asked to configure out-of-band management and then enter the mgmt0 IPv4 address and subnet mask.
Continue with Out-of-band (mgmt0) management configuration? [yes/no] [y]: **yes**
Mgmt0 IPv4 address: **172.28.15.152**
Mgmt0 IPv4 netmask: **255.255.255.0**
- Step 11** Enter yes when asked to configure the default gateway.
Configure the default-gateway: (yes/no) [y]: **yes**

IPv4 address of the default gateway : 172.23.233.1
- Step 12** Enter yes when asked to enable the Telnet service.
Enable the telnet service? (yes/no) [y]: **yes**
- Step 13** Enter yes when asked to enable the SSH service, and then enter the key type and number of key bits. For more information, see the *Cisco Nexus 1000V Security Configuration Guide, Release 4.2(1)SV1(4)*.
Enable the ssh service? (yes/no) [y]: **yes**
Type of ssh key you would like to generate (dsa/rsa) : **rsa**
Number of key bits <768-2048> : **1024**
- Step 14** Enter yes when asked to enable the HTTP server.
Enable the http-server? (yes/no) **yes**
- Step 15** Enter no when asked to configure the NTP server.
Configure NTP server? (yes/no) [n]: **no**
- Step 16** Enter no when asked to configure the VEM feature level.
Vem feature level will be set to 4.2(1)SV1(4a), Do you want to reconfigure? (yes/no) [n]
no
- Step 17** Enter no when asked to configure the svs domain parameters.
Configure svs domain parameters? (yes/no) [y]: **no**

The system now summarizes the complete configuration and prompts you to edit it.

The following configuration will be applied:

```
interface Mgmt0
```

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

```
ip address 172.28.15.152 255.255.255.0
no shutdown
vrf context management
ip route 0.0.0.0/0 10.78.111.11
no telnet server enable
ssh key rsa 1024 force
ssh server enable
feature http-server
svs-domain
  svl mode L2
  control vlan 1
  packet vlan 1
  domain id 1
```

Step 18 Enter no when asked if you would like to edit the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: no
```

```
Enter SVS Control mode (L2 / L3) : L2
Enter control vlan <1-3967, 4048-4093> : 100
Enter packet vlan <1-3967, 4048-4093> : 101
```

Step 19 Enter yes when asked to use and save this configuration.



Caution

If you do not save the configuration now, none of your changes are part of the configuration the next time the switch is rebooted. Enter **yes** to save the new configuration. This ensures that the kickstart and system images are configured automatically.

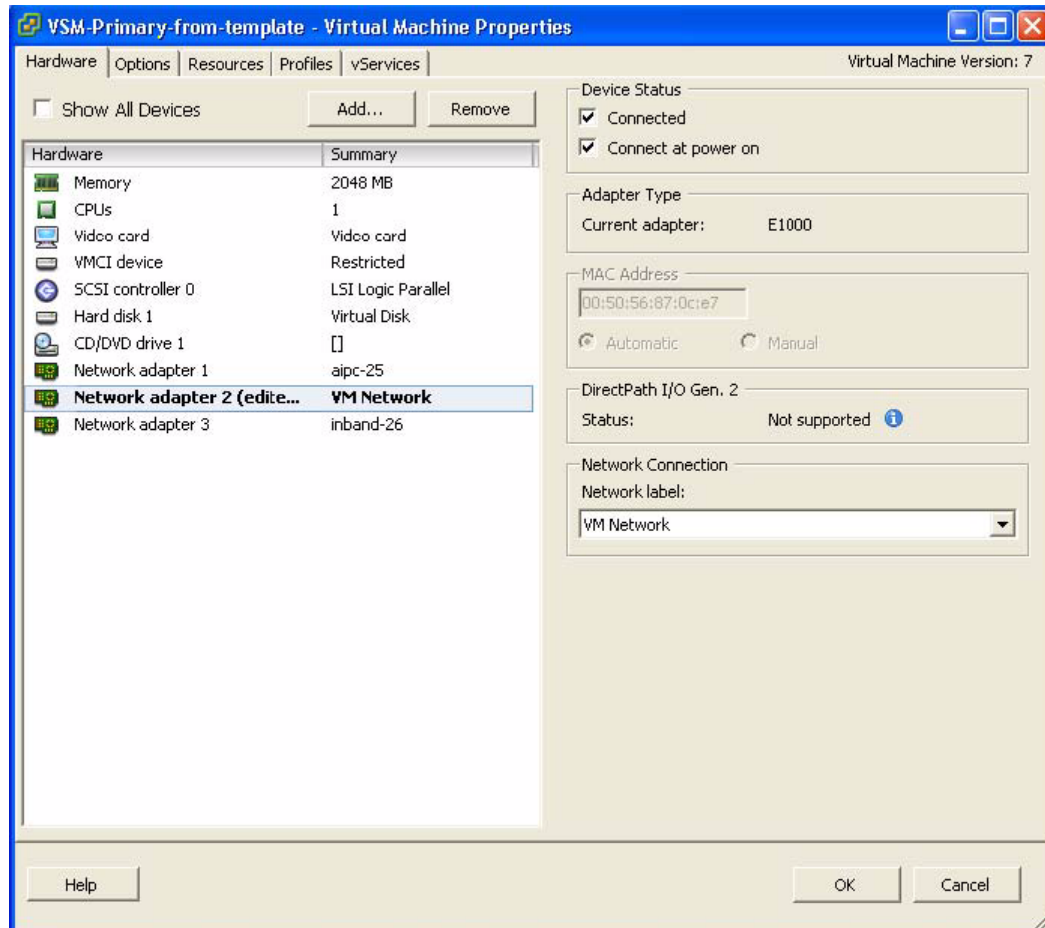
```
Use this configuration and save it? (yes/no) [y]: yes
[#####] 100%
```

Step 20 In the vSphere Client, right-click the **VSM** and choose **Edit Settings**.

The VSM Virtual Machine Properties window opens. See [Figure 14-16](#).

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

Figure 14-16 VSM Virtual Machine Properties Window



- Step 21** In the Hardware / Summary pane, choose **Network adapter 2**.
- Step 22** Check the **Connect at power on** check box.
- Step 23** Log in to the VSM.
- Step 24** Copy the backup configuration to the VSM bootflash by entering the following command:

```
switch# copy scp://root@10.78.19.15/tftpboot/backup/VSM-Backup-running-config
bootflash:
Enter vrf (If no input, current vrf 'default' is considered):
The authenticity of host '10.78.19.15 (10.78.19.15)' can't be established.
RSA key fingerprint is 29:bc:4c:26:e3:6f:53:91:d4:b9:fe:d8:68:4a:b4:a3.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.78.19.15' (RSA) to the list of known hosts.
root@10.78.19.15's password:
switch-running-config 100%
6090 6.0KB/s 00:00
switch#
```

- Step 25** Copy the backup configuration to the running configuration by entering the following command:

```
switch# copy bootflash:VSM-Backup-running-config running-config
Disabling ssh: as its enabled right now:
Can't disable ssh for key generation:Current user is logged in through ssh
Please do a "copy running startup" to ensure the new setting takes effect
on next reboot
```

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

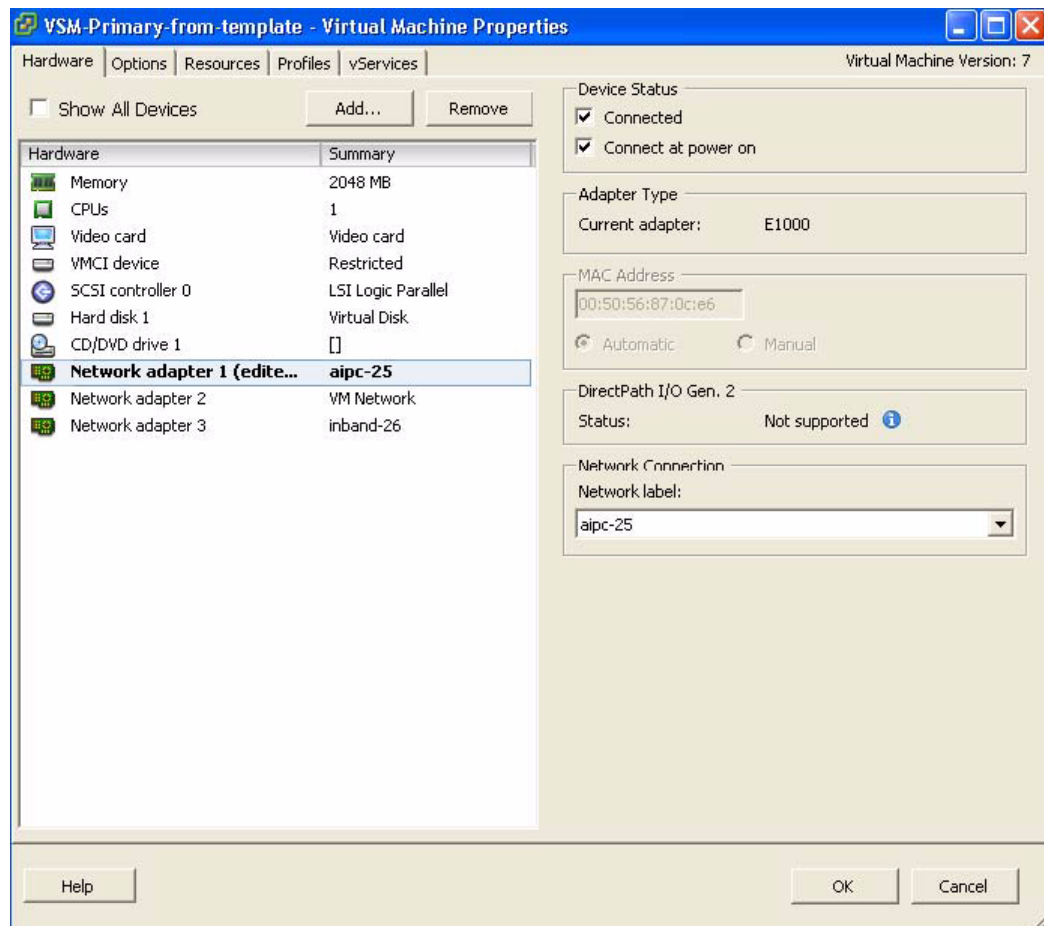
```
LACP Offload Status can be verified using "show lacp offload status"
Change in LACP Offload Status takes effect only on the next VSM Reboot
This can potentially cause modules with LACP uplinks to flap
Syntax error while parsing 'limit-resource m4route-mem minimum 58 maximum 58'
Syntax error while parsing 'limit-resource m6route-mem minimum 8 maximum 8'
Syntax error while parsing 'interface Ethernet3/2'
Syntax error while parsing 'inherit port-profile uplink-cdp'
Warning: Config saved but not pushed to vCenter Server due to inactive connection!
Warning: Config saved but not pushed to vCenter Server due to inactive connection!
Warning: Config saved but not pushed to vCenter Server due to inactive connection!
command failed. Invalid ip address.
Syntax error while parsing 'log-level '
Syntax error while parsing 'no ip dhcp relay'
switch#
```



Note

You might see syntax errors. You can ignore them.

Figure 14-17 Virtual Machine Properties Window



Step 26 In the Hardware / Summary pane, choose **Network adapter 1**.

Step 27 In the Device Status area, check the **Connect at power on** check box.

Step 28 Confirm that the VEMs are attached to the VSM by entering the following command:

```
switch# show module
```

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

```

Mod Ports Module-Type Model Status
-----
1 0 Virtual Supervisor Module Nexus1000V active *
3 248 Virtual Ethernet Module NA ok
Mod Sw Hw
-----
1 4.2(1)SV1(4a) 0.0
3 4.2(1)SV1(4a) VMware ESXi 4.0.0 Releasebuild-261974 (1.20)
Mod MAC-Address(es) Serial-Num
-----
1 00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8 NA
3 02-00-0c-00-03-00 to 02-00-0c-00-03-80 NA
Mod Server-IP Server-UUID Server-Name
-----
1 10.78.111.20 NA NA
3 10.78.111.186 0e973f80-e804-11de-956e-4bc311a28ede VEM-186-KLU2
* this terminal session
switch#

```

Step 29 Copy the backup configuration to the running configuration after all the VEMs attach by entering the following command:

```

switch# copy bootflash:VSM-Backup-running-config running-config
Disabling ssh: as its enabled right now:
Can't disable ssh for key generation:Current user is logged in through ssh
2011 Apr 26 12:21:22 switch %KERN-3-SYSTEM_MSG: redun_platform_ioctl :
Entered - kernel
2011 Apr 26 12:21:22 switch %KERN-3-SYSTEM_MSG: redun_platform_ioctl : Host
name is set switch - kernel
2011 Apr 26 12:21:22 switch %KERN-3-SYSTEM_MSG: redun_platform_ioctl :
Entered - kernel
2011 Apr 26 12:21:22 switch %KERN-3-SYSTEM_MSG: redun_platform_ioctl : Host
name is set switch - kernel
ERROR: Flow Record: Record is in use. Remove from all clients before modifying.
ERROR: Flow Record: Record is in use. Remove from all clients before modifying.
ERROR: Flow Record: Record is in use. Remove from all clients before modifying.
Please do a "copy running startup" to ensure the new setting takes effect
on next reboot
LACP Offload Status can be verified using "show lacp offload status"
Change in LACP Offload Status takes effect only on the next VSM Reboot
This can potentially cause modules with LACP uplinks to flap
2011 Apr 26 12:21:23 switch %VMS-5-DVS_NAME_CHANGE: Changed dvs switch
name to 'switch' on the vCenter Server.
Syntax error while parsing 'limit-resource m4route-mem minimum 58 maximum 58'
Syntax error while parsing 'limit-resource m6route-mem minimum 8 maximum 8'
ERROR: Port-channel interface has non-zero members!
2011 Apr 26 12:21:34 switch %MSP-5-DOMAIN_CFG_SYNC_DONE: Domain config
successfully pushed to the management server.
ERROR: Cannot change connection configuration in 'Enabled' state.
ERROR: Cannot change connection configuration in 'Enabled' state.
ERROR: Cannot change the data-center name in connected state.
command failed. Invalid ip address.
Syntax error while parsing 'log-level '
Syntax error while parsing 'no ip dhcp relay'
switch# 2011 Apr 26 12:21:35 switch last message repeated 3 times
2011 Apr 26 12:21:35 switch %ETHPORT-5-SPEED: Interface port-channel1,
operational speed changed to 1 Gbps
2011 Apr 26 12:21:35 switch %ETHPORT-5-IF_DUPLEX: Interface port-channel1,
operational duplex mode changed to Full
2011 Apr 26 12:21:35 switch %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface portchannel1,
operational Receive Flow Control state changed to on
2011 Apr 26 12:21:35 switch %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface portchannel1,
operational Transmit Flow Control state changed to on
VSM backup and Recovery Procedure EDCS-1017832

```

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

```
Cisco Systems Pvt Ltd Internal Document April-27-2011
2011 Apr 26 12:21:35 switch %ETH_PORT_CHANNEL-5-PORT_UP: port-channel1:
Ethernet3/2 is up
2011 Apr 26 12:21:35 switch %ETH_PORT_CHANNEL-5-FOP_CHANGED: portchannel1:
first operational port changed from none to Ethernet3/2
2011 Apr 26 12:21:35 switch %ETHPORT-5-IF_UP: Interface Ethernet3/2 is up in
mode trunk
2011 Apr 26 12:21:35 switch %ETHPORT-5-IF_UP: Interface port-channel1 is up in
mode trunk
switch#
```



Note

This step is necessary if features are configured directly through the interface configuration mode for Ethernet interfaces and for features like ERSPAN/NFM.

Step 30 Copy the running-configuration to the startup-configuration by entering the following command:

