



Cisco Nexus 1000V System Management Configuration Guide, Release 4.2(1)SV2(1.1)

First Published: October 22, 2012

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-27748-01

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

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

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

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

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

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

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.

© 2009-2012 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface **xiii**

Audience **xiii**

Document Conventions **xiii**

Related Documentation for Nexus 1000V Series NX-OS Software **xv**

Documentation Feedback **xvi**

Obtaining Documentation and Submitting a Service Request **xvi**

CHAPTER 1

New and Changed Information for this Release **1**

New and Changed Information **1**

CHAPTER 2

Overview **3**

System Management Overview **3**

CDP **3**

Domains **4**

Server Connections **4**

Configuration Management **4**

File Management **4**

User Management **4**

NTP **4**

Local SPAN and ERSPAN **5**

SNMP **5**

NetFlow **5**

System Messages **5**

iSCSI Multipath **5**

Troubleshooting **5**

CHAPTER 3

Configuring CDP **7**

Information About CDP	7
High Availability	8
Guidelines and Limitations	8
Default Settings	8
Configuring CDP	9
CDP Global Configuration	9
Enabling or Disabling CDP Globally	9
Advertising a CDP Version	9
Configuring CDP Options	10
CDP Interface Configuration	12
Enabling CDP on an Interface	12
Disabling CDP on an Interface	13
Monitoring CDP	14
Clearing CDP Statistics	14
Verifying the CDP Configuration	14
Configuration Example for CDP	15
Feature History for CDP	15

CHAPTER 4**Configuring the Domain 17**

Information About the Domain	17
Layer 3 Control	17
Guidelines and Limitations	18
Default Settings	19
Configuring the Domain	19
Creating a Domain	19
Changing to Layer 3 Transport	21
Changing to Layer 2 Transport	23
Creating a Port Profile for Layer 3 Control	24
Creating a Control VLAN	26
Creating a Packet VLAN	27
Feature History for the VSM Domain	28

CHAPTER 5**Managing Server Connections 31**

Information About Server Connections	31
Guidelines and Limitations	32

Connecting to the vCenter Server	32
Disconnecting From the vCenter Server	33
Removing the DVS from the vCenter Server	34
Removing the DVS from the vCenter Server When the VSM Is Not Connected	35
Configuring the Admin User or Admin Group	35
Removing the DVS from the vCenter Server Using the Graphical User Interface	36
Configuring Host Mapping	36
Information about Host Mapping	36
Removing Host Mapping from a Module	37
Mapping to a New Host	37
Viewing Host Mapping	38
Verifying Connections	38
Verifying the Domain	39
Verifying the Configuration	40
Verifying Module Information	40
Feature History for Server Connections	42

CHAPTER 6**Managing the Configuration 43**

Information About Configuration Management	43
Changing the Switch Name	43
Configuring a Message of the Day	44
Verifying the Configuration	45
Verifying the Software and Hardware Versions	45
Verifying the Running Configuration	46
Comparing the Startup and Running Configurations	47
Verifying the Interface Configuration	48
Verifying the Interface Configuration in a Brief Version	49
Verifying an Interface Configuration in a Detailed Version	49
Verifying All Interfaces in a Brief Version	50
Verifying the Running Configuration for all Interfaces	50
Saving a Configuration	51
Erasing a Configuration	51
Feature History for Configuration Management	52

CHAPTER 7**Working with Files 53**

Information About Files	53
Navigating the File System	54
Specifying File Systems	54
Identifying the Directory You are Working From	54
Changing Your Directory	55
Listing the Files in a File System	56
Identifying Available File Systems for Copying Files	56
Using Tab Completion	57
Copying and Backing Up Files	58
Creating a Directory	59
Removing an Existing Directory	60
Moving Files	60
Deleting Files or Directories	61
Compressing Files	62
Uncompressing Files	63
Directing Command Output to a File	63
Verifying a Configuration File before Loading	64
Rolling Back to a Previous Configuration	64
Displaying Files	65
Displaying File Contents	65
Displaying Directory Contents	66
Displaying File Checksums	66
Displaying the Last Lines in a File	67
Feature History for File Management	67

CHAPTER 8**Managing users 69**

Information About User Management	69
Displaying Current User Access	69
Sending a Message to Users	70
Feature History for User Management	70

CHAPTER 9**Configuring NTP 71**

Information about NTP	71
NTP Peers	72
High Availability	72

Prerequisites for NTP	72
Guidelines and Limitations for NTP	73
Default Settings for NTP	73
Configuring an NTP Server and Peer	73
Clearing NTP Sessions	74
Clearing NTP Statistics	74
Verifying the NTP Configuration	74
NTP Example Configuration	74
Feature History for NTP	75

CHAPTER 10

Configuring Local SPAN and ERSPAN	77
Information About SPAN and ERSPAN	77
SPAN Sources	77
Characteristics of SPAN Sources	78
SPAN Destinations	78
Characteristics of Local SPAN Destinations	78
Characteristics of ERSPAN Destinations	78
Local SPAN	79
Encapsulated Remote SPAN	79
Network Analysis Module	80
SPAN Sessions	80
Guidelines and Limitations for SPAN	81
Default Settings for SPAN	82
Configuring SPAN	82
Configuring a Local SPAN Session	82
Configuring an ERSPAN Port Profile	85
Configuring an ERSPAN Session	87
Shutting Down a SPAN Session from Global Configuration Mode	90
Shutting Down a SPAN Session from Monitor Configuration Mode	91
Resuming a SPAN Session from Global Configuration Mode	92
Resuming a SPAN Session from Monitor Configuration Mode	93
Configuring the Allowable ERSPAN Flow IDs	94
Verifying the SPAN Configuration	95
Configuration Example for an ERSPAN Session	95
Example of Configuring a SPAN Session	96

Example of a Configuration to Enable SPAN Monitoring	97
Feature History for SPAN and ERSPAN	97

CHAPTER 11**Configuring SNMP 99**

Information About SNMP	99
SNMP Functional Overview	99
SNMP Notifications	100
SNMPv3	100
Security Models and Levels for SNMPv1, v2, v3	100
User-Based Security Model	101
CLI and SNMP User Synchronization	102
Group-Based SNMP Access	102
High Availability	103
Guidelines and Limitations for SNMP	103
Default Settings for SNMP	103
Configuring SNMP	103
Configuring SNMP Users	104
Enforcing SNMP Message Encryption for All Users	105
Creating SNMP Communities	105
Configuring SNMP Notification Receivers	105
Configuring the Notification Target User	105
Enabling SNMP Notifications	106
Disabling LinkUp/LinkDown Notifications on an Interface	107
Enabling a One-time Authentication for SNMP over TCP	108
Assigning the SNMP Switch Contact and Location Information	108
Configuring a Host Receiver for SNMPv3 Traps or Informs	109
Disabling SNMP	109
Modifying the AAA Synchronization Time	110
Verifying the SNMP Configuration	110
Configuration Example for SNMP	111
Related Documents for SNMP	111
MIBs	112
Feature History for SNMP	113

CHAPTER 12**Configuring NetFlow 115**

Information about NetFlow	115
What is a Flow	116
Flow Record Definition	117
Predefined Flow Records	118
Accessing NetFlow Data	119
Command Line Interface for NetFlow	119
Flow Monitor	120
Flow Exporter	120
NetFlow Collector	120
Exporting Flows to the NetFlow Collector Server	121
What NetFlow Data Looks Like	122
Network Analysis Module	122
High Availability for NetFlow	122
Prerequisites for NetFlow	122
Configuration Guidelines and Limitations for NetFlow	123
Default Settings for NetFlow	123
Enabling the NetFlow Feature	124
Configuring Netflow	125
Defining a Flow Record	125
Defining a Flow Exporter	126
Defining a Flow Monitor	128
Assigning a Flow Monitor to an Interface	130
Adding a Flow Monitor to a Port Profile	131
Verifying the NetFlow Configuration	132
Netflow Example Configuration	135
Related Documents for NetFlow	136
Feature History for NetFlow	136
<hr/>	
CHAPTER 13	Configuring System Message Logging 137
	Information about System Message Logging 137
	System Message Logging Facilities 138
	Guidelines and Limitations for System Message Logging 142
	Default System Message Logging Settings 142
	Configuring System Message Logging 143
	Configuring System Message Logging to Terminal Sessions 143

Restoring System Message Logging Defaults for Terminal Sessions	144
Configuring System Message Logging for Modules	144
Restoring System Message Logging Defaults for Modules	145
Configuring System Message Logging for Facilities	146
Restoring System Message Logging Defaults for Facilities	146
Configuring syslog Servers	147
Restoring System Message Logging Defaults for Servers	147
Using a UNIX or Linux System to Configure Logging	148
Displaying Log Files	148
Verifying the System Message Logging Configuration	149
System Message Logging Example Configuration	152
Feature History for System Message Logging	152

CHAPTER 14
Configuring iSCSI Multipath 153

Information About iSCSI Multipath	153
Overview	154
Supported iSCSI Adapters	154
iSCSI Multipath Setup on the VMware Switch	155
Guidelines and Limitations	157
Pre-requisites	158
Default Settings	158
Configuring iSCSI Multipath	158
Uplink Pinning and Storage Binding	159
Process for Uplink Pinning and Storage Binding	159
Creating a Port Profile for a VMkernel NIC	159
Creating VMkernel NICs and Attaching the Port Profile	161
Manually Pinning the NICs	162
Identifying the iSCSI Adapters for the Physical NICs	164
Identifying iSCSI Adapters on the vSphere Client	164
Identifying iSCSI Adapters on the Host Server	164
Binding the VMkernel NICs to the iSCSI Adapter	165
Converting to a Hardware iSCSI Configuration	166
Converting to a Hardware iSCSI Configuration	166
Removing the Binding to the Software iSCSI Adapter	167
Adding the Hardware NICs to the DVS	167

Changing the VMkernel NIC Access VLAN	168
Process for Changing the Access VLAN	168
Changing the Access VLAN	169
Verifying the iSCSI Multipath Configuration	171
Managing Storage Loss Detection	172
Related Documents	174
Feature History for iSCSI Multipath	174

CHAPTER 15

Configuring VSM Backup and Recovery	175
Information About VSM Backup and Recovery	175
Guidelines and Limitations	175
Configuring VSM Backup and Recovery	176
Backing Up the VSM	176
Performing a Backup of the VSM	176
Performing a Periodic Backup	182
Recovering the VSM	182
Deploying the Backup VSM VM	182
Erasing the Old Configuration	190
Restoring the Backup Configuration on the VSM	191
Feature History for VSM Backup and Recovery	198

CHAPTER 16

Enabling vTracker	199
Information About vTracker	200
Guidelines and Limitations	201
Default Settings for vTracker Parameters	201
Enabling vTracker Globally	201
Upstream View	203
Upstream View Overview	203
Displaying Upstream View	204
Upstream View Field Description	204
Virtual Machine (VM) View	205
Virtual Machine (VM) View Overview	205
Displaying the VM vNIC View	206
VM vNIC View Field Description	207
Displaying the VM Info View	208

VM Info View Field Description	209
Module pNIC View	211
Module pNIC View Overview	211
Displaying the Module pNIC View	211
Module pNIC View Field Description	212
VLAN View	212
VLAN View Overview	212
Displaying the VLAN View	213
VLAN View Field Description	213
VMotion View	214
VMotion View Overview	214
Displaying the VMotion View	214
VMotion View Field Description	215
Feature History for vTracker	216

CHAPTER 17
Configuring Virtualized Workload Mobility 217

Information About Virtualized Workload Mobility (DC to DC vMotion)	217
Stretched Cluster	217
Split Cluster	218
Prerequisites for Virtualized Workload Mobility (DC to DC vMotion)	218
Guidelines and Limitations	218
Physical Site Considerations	218
Handling Inter-Site Link Failures	219
Headless Mode of Operation	219
Handling Additional Distance/Latency Between the VSM and VEM	219
Migrating a VSM	219
Migrating a VSM Hosted on an ESX	220
Verifying and Monitoring the Virtualized Workload Mobility (DC to DC vMotion)	
Configuration	220
Feature History for Virtualized Workload Mobility (DC to DC vMotion)	221



Preface

This preface contains the following sections:

- [Audience, page xiii](#)
- [Document Conventions, page xiii](#)
- [Related Documentation for Nexus 1000V Series NX-OS Software , page xv](#)
- [Documentation Feedback , page xvi](#)
- [Obtaining Documentation and Submitting a Service Request, page xvi](#)

Audience

This publication is for experienced network administrators who configure and maintain Cisco Nexus devices .

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

- An understanding of virtualization
- Using VMware software to create a virtual machine and configure a VMware vSwitch



Note

Knowledge of VMware vNetwork Distributed Switch is not required.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.

Convention	Description
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element(keyword or argument).
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
variable	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
screen font	Terminal sessions and information the switch 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:



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.

Related Documentation for Nexus 1000V Series NX-OS Software

This section lists the documents used with the Cisco Nexus 1000V and available on 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

Cisco Nexus 1000V Release Notes

Cisco Nexus 1000V and VMware Compatibility Information

Install and Upgrade

Cisco Nexus 1000V Installation and Upgrade Guide

Configuration Guides

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide

Cisco Nexus 1000V Interface Configuration Guide

Cisco Nexus 1000V Layer 2 Switching Configuration Guide

Cisco Nexus 1000V License Configuration Guide

Cisco Nexus 1000V Network Segmentation Manager Configuration Guide

Cisco Nexus 1000V Port Profile Configuration Guide

Cisco Nexus 1000V Quality of Service Configuration Guide

Cisco Nexus 1000V Security Configuration Guide

Cisco Nexus 1000V System Management Configuration Guide

Cisco Nexus 1000V vCenter Plugin Configuration Guide

Cisco Nexus 1000V VXLAN Configuration Guide

Programming Guide

Cisco Nexus 1000V XML API Configuration Guide

Reference Guides

Cisco Nexus 1000V Command Reference

Cisco Nexus 1000V MIB Quick Reference

Cisco Nexus 1000V Resource Availability Reference

Troubleshooting and Alerts

Cisco Nexus 1000V Troubleshooting Guide

Cisco Nexus 1000V Password Recovery Procedure

Cisco NX-OS System Messages Reference

Virtual Services Appliance Documentation

The *Cisco Nexus Virtual Services Appliance* documentation is available at http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html.

Virtual Security Gateway Documentation

The *Cisco Virtual Security Gateway for Nexus 1000V Series Switch* documentation is available at http://www.cisco.com/en/US/products/ps11208/tsd_products_support_model_home.html.

Virtual Wide Area Application Services (vWAAS) Documentation

The *Virtual Wide Area Application Services* documentation is available at http://www.cisco.com/en/US/products/ps6870/tsd_products_support_series_home.html.

ASA 1000V Cloud Firewall Documentation

The *ASA 1000V Cloud Firewall* documentation is available at http://www.cisco.com/en/US/products/ps12233/tsd_products_support_series_home.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus1k-docfeedback@cisco.com. We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

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

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

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



CHAPTER 1

New and Changed Information for this Release

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to the configuration guides or of the new features in this release.

- [New and Changed Information, page 1](#)

New and Changed Information

Table 1: New and Changed Features

Content	Description	Changed in Release	Where Documented
vTracker	This feature is introduced.	4.2(1)SV2(1.1)	Enabling vTracker, on page 199
Virtualized Workload Mobility (DC to DC vMotion)	Addresses the Cisco Nexus 1000 across two physical data centers.	4.2(1)SV1(4a)	Configuring Virtualized Workload Mobility, on page 217
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, on page 31
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, on page 175
Enable NetFlow feature	You can enable/disable the NetFlow feature.	4.2(1)SV1(4)	Configuring NetFlow, on page 115
Add port profile as Local SPAN source	Allows you to use a port profile as a source for Local SPAN monitor traffic.	4.2(1)SV1(4)	Configuring Local SPAN and ERSPAN, on page 77
Add port profile as ERSPAN source	Allows you to use a port profile as a source for ERSPAN monitor traffic.	4.2(1)SV1(4)	Configuring Local SPAN and ERSPAN, on page 77

Content	Description	Changed in Release	Where Documented
Hardware iSCSI Multipath	Allows you to use a hardware iSCSI adapter for multipathing.	4.2(1)SV1(4)	Configuring iSCSI Multipath, on page 153
SNMP MIBs added	Added list of supported MIBs.	4.2(1)SV1(4)	Configuring SNMP, on page 99
Network Analysis Module (NAM)	NAM support for NetFlow data sources	4.2(1)SV1(4)	Configuring NetFlow, on page 115
	NAM support for ERSPAN data sources	4.0(4)SV1(3)	Configuring Local SPAN and ERSPAN, on page 77
ERSPAN Type-III header	Provides the ERSPAN Type-III extended format header frame that enhances support for network management, intrusion detection, and lawful intercept	4.0(4)SV1(3)	Configuring Local SPAN and ERSPAN, on page 77
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, on page 17
iSCSI Multipath	Allows multiple routes between a server and its storage devices.	4.0(4)SV1(2)	Configuring iSCSI Multipath, on page 153



CHAPTER 2

Overview

This chapter contains the following sections:

- [System Management Overview](#) , page 3

System Management Overview

This chapter describes the following system management features:

- CDP
- Domains
- Server Connections
- Configuration Management
- File Management
- User Management
- NTP
- Local SPAN and ERSPAN
- SNMP System Messages
- NetFlow
- System Messages
- iSCSI Multipath
- Troubleshooting

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 [Configuring CDP](#) section.