```
switch# copy running-config startup-config
[#####] 100%
switch#
```

Step 31 Create the standby VSM by using the OVA/OVF files to form an HA pair. See the “Installing the Software from an OVA or OVF Image” section in the *Cisco Nexus 1000V Software Installation Guide, Release 4.2(1)SV1(4a)*.

- For release 4.2(1)SV1(4) and later releases, deploy the OVF template from the VMware vSphere Client and choose **Nexus 1000V Secondary** from the Configuration drop-down list.
- For release 4.0(4)SV1(2) through release 4.0(4)SV1(3d), choose **Manual Install of Nexus 1000V** from the Configuration drop-down list and assign the HA role of secondary in the System Admin Setup of the VSM.

The recovery is complete.

Additional References

For additional information related to implementing system message logging, see the following sections:

- [Related Documents, page 14-22](#)
- [Standards, page 14-23](#)

Related Documents

Related Topic	Document Title
System messages	<i>Cisco NX-OS System Messages Reference</i>
Complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Nexus 1000V Command Reference, Release 4.2(1)SV1(4a)</i>

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

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.	—

Feature History for VSM Backup and Recovery

This section provides the VSM backup and Recovery feature release history.

Feature Name	Releases	Feature Information
VSM Backup and Recovery	4.2(1)SV1(4a)	This feature was introduced.

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



CHAPTER 15

Virtualized Workload Mobility (DC to DC vMotion)

This chapter describes the Virtualized Workload Mobility (DC to DC vMotion) feature of the Cisco Nexus 1000V. This feature addresses Cisco Nexus 1000V across two physical data centers.

This chapter includes the following sections:

- [Information About Virtualized Workload Mobility \(DC to DC vMotion\), page 15-1](#)
- [Prerequisites for Virtualized Workload Mobility \(DC to DC vMotion\), page 15-2](#)
- [Guidelines and Limitations, page 15-2](#)
- [Verifying the Virtualized Workload Mobility \(DC to DC vMotion\) Configuration, page 15-4](#)
- [Monitoring Virtualized Workload Mobility \(DC to DC vMotion\), page 15-4](#)
- [Configuration Limits, page 15-4](#)
- [Feature History for Virtualized Workload Mobility \(DC to DC vMotion\), page 15-5](#)

Information About Virtualized Workload Mobility (DC to DC vMotion)

This section describes the Virtualized Workload Mobility (DC to DC vMotion) configurations and includes the following topics:

- [Stretched Cluster, page 15-1](#)
- [Split Cluster, page 15-2](#)
- [Physical Site Considerations, page 15-2](#)

Stretched Cluster



Note

A stretched cluster is a cluster with ESX/ESXi hosts in different physical locations.

In an environment where the same Cisco Nexus 1000V instance spans two data centers, this configuration allows you to have Virtual Ethernet Modules (VEMs) in different data centers be part of the same vCenter Server cluster.

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

By choosing this configuration, you ensure that the VEMs in either data center (in a two data center environment) are a part of the same Dynamic Resource Scheduling (DRS) / VMware High Availability (VMW HA) / Fault Tolerance (FT) domain that allows for multiple parallel virtual machine (VM) migration events.

Split Cluster

The Split Cluster configuration is an alternate to the Stretched Cluster deployment. With this configuration, the deployment consists of one or more clusters on either physical site with no cluster that contains VEMs in multiple data centers. While this configuration allows for VM migration between physical data centers, these events are not automatically scheduled by DRS.

Prerequisites for Virtualized Workload Mobility (DC to DC vMotion)

Virtualized Workload Mobility (DC to DC vMotion) has the following prerequisites:

- You must set up your DRS affinity rules to ensure that the VSM pair is restricted to one site.
- Layer 2 extension between the two physical data centers over the DCI link.

Guidelines and Limitations

Virtualized Workload Mobility (DC to DC vMotion) has the following guidelines and limitations:

- The VSM HA pair must be located in the same site as their storage and the active vCenter Server.
- Layer 3 control mode is preferred.
- If you are using Link Aggregation Control Protocol (LACP) on the VEM, use LACP offload.
- Quality of Service bandwidth guarantees for control traffic over the DCI link.
- Limit the number of physical data centers to two.
- A maximum latency of 5 ms is supported for VSM-VEM control traffic.

For configuration limits in a two data center environment, see the [Chapter 16, “Configuration Limits”](#).

Physical Site Considerations

When you are designing a physical site, follow these guidelines:

- Check the average and maximum latency between a Virtual Supervisor Module (VSM) and VEM.
- Follow the procedures to perform actions you would intend to do in normal operation. For example, VSM migration.
- Design the system to handle the high probability of VSM-VEM communication failures where a VEM must function in headless mode due to data center interconnect (DCI) link failures.

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

Handling Inter-Site Link Failures

If the DCI link or Layer 2 extension mechanism fails, a set of VEM modules might run with their last known configuration for a period of time.

Headless Mode of Operation

For the period of time that the VSM and VEM cannot communicate, the VEM continues to operate with its last known configuration. Once the DCI link connectivity is restored and the VSM-VEM communication is reestablished, the system should come back to its previous operational state.

This mode type is no different than the headless mode of operation within a data center and has the following limitations for the headless VEM:

- No new ports can be brought up on the headless VEM (new VMs coming up or VMs coming up after vMotion).
- No NetFlow data exports.
- Ports shut down because DHCP/DAI rate limits are not automatically brought up until the VSM reconnects.
- Port security options, such as aging or learning secure MAC addresses and shutting down/recovering from port-security violations, are not available until the VSM reconnects.
- The Cisco Discovery Protocol (CDP) does not function for the disconnected VEM.
- IGMP joins/leaves are not processed until the VSM reconnects.
- Queries on BRIDGE and IF-MIB processed at the VSM give the last known status for the hosts in headless mode.

Handling Additional Distance/Latency Between the VSM and VEM

In a network where there is a considerable distance between the VSM and VEM, latency becomes a critical factor.

Because the control traffic between the VSM and VEM faces a sub-millisecond latency within a data center, latency can increase to a few milliseconds depending on the distance.

With an increased round-trip time, communication between the VSM and VEM takes longer. As you add VEMs and vEthernet interfaces, the time it takes to perform actions such as configuration commands, module insertions, port bring-up, and **show** commands increase because that many tasks are serialized.

Migrating a VSM

This section describes how migrate a VSM from one physical site to another.



Note

If you are migrating a VSM on a Cisco Nexus 1010, see the *Cisco Nexus 1010 Software Configuration Guide, Release 4.2(1)SP1(3)*.

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

Migrating a VSM Hosted on an ESX or ESXi Host

Use the following procedure to migrate a VSM that is hosted on an ESX or ESXi host from the local data center to the remote data center.

BEFORE YOU BEGIN

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

- Reduce the amount of time where the VSM runs with remote storage in another data center.
- Do not bring up any new VMs or vMotion VMs that are hosted on any VEMs corresponding to the VSM that is being migrated.



Note

For information on vMotion or storage vMotion, see the VMware documentation.

PROCEDURE

-
- Step 1** Migrate the standby VSM to the backup site.
- Step 2** Perform a storage vMotion for the standby VSM storage.
- Step 3** Enter the **system switchover** command.
- ```
switch # system switchover
```
- Step 4** Migrate the original active VSM to the backup site.
- Step 5** Perform a storage vMotion for the original active VSM storage.
- 

## Verifying the Virtualized Workload Mobility (DC to DC vMotion) Configuration