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 1000Vwhen 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.

Server Connections

In order to connect to vCenter Server or an ESX server, you must first define the connection in the Cisco Nexus 1000V. The Managing Server Connections section 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 Managing the Configuration section.

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 Working with Files section.

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 Managing Users section.

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 Configuring NTP section.

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 [Configuring Local SPAN and ERSPAN](#) section. 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](#).

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 [Configuring SNMP](#) section.

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 [Configuring NetFlow](#) section. 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](#).

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 [Configuring System Message Logging](#) section.

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](#) section.

Troubleshooting

Ping and trace route are among the available troubleshooting tools. For more information, see the [Cisco Nexus 1000V Troubleshooting Guide](#).



CHAPTER 3

Configuring CDP

This chapter contains the following sections:

- [Information About CDP, page 7](#)
- [Guidelines and Limitations, page 8](#)
- [Default Settings, page 8](#)
- [Configuring CDP, page 9](#)

Information About CDP

The Cisco Discovery Protocol (CDP), which runs over the data link layer, 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.

Each device that 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
- 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.

High Availability

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

Guidelines and Limitations

- CDP gathers protocol addresses of neighboring devices and discovers the platform of those devices. CDP runs over the data link layer only. With CDP, two systems that support different Layer 3 protocols can learn about each other.
- 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.
- You can configure CDP on physical interfaces and port channels only.

Default Settings

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

Configuring CDP

This section includes the following topics:

- CDP Global Configuration
- Enabling CDP on an Interface
- Disabling CDP on an Interface

CDP Global Configuration

This section includes the following topics:

- Enabling or Disabling CDP Globally
- Advertising a CDP Version
- Configuring CDP Options

Enabling or Disabling CDP Globally

Be sure you understand that when you globally disable the CDP feature, all CDP configurations are removed.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# [no] cdp enable	Enables or disables the CDP feature globally.

```
switch# config t
switch(config)# no cdp enable
```

Advertising a CDP Version

Before beginning this procedure, be sure you have know the following information:

- 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.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# cdp advertise {v1 v2}	Assigns the CDP version to advertise: <ul style="list-style-type: none"> • CDP Version 1 • CDP Version 2
Step 3	switch(config)# show cdp global	(Optional) Displays the CDP version that is being advertised or sent to other devices.
Step 4	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch# config t
switch(config)# cdp advertise v1
switch(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
switch(config)# copy running-config startup-config

```

Configuring CDP Options

You can 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



Note You can view output from the upstream Catalyst 6500 Series switch by using the **show cdp neighbor** command.

Before You Begin

Before beginning this procedure, be sure you know the following information:

- How long you want CDP to retain neighbor information if you are setting the holdtime.
- How often you want CDP to advertise if you are setting the CDP timer.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# cdp format device-id system-name	(Optional) Specifies that CDP uses the system name for the device ID format.
Step 3	switch(config)# show cdp neighbors	Displays your device from the upstream device.
Step 4	switch(config)# show cdp neighbors	Displays the upstream device from your device.
Step 5	switch(config)# cdp holdtime seconds	(Optional) Sets the maximum amount of time that CDP holds onto neighbor information before discarding it. <ul style="list-style-type: none"> • The range for the <i>seconds</i> argument is from 10 to 255 seconds. • The default is 180 seconds.
Step 6	switch(config)# cdp timer seconds	(Optional) Sets the refresh time for CDP to send advertisements to neighbors. <ul style="list-style-type: none"> • The range for the <i>seconds</i> argument is from 5 to 254 seconds.
Step 7	switch(config)# show cdp global	(Optional) Displays the CDP version that is being advertised or sent to other devices.
Step 8	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch# config t
switch(config)# cdp format device-id system-name
switch# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID      Local Infrfce  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
switch(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

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
switch(config)# cdp holdtime 10
switch(config)# cdp timer 5
switch(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
switch(config-if)# copy running-config startup-config

```

CDP Interface Configuration

This section includes the following procedures:

- Enabling CDP on an Interface
- Disabling CDP on an Interface

Enabling CDP on an Interface

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

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# interface <i>interface-type number</i>	Places you in interface configuration mode for the specific interface.
Step 3	switch(config-if)# no cdp enable	Disables CDP on this interface.
Step 4	switch(config-if)# cdp enable	Enables CDP on this interface.
Step 5	switch(config-if)# show cdp interface <i>interface-type number</i>	(Optional) Displays CDP information for the specified interface.

	Command or Action	Purpose
Step 6	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```
switch# config t
switch(config)# interface port-channel 2
switch(config-if)# no cdp enable
switch(config-if)# cdp enable
switch(config-if)# show cdp interface mgmt0
mgmt0 is up
      CDP disabled on interface
      Sending CDP packets every 60 seconds
      Holdtime is 180 seconds
switch(config)# copy running-config startup-config
```

Disabling CDP on an Interface

Before You Begin

Before beginning this procedure, be sure of the following:

- CDP is currently enabled on the device.



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

- CDP is currently enabled on the specific interface you want to configure.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# interface <i>interface-type number</i>	Places you in interface configuration mode for the specific interface.
Step 3	switch(config-if)# no cdp enable	Disables CDP on this interface.
Step 4	switch(config-if)# show cdp interface <i>interface-type number</i>	(Optional) Displays CDP information for the specified interface.
Step 5	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

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

```

Monitoring CDP

Command	Purpose
show cdp traffic interface <i>interface-type slot/port</i>	Displays the CDP traffic statistics on an interface.

Clearing CDP Statistics

Use one of the following commands to clear CDP statistics:

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

Verifying the CDP Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show cdp all	Displays all interfaces that have CDP enabled.
show cdp entry { all name <i>entry-name</i> }	Displays the CDP database entries.
show cdp global	Displays the CDP global parameters.
show cdp interface <i>interface-type slot/port</i>	Displays the CDP interface status.
show cdp neighbors { detail interface <i>interface-type slot/port</i> }	Displays the CDP neighbor status.

Configuration Example for CDP

This example shows how to enable the CDP feature and configures the refresh and hold timers:

```
switch# config t
switch(config)# cdp enable
switch(config)# cdp timer 50
switch(config)# cdp holdtime 100
```

Feature History for CDP

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



CHAPTER 4

Configuring the Domain

This chapter contains the following sections:

- [Information About the Domain, page 17](#)
- [Guidelines and Limitations, page 18](#)
- [Default Settings, page 19](#)
- [Configuring the Domain, page 19](#)
- [Feature History for the VSM Domain, page 28](#)

Information About the Domain

You must create a domain name for the 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 by using the **setup** command or the procedures described in this chapter.

Layer 3 Control

Layer 3 control, or IP connectivity, is supported between the Virtual Supervisor Module (VSM) and the Virtual Ethernet Module (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. In the Layer 3 mode, all the VEMs (hosts) managed by VSM and the VSM can be in different networks.


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.
- Configure a port profile.
- 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.

In the following diagram, VSM0 controls VEM_0_1, VEM_0_1 hosts VSM1 and VSM2, and VSM1 and VSM2 control VEMs in other Layer 2 networks.

Figure 1: Example of Layer 3 Control IP Connectivity



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, make sure that 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 might 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.
- 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 network file system (NFS) mount.
- You must configure control VLANs, packet VLANs, and management VLANs as regular VLANs and not as private VLANs.

Default Settings

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
- Changing to Layer 3 Transport
- Changing to Layer 2 Transport
- Creating a Port Profile for Layer 3 Control
- Creating a Control VLAN
- Creating a Packet VLAN

Creating a Domain

You can 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 by using this procedure.

**Note**

We recommend the following:

- Use one VLAN for control traffic and a different VLAN for packet traffic.
- Use a distinct VLAN for each instances of Cisco Nexus 1000V (different domains)

Before You Begin

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

You must know the following information:

- If two or more VSMs share the same control and/or packet VLAN, the domain helps identify the VEMs managed by each VSM.
- A unique domain ID for this Cisco Nexus 1000V instance.
- Identity of the VLANs to be used for control and packet traffic.
- The **svs mode** command in the SVS Domain Configuration mode is not used and has no effect on a configuration.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# svs-domain	Places you in SVS domain configuration mode.
Step 3	switch(config-svs-domain)# domain id number	Creates the domain ID for this Cisco Nexus 1000V instance.
Step 4	switch(config-svs-domain)# control vlan number	Assigns the control VLAN for this domain.
Step 5	switch(config-svs-domain)# packet vlan number	Assigns the packet VLAN for this domain.
Step 6	switch(config-vlan)# show svs domain	(Optional) Displays the domain configuration.
Step 7	switch(config-vlan)# exit	Returns you to global configuration mode.
Step 8	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```
switch# config t
switch(config)# svs-domain
switch(config-svs-domain)# domain id 100
```

```

switch(config-svs-domain)# control vlan 190
switch(config-svs-domain)# packet vlan 191
switch(config-vlan)# exit

switch(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.

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

```

Changing to Layer 3 Transport

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.

Before You Begin

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

You have already configured the Layer 3 interface (mgmt 0 or control 0) and assigned an IP address.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# show svcs domain	Displays the existing domain configuration, including control and packet VLAN IDs.
Step 2	switch# config t	Places you in global configuration mode.
Step 3	switch(config)# svcs-domain	Places you in SVS domain configuration mode.
Step 4	switch(config-svs-domain)# no packet vlan	Removes the packet VLAN configuration.
Step 5	switch(config-svs-domain)# no control vlan	Removes the control VLAN configuration.
Step 6	switch(config-svs-domain)# show svcs domain	(Optional) Displays the domain configuration.
Step 7	switch(config-svs-domain)# svcs mode L3 interface { mgmt0 control0 }	Configures Layer 3 transport mode for the VSM domain. If configuring Layer 3 transport, then you must designate which interface to use; and the interface must already have an IP address configured.

	Command or Action	Purpose
Step 8	switch(config-vlan)# show svcs domain	(Optional) Displays the new Layer 3 control mode configuration for this VSM domain.
Step 9	switch(config-svs-domain)# [no] control type multicast	Configures the control type multicast in Layer 3 mode on the VSM.
Step 10	switch(config-vlan)# show svcs domain	(Optional) Displays the control type multicast status in Layer 3 mode on the VSM.
Step 11	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch(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.
switch# config t
switch(config)# svcs-domain
switch(config-svs-domain)# no packet vlan
switch(config-svs-domain)# no control vlan
switch(config)# show svcs 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)# svcs mode l3 interface mgmt0
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.
switch(config-svs-domain)# show svcs domain

switch(config-svs-domain)# control type multicast
switch(config)# show svcs domain
SVS domain config:
  Domain id:    343
  Control vlan: NA
  Packet vlan: NA
  L2/L3 Control mode: L3
  L3 control interface: mgmt0
  Status: Config push to VC successful.
  Control type multicast: Yes

switch(config-svs-domain)# no control type multicast
switch(config)# show svcs domain
SVS domain config:
  Domain id:    343
  Control vlan: NA

```

```

Packet vlan: NA
L2/L3 Control mode: L3
L3 control interface: mgmt0
Status: Config push to VC in progress.
Control type multicast: No
Limitation : Control type multicast is configured. It is not applicable in svcs L2 mode.

switch(config-svs-domain)# copy running-config startup-config
[#####] 100%
switch(config-svs-domain)#

```

Changing to Layer 2 Transport

You can 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.

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.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch(config)# show svcs domain	Displays the existing domain configuration, including control and packet VLAN IDs and the Layer 3 interface configuration.
Step 2	switch# config t	Places you in global configuration mode.
Step 3	switch(config)# svcs-domain	Places you in SVS domain configuration mode.
Step 4	switch(config-svs-domain)# svcs mode L2	Configures Layer 2 transport mode for the VSM domain.
Step 5	switch(config-svs-domain)# control vlan vlanID	Configures the specified VLAN ID as the control VLAN for the VSM domain.
Step 6	switch(config-svs-domain)# packet vlanvlanID	Configures the specified VLAN ID as the packet VLAN for the VSM domain.
Step 7	switch(config-vlan)# show svcs domain	(Optional) Displays the new Layer 2 control mode configuration for this VSM domain.
Step 8	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch# show svcs domain
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.
switch# config t
switch(config)# svcs-domain
switch(config-svs-domain)# svcs mode l2
switch(config-svs-domain)# control vlan 100
switch(config-svs-domain)# packet vlan 101
switch(config-svs-domain)# 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.
switch(config-svs-domain)# copy running-config startup-config
[#####] 100%

```

Creating a Port Profile for Layer 3 Control

You can allow the VSM and VEM to communicate over IP for control and packet traffic.

Before You Begin

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

You must know the following information:

- The transport mode for the VSM domain has already been configured as Layer 3.
- All VEMs belong to the same Layer 2 domain.
- The VEM VM kernel NIC connects to this Layer 3 control port profile when you add 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.
- 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.
 - 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 with the **capability L3 control** command.
- Different hosts can use different VLANs for Layer 3 control.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.

	Command or Action	Purpose
Step 2	<code>switch(config)# port-profile name</code>	Creates a port profile and places you into Port Profile Configuration mode for the named port profile. The <i>name</i> argument can be up to 80 characters and must be unique for each port profile on the Cisco Nexus 1000V.
Step 3	<code>switch(config-port-prof)# capability l3control</code>	Allows the port to be used for IP connectivity.
Step 4	<code>switch(config-port-prof)# vmware port-group [name]</code>	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. 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	<code>switch(config-port-prof)# switchport mode access</code>	Designates that the interfaces are switch access ports (the default).
Step 6	<code>switch(config-port-prof)# switchport access vlan vlanID</code>	Assigns the system VLAN ID to the access port for this Layer 3 control port profile.
Step 7	<code>switch(config-port-prof)# no shutdown</code>	Administratively enables all ports in the profile.
Step 8	<code>switch(config-port-prof)# system vlan vlanID</code>	Adds the system VLAN to this Layer 3 control port profile. This command ensures that, when the host is added for the first time or rebooted later, the VEM can reach the VSM. One of the uplink ports must have this VLAN in its system VLAN range.
Step 9	<code>switch(config-port-prof)# state enabled</code>	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 vCenter Server.
Step 10	<code>switch(config-port-prof)# show port-profile name name</code>	(Optional) Displays the current configuration for the port profile.
Step 11	<code>switch(config)# copy running-config startup-config</code>	(Optional) Copies the running configuration to the startup configuration.

```

switch# config t
switch(config)# port-profile l3control-150
switch(config-port-prof)# capability l3control
switch(config-port-prof)# vmware port-group
switch(config-port-prof)# switchport mode access
switch(config-port-prof)# switchport access vlan 150

```

```

switch(config-port-prof)# no shutdown
switch(config-port-prof)# system vlan 150
switch(config-port-prof)# state enabled
switch(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:
switch(config-port-prof)# copy running-config startup-config

```

Creating a Control VLAN

Before You Begin

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

Be sure you have already configured and enabled the required switched virtual interface (SVI) using the document, *Cisco Nexus 1000V Interface Configuration Guide*. The SVI is also called the VLAN interface and provides communication between VLANs.

You must know the following:

- If Layer 3 Control is configured on your VSM, you cannot create a control VLAN. You must first disable Layer 3 Control.
- How VLANs are numbered.
- That newly created VLANs remain unused until Layer 2 ports are assigned to them.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.
Step 2	switch(config)# vlan 30	Creates VLAN ID 30 for control traffic and places you in 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	switch(config-vlan)# name cp_control	Adds the descriptive name, cp_control, to this VLAN.

	Command or Action	Purpose
Step 4	switch(config-vlan)# state active	Changes the operational state of the VLAN to active.
Step 5	switch(config-vlan)# show vlan id 30	(Optional) Displays the configuration for VLAN ID 30.
Step 6	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch# config t
switch(config)# vlan 30
switch(config-vlan)# name cp_control
switch(config-vlan)# state active
switch(config-vlan)# show vlan id 30
VLAN Name                Status      Ports
-----
30    cp_control            active
VLAN Type MTU
----
5     enet 1500
Remote SPAN VLAN
-----
Disabled
Primary  Secondary  Type          Ports
-----
switch(config-vlan)# copy running-config startup-config

```

Creating a Packet VLAN

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode
- Configured and enabled the required switched virtual interface (SVI)
- Familiarized yourself with how VLANs are numbered.



Note

Newly created VLANs remain unused until Layer 2 ports are assigned to them.