```
switch# show module
switch# show interface
```

## Monitoring Virtualized Workload Mobility (DC to DC vMotion)

```
switch# show module
switch# show interface
```

## Configuration Limits

For information about Cisco Nexus 1000V configuration limits, refer to [Chapter 16, “Configuration Limits”](#).

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

## Feature History for Virtualized Workload Mobility (DC to DC vMotion)

Table 15-1 lists the release history for this feature.

**Table 15-1** Feature History for Virtualized Workload Mobility (DC to DC vMotion)

| Feature Name                                     | Releases      | Feature Information          |
|--------------------------------------------------|---------------|------------------------------|
| Virtualized Workflow Mobility (DC to DC vMotion) | 4.2(1)SV1(4a) | This feature was introduced. |

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



# CHAPTER 16

## Configuration Limits

**Table 16-1** Configuration Limits for Cisco Nexus 1000V

| Component                                        | Supported Limits for Cisco Nexus 1000V in the Same Datacenter  |          | Supported Limits for Cisco Nexus 1000V Across Two Datacenters  |          |
|--------------------------------------------------|----------------------------------------------------------------|----------|----------------------------------------------------------------|----------|
|                                                  | Per DVS                                                        | Per Host | Per DVS                                                        | Per Host |
| Maximum Modules                                  | 66                                                             |          | 34                                                             |          |
| Virtual Ethernet Module(VEM)                     | 64                                                             |          | 32                                                             |          |
| Virtual Supervisor Module (VSM)                  | 2 in an HA Pair (active-standby hosted in the same datacenter) |          | 2 in an HA Pair (active-standby hosted in the same datacenter) |          |
| vCenter Server Datacenters per VSM               | 1                                                              |          | 1                                                              |          |
| Hosts                                            | 64                                                             |          | 32                                                             |          |
| Active VLANs across all VEMs                     | 2048                                                           |          | 1024                                                           |          |
| MACs per VEM                                     | 32000                                                          |          | 32000                                                          |          |
| MACs per VLAN per VEM                            | 4000                                                           |          | 4000                                                           |          |
| vEthernet interfaces per port profile            | 1024                                                           |          | 1024                                                           |          |
| PVLAN                                            | 512                                                            |          | 128                                                            |          |
| Distributed Virtual Switches (DVSEs) per vCenter | 12                                                             |          | 12                                                             |          |
| vCenter Server connections                       | 1 per VSM HA Pair <sup>1</sup>                                 |          | 1 per VSM HA Pair <sup>1</sup>                                 |          |
| Maximum latency between VSMs and VEMs            | —                                                              |          | 5 ms                                                           |          |
|                                                  | Per DVS                                                        | Per Host | Per DVS                                                        | Per Host |
| Virtual Service Domains (VSDs)                   | 64                                                             | 6        | 32                                                             | 3        |
| VSD interfaces                                   | 2048                                                           | 216      | 1024                                                           | 108      |
| vEthernet interfaces                             | 2048                                                           | 216      | 1024                                                           | 108      |
| Port profiles                                    | 2048                                                           | —        | 1024                                                           | —        |
| System port profiles                             | 32                                                             | 32       | 16                                                             | 16       |
| Port channels                                    | 256                                                            | 8        | 128                                                            | 4        |
| Physical trunks                                  | 512                                                            | —        | 256                                                            | —        |
| Physical NICs                                    | —                                                              | 32       | —                                                              | 16       |

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

**Table 16-1 Configuration Limits for Cisco Nexus 1000V (continued)**

| Component                                     | Supported Limits for Cisco Nexus 1000V in the Same Datacenter |                  | Supported Limits for Cisco Nexus 1000V Across Two Datacenters |                 |
|-----------------------------------------------|---------------------------------------------------------------|------------------|---------------------------------------------------------------|-----------------|
|                                               |                                                               |                  |                                                               |                 |
| vEthernet trunks                              | 256                                                           | 8                | 128                                                           | 4               |
| ACLs                                          | 128                                                           | 16 <sup>2</sup>  | 64                                                            | 8 <sup>2</sup>  |
| ACEs per ACL                                  | 128                                                           | 128 <sup>2</sup> | 64                                                            | 64 <sup>2</sup> |
| ACL interfaces                                | 2048                                                          | 256              | 1024                                                          | 128             |
| NetFlow policies                              | 32                                                            | 8                | 16                                                            | 4               |
| NetFlow interfaces                            | 256                                                           | 32               | 128                                                           | 16              |
| SPAN/ERSPAN sessions                          | 64                                                            | 64               | 32                                                            | 32              |
| QoS policy maps                               | 128                                                           | 128              | 64                                                            | 64              |
| QoS class maps                                | 1024                                                          | 1024             | 512                                                           | 512             |
| QoS interfaces                                | 2048                                                          | 256              | 1024                                                          | 128             |
| Port security                                 | 2048                                                          | 216              | 1024                                                          | 108             |
| Multicast groups                              | 512                                                           | 512              | 256                                                           | 256             |
| DHCP snoop binding entries (static + dynamic) | 2048                                                          | 2048             | 1024                                                          | 1024            |

1. Only one connection to vCenter server is permitted at a time.
2. This number can be exceeded if VEM has available memory.



## INDEX

---

### B

- backing up files [6-6](#)
- banner message
  - configuring [5-2](#)

---

### C

- capability, Layer 3 control [3-10, 13-8](#)
- capability command [3-10, 13-8](#)
- class-map limits [16-1](#)
- command
  - directing output to a file [6-12](#)
- configuration
  - clearing [5-10](#)
  - displaying [5-3](#)
  - rolling back to previous [6-13](#)
  - saving [5-10](#)
- configuration, viewing [4-12](#)
- configuration files
  - backing up [6-6](#)
  - copying [6-6](#)
  - deleting [6-9](#)
  - downloading [6-6](#)
- configuration limits [16-1](#)
- configured domain
  - viewing [4-12](#)