Procedure

	Command or Action	Purpose
Step 1	switch# config t	Places you in global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# vlan <i>vlan-id</i>	Creates VLAN ID for packet traffic and places you in 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	switch(config-vlan)# name <i>vlan-name</i>	Adds the descriptive name to this VLAN.
Step 4	switch(config-vlan)# state <i>vlan-state</i>	Changes the operational state of the VLAN to active or suspend.
Step 5	switch(config-vlan)# show vlan id <i>vlan-id</i>	(Optional) Displays the configuration for the VLAN ID.
Step 6	switch(config-vlan)# exit	Returns you to global configuration mode.
Step 7	switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

```

switch# config t
switch(config)# vlan 31
switch(config-vlan)# name cp_packet
switch(config-vlan)# state active
switch(config-vlan)# exit
switch(config)# show vlan id 31

VLAN Name                Status    Ports
-----
31  cp_packet              active

VLAN Type MTU
-----
5   enet 1500

Remote SPAN VLAN
-----
Disabled

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

```

Feature History for the VSM Domain

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

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• Guidelines and Limitations• Changing to Layer 2 Transport• Changing to Layer 3 Transport• Creating a Port Profile for Layer 3 Control
VSM Domain	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 5

Managing Server Connections

This chapter contains the following sections:

- [Information About Server Connections, page 31](#)
- [Guidelines and Limitations, page 32](#)
- [Connecting to the vCenter Server, page 32](#)
- [Disconnecting From the vCenter Server, page 33](#)
- [Removing the DVS from the vCenter Server, page 34](#)
- [Removing the DVS from the vCenter Server When the VSM Is Not Connected, page 35](#)
- [Configuring Host Mapping, page 36](#)
- [Verifying Connections, page 38](#)
- [Verifying the Domain, page 39](#)
- [Verifying the Configuration, page 40](#)
- [Verifying Module Information, page 40](#)
- [Feature History for Server Connections, page 42](#)

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
- All communication with vCenter Server is secured by the Transport Layer Security (TLS) protocol.

Guidelines and Limitations

A single Virtual Supervisor Module (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

Before You Begin

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

You must know the following:

- The datacenter name
- The vCenter Server IP address or hostname.

You must be sure the following is set up:

- The vCenter Server management station is installed and running.
- The ESX servers are installed and running.
- The Cisco Nexus 1000V appliance is installed.
- The management port is configured.
- The DNS is already configured if you are configuring a connection using a hostname.
- An extension with vCenter Server has been registered. The extension includes the extension key and public certificate for the VSM. vCenter Server uses the extension 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 Installation and Upgrade Guide*.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# svs connection name	Places you in connection configuration mode for adding this connection between the Cisco Nexus 1000V and either a particular ESX server or vCenter Server. By using a name, information for multiple connections can be stored in the configuration.
Step 3	switch(config-svs-conn)# protocol vmware-vim [http]	Use the http keyword to specify that this connection uses the VIM protocol. This command is stored locally. http : Specifies that the VIM protocol runs over HTTP. The default is to use HTTP over SSL (HTTPS).
Step 4	Do one of the following:	<ul style="list-style-type: none"> • If you are configuring an IP address, go to Step 5.

	Command or Action	Purpose
		<ul style="list-style-type: none"> If you are configuring a hostname, go to Step 6.
Step 5	switch(config-svs-conn)# remote ip address <i>ipaddress</i>	Specifies the IP address of the ESX server or vCenter Server for this connection. This command is stored locally. Go to step 7 to configure the datacenter name.
Step 6	switch(config-svs-conn)# remote hostname <i>hostname</i>	Specifies the DNS name of the ESX server or vCenter Server for this connection. This command is stored locally. Note DNS is already configured.
Step 7	switch(config-svs-conn)# vmware dvs datacenter-name <i>name</i>	Identifies the datacenter name in the vCenter Server where the 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.
Step 8	switch(config-svs-conn)# connect	Initiates the connection. If the username and password have not been configured for this connection, the you are prompted for a username and password. 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 appears and the command is rejected until you close the previous connection by entering no connect.

```

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

Disconnecting From the vCenter Server

You can disconnect from the vCenter Server, for example, after correcting a vCenter Server configuration.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the Cisco Nexus 1000V in EXEC mode.

- Configured an Cisco Nexus 1000V connection
- Connected the Cisco Nexus 1000V to vCenter Server/ESX.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# svs connection name	Places you in a global configuration submode for the connection to vCenter Server.
Step 3	switch(config-svs-conn)# no connect	Closes the connection.

```
switch# config t
switch# (config)# svs connection vcWest
switch# (config-svs-conn)# no connect
```

Removing the DVS from the vCenter Server

Use this procedure to remove the Distributed Virtual Switch (DVS) from the vCenter Server.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the Cisco Nexus 1000V in EXEC mode
- Configured a connection to the vCenter Server
- Connected the Cisco Nexus 1000V to vCenter Server/ESX
- Checked that the server administrator has removed all of the hosts that are connected to the Cisco Nexus 1000V from the VI client. For more information, see the VMware documentation.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# svs connection name	Places you in a global configuration submode for the connection to vCenter Server.
Step 3	switch(config-svs-conn)# no vmware dvs	Removes the DVS associated with the specified connection from the vCenter Server.

```
switch# config t
switch(config)# svs connection vcWest
switch(config-svs-conn)# no vmware dvs
```

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

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

Procedure

-
- Step 1** Configure the admin user or group. See the Configuring the Admin User or Admin Group section.
 - Step 2** Remove the DVS from the vCenter Server. See the Removing the DVS from the vCenter Server section
-

Configuring the Admin User or Admin Group

Before You Begin

Before beginning this procedure, ensure that the system administrator has created an admin user or admin group on vCenter Server to manage and delete the DVS. This user should not be given any other permissions such as deploying VMs or hosts, and so on. The admin user name configured on the VSM should be the same as the username on vCenter Server.

Procedure

-
- Step 1** Determine the name of the DVS.

Example:

```
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.

Example:

```
switch# config t
switch(config)# svcs connection VC
switch(config-svcs-conn) # admin user N1K-DC
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.

Example:

```
switch# show vs connections

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

Removing the DVS from the vCenter Server Using the Graphical User Interface

Procedure

- Step 1** Log in to the vCenter Server through the VMware vSphere Client with the admin user account
 - 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
- Removing Host Mapping from a Module
- Mapping to a New Host
- Viewing Host Mapping

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 the universally unique identifier (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

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the Cisco Nexus 1000V in EXEC mode.
- Removed the host from the Cisco Nexus 1000V DVS on vCenter

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# no vem <i>module-number</i>	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	switch(config)# show module vem mapping	(Optional) Displays the mapping of modules to host servers.
Step 4	switch(config)# copy running-config startup-config	Copies the running configuration to the startup configuration.

```
switch# configure terminal
switch(config)# no vem 4
switch(config)# no vem 3
cannot modify slot 3: host module is inserted
switch(config)# show module vem mapping
Mod      Status          UUID                               License Status
-----
  3      powered-up      93312881-309e-11db-afaf-0015170f51a8  licensed
switch(config-vem-slot)# copy running-config startup-config
```

Mapping to a New Host

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode
- Removed the host from the Cisco Nexus 1000V DVS on vCenter



Note

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# vem module number	Places you in VEM slot configuration mode.
Step 3	switch(config-vem-slot)# host vmware id server-bios-uuid	Assigns a different host server UUID to the specified module.
Step 4	switch(config-vem-slot)# show module vem mapping	(Optional) Displays the mapping of modules to host servers.
Step 5	switch(config-vem-slot)# copy running-config startup-config	Copies the running configuration to the startup configuration.

```
switch# config t
switch(config)# vem 3
switch(config-vem-slot)# host vmware id 6dd6c3e3-7379-11db-abcd-000bab086eb6
switch(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

switch(config-vem-slot)# copy running-config startup-config
```

Viewing Host Mapping

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

Procedure

Display the mapping on modules to host servers by entering the following command: **show module vem mapping**

Example:

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

Procedure

show svcs connections [*name*]

Example:

```
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#
```

Displays the current connections to the Cisco Nexus1000V.

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 **no connect**

```
connect
```

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.

Procedure

show svcs domain

Example:

```
n1000v# show svcs domain
SVS domain config:
Domain id: 98
Control vlan: 70
Packet vlan: 71
Sync state: -
n1000v#
```

Display the domain configured on the Cisco Nexus 1000V.

Verifying the Configuration

Use one of the following commands to verify the configuration:

Command	Description
show running-config	Displays 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.
show svcs connections [<i>name</i>]	Displays the current connections to the Cisco Nexus 1000V. 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: no connect connect
show svcs domain	Displays the domain configured on the Cisco Nexus 1000V.
show module	Displays module information.
show server_info	Displays server information.
show interface brief	Displays interface information, including the uplinks to vCenter Server.
show interface virtual	Displays virtual interface information.
show module vem mapping	Displays the mapping of modules to host servers.

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.
- 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.

Procedure

Step 1 show module

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 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
```

Displays module information.

Step 2 show server_info

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#
```

Displays server information.

Step 3 show interface brief

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#
Displays interface information, including the uplinks to vCenter Server.

```

Step 4 show interface virtual**Example:**

```

n1000v# show interface virtual
-----
Port Adapter Owner Mod Host
-----
Veth49 R-VM-1 2 mcs-srvr35
Displays virtual interface information.

```

Feature History for Server Connections

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.



CHAPTER 6

Managing the Configuration

This chapter contains the following sections:

- [Information About Configuration Management, page 43](#)
- [Changing the Switch Name, page 43](#)
- [Configuring a Message of the Day, page 44](#)
- [Verifying the Configuration, page 45](#)
- [Verifying the Interface Configuration, page 48](#)
- [Saving a Configuration, page 51](#)
- [Erasing a Configuration, page 51](#)
- [Feature History for Configuration Management, page 52](#)

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.

If the VSM is connected to vCenter Server then this procedure also changes the Dynamic Vectoring and Streaming (DVS) engine that the VSM is managing. If you make an error when renaming the DVS, a syslog is generated and the DVS on vCenter Server continues to use the old DVS name.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch(config)# switchname	Changes the switch prompt.

```
switch(config)# switchname metro
metro(config)# exit
metro#
```

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.

- 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.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch(config)# banner motd [<i>delimiting-character message</i> <i>delimiting-character</i>]	Configures a banner message of the day with the following features: <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	switch(config)# show banner motd	Displays the configured banner message.

```
switch(config)# banner motd #April 16, 2011 Welcome to the svr#
switch(config)# show banner motd
April 16, 2011 Welcome to the Switch
```

Verifying the Configuration

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

- Verifying the Software and Hardware Versions
- Verifying the Running Configuration
- Comparing the Startup and Running Configurations
- Verifying the Interface Configuration

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

Procedure

	Command or Action	Purpose
Step 1	swtich# show version	Displays the versions of system software and hardware that are currently running on the switch,

```
switch# 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")
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 procedure to view the configuration currently running on the system.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show running-config	Displays the versions of system software and hardware that are currently running on the switch.

```

switch# 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 switch 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 switch
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
 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

switch#
```

Comparing the Startup and Running Configurations

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	swtich# show running-config diff	Displays the difference between the startup configuration and the running configuration currently on the switch.

```

switch# 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

+ 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

switch#

```

Verifying the Interface Configuration

This section includes the following procedures:

- Verifying a Brief Version of an Interface Configuration
- Verifying a Detailed Version of an Interface Configuration
- Verifying a Brief Version of all Interfaces

- Verifying the Running Configuration for all Interfaces

Verifying the Interface Configuration in a Brief Version

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show interface <i>{type}</i> <i>{name}</i> brief	Displays a brief version of information about the specified interface configuration.

```
switch# show interface mgmt 0 brief
```

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

Verifying an Interface Configuration in a Detailed Version

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show interface <i>{type}</i> <i>{name}</i>	Displays details about the specified interface configuration.

```
switch# 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
```

Verifying All Interfaces in a Brief Version

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show interface brief	Displays a brief version of all interface configurations on your system,

```
switch# 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                                Speed  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
switch#
```

Verifying the Running Configuration for all Interfaces

The output for the command, **show running-config interface** differs from the output of the **show interface** command.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show running-config interface	Displays the running configuration for all interfaces on your system.

```
switch# 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

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

switch#
    
```

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 beginning this procedure, you must be logged in to the CLI in any command mode.

Procedure

	Command or Action	Purpose
Step 1	switch# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```

switch# copy run start
[#####] 100%
switch#
    
```

Erasing a Configuration

Use this procedure to erase a startup configuration.

**Caution**

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

Before You Begin

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

•

Procedure

	Command or Action	Purpose
Step 1	switch# write erase [boot debug]	<p>The existing startup configuration is completely erased and all settings revert to their factory defaults.</p> <p>The running configuration is not affected.</p> <p>The following parameters are used with this command:</p> <ul style="list-style-type: none"> • boot: Erases the boot variables and the mgmt0 IP configuration. • debug: Erases the debug configuration.

```
switch# write erase debug
```

Feature History for Configuration Management

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



CHAPTER 7

Working with Files

This chapter contains the following sections:

- [Information About Files, page 53](#)
- [Navigating the File System, page 54](#)
- [Copying and Backing Up Files, page 58](#)
- [Creating a Directory, page 59](#)
- [Removing an Existing Directory, page 60](#)
- [Moving Files, page 60](#)
- [Deleting Files or Directories, page 61](#)
- [Compressing Files, page 62](#)
- [Uncompressing Files, page 63](#)
- [Directing Command Output to a File, page 63](#)
- [Verifying a Configuration File before Loading, page 64](#)
- [Rolling Back to a Previous Configuration , page 64](#)
- [Displaying Files, page 65](#)
- [Feature History for File Management, page 67](#)

Information About Files

The Cisco Nexus 1000V file system provides a single interface to all the file systems that the Cisco Nexus 1000V 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)

Navigating the File System

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

- Specifying File Systems
- Identifying the Directory You are Working From
- Changing Your Directory
- Listing the Files in a File System
- Identifying Available File Systems for Copying Files
- Using Tab Completion

Specifying File Systems

The syntax for specifying a file system is `<file system name>:[//server/]`. The following table describes file system syntax.

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

You can display the directory name of your current CLI location.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# pwd	Displays the present working directory.

```
switch# pwd
bootflash:
```

Changing Your Directory

You can change your location in the CLI, from one directory or file system to another.

Cisco Nexus 1000V CLI defaults to the bootflash: file system.

**Note**

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

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# pwd	Displays the directory name of your current CLI location.
Step 2	switch# cd directory name <ul style="list-style-type: none"> • switch# cd bootflash: Changes your CLI location to the root directory on the bootflash: file system. • switch# cd bootflash:mydir Changes your CLI location to the mydir directory that resides in the bootflash: file system. • switch# cd mystorage Changes your CLI location to the mystorage directory that resides within the current directory. If the current directory is bootflash: mydir, this command changes the current directory to bootflash: mydir/mystorage. 	Changes your CLI location to the root directory on the bootflash: file system.

```

switch# pwd
volatile:
switch# cd bootflash:

switch# pwd
volatile:
switch# cd bootflash:mydir

switch# pwd
volatile:
switch# cd mystorage

```

Listing the Files in a File System

Procedure

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

```

switch# 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
switch#

```

Identifying Available File Systems for Copying Files

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# copy ?	Displays the source file systems available to the copy command.

	Command or Action	Purpose
Step 2	switch# copy filename ?	Displays the destination file systems available to the copy command for a specific file.

```
switch# 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
```

Using Tab Completion

You can have the CLI complete a partial file name in a command.

Procedure

	Command or Action	Purpose
Step 1	switch# show file filesystem <i>name: partial filename</i> <Tab>	Completes the filename when you type a partial filename and then press Tab and if the characters you typed are unique to a single file. If not, the CLI lists a selection of file names that match the characters that you typed. You can then retype enough characters to make the file name unique; and CLI completes the filename for you.
Step 2	switch# show file bootflash:c <Tab>	Completes the file name for you

```
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.bin
bootflash:nexus-1000v-kickstart-mzg.4.0.4.SV1.0.42.bin
n1000v# show file bootflash:c<Tab>
-----BEGIN RSA PRIVATE KEY-----
MIICXgIBAAKBgQDSq93BrlHcg3bX1jXDMY5c9+yZSST3VhuQBqogvCPDGeLecA+j
...
...
n1000v#
```

Copying and Backing Up Files

You can 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.



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.

Before You Begin

Before beginning this procedure, you must be of the following:

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

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# copy [source filesystem:] filename [destination filesystem:] filename</pre> <ul style="list-style-type: none"> • switch# copy system:running-config system run.cfg Saves a copy of the running configuration to a remote switch. • switch# copy bootflash: system_image bootflash://sup-standby/system_image Copies a file from bootflash in the active supervisor module to bootflash in the standby supervisor module. • switch# copy system:running-config bootflash:config Copies a running configuration to the bootflash: file system. • switch# copy scp://[username@]server[/path]/filename Copies a source or destination URL for a network server that supports Secure Shell (SSH) and accepts copies of files using the secure copy protocol (scp). • switch# copy sftp://[username@]server[/path]/filename/// Copies a source or destination URL for an SSH FTP (SFTP) network server • switch# copy system:running-config bootflash:my-config 	Copies a file from the specified source location to the specified destination location.

	Command or Action	Purpose
	<p>Places a back up copy of the running configuration on the bootflash: file system (ASCII file).</p> <ul style="list-style-type: none"> • switch# copy bootflash: filename bootflash:directory/filename Copies the specified file from the root directory of the bootflash: file system to the specified directory. • switch# copy filename directory/filename Copies a file within the current file system. • switch# copy tftp://server[:port]][/path]/filename Copies the source file to the running configuration on the switch, and configures the switch as the file is parsed line by line. 	

```

switch# copy system:running-config tftp://10.10.1.1/home/configs/switch3-run.cfg
switch# copy bootflash:system_image bootflash://sup-2/system_image
switch# copy system:running-config bootflash:my-config
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
switch# copy sftp://172.16.10.100/myscript.txt volatile:myscript.txt
switch# copy system:running-config bootflash:my-config
switch# copy bootflash:samplefile bootflash:mystorage/samplefile

switch# copy samplefile mystorage/samplefile
switch# copy tftp://10.10.1.1/home/configs/switch3-run.cfg system:running-config

```

Creating a Directory

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# mkdir <i>directory name</i></pre> <ul style="list-style-type: none"> • mkdir {bootflash: debug: volatile:} Specifies the directory name you choose: <ul style="list-style-type: none"> ◦ bootflash: ◦ debug: ◦ volatile: • switch# mkdir bootflash:directory name Creates a directory that you name in the bootflash: directory. 	Creates a directory at the current directory level.

```
switch# mkdir test
switch# mkdir bootflash:test
```

Removing an Existing Directory

This command is valid only on Flash file systems.

Before You Begin

Before beginning this procedure, be sure of the following:

- You are logged in to the CLI.
- The directory you want to remove is empty.

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# rmdir [filesystem:[//module/]]directory</pre> <ul style="list-style-type: none"> • <code>switch# rmdir directory</code> Removes the specified directory at the current directory level. • <code>switch# rmdir {bootflash: debug: volatile:} directory</code> Removes a directory from the file system. 	<p>Removes a directory.</p> <p>The directory name is case sensitive.</p>

```
switch# rmdir test
switch# rmdir bootflash:test
```

Moving Files



Caution

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

The move will not complete if there is not enough space in the destination directory.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# move {source path and filename} {destination path and filename}</pre> <ul style="list-style-type: none"> switch# move filename path/filename Moves the file from one directory to another in the current file system. 	Moves the file from one directory to another in the same file system (bootflash:).

```
switch# move bootflash:samplefile bootflash:mystorage/samplefile
switch# move samplefile mystorage/samplefile
```

Deleting Files or Directories

You can delete files or directories on a Flash Memory device.

**Caution**

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

Before You Begin

You must understand the following information:

- 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.

Procedure

	Command or Action	Purpose
Step 1	<pre>switch# delete [bootflash: debug: log: volatile:] filename or directory name</pre> <ul style="list-style-type: none"> switch# delete filename Deletes the named file from the current working directory. switch# delete bootflash:directory name Deletes the named directory and its contents. 	Deletes a specified file or directory.

```
switch# delete bootflash:dns_config.cfg
switch# delete dns_config.cfg
```

Compressing Files

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show command > [path] filename	Directs the show command output to a file.
Step 2	switch# dir	Displays the contents of the current directory, including the new file created in the first step.
Step 3	switch# gzip [path] filename	Compresses the specified file
Step 4	switch# 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.

```
switch# show system internal l2fm event-history errors >errorsfile
switch# 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
switch# gzip bootflash:errorsfile
switch# 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
```

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

Uncompressing Files

You can uncompress (unzip) a specified file that is compressed using LZ77 coding.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# gunzip [<i>path</i>] <i>filename</i>	Uncompresses the specified file. The filename is case sensitive .
Step 2	switch# dir	Displays the contents of a directory, including the newly uncompresssed file.

```
switch# gunzip bootflash:errorsfile.gz
switch# 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#
```

Directing Command Output to a File

Procedure

	Command or Action	Purpose
Step 1	switch# show running-config > [<i>path</i> <i>filename</i>] <ul style="list-style-type: none"> • switch# show running-config > volatile:<i>filename</i> Directs the output of the command, show running-config, to the specified filename on the volatile file system. • switch# show running-config > bootflash:<i>filename</i> Directs the output of the command, show running-config, to the specified file in bootflash. • switch# show running-config > tftp:// <i>ipaddress</i>/<i>filename</i> 	Directs the output of the command, show running-config , to a path and filename.

	Command or Action	Purpose
	Directs the output of the command, show running-config , to the specified file on a TFTP server. • switch# show interface > filename Directs the output of the command, show interface , to the specified file at the same directory level, for example, in bootflash.	

```
switch# show running-config > volatile:switch1-run.cfg
switch# show running-config > bootflash:switch2-run.cfg
switch# show running-config > tftp://10.10.1.1/home/configs/switch3-run.cfg
switch# show interface > samplefile
```

Verifying a Configuration File before Loading

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

Procedure

	Command or Action	Purpose
Step 1	switch# copy source path and file 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.
Step 2	switch# show version image [bootflash: modflash: volatile:]	Validates the specified image. bootflash:—specifies bootflash as the directory name. volatile:—Specifies volatile as the directory name. modflash:—Specifies modflash as the directory name.

```
switch# copy tftp://10.10.1.1/home/configs/switch3-run.cfg system:running-config
switch# 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]
```

Rolling Back to a Previous Configuration

You can recover your configuration from a previously saved version.

**Note**

Each time you use a **copy running-config startup-config** command, 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.

Procedure

	Command or Action	Purpose
Step 1	switch# copy running-config bootflash: <i>{filename}</i>	Reverts to a snapshot copy of a previously saved running configuration (binary file).
Step 2	switch# copy bootflash: <i>{filename}</i> startup-config	Reverts to a configuration copy that was previously saved in the bootflash: file system (ASCII file).

```
switch# copy running-config bootflash:June03-Running
switch# copy bootflash:my-config startup-config
```

Displaying Files

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

- Displaying File Contents
- Displaying Directory Contents
- Displaying File Checksums
- Displaying the Last Lines in a File

Displaying File Contents

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# show file [bootflash: debug: volatile:] <i>filename</i>	Displays the contents of the specified file.

```
switch# show file bootflash:sample_test.txt
config t
Int veth1/1
```

```
no shut
end
show int veth1/1

switch#
```

Displaying Directory Contents

You can display the contents of a directory or file system.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# pwd	Displays the present working directory.
Step 2	switch# dir	Displays the contents of the directory.

```
switch# pwd
bootflash:
switch# dir

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

Displaying File Checksums

You can display checksums for checking file integrity.

Procedure

	Command or Action	Purpose
Step 1	switch# show file <i>filename</i> [cksum md5sum] show file { bootflash: volatile: debug: } <i>filename</i> [cksum md5sum]	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.

```
switch# show file bootflash:cisco_svs_certificate.pem cksum
266988670

switch# show file bootflash:cisco_svs_certificate.pem md5sum
d3013f73aea3fda329f7ea5851ae81ff
```

Displaying the Last Lines in a File

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# tail <i>{path}[filename] {Number of lines}</i>	Displays the requested number of lines from the end of the specified file. The range for the number of lines is from 0 to 80.

```
switch# 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
```

Feature History for File Management

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



CHAPTER 8

Managing users

This chapter contains the following sections:

- [Information About User Management, page 69](#)
- [Displaying Current User Access, page 69](#)
- [Sending a Message to Users, page 70](#)
- [Feature History for User Management, page 70](#)

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*.

Displaying Current User Access

You can display all users currently accessing the switch.

•

Before You Begin

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

•

Procedure

	Command or Action	Purpose
Step 1	switch# show users	Displays a list of users who are currently accessing the system.

```

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

```

Sending a Message to Users

You can send a message to all active CLI users currently using the system.

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# send <i>{session device}</i> <i>line</i>	Sends a message to users currently logged in to the system. <ul style="list-style-type: none"> The <i>session</i> argument sends the message to a specified pts/tty device type. The <i>device</i> argument specifies the device type. The <i>line</i> argument is a message of up to 80 alphanumeric characters in length.

```

switch# 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.

switch#

```

Feature History for User Management

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



CHAPTER 9

Configuring NTP

This chapter contains the following sections:

- [Information about NTP, page 71](#)
- [Prerequisites for NTP, page 72](#)
- [Guidelines and Limitations for NTP, page 73](#)
- [Default Settings for NTP, page 73](#)
- [Configuring an NTP Server and Peer, page 73](#)
- [Verifying the NTP Configuration, page 74](#)
- [NTP Example Configuration, page 74](#)
- [Feature History for NTP, page 75](#)

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.

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 by 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.

The following diagram shows a network with two NTP stratum 2 servers and two switches.

Figure 2: 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

You must have connectivity to at least one server that is running NTP.

Guidelines and Limitations for NTP

- 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 a 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 a 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 for NTP

Parameter	Default
NTP	Enabled

Configuring an NTP Server and Peer

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

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# ntp server { <i>ip-address</i> <i>dns-name</i> }	Forms an association with a server.
Step 3	switch(config)# ntp peer { <i>ip-address</i> <i>dns-name</i> }	Forms an association with a peer. You can specify multiple peer associations.
Step 4	switch(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.

	Command or Action	Purpose
Step 5	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
switch(config)# ntp server 192.0.2.10
switch(config)# ntp peer 2001:0db8::4101
```

Clearing NTP Sessions

Command	Purpose
clear ntp session	Clears the NTP sessions.

Clearing NTP Statistics

Command	Purpose
clear ntp statistics	Clears the NTP sessions.

Verifying the NTP Configuration

Use one of the following commands to verify the configuration:

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

NTP Example Configuration

This example configures an NTP server:

Procedure

- Step 1** `switch# configure terminal`
Enters global configuration mode.
- Step 2** `ntp server 192.0.2.10`
Configures an NTP server.
-

Feature History for NTP

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



CHAPTER 10

Configuring Local SPAN and ERSPAN

This chapter contains the following sections:

- [Information About SPAN and ERSPAN, page 77](#)
- [Guidelines and Limitations for SPAN, page 81](#)
- [Default Settings for SPAN, page 82](#)
- [Configuring SPAN, page 82](#)
- [Verifying the SPAN Configuration, page 95](#)
- [Configuration Example for an ERSPAN Session, page 95](#)
- [Feature History for SPAN and ERSPAN, page 97](#)

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 allows you to 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.

SPAN Sources

The interfaces from which traffic can be monitored are called SPAN sources. These sources 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 that 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 as described by the following:

- **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 Virtual Ethernet Module (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.

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)
- 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.

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.

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 a port profile. The destination can be a port or port profile.

The diagram 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 3: Local SPAN




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 the following diagram, 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

decapsulated and forwarded to the attached network analyzer. The destination may also be on the same Layer 2 network as the source.

Figure 4: ERSPAN Example



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.

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.

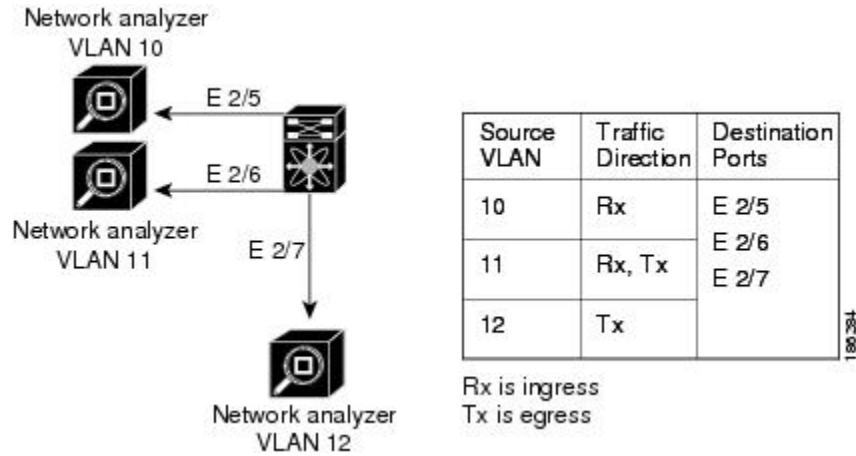
When trunk ports are configured as SPAN sources and destinations, you can filter VLANs to send to the destination ports from among those allowed. Both sources and destinations must be configured to allow the VLANs.

The following diagram 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 the diagram, 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 5: VLAN-based SPAN Configuration Example



Guidelines and Limitations for SPAN

- A maximum of 64 SPAN sessions (Local SPAN plus ERSPAN) can be configured on the Virtual Supervisor Module (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.



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 four 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 for SPAN

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
- Configuring an ERSPAN Port Profile
- Configuring an ERSPAN Session
- Shutting Down a SPAN Session
- Resuming a SPAN Session
- Verifying the SPAN Configuration

Configuring a Local SPAN Session

This procedure involves creating the SPAN session in monitor configuration mode, and then, optionally, configuring allowed VLANs in interface configuration mode.

It is important to know the following information about SPAN:

- 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. This procedure includes how to do this.
- 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*.

Before You Begin

Before beginning this procedure, you must be logged in to the CLI in EXEC mode and be sure you know the number of the SPAN session you are going to configure.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# no monitor session <i>session-number</i>	Clears the specified session.
Step 3	switch(config)# monitor session <i>session-number</i>	Creates a session with the given session number and places you in monitor configuration mode to further configure the session.
Step 4	switch(config-monitor)# description <i>description</i>	Adds a description for the specified SPAN session. The <i>description</i> can be up to 32 alphanumeric characters. The default is blank (no description)
Step 5	switch(config-monitor)# source {interface {type} {id} vlan {id range} port-profile {name}} [rx tx both]	For the specified session, configures the sources and the direction of traffic to monitor. <ul style="list-style-type: none"> For the <i>type</i> argument, specify the interface type—Ethernet or vEthernet. For the <i>id</i> argument, specify the vEthernet number, the Ethernet slot/port, or the VLAN ID to monitor. For the <i>range</i> argument, specify the VLAN range to monitor. For the <i>name</i> argument, 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 For the traffic direction keywords, specify as follows: <ul style="list-style-type: none"> rx which is the VLAN default indicates receive. tx indicates transmit. both is the default keyword
Step 6	Repeat Step 5 to configure additional SPAN sources.	(Optional)

	Command or Action	Purpose
Step 7	switch(config-monitor)# filter vlan { <i>id</i> <i>range</i> }	(Optional) For the specified SPAN session, configures the filter from among the source VLANs.
Step 8	Repeat Step 7 to configure all source VLANs to filter.	(Optional)
Step 9	switch(config-monitor)# destination { interface { <i>type</i> } { <i>id</i> <i>range</i> } port-profile { <i>name</i> }}	For the specified SPAN session, configures the destination(s) for copied source packets. <ul style="list-style-type: none"> • For the <i>type</i> argument, specify the interface type—Ethernet or vEthernet. • For the <i>id</i> argument, specify the vEthernet number or the Ethernet slot/port to monitor. • For the <i>name</i> argument specify the name of the port profile to monitor.
Step 10	Repeat Step 9 to configure all SPAN destination ports.	(Optional)
Step 11	switch(config-monitor)# no shut	Enables the SPAN session. By default, the session is created in the shut state.
Step 12	switch(config-monitor)# exit	(Optional) Exits monitor configuration mode and places you in interface configuration mode.
Step 13	switch(config-if)# show monitor session <i>session-number</i>	(Optional) Displays the configured monitor session.
Step 14	switch(config-if)# show interface { <i>type</i> } { <i>id</i> } switchport	Displays the configured port including allowed VLANs. <ul style="list-style-type: none"> • For the <i>type</i> argument, specify the interface type—Ethernet or vEthernet. • For the <i>id</i> argument, specify the vEthernet number or the Ethernet slot/port to monitor.
Step 15	switch(config-if)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```

switch# configure terminal
switch(config)# no monitor session 3
switch(config)# monitor session 3
switch(config-monitor)# description my_span_session_3
switch(config-monitor)# source interface ethernet 2/1-3, ethernet 3/1 rx
switch(config-monitor)# filter vlan 3-5, 7
switch(config-monitor)# destination interface ethernet 2/5, ethernet 3/7
switch(config-monitor)# no shut

```

```
switch(config-monitor)# exit
switch(config-if)# show monitor session 3
switch(config-if)# show interface ethernet 2/5 switchport
switch(config-if)# copy running-config startup-config
```

Configuring an ERSPAN Port Profile

You can configure a port profile on the VSM to carry ERSPAN packets through the IP network to a remote destination analyzer.

You must complete this configuration for all hosts in vCenter Server.

This procedure includes steps to configure the port profile for the following requirements:

- ERSPAN for Layer 3 control.
- An access port profile. It cannot be a trunk port profile.

Only one VMKNIC can be assigned to this Layer 3 control port profile per host as follows:

- If more than one VMKNIC is assigned to a host, the first one assigned takes effect. The second one is not considered a Layer 3 control VMKNIC.
- If more than one VMKNIC is assigned to a host, and you remove the second assigned one, the VEM does not use the first assigned one. Instead, you must remove both VMKNICs and then add one back.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode
- Established the name to be used for this port profile



Note The port profile name is used to configure the VM Kernal NIC (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.

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

For more information about system port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide*.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# port-profile <i>port_profile_name</i>	Creates the port profile and places you in global configuration mode for the specified port profile. This command saves the port profile in the running configuration. 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	switch(config-prot-prof)# capability l3control	Configures the port profile to carry ERSPAN traffic and saves the port profile in the running configuration.
Step 4	switch(config-prot-prof)# vmware port-group <i>name</i>	Designates the port profile as a VMware port group and adds the name of the VMware port group to which this profile maps. This command saves the settings in the running configuration. 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. The <i>name</i> argument is the same as the port profile name if you do not specify a port group 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	switch(config-prot-prof)# switchport mode access	Designates the interfaces as switch access ports (the default).
Step 6	switch(config-prot-prof)# switchport access <i>vlan id</i>	Assigns a VLAN ID to the access port for this port profile and saves the setting in the running configuration. This VLAN is used to send IP traffic to the ERSPAN destination.
Step 7	switch(config-prot-prof)# no shutdown	Enables the interface in the running configuration.
Step 8	switch(config-prot-prof)# system vlan <i>id</i>	Associates the system VLAN ID with the port profile and saves it in the running configuration. The ID must match the VLAN ID that is assigned to the access port. If it does not match, then the following error message is generated: ERROR: System vlan being set does not match the switchport access vlan 2
Step 9	switch(config-prot-prof)# state enabled	Enables the port profile in the running configuration. This port profile is now ready to send out ERSPAN packets on all ESX hosts with ERSPAN sources.

	Command or Action	Purpose
Step 10	switch(config-prot-prof)# show port-profile name <i>port_profile_name</i>	(Optional) Displays the configuration for the specified port profile as it exists in the running configuration.
Step 11	switch(config-port-prof)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.
Step 12	Using the VMware documentation, go to vSphere Client and configure a VMKNIC on each ESX Host for sending ERSPAN-encapsulated packets. Make sure that the VMKNIC points to this port profile as a new virtual adapter. This VMKNIC must have IP connectivity to the ERSPAN destination IP address.	

```

switch# configure terminal
switch(config)# port-profile erspan_profile
switch(config-port-prof)# capability l3control
switch(config-port-prof)# vmware port-group erspan
switch(config-port-prof)# switchport mode access
switch(config-port-prof)# switchport access vlan 2
switch(config-port-prof)# no shutdown
switch(config-port-prof)# system vlan 2
switch(config-port-prof)# state enabled
switch(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:
n1000v(config-port-prof)# copy running-config startup-config

```

Configuring an ERSPAN Session

This procedure involves creating the SPAN session in ERSPAN source configuration mode (config-erspan-source).

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. The step to do this is included in the procedure.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode
- Obtained the number of the SPAN session that you are going to configure
- Configured an ERSPAN-capable port profile on the VSM
- Using the VMware documentation for adding a new virtual adapter, you have already configured the required VMKNIC on each ESX host. The VMKNIC must have IP connectivity to the ERSPAN destination IP address for sending ERSPAN-encapsulated packets.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# no monitor session <i>session-number</i>	Clears the specified session.
Step 3	switch(config)# monitor session <i>session-number</i> type erspan-source	Creates a session with the given session number and places you in ERSPAN source configuration mode. This configuration is saved in the running configuration.
Step 4	switch(config-erspan-src)# description <i>description</i>	For the specified ERSPAN session, adds a description and saves it in the running configuration. The <i>description</i> can be up to 32 alphanumeric characters The default is blank (no description)
Step 5	switch(config-erspan-src)# source {interface <i>type</i> { <i>number</i> <i>range</i> } vlan { <i>number</i> <i>range</i> } port-profile { <i>name</i> } } [rx tx both]	For the specified session, configures the sources and the direction of traffic to monitor and saves them in the running configuration. <ul style="list-style-type: none"> • For the <i>type</i> argument, specify the interface type—ethernet, port-channel, vethernet. • For the <i>number</i> argument, specify the interface slot/port or range; or the VLAN number or range to monitor. • For the <i>name</i> argument, specify the name of the existing port profile. • For the traffic direction keywords, specify as follows: <ul style="list-style-type: none"> ◦ rx which is the VLAN default indicates receive. ◦ tx indicates transmit. ◦ both is the default keyword

	Command or Action	Purpose
Step 6	Repeat Step 5 to configure additional ERSPAN sources.	(Optional)
Step 7	switch(config-erspan-src)# filter vlan {number range}	(Optional) For the specified ERSPAN session, configures the VLANs, VLAN lists, or VLAN ranges to be monitored; and saves the VLAN arguments to the running configuration. On the monitor port, only the traffic from the VLANs that match the VLAN filter list are replicated to the destination.
Step 8	Repeat Step 7 to configure all source VLANs to filter.	(Optional)
Step 9	switch(config-erspan-src)# destination ip ip_address	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	switch(config-erspan-src)# ip ttl ttl_value	(Optional) Specifies the IP time-to-live value, from 1 to 255, for ERSPAN packets in this monitor session and saves it in the running configuration.
Step 11	switch(config-erspan-src)# ip prec precedence_value	(Optional) Specifies the IP precedence value, from 0 to 7, for the ERSPAN packets in this monitor session and saves it in the running configuration. The default value is 0.
Step 12	switch(config-erspan-src)# ip dscp dscp_value	(Optional) Specifies the IP DSCP value, from 0 to 63. for the ERSPAN packets in this monitor session and saves it in the running configuration. The default is 0.
Step 13	switch(config-erspan-src)# mtu mtu_value	(Optional) Specifies an MTU size (from 50 to 1500) for ERSPAN 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. The default is 1500. 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.
Step 14	switch(config-erspan-src)# header-type value	Specifies the ERSPAN header type (2 or 3) used for ERSPAN encapsulation for this monitor session as follows: <ul style="list-style-type: none"> • 2 is the ERSPANv2 header type (the default)

	Command or Action	Purpose
		<ul style="list-style-type: none"> • 3 is the ERSPANv3 header type (Used with NAM setups. Any other type of destination works only with the default v2 headers.)
Step 15	switch(config-erspan-src)# erspan-id <i>flow_id</i>	<p>Adds an ERSPAN ID from 1 to 1023 to the session configuration and saves it in the running configuration.</p> <p>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.</p>
Step 16	switch(config-erspan-src)# no shut	<p>Enables the ERSPAN session and saves it in the running configuration.</p> <p>By default, the session is created in the shut state.</p>
Step 17	switch(config-erspan-src)# show monitor session <i>session_id</i>	<p>(Optional)</p> <p>Displays the ERSPAN session configuration as it exists in the running configuration</p>
Step 18	switch(config-erspan-src)# copy running-config startup-config	<p>(Optional)</p> <p>Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.</p>

```

switch# configure terminal
switch(config)# no monitor session 3
switch(config)# monitor session 3 type erspan
switch(config-erspan-src)# description my_erspan_session_3
switch(config-erspan-src)# source interface ethernet 2/1-3, ethernet 3/1 rx
switch(config-erspan-src)# filter vlan 3-5, 7
switch(config-erspan-src)# destination ip 10.54.54.1
switch(config-erspan-src)# ip ttl 64
switch(config-erspan-src)# ip prec 1
switch(config-erspan-src)# ip dscp 24
switch(config-erspan-src)# mtu 1000
switch(config-erspan-src)# header-type 2
switch(config-erspan-src)# erspan-id 51
switch(config-erspan-src)# no shut
switch(config-erspan-src)# show monitor session 3
switch(config-erspan-src)# copy running-config startup-config

```

Shutting Down a SPAN Session from Global Configuration Mode

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode.
- Determined which session you want to shutdown

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# monitor session { <i>session-number</i> <i>session-range</i> all } shut	Shuts down the specified SPAN monitor session(s) from global configuration mode. <ul style="list-style-type: none"> • The <i>session-number</i> argument specifies a particular SPAN session number. • The <i>session-range</i> argument specifies a range of SPAN sessions from 1 to 64. • The all keyword specifies all SPAN monitor sessions.
Step 3	switch(config)# show monitor	(Optional) Displays the status of the SPAN sessions.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
switch(config)# monitor session 3 shut
switch(config)# show monitor
switch(config)# copy running-config startup-config
```

Shutting Down a SPAN Session from Monitor Configuration Mode

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode.
- Determined which session you want to shutdown

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# monitor session { <i>session-number</i> <i>session-range</i> all } [type erspan-source]	Specifies the SPAN monitor session(s) you want to shut down from monitor-configuration mode. <ul style="list-style-type: none"> • The <i>session-number</i> argument specifies a particular SPAN session number.

	Command or Action	Purpose
		<ul style="list-style-type: none"> The <i>session-range</i> argument specifies a range of SPAN sessions from 1 to 64. The all keyword specifies all SPAN monitor sessions.
Step 3	switch(config)# shut	Shuts down the specified SPAN monitor session(s) from monitor configuration mode.
Step 4	switch(config-monitor)# show monitor	(Optional) Displays the status of the SPAN sessions.
Step 5	switch(config-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
switch(config)# monitor session 3
switch(config-monitor)# shut
switch(config-monitor)# show monitor
switch(config-monitor)# copy running-config startup-config
```

Resuming a SPAN Session from Global Configuration Mode

You can discontinue copying packets from one source and destination and then resume from another source and destination in global configuration mode.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode.
- Determined which SPAN session that you want to configure.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no]monitor session {session-number session-range all} shut	Shuts down the specified SPAN monitor session(s) from global configuration mode. <ul style="list-style-type: none"> The <i>session-number</i> argument specifies a particular SPAN session number. The <i>session-range</i> argument specifies a range of SPAN sessions from 1 to 64.

	Command or Action	Purpose
		<ul style="list-style-type: none"> The all keyword specifies all SPAN monitor sessions.
Step 3	switch(config)# show monitor	(Optional) Displays the status of the SPAN sessions.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
switch(config)# no monitor session 3 shut
switch(config)# show monitor
switch(config)# copy running-config startup-config
```

Resuming a SPAN Session from Monitor Configuration Mode

You can discontinue copying packets from one source and destination and then resume from another source and destination in monitor configuration mode.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode.
- Determined which SPAN session that you want to configure.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] monitor session { <i>session-number</i> <i>session-range</i> all } shut	Shuts down the specified SPAN monitor session(s) from monitor configuration mode. <ul style="list-style-type: none"> The <i>session-number</i> argument specifies a particular SPAN session number. The <i>session-range</i> argument specifies a range of SPAN sessions from 1 to 64. The all keyword specifies all SPAN monitor sessions.
Step 3	switch(config-monitor)# show monitor	(Optional) Displays the status of the SPAN sessions.

	Command or Action	Purpose
Step 4	switch(config-monitor)# show monitor session <i>session-id</i>	(Optional) Displays detailed configuration and status of a specific SPAN session for verification.
Step 5	switch(config-monitor)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
switch(config)# monitor session 3
switch(config-monitor)# no shut
switch(config-monitor)# show monitor
switch(config-monitor)# show monitor session 3
switch(config-monitor)# copy running-config startup-config
```

Configuring the Allowable ERSPAN Flow IDs

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

The available ERSPAN flow IDs are from 1 to 1023.

Before You Begin

Before beginning this procedure, be sure you have done the following:

- Logged in to the CLI in EXEC mode.
- Determined the restricted range of ERSPAN flow IDs that you want to designate.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] limit-resource erspan-flow-id minimum <i>min_val</i> maximum <i>max_val</i>	Restricts the allowable range of ERSPAN flow IDs that can be assigned. The allowable range is from 1 to 1023. The defaults are as follows: The minimum value = 1 The maximum value = 1023 The no form of this command removes any configured values and restores default values.

	Command or Action	Purpose
Step 3	switch(config)# show running monitor	(Optional) Displays changes to the default limit-resource erspan-flow-id values for verification
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```
switch# configure terminal
switch(config)# limit-resource erspan-flow-id minimum 20 maximum 40
switch(config)# show monitor
switch(config)# show running monitor
switch(config)# copy running-config startup-config
```

Verifying the SPAN Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show monitor session {all <i>session-number</i> range <i>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.

Configuration Example 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.CSCtn56340 Packets arriving at the destination IP are identified by the ID 999 in their header.

```
switch# monitor session 2 type erspan-source
switch(config-erspan-src)# source interface ethernet 3/3
switch(config-erspan-src)# source port-profile my_profile_src
switch(config-erspan-src)# destination ip 10.54.54.1
switch(config-erspan-src)# erspan-id 999
switch(config-erspan-src)# mtu 1000
switch(config-erspan-src)# no shut

switch(config-erspan-src)# 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
ERSPAN TTL           : 64
ERSPAN IP Prec.     : 0
ERSPAN DSCP          : 0
ERSPAN MTU           : 1000
ERSPAN Header Type  : 2

switch(config-erspan-src)# 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

```

Example of Configuring a SPAN Session

```

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

switch(config)# show monitor session 1
  session 1
  -----
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

switch# 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 of a Configuration to Enable SPAN Monitoring

This example shows how to configure destination ports in access or trunk mode, and enable SPAN monitoring.

```
switch# configure terminal
switch(config)# interface ethernet 2/5
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# no shut
switch(config-if)# exit
switch(config)#
```

Feature History for SPAN and ERSPAN

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.



CHAPTER 11

Configuring SNMP

This chapter contains the following sections:

- [Information About SNMP, page 99](#)
- [Guidelines and Limitations for SNMP, page 103](#)
- [Default Settings for SNMP, page 103](#)
- [Configuring SNMP, page 103](#)
- [Verifying the SNMP Configuration, page 110](#)
- [Configuration Example for SNMP, page 111](#)
- [Related Documents for SNMP, page 111](#)
- [MIBs, page 112](#)
- [Feature History for SNMP, page 113](#)

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.

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.
- 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. The 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 the 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 [Configuring SNMP Notification Receivers](#) 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.

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.

The following table identifies what the combinations of security models and levels mean.

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.
- 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

The 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, on page 110](#) section for information on how to modify this default value.

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 for SNMP

- 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

Default Settings for SNMP

Parameters	Default
license notifications	enabled

Configuring SNMP

This section includes the following topics:

- Configuring SNMP
- Users Enforcing SNMP Message Encryption
- Creating SNMP Communities
- Configuring SNMP Notification Receivers
- Configuring the Notification Target User
- Enabling SNMP Notifications
- Disabling LinkUp/LinkDown Notifications on an Interface
- Enabling a One-time Authentication for SNMP over TCP
- Assigning the SNMP Switch Contact and Location Information

- Disabling SNMP
- Modifying the AAA Synchronization Time

Configuring SNMP Users

Before You Begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# snmp-server user name [auth {md5 sha} passphrase [auto] [priv [aes-128] passphrase] [engineID id] [localizedkey]]	<p>Configures an SNMP user with authentication and privacy parameters. The <i>passphrase</i> can be any case-sensitive, alphanumeric string up to 64 characters. If you use the localizedkey keyword, the <i>passphrase</i> can be any case-sensitive, alphanumeric string up to 130 characters.</p> <p>The <i>name</i> argument is the name of a user who can access the SNMP engine.</p> <p>The auth keyword enables one-time authentication for SNMP over a TCP session. It is optional.</p> <p>The md5 keyword specifies HMAC MD5 algorithm for authentication. It is optional.</p> <p>The sha keyword specifies HMAC SHA algorithm for authentication. It is optional.</p> <p>The priv keyword specifies encryption parameters for the user. It is optional.</p> <p>The aes-128 keyword specifies a 128-byte AES algorithm for privacy. It is optional.</p> <p>The engineID keyword specifies the engineID for configuring the notification target user (for V3 informs). It is optional.</p> <p>The <i>id</i> is a 12-digit colon-separated decimal number.</p>
Step 3	switch(config-callhome)# show snmp user	(Optional) Displays information about one or more SNMP users.
Step 4	switch(config-callhome)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

```
switch# configure terminal
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 for All Users

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server globalEnforcePriv	Enforces SNMP message encryption for all users.

```
switch(config)# snmp-server globalEnforcePriv
```

Creating SNMP Communities

You can create SNMP communities for SNMPv1 or SNMPv2c.

Before You Begin

You must be in global configuration mode.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server community <i>name</i> {ro rw}	Creates an SNMP community string.

```
switch(config)# snmp-server community public ro
```

Configuring SNMP Notification Receivers

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

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.

The 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

Before You Begin

You must be in global configuration mode.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server user <i>name</i> [auth { md5 sha } <i>passphrase</i> [auto] [priv [aes-128] <i>passphrase</i>] [engineID <i>id</i>]	Configures the notification target user with the specified engine ID for notification host receiver. The <i>id</i> is a 12-digit colon-separated decimal number.
	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
```

Enabling SNMP Notifications

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

The following table 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.

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

The license notifications are enabled by default. All other notifications are disabled by default.

Before You Begin

You must be in global configuration mode to enable the specified notification

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server enable traps	Enables all SNMP notifications.
Step 2	switch(config)# snmp-server enable traps aaa [server-state-change]	Enables the AAA SNMP notifications.
Step 3	switch(config)# snmp-server enable traps entity [fru]	Enables the ENTITY-MIB SNMP notifications.
Step 4	switch(config)# snmp-server enable traps license	Enables the license SNMP notification.
Step 5	switch(config)# snmp-server enable traps link	Enables the link SNMP notifications.
Step 6	switch(config)# snmp-server enable traps port-security	Enables the port security SNMP notifications
Step 7	switch(config)# snmp-server enable traps snmp [authentication]	Enables the SNMP agent notifications.

```
switch(config)# snmp-server enable traps

switch(config)# snmp-server enable traps aaa
switch(config)# snmp-server enable traps entity
switch(config)# snmp-server enable traps license
switch(config)# snmp-server enable traps port-security
switch(config)# snmp-server enable traps link
switch(config)# snmp-server enable traps snmp
```

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).

Before You Begin

You must be in interface configuration mode to disable linkUp/linkDown notifications for the interface.

•

Procedure

	Command or Action	Purpose
Step 1	switch(config-if)# no snmp trap link-status	Disables SNMP link-state traps for the interface. This command is enabled by default.

```
switch(config-if)# no snmp trap link-status
```

Enabling a One-time Authentication for SNMP over TCP

Before You Begin

You must be in global configuration mode to enable one-time authentication for SNMP over TCP

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server tcp-session [auth]	Enables a one-time authentication for SNMP over a TCP session. The default is disabled.

```
switch(config)# snmp-server tcp-session
```

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

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# snmp-server contact <i>name</i>	Configures sysContact, which is the SNMP contact name.
Step 3	switch(config)# snmp-server location <i>name</i>	Configures sysLocation, which is the SNMP location.

	Command or Action	Purpose
Step 4	switch(config)# show snmp	(Optional) Displays information about one or more destination profiles.
Step 5	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# snmp contact Admin
switch(config)# snmp location Lab-7
switch(config)# show snmp
switch(config)# copy running-config startup-config
```

Configuring a Host Receiver for SNMPv3 Traps or Informs



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

Before You Begin

You must be in global configuration mode.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server host ip-address {traps informs} version 3 {auth noauth priv} username [udp_port number]	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.

```
switch(config)# snmp-server host 192.0.2.1 informs version 3 auth NMS
```

Disabling SNMP

Before You Begin

You must be in global configuration mode to disable the SNMP protocol on a device.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# no snmp-server protocol enable	Disables the SNMP protocol. This command is enabled by default.

```
switch(config)# no snmp-server protocol enable
```

Modifying the AAA Synchronization Time

You can modify how long Cisco NX-OS holds the synchronized user configuration.

.

Before You Begin

You must be in global configuration mode.

Procedure

	Command or Action	Purpose
Step 1	switch(config)# snmp-server aaa-user cache-timeout seconds	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.

```
switch(config)# snmp-server aaa-user cache-timeout 1200
```

Verifying the SNMP Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show running-config snmp [all]	Displays the SNMP running configuration.
show snmp	Displays the SNMP status.
show snmp community	Displays the SNMP community strings.
show snmp context	Displays the SNMP context mapping.
show snmp engineID	Displays the SNMP engineID.
show snmp group	Displays SNMP roles.

Command	Purpose
show snmp session	Displays SNMP sessions.
show snmp trap	Displays the SNMP notifications enabled or disabled.
show snmp user	Displays SNMPv3 users.

Configuration Example for SNMP

This example shows how to configure sending the Cisco linkUp/Down notifications to one notification host receiver using the Blue VRF and define two SNMP users, Admin and NMS

```
switch# configure terminal
switch(config)# snmp-server contact Admin@company.com
switch(config)# snmp-server user Admin auth sha abcd1234 priv abcdefgh
switch(config)# snmp-server user NMS auth sha abcd1234 priv abcdefgh engineID
00:00:00:63:00:01:00:22:32:15:10:03
switch(config)# snmp-server host 192.0.2.1 informs version 3 auth NMS
switch(config)# snmp-server host 192.0.2.1 use-vrf Blue
switch(config)# snmp-server enable traps link cisco
```

Related Documents for SNMP

Related Topic	Document Title
MIBs	http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

MIBs

<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:</p> <p>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 	
--	--

Feature History for SNMP

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



CHAPTER 12

Configuring NetFlow

This chapter contains the following sections:

- [Information about NetFlow, page 115](#)
- [Prerequisites for NetFlow, page 122](#)
- [Configuration Guidelines and Limitations for NetFlow, page 123](#)
- [Default Settings for NetFlow, page 123](#)
- [Enabling the NetFlow Feature, page 124](#)
- [Configuring Netflow, page 125](#)
- [Verifying the NetFlow Configuration, page 132](#)
- [Netflow Example Configuration, page 135](#)
- [Related Documents for NetFlow, page 136](#)
- [Feature History for NetFlow, page 136](#)

Information about NetFlow


NetFlow lets you evaluate IP traffic and understand how and where it flows. 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. NetFlow gathers data that can be used in accounting, network monitoring, and network planning.

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 6: NetFlow Cache Example



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.

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.

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
- Talled 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.

The following table describes the criteria defined in a flow record.

Table 2: Flow record criteria

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 Predefined Flow Record: Netflow-Original

```
switch# show flow record netflow-original
Flow record netflow-original:
  Description: Traditional IPv4 input NetFlow with origin ASs
  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#
```



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 effect on the configuration.

```
collect routing source as
collect routing destination as
collect routing next-hop address ipv4
```

Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Input

```
switch# 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
switch#
```

Cisco Nexus 1000V Predefined Flow Record: Netflow IPv4 Original-Output

```
switch# show flow record netflow ipv4 original-output
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#
```

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)
- NetFlow Collector

Command Line Interface for NetFlow

Use the Command Line Interface (CLI) to access NetFlow data, and to view what is happening in your network now.

The CLI uses the Flow Monitor and Flow Exporter to capture and export flow records to the Netflow Collector. Cisco Nexus 1000V supports the NetFlow Version 9 export format.

**Note**

Cisco Nexus 1000V supports UDP as the transport protocol for exporting data to up to two exporters per monitor.

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.

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.

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. 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 7: Exporting Flows to the NetFlow Collector Server

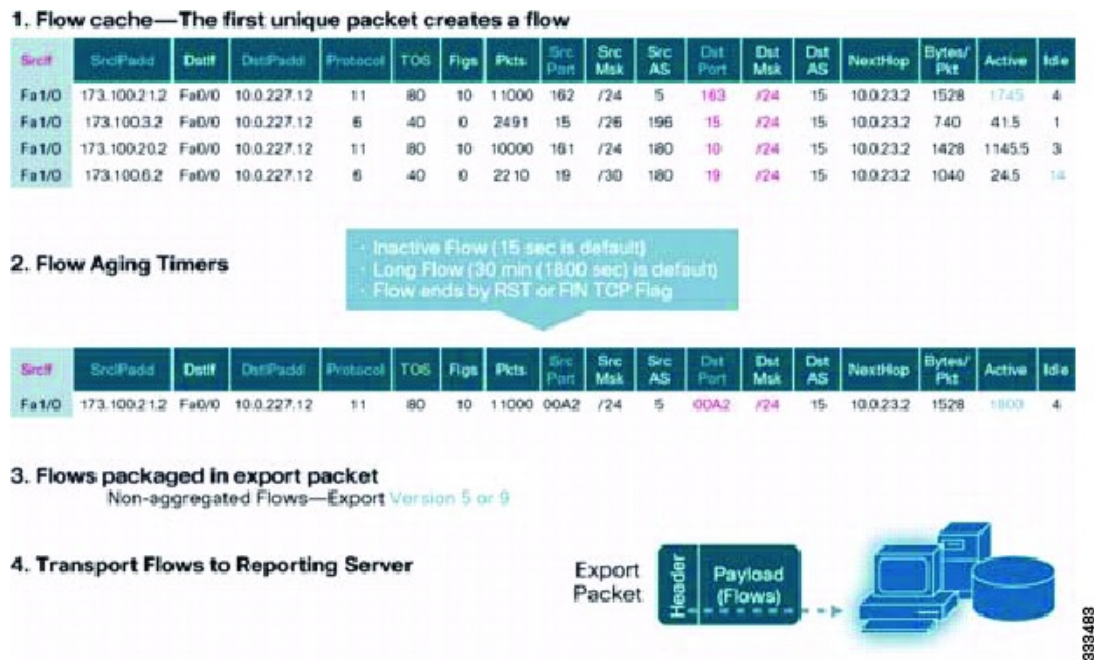


33-3482

What NetFlow Data Looks Like

The following figure shows an example of NetFlow data.

Figure 8: 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 Prime Network Analysis Module (NAM) for Nexus 1010 Installation and Configuration Guide*.

High Availability for NetFlow

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.

Configuration Guidelines and Limitations for NetFlow

- 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.

- 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 on port channels or interfaces in a port-channel.

Default Settings for NetFlow

Table 3: Default NetFlow Parameters

Parameters	Default
NetFlow version	9
source interface	mgmt0

Parameters	Default
match	direction and interface (incoming/outgoing)
flow monitor active timeout	1800
flow monitor inactive timeout	15
flow monitor cache size	4096
flow exporter UDP port transport udp command	9995
DSCP	default/best-effort (0)
VRF	default

Enabling the NetFlow Feature

Before You Begin

You are logged in to the CLI in EXEC mode.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# feature netflow	Enables the NetFlow feature.
Step 3	switch(config)# show feature	(Optional) Displays the available features and whether or not they are enabled.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to enable the NetFlow feature:

```
switch# configure terminal
switch(config)# feature netflow
switch(config)#
```

Configuring Netflow

Defining a Flow Record

Before You Begin

- You know which of the options you want this flow record to match.
- You know which options you want this flow record to collect.



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
```

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# flow record name	Creates a Flow Record by name, and places you in the CLI Flow Record Configuration mode for that specific record.
Step 3	switch(config)# description string	(Optional) Adds a description of up to 63 characters to the Flow Record and saves it to the running configuration.
Step 4	switch(config)# match {ip {protocol tos} ipv4 {destination address source address} transport {destination-port source-port}}	<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

	Command or Action	Purpose
Step 5	switch(config)# collect { counter { bytes [long] packets [long]} timestamp sys-uptime transport tcp flags }	<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	switch(config)# show flow record <i>name</i>	(Optional) Displays information about Flow Records.
Step 7	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to create a flow record:

```
switch# configure terminal
switch(config)# flow record RecordTest
switch(config-flow-record)# description Ipv4flow
switch(config-flow-record)# match ipv4 destination address
switch(config-flow-record)# collect counter packets
switch(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
switch(config-flow-record)#
```

Defining a Flow Exporter

A Flow Exporter defines where and how Flow Records are exported to the NetFlow Collector Server.

- Export format version 9 is supported.
- A maximum of two flow exporters per monitor are permitted.

Before You Begin

- You know the 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.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# flow exporter <i>name</i>	Creates a Flow Exporter, saves it in the running configuration, and then places you in CLI Flow Exporter Configuration mode.
Step 3	switch(config-flow-exporter)# description <i>string</i>	Adds a description of up to 63 characters to this Flow Exporter and saves it in the running configuration.
Step 4	switch(config-flow-exporter)# destination { <i>ipv4-address</i> <i>ipv6-address</i> }	Specifies the IP address of the destination interface for this Flow Exporter and saves it in the running configuration.
Step 5	switch(config-flow-exporter)# dscp <i>value</i>	Specifies the differentiated services codepoint value for this Flow Exporter, between 0 and 63, and saves it in the running configuration.
Step 6	switch(config-flow-exporter)# source <i>mgmt interface_number</i>	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	switch(config-flow-exporter)# transport udp <i>port-number</i>	Specifies the destination UDP port, between 0 and 65535, used to reach the NetFlow collector, and saves it in the running configuration.
Step 8	switch(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	switch(config-flow-exporter-version-9)# option { exporter-stats interface-table sampler-table } timeout <i>value</i>	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	switch(config-flow-exporter-version-9)# template data timeout <i>seconds</i>	Sets the template data resend timer and its value, between 1 and 86400 seconds, and saves it in the running configuration.

	Command or Action	Purpose
Step 11	switch(config-flow-exporter-version-9)# show flow exporter [name]	(Optional) Displays information about the Flow Exporter.
Step 12	switch(config-flow-exporter-version-9)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to create a flow exporter:

```
switch# configure terminal
switch(config)# flow exporter ExportTest
switch(config-flow-exporter)# description ExportHamilton
switch(config-flow-exporter)# destination 192.0.2.1
switch(config-flow-exporter)# dscp 2
switch(config-flow-exporter)# source mgmt 0
switch(config-flow-exporter)# transport udp 200
switch(config-flow-exporter)# version 9
switch(config-flow-exporter-version-9)# option exporter-stats timeout 1200
switch(config-flow-exporter-version-9)# template data timeout 1200
switch(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
switch(config-flow-exporter-version-9)# copy running-config startup-config
switch(config-flow-exporter-version-9)#
```

Defining a Flow Monitor

A Flow Monitor is associated with a Flow Record and a Flow Exporter.

A maximum of one flow monitor per interface per direction is permitted.

Before You Begin

- You know the name of an existing Flow Exporter to associate with this flow monitor.
- 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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# flow monitor <i>name</i>	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	switch(config-flow-monitor)# description <i>string</i>	(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	switch(config-flow-monitor)# exporter <i>name</i>	For the specified flow monitor, adds an existing flow exporter and saves it in the running configuration.
Step 5	switch(config-flow-monitor)# record { <i>name</i> netflow { ipv4 }}	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 ipv4: Traditional IPv4 NetFlow collection schemes
Step 6	switch(config-flow-monitor)# timeout { active <i>value</i> inactive <i>value</i> }	(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	switch(config-flow-monitor)# cache { size <i>value</i> }	(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. This option is used to limit the impact of the monitor cache on memory and performance.
Step 8	switch(config-flow-monitor)# show flow monitor [<i>name</i>]	(Optional) Displays information about existing flow monitors.

	Command or Action	Purpose
Step 9	switch(config-flow-monitor)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to create a flow exporter:

```
switch# configure terminal
switch(config)# flow monitor MonitorTest
switch(config-flow-monitor)# description Ipv4Monitor
switch(config-flow-monitor)# exporter ExportTest
switch(config-flow-monitor)# record RecordTest
switch(config-flow-monitor)# cache size 15000
switch(config-flow-monitor)# timeout inactive 600
switch(config-flow-monitor)# show flow monitor MonitorTest
Flow Monitor monitorTest:
  Use count: 0
  Inactive timeout: 600
  Active timeout: 1800
  Cache Size: 15000
switch(config-flow-monitor)#
```

Assigning 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.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type</i> <i>interface-number</i>	Places you in the CLI Interface Configuration mode for the specified interface.
Step 3	switch(config)# ip flow monitor <i>name</i> { input output }	For the specified interface, assigns a flow monitor for input or output packets and saves it in the running configuration.
Step 4	switch(config)# show flow <i>interface-type</i> <i>interface-number</i>	(Optional) For the specified interface, displays the NetFlow configuration.
Step 5	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to assign a flow monitor to an interface:

```
switch# configure terminal
switch(config)# interface veth 2
switch(config-if)# ip flow monitor MonitorTest output
switch(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
  Direction: Output
switch(config-if)#
```

Adding a Flow Monitor to a Port Profile

Before You Begin

- You are logged in to the CLI in EXEC mode.
- You have already created the flow monitor.
- 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.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# port-profile [type { ethernet vethernet }] <i>name</i>	Enters port profile configuration mode for the named port profile.
Step 3	switch(config-port-prof)# ip flow monitor <i>name</i> { input output }	Applies a named flow monitor to the port profile for either incoming (input) or outgoing (output) traffic.
Step 4	switch(config-port-prof)# show port-profile [brief expand-interface usage] [name <i>profile-name</i>]	(Optional) Displays the configuration for verification.
Step 5	switch(config-port-prof)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to add a flow monitor to a port profile:

```
switch# configure terminal
switch(config)# port-profile AccessProf
switch(config-port-prof)# ip flow monitor allaces4 output
switch(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:
switch(config-port-prof) #

```

Verifying the NetFlow Configuration

Use one of the following commands to verify the configuration:

Table 4: Verifying the NetFlow Configuration

Command	Purpose
show flow exporter [<i>name</i>]	Displays information about NetFlow flow exporter maps.
show flow interface [<i>interface-type number</i>]	Displays information about NetFlow interfaces.
show flow monitor [<i>name</i> [<i>cache module number</i> <i>statistics module number</i>]]	Displays information about NetFlow flow monitors. Note The show flow monitor cache module command differs from the show flow monitor statistics module 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.
show flow record [<i>name</i>]	Displays information about NetFlow flow records.

Example: show flow exporter

```

switch(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
switch(config-flow-exporter-version-9)#
    
```

Example: show flow interface

```

switch(config-if)# show flow interface VEth2
Interface veth2:
  Monitor: MonitorTest
  Direction: Output
switch(config-if)#
    
```

Example: show flow monitor

```

switch(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
switch(config)#
    
```

Example: show flow monitor cache module

```

switch# 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
0.0.0.0	255.255.255.255	17	Veth1		
7.192.192.10	7.192.192.2	1	Veth1	Eth5/2	

```

Cache entries on Processor1
- Active Flows:          0
- Free Flows:           4096

Cache entries on Processor2
- Active Flows:          1
- 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
switch#

```

Example: show flow monitor statistics module

```

switch# 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

```

```

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
switch#

```

Example: show flow record

```

switch(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
switch(config-flow-record)#

```

Netflow Example Configuration

The following example shows how to configure flow monitor using a new flow record and apply it to an interface:

```

switch# configure terminal
switch(config)# flow record RecordTest
switch(config-flow-record)# description Ipv4flow
switch(config-flow-record)# match ipv4 destination address
switch(config-flow-record)# collect counter packets
switch(config-flow-record)# exit
switch(config)# flow exporter ExportTest
switch(config-flow-exporter)# description ExportHamilton
switch(config-flow-exporter)# destination 192.0.2.1
switch(config-flow-exporter)# dscp 2
switch(config-flow-exporter)# source mgmt 0
switch(config-flow-exporter)# transport udp 200
switch(config-flow-exporter)# version 9
switch(config-flow-exporter-version-9)# option exporter-stats timeout 1200
switch(config-flow-exporter-version-9)# template data timeout 1200
switch(config-flow-exporter-version-9)# exit
switch(config-flow-exporter)# exit
switch(config)# flow monitor MonitorTest
switch(config-flow-monitor)# description Ipv4Monitor
switch(config-flow-monitor)# exporter ExportTest
switch(config-flow-monitor)# record RecordTest
switch(config-flow-monitor)# exit
switch(config)# interface veth 2/1
switch(config-if)# ip flow monitor MonitorTest output
switch(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
  Direction: Output
switch(config-if)#

```

The following example shows how to configure flow monitor using a pre-defined record and apply it to an interface:

```
switch# configure terminal
switch(config)# flow exporter ExportTest
switch(config-flow-exporter)# description ExportHamilton
switch(config-flow-exporter)# destination 192.0.2.1
switch(config-flow-exporter)# dscp 2
switch(config-flow-exporter)# source mgmt 0
switch(config-flow-exporter)# transport udp 200
switch(config-flow-exporter)# version 9
switch(config-flow-exporter-version-9)# option exporter-stats timeout 1200
switch(config-flow-exporter-version-9)# template data timeout 1200
switch(config-flow-exporter-version-9)# exit
switch(config-flow-exporter)# exit
switch(config)# flow monitor MonitorTest
switch(config-flow-monitor)# description Ipv4Monitor
switch(config-flow-monitor)# exporter ExportTest
switch(config-flow-monitor)# record netflow-original
switch(config-flow-monitor)# exit
switch(config)# interface veth 2/1
switch(config-if)# ip flow monitor MonitorTest output
switch(config-if)# show flow interface veth 2
Interface veth 2:
  Monitor: MonitorTest
  Direction: Output
switch(config-if)#
```

Related Documents for NetFlow

Related Topic	Document Title
Cisco NetFlow Overview	http://cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html

Feature History for NetFlow

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

Table 5: Feature 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.



CHAPTER 13

Configuring System Message Logging

This chapter contains the following sections:

- [Information about System Message Logging, page 137](#)
- [System Message Logging Facilities, page 138](#)
- [Guidelines and Limitations for System Message Logging, page 142](#)
- [Default System Message Logging Settings, page 142](#)
- [Configuring System Message Logging, page 143](#)
- [Verifying the System Message Logging Configuration, page 149](#)
- [System Message Logging Example Configuration, page 152](#)
- [Feature History for System Message Logging, page 152](#)

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.

The following table describes the severity levels used in system messages. When you configure the severity level, the system outputs messages at that level and lower.

Table 6: System Message Severity Levels

Level	Description
0 – emergency	System unusable
1 – alert	Immediate action needed

Level	Description
2 – critical	Critical condition
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.

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.



Note

When the device first initializes, messages are sent to syslog servers only after the network is initialized.

System Message Logging Facilities

The following table lists the facilities that you can use in system message logging configuration

Table 7: 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

Facility	Description
bootvar	Bootvar
callhome	Call home manager
capability	MIG utilities daemon
cdp	CDP manager
cert-enroll	Certificate enroll daemon
cfs	CFS manager
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

Facility	Description
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
mail	Mail system
mfdm	MFDM manager
module	Module manager
monitor	Ethernet SPAN manager
mrib	MRIB manager
mvsh	MVSH manager

Facility	Description
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

Facility	Description
tcpudp	TCP and UDP manager
u2	U2 manager
u6rib	U6RIB manager
ufdm	UFDM manager
urib	URIB manager
user	User process
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 for System Message Logging

System messages are logged to the console and the logfile by default.

Default System Message Logging Settings

Table 8: System Message Logging Defaults

Parameter	Default
Console logging	Enabled at severity level 2
Monitor logging	Enabled at severity level 5

Parameter	Default
Log file logging	Enabled to log messages at severity level 5
Module logging	Enabled at severity level 5
Facility logging	Enabled
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
- Restoring System Message Logging Defaults for Terminal Sessions
- Configuring System Message Logging for Modules
- Restoring System Message Logging Defaults for Modules
- Configuring System Message Logging for Facilities
- Restoring System Message Logging Defaults for Facilities
- Configuring syslog Servers
- Restoring System Message Logging Defaults for Servers
- Using a UNIX or Linux System to Configure Logging
- Displaying Log Files

Configuring System Message Logging to Terminal Sessions

You can log messages by severity level to console, telnet, and SSH sessions. By default, logging is enabled for terminal sessions.

Procedure

	Command or Action	Purpose
Step 1	switch# terminal monitor	Enables the device to log messages to the console.
Step 2	switch# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 3	switch(config)# logging console [<i>severity-level</i>]	Configures the device to log messages to the console session based on a specified severity level or higher. The default severity level is 2.
Step 4	switch(config)# show logging console	(Optional) Displays the console logging configuration.
Step 5	switch(config)# logging monitor [<i>severity-level</i>]	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. The default severity level is 2.
Step 6	switch(config)# show logging monitor	(Optional) Displays the monitor logging configuration.
Step 7	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

```
switch# terminal monitor
switch# configure terminal
switch(config)# logging console 2
switch(config)# show logging console
Logging console:                enabled (Severity: critical)
switch(config)# logging monitor 3
switch(config)# show logging monitor
Logging monitor:                enabled (Severity: errors)
switch(config)# copy running-config startup-config
switch(config)#
```

Restoring System Message Logging Defaults for Terminal Sessions

You can use the following commands in the CLI Global Configuration mode to restore default settings for system message logging for terminal sessions.

Table 9: Restoring System Message Logging Defaults for Terminal Sessions

Command	Description
no logging console [<i>severity-level</i>]	Disables the device from logging messages to the console.
no logging monitor [<i>severity-level</i>]	Disables logging messages to telnet and SSH sessions.

Configuring System Message Logging for Modules

You can configure the severity level and time-stamp units of messages logged by modules.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# logging module [severity-level]	Enables module log messages that have the specified severity level or higher. If the severity level is not specified, the default of 5 is used.
Step 3	switch(config)# show logging module	
Step 4	switch(config)# logging timestamp {microseconds milliseconds seconds}	Sets the logging time-stamp units. The default unit is seconds.
Step 5	switch(config)# show logging timestamp	(Optional) Displays the logging time-stamp units configured.
Step 6	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to configure system message logging for modules.

```
switch# configure terminal
switch(config)# logging module 3
switch(config)# show logging module
Logging linecard:          enabled (Severity: errors)
switch(config)# logging timestamp microseconds
switch(config)# show logging timestamp
Logging timestamp:        Microseconds
switch(config)# copy running-config startup-config
switch(config)#
```

Restoring System Message Logging Defaults for Modules

You can use the following commands in the CLI Global Configuration mode to restore default settings for system message logging for modules.

Table 10: Restoring System Message Logging Defaults for Modules

Command	Description
no logging module [severity-level]	Restores the default severity level for logging module system messages.
no logging timestamp {microseconds milliseconds seconds}	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.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# logging module [<i>severity-level</i>]	Enables module log messages that have the specified severity level or higher. If the severity level is not specified, the default of 5 is used.
Step 3	switch(config)# show logging module	(Optional) Displays the module logging configuration.
Step 4	switch(config)# logging timestamp { microseconds milliseconds seconds }	Sets the logging time-stamp units. The default unit is seconds.
Step 5	switch(config)# show logging timestamp	(Optional) Copies the running configuration to the startup configuration.
Step 6	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to configure system message logging for modules.

```
switch# configure terminal
switch(config)# logging module 3
switch(config)# show logging module
Logging linecard:          enabled (Severity: errors)
switch(config)# logging timestamp microseconds
switch(config)# show logging timestamp
Logging timestamp:        Microseconds
switch(config)# copy running-config startup-config
switch(config)#
```

Restoring System Message Logging Defaults for Facilities

You can use the following commands to restore system message logging defaults for facilities.

Table 11: Restoring System Message Logging Defaults for Facilities

Command	Description
no logging level [<i>facility severity-level</i>]	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.

Command	Description
no logging timestamp {microseconds milliseconds seconds}	Resets the logging time-stamp unit to the default (seconds).

Configuring syslog Servers

Use this procedure to configure syslog servers for system message logging.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# logging server <i>host</i> [<i>severity-level</i> [use-vrf <i>vrf-name</i>]]	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 range from 0 to 7. The default outgoing facility is local7.
Step 3	switch(config)# show logging server	(Optional) Displays the syslog server configuration.
Step 4	switch(config)# copy running-config startup-config	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

The following example shows how to forward all messages on facility local7.

```
switch# configure terminal
switch(config)# logging server 10.10.2.2 7
switch(config)# show logging server
Logging server:                enabled
{10.10.2.2}
    server severity:           debugging
    server facility:           local7
switch(config)# copy running-config startup-config
switch(config)#
```

Restoring System Message Logging Defaults for Servers

You can use the following command to restore server system message logging default.

Table 12: Restoring System Message Logging Defaults for Servers

Command	Description
no logging server <i>host</i>	Removes the logging server for the specified host.

Using a UNIX or Linux System to Configure Logging

Before You Begin

The following UNIX or Linux fields must be configured 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.

Procedure

-
- 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
- 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.

Procedure

	Command or Action	Purpose
Step 1	show logging last <i>number-lines</i>	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.

```
switch# 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
switch#
```

Verifying the System Message Logging Configuration

Use one of the following commands to verify the configuration:

Table 13: Verifying the System Message Logging Configuration

Command	Purpose
show logging console	Displays the console logging configuration.
show logging info	Displays the logging configuration.
show logging last <i>number-lines</i>	Displays the last number of lines of the log file.
show logging level [<i>facility</i>]	show logging level [<i>facility</i>]
show logging module	Displays the module logging configuration.
show logging monitor	Displays the monitor logging configuration.
show logging server	Displays the syslog server configuration.
show logging session	Displays the logging session status.
show logging status	Displays the logging status.
show logging timestamp	Displays the logging time-stamp units configuration.

Example: show logging console

```
switch# show logging console
Logging console:          disabled
switch#
```

Example: show logging info

```
switch# 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
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)
switch#
    
```

Example: show logging last

```

switch# 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:
clis
switch#
    
```

Example: show logging level aaa

```

switch# show logging level aaa
Facility           Default Severity      Current Session Severity
-----
aaa                2                      2

0(emergencies)    1(alerts)      2(critical)
3(errors)         4(warnings)    5(notifications)
6(information)   7(debugging)
switch#
    
```

Example: show logging module

```

switch# show logging module
Logging linecard:          enabled (Severity: notifications)
switch#
    
```

Example: show logging monitor

```

switch# show logging monitor
Logging monitor:          enabled (Severity: errors)
switch#
    
```

Example: show logging server

```

switch# show logging server
Logging server:          enabled
{10.10.2.2}
server severity:        debugging
server facility:        local7
switch#
    
```

Example: show logging session status

```
switch# 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
switch#
```

Example: show logging status

```
switch# show logging status
Fabric Distribute          : Enabled
Session State              : IDLE
switch#
```

Example: show logging timestamp

```
switch# show logging timestamp
Logging timestamp:          Seconds
switch#
```

System MESSage Logging Example Configuration

The following example shows how to configure system message logging:

```
switch# configure terminal
switch(config)# logging console 3
switch(config)# logging monitor 3
switch(config)# logging logfile my_log 6
switch(config)# logging module 3
switch(config)# logging level aaa 2
switch(config)# logging timestamp milliseconds
switch(config)# logging distribute
switch(config)# logging server 172.28.254.253
switch(config)# logging server 172.28.254.254 5 local3
switch(config)# logging commit
switch(config)# copy running-config startup-config
switch(config)#
```

Feature History for System Message Logging

Feature Name	Releases	Feature Information
System Message Logging	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 14

Configuring iSCSI Multipath

This chapter contains the following sections:

- [Information About iSCSI Multipath, page 153](#)
- [Guidelines and Limitations, page 157](#)
- [Pre-requisites, page 158](#)
- [Default Settings, page 158](#)
- [Configuring iSCSI Multipath, page 158](#)
- [Uplink Pinning and Storage Binding, page 159](#)
- [Converting to a Hardware iSCSI Configuration, page 166](#)
- [Changing the VMkernel NIC Access VLAN, page 168](#)
- [Verifying the iSCSI Multipath Configuration, page 171](#)
- [Managing Storage Loss Detection, page 172](#)
- [Related Documents, page 174](#)
- [Feature History for iSCSI Multipath, page 174](#)

Information About iSCSI Multipath

This section includes the following topics:

- Overview
- Supported iSCSI Adapters
- iSCSI Multipath Setup on the VMware Switch

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.

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

The default settings in the iSCSI Multipath configuration are listed in the following table.

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

Parameter	Default
State (port-profile)	Disabled

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 the following illustration, 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 the following illustration, vmk1 is bound to

VMHBA33, the iSCSI adapter associated with vmnic1 and to which vmk1 is pinned. Similarly vmk2 is bound to VMHBA34.

Figure 9: iSCSI Multipath on VMware Virtual Switch



The illustration now is the part

Two corrections

The vmhba in host2 (on the right) is vmhba33 (on the left) and vmhba34 (on the right)

The adapter label in host2 (on the right) is vmhba33 (on the left) and vmhba34 (on the right)

The following are the adapters and NICs used in the hardware and software iSCSI multipathing configuration shown in the iSCSI Multipath on VMware Virtual Switch illustration.

Software HBA	VMkernel NIC	VM NIC
VMHBA35	1	1
	2	2
Hardware HBA		
VMHBA33	1	1
VMHBA34	2	2

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 I3 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** 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 sekriahn-bl-vsm %VEM_MGR-SLOT8-1-VEM_SYSLOG_ALERT: sfpport :
  Removing Uplink Port Eth8/3 (l1l 19), when vmknic lveth8/1 (l1l 49) is pinned to this
  port for iSCSI Multipathing
```
- Hardware iSCSI is new in Cisco Nexus 1000V Release 4.2(1)SV1(5.1). 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.

- An iSCSI target and initiator should be in the same subnet.

Pre-requisites

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

Parameters	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
- Converting to a Hardware iSCSI Configuration procedure
- Changing the VMkernel NIC Access VLAN procedure

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.

Process for Uplink Pinning and Storage Binding

Refer to the following process for Uplink Pinning and Storage Binding:

- Creating a Port Profile for a VMkernel NIC procedure.
- Creating VMkernel NICs and Attaching the Port Profile procedure.

Do one of the following:

- If you want to override the automatic pinning of NICS, go to Manually Pinning the NICs procedure.
- If not, continue with storage binding.
- You have completed uplink pinning. Continue with the next step for storage binding.
- Identifying the iSCSI Adapters for the Physical NICs procedure.
- Binding the VMkernel NICs to the iSCSI Adapter procedure.
- Verifying the iSCSI Multipath Configuration procedure.

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.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Places you in global configuration mode.
Step 2	switch(config)# port-profile type vethernet <i>name</i>	Places you into the CLI Port Profile Configuration mode for the specified port profile. type: Defines the port-profile as Ethernet or vEthernet type. Once configured, this setting cannot be changed. The default is vEthernet type. If a port profile is configured as an Ethernet type, then it cannot be used to configure VMware virtual ports. name: 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	switch(config)# description <i>profile description</i>	Adds a description to the port profile. This description is automatically pushed to the vCenter Server. profile description: up to 80 ASCII characters. If the description includes spaces, it must be surrounded by quotations.
Step 4	switch(config)# vmware port-group <i>name</i>	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: 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	switch(config)# switchport mode access	Designates that the interfaces are switch access ports (the default).
Step 6	switch(config)# switchport access vlan <i>vlanID</i>	Assigns the system VLAN ID to the access port for this port profile. The VLAN assigned to this iSCSI port profile must be a system VLAN.
Step 7	switch(config)# no shutdown	Administratively enables all ports in the profile.
Step 8	switch(config)# system vlan <i>vlanID</i>	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	switch(config)# capability iscsi-multipath	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	switch(config)# state enabled	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.

	Command or Action	Purpose
Step 11	switch(config)# show port-profile name <i>name</i>	(Optional) Displays the current configuration for the port profile.
Step 12	switch(config)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

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.

- 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.
- 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.

Before You Begin

Before starting this procedure, you must know or do the following

Procedure

-
- 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, use the command `# vemcmd show iscsi pinning` to display the auto pinning configuration for verification.

Example:

```
Example:~ # vemcmd show iscsi pinning
Vmknick  LTL      Pinned_Uplink  LTL
vmk6     49        vmnic2         19
vmk5     50        vmnic1         18
```

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.
- 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.

Procedure

Step 1 List the binding for each VMHBA to identify the binding to remove (iSCSI VMkernel NIC to VMHBA) with the command `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:

```

Example:
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk6
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk5

```

If active iSCSI sessions exist between the host and targets, the iSCSI port cannot be disconnected.

Step 3 From the ESX host, display the auto pinning configuration with the command **# vemcmd show iscsi pinning**.

Example:

```

Example:
~ # vemcmd show iscsi pinning
Vmknuc   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 with the command **# vemcmd set iscsi pinningvmk-ltl vmnic-ltl**.

Example:

```

Example:
~ # vemcmd set iscsi pinning 50 20

```

Step 5 Manually pin the VMkernel NIC to the physical NIC, overriding the auto pinning configuration with the command **# vemcmd set iscsi pinningvmk-ltl vmnic-ltl**.

Example:

```

Example:
~ # vemcmd set iscsi pinning 50 20

```

Step 6 You have completed this procedure. Return to the [Process for Uplink Pinning and Storage Binding](#), on page 159 section.

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
- Identifying iSCSI Adapters on the Host Server procedure

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.

Procedure

- 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.
- 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.
-

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

Procedure

Step 1 Use the command `esxcfg-scsidevs -a` to list the storage adapters on the server.

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 using the command `esxcli swiscsi vmnic list -dadapter-name` to list the storage adapters on the server.

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.

You have completed this procedure.

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.

Procedure

Step 1 Find the physical NICs to which the VEM has pinned the VMkernel NICs.

Example:

```
Vmknick LTL Pinned_Uplink LTL
vmk2    48  vmnic2          18
vmk3    49  vmnic3          19
```

Step 2 Bind the physical NIC to the iSCSI adapter.

Example:

```
Example:
esxcli swiscsi nic add --adapter vmhba33 --nic vmk2
```

```
Example:
esxcli swiscsi nic add --adapter vmhba34 --nic vmk3
```

For more information, refer to Identifying the iSCSI Adapters for the Physical NICs procedure.

You have completed this procedure.

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.

Procedure

- 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.
 - 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** 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, refer to the Adding the Hardware NICs to the DVS procedure.
 - 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.
 - Step 9** Bind the NIC to the adapter, using the Binding the VMkernel NICs to the iSCSI Adapter procedure.
 - Step 10** Verify the iSCSI multipathing configuration, using the Verifying the iSCSI Multipath Configuration procedure.
-

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.

Procedure

- 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.
 - 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** 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, refer to the Adding the Hardware NICs to the DVS procedure.
 - 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.
 - Step 9** Bind the NIC to the adapter, using the Binding the VMkernel NICs to the iSCSI Adapter procedure.
 - Step 10** Verify the iSCSI multipathing configuration, using the Verifying the iSCSI Multipath Configuration procedure.
-

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.

Procedure

Remove the iSCSI VMkernel NIC binding to the VMHBA.

Example:

```
Example:  
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk6  
esxcli swiscsi nic remove --adapter vmhba33 --nic vmk5
```

You have completed this procedure. Return to the Process for Converting to a Hardware iSCSI Configuration section.

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.

Procedure

- 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.
-

What to Do Next

You have completed this procedure. Return to the Process for Converting to a Hardware iSCSI Configuration section.

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:

Procedure

- 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.
 - 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.
 - 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.
 - Step 8** Bind the NIC to the adapter, using the Binding the VMkernel NICs to the iSCSI Adapter procedure.
 - Step 9** Verify the iSCSI multipathing configuration, using the [Verifying the iSCSI Multipath Configuration, on page 171](#) procedure.
-

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.

Procedure

- 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).

Example:

```
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 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.

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.

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.

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.

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.

Example:

```

n1000v# config t

```

```
n1000v(config)# port-profile type vethernet VMK-port-profile
n1000v(config-port-prof)# capability iscsi-multipath
```

What to Do Next

You have completed this procedure.

-

Verifying the iSCSI Multipath Configuration

Refer to the following commands and the examples.

Before You Begin

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.
<code>esxcli swiscsi nic list -d vmhba33</code>	Displays the iSCSI adapter binding of VMkernel NICs. See Example 13-2.
<code>show port-profile [brief expand-interface usage] name [profile-name]</code>	Displays the port profile configuration. See Example.

Procedure

Step 1 ~ # vemcmd show iscsi pinning

Example:

```
~ # vemcmd show iscsi pinning
Vmknict LTL Pinned_Uplink LTL
vmk6 49 vmnic2 19
vmk5 50 vmnic1 18
```

Step 2 esxcli swiscsi nic list -d vmhba33

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

```

Step 3 show port-profile name iscsi-profile

Example:

```

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:
system vlans: 254
capability l3control: no
capability iscsi-multipath: yes
port-profile role: none
port-binding: static
n1000v#

```

Managing Storage Loss Detection

This section describes the command that provides the configuration to detect storage connectivity losses and provides support when storage loss is detected. When VSMS are hosted on remote storage systems such as NFS or iSCSI, storage connectivity can be lost. This connectivity loss can cause unexpected VSM behavior.

Use the following command syntax to configure storage loss detection: **system storage-loss** { *log* | *reboot* } [*time* <*interval*>] | **no system storage-loss** [{ *log* | *reboot* }] [*time* <*interval*>]

The time interval value is the intervals at which the VSM checks for storage connectivity status. If a storage loss is detected, the syslog displays. The default interval is 30 seconds. You can configure the intervals from 30 seconds to 600 seconds. The default configuration for this command is: **system storage-loss log time 30**



Note

Configure this command in EXEC mode. Do not use configuration mode.

The following describes how this command manages storage loss detection:

- **Log only:** A syslog message is displayed stating that a storage loss has occurred. The administrator must take action immediately to avoid an unexpected VSM state.
- **Reset:** The VSM on which the storage loss is detected is reloaded automatically to avoid an unexpected VSM state.
 - **Storage loss on the active VSM:** The active VSM is reloaded. The standby VSM becomes active and takes control of the hosts.

- Storage loss on the standby VSM: The standby VSM is reloaded. The active VSM continues to control the hosts.



Note Do not keep both the active and standby VSMs on the same remote storage, so that storage losses do not affect the VSM operations.

Before You Begin

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

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

Procedure

Step 1 `system storage-loss log time 30`

Example:

```
n1000v# system storage-loss log time 30
n1000v#
```

Sets the time interval in seconds to check storage connectivity and log the status. Thirty seconds is the default interval.

Step 2 `copy running-config startup-config`

Example:

```
n1000v# copy run start
n1000v#
```

Example:

The following command disables the storage-loss checking. Whenever the VSMs are installed on local storage, this is the configuration we recommend.

Note Disable storage loss checking if both VSMs are in local storage.

```
n1000v# no system storage-loss
```

The following command enables storage loss detection time intervals on an active or standby VSM and displays a syslog message about the storage loss. In this example, the VSM is checked for storage loss every 60 seconds. The administrator has to take action to recover the storage and avoid an inconsistent VSM state:

```
n1000v# system storage-loss log time 60
```

The following example shows the commands you use to configure the VSM to reboot when storage loss is detected:

```
n1000v# system storage-loss reboot time 60
n1000v# copy run start
```

The following example shows the commands you use to disable storage loss checking:

```
n1000v# no system storage-loss
n1000v# copy run start
```

Saves configuration changes in the running configuration to the startup configuration in persistent memory.

What to Do Next

-

Related Documents

Related Topic	Document Title
VMware SAN Configuration	VMware SAN Configuration Guide

Feature History for iSCSI Multipath

Feature	Releases	Feature Information
Hardware iSCSI Multipath	4.2(1)SV1(4)	Added support for hardware iSCSI Multipath.
Configuring iSCSI Multipath	4.0(4)SV1(1)	This feature was introduced.



CHAPTER 15

Configuring VSM Backup and Recovery

This chapter contains the following sections:

- [Information About VSM Backup and Recovery, page 175](#)
- [Guidelines and Limitations, page 175](#)
- [Configuring VSM Backup and Recovery, page 176](#)

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 for more information.

Guidelines and Limitations

VSM backup and recovery has the following configuration guidelines and limitations:

- Backing up the VSM is a onetime task.
- Backing up the VSM 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.

Configuring VSM Backup and Recovery

This section includes the following topics:

- Performing a Backup of the VSM
- Performing a Periodic Backup
- Recovering the VSM

**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
- Performing a Periodic Backup

Performing a Backup of the VSM

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.

Procedure

Step 1 Open the vSphere Client.

The vSphere Client window opens as displayed in the following illustration.


Figure 10: vSphere Client Window



Step 2 In the left navigation pane, right-click the standby VSM. A drop-down list is displayed.

Step 3 Choose **Power > Power Off**.
The action is displayed in the Clone to Template Window.

Figure 11: Clone to Template Window




Step 4 In the left navigation pane, right-click the standby VSM.
A drop-down list is displayed.

Step 5 Choose **Template > Clone to Template**.

The Clone Virtual Machine to Template window opens.

Figure 12: Clone Virtual Machine to Template Window




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 Window opens.

Figure 13: Host Window




Step 9 Choose the host on which the template will be stored.

Step 10 Click Next.

The Choosing a Datastore window opens.

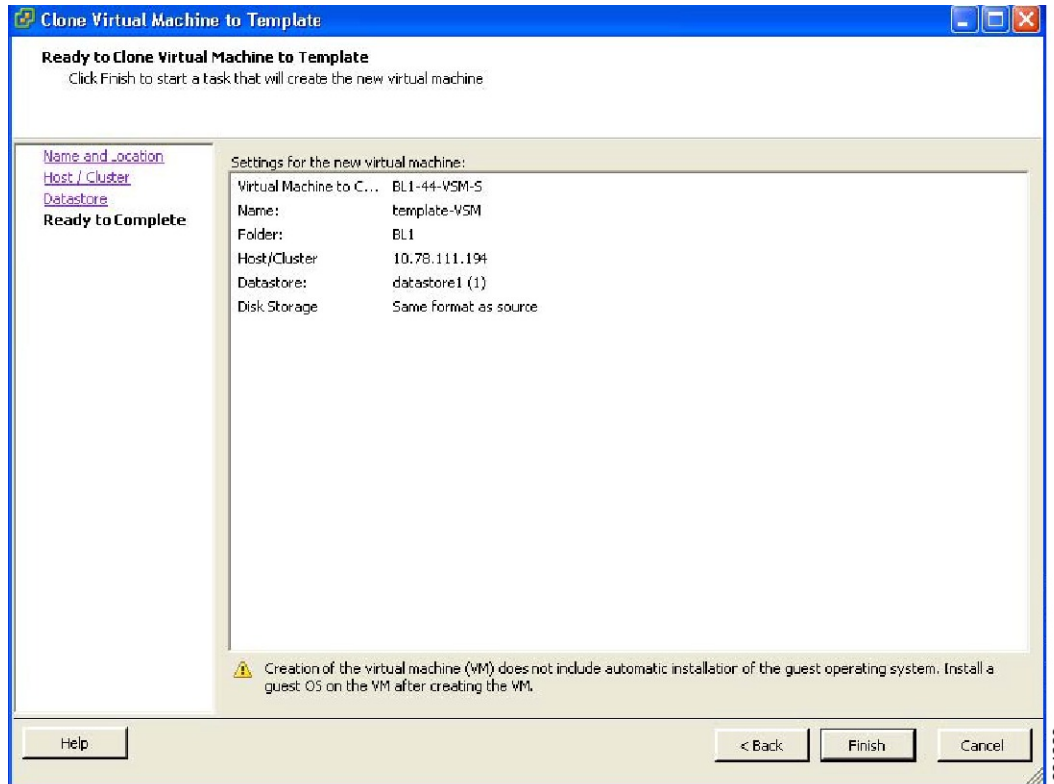
Figure 14: Choosing a Datastore Window



- 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.

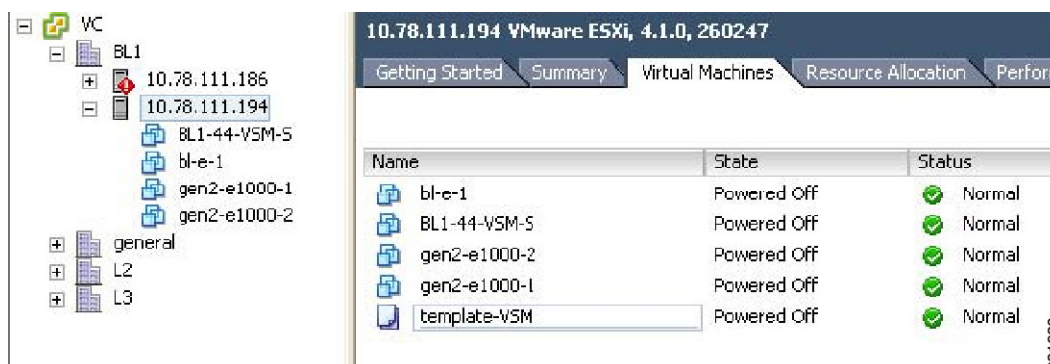
The Confirming the Settings window opens.

Figure 15: Confirming the Settings Window



- 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.
- Step 15** The Template Virtual Machine window opens.
The template creation is complete.

Figure 16: Template Virtual Machine Window



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.

Before You Begin

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

Enter the command `copy running-config scp://root@10.78.19.15/tftpboot/config/` to back up the VSM.

Example:

```
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
- Erasing the Old Configuration
- Restoring the Backup Configuration on the VSM

Deploying the Backup VSM VM

This section describes how to deploy the backup VSM VM when the primary and secondary VSMs are not present.




Note

While deploying the VSM VM, do not power it on.

Procedure

- Step 1** Open the vSphere Client.
The vSphere Client window opens as displayed in the following illustration.

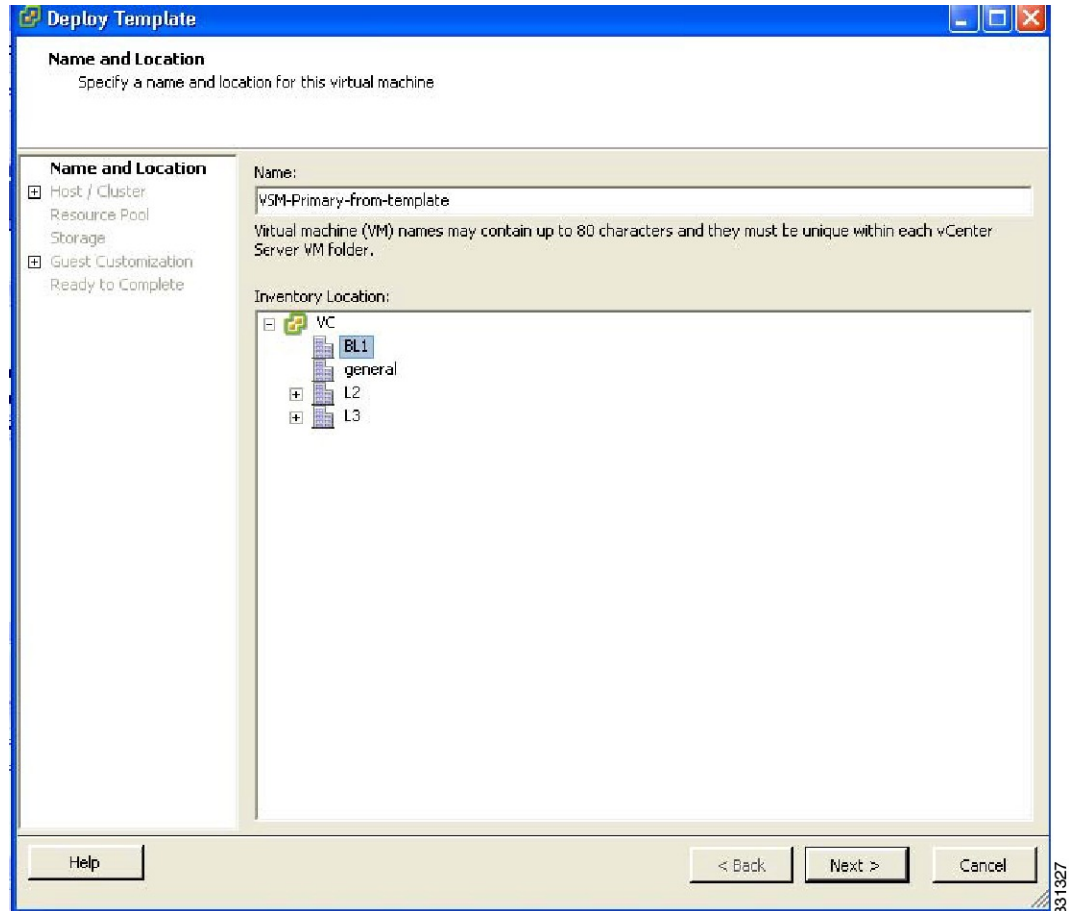
Figure 17: vSphere Client Window



- 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 window opens.


Figure