- connections, viewing [4-11](#)
- connect to vCenter Server [4-2](#)
- control VLAN
  - CLI setup [14-18](#)
- copying files [6-6](#)
- creating VLANs

- default state [3-12, 3-14](#)

- current directory
  - changing [6-3](#)
  - displaying [6-2](#)

---

### D

- default gateway
  - CLI setup [14-17](#)
- default settings
  - SNMP [3-3, 10-5, 13-5](#)
- description command [13-7](#)
- directories
  - creating [6-7](#)
  - deleting [6-8, 6-9](#)
  - display current [6-2](#)
  - listing files [6-4](#)
  - moving files [6-8](#)
- disconnect from vCenter Server [4-4](#)
- display switch configuration [4-12](#)
- documentation
  - additional publications [i-xvii](#)
- domain
  - CLI setup [14-18](#)
- domain ID, VSM
  - CLI setup [14-16](#)

---

### E

- enable
  - port profile [3-11, 13-8](#)
  - ports in the profile [3-11, 13-8](#)
- ERSPAN

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

about [9-4](#)  
 configuring a session [9-13](#)  
 implementation [9-4](#)  
 NAM monitoring [9-4](#)

---

## F

### Fibre Channel interfaces

default settings [2-2, 8-3, 9-6, 11-9, 12-5](#)

### files

compressing [6-10](#)  
 copying or backing up [6-6](#)  
 deleting [6-9](#)  
 displaying checksums [6-15](#)  
 displaying contents [6-13](#)  
 displaying last lines [6-15](#)  
 moving [6-8](#)  
 uncompressing [6-10](#)

### file systems

changing directories [6-3](#)  
 creating directories [6-7](#)  
 deleting directories [6-8](#)  
 displaying current directory [6-2](#)  
 listing files [6-4](#)  
 specifying [6-2](#)

flow exporter [11-6](#)

flow monitor [11-6](#)

---

## H

### HA role

CLI setup [14-16](#)

### high availability

SNMP [10-5](#)

### HTTP

CLI setup [14-17](#)

---

## I

### interfaces

default settings [2-2, 8-3, 9-6, 11-9, 12-5](#)

IP connectivity [3-9](#)

ip flow monitor command [11-20](#)

---

## L

Layer 3 connectivity [3-9](#)

limits, configuration [16-1](#)

### Local SPAN

about [9-3](#)

configuring a session [9-7](#)

implementation [9-3](#)

---

## M

### management interfaces

default settings [2-2, 8-3, 9-6, 11-9, 12-5](#)

match criteria limit [16-1](#)

### mgmt0

CLI setup [14-17](#)

### mgmt0 interfaces

default settings [2-2, 8-3, 9-6, 11-9, 12-5](#)

### MIBs

description [10-2](#)

location to download [10-15](#)

SNMP [10-15](#)

### modifying VLANs

allowed parameters [3-12, 3-14](#)

modules, displaying [4-15](#)

---

## N

### NAM

ERSPAN data sources [9-4](#)

NetFlow data source [11-8](#)



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

## NetFlow 11-20

- exporter 11-6
- monitor 11-6
- NAM monitoring 11-8

## Network Analysis Module

- ERSPAN data source 9-4
- NetFlow data source 11-8

no shutdown command 3-11, 13-8

## NTP

- CLI setup 14-17
- configuring 2-3 to 2-10

---

## P

### packet VLAN

- CLI setup 14-18

pg-name option 3-11, 9-11, 13-7

policy map limits 16-1

### Port Profile

- IP connectivity 3-9
- Layer 3 control 3-9

port-profile command 3-10

### port profiles

- NetFlow 11-20

---

## R

related documents i-xvii

remove Nexus1000V from vCenter Server 4-5

---

## S

service policy limits 16-1

### show commands

- show interface brief 4-15
- show interface virtual 4-15
- show module 4-15
- show running-config 4-12

show server-info 4-15

show svcs connections 4-11

show svcs domain 4-12

Simple Network Management Protocol. See SNMP

## SNMP

agent 10-2

assigning contact 10-11

assigning location 10-11

authentication 10-4

configuring a user 10-6

creating communities 10-8

default settings 3-3, 10-5, 13-5

description 10-1 to ??

disabling protocol 10-12

enabling one-time authentication 10-11

enforcing encryption 10-7

engine ID format 10-7

example configuration 10-13

group-based access 10-5

guidelines 10-5

high availability 10-5

limitations 10-5

manager 10-1

MIBs 10-2

MIBs supported 10-15

### notifications

configuring LinkUp/LinkDown  
notifications 10-11

configuring notification receivers 10-8

configuring the notification target user 10-9

description 10-2

enabling individual notifications 10-9

informs 10-2

trap 10-2

RFCs 10-2

user synchronization with CLI 10-4

verifying configuration 10-13

### versions

security models and levels 10-3

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

SNMPv3 [10-2](#)

USM [10-3](#)

## SPAN

egress sources [9-2](#)

## SPAN sessions

description [9-5](#)

resuming [9-17, 9-19](#)

shutting down [9-16](#)

## SPAN sources

egress [9-2](#)

## SSH

CLI setup [14-17](#)

state enabled command [3-11, 13-8](#)

## SVIs

VLAN interfaces [3-12, 3-14](#)

## switch name

CLI setup [14-17](#)

switchport access vlan, command [3-11, 13-8](#)

switchport mode command [3-11, 13-7](#)

system vlan command [3-11, 13-8](#)

remove Nexus 1000V from [4-5](#)

## VEM

feature level

CLI setup [14-17](#)

## VLAN interfaces

communicating between VLANs [3-12, 3-14](#)

## VLANs

SVIs [3-12, 3-14](#)

vmware port-group command [3-11](#)

## volatile:

switch reboots [6-3](#)

## VSM

credentials

CLI setup [14-16](#)

---

## T

### Telnet

CLI setup [14-17](#)

trap. See SNMP

---

## U

### users

displaying [7-1](#)

sending messages [7-2](#)

---

## V

### vCenter Server

connect to [4-2](#)

disconnect from [4-4](#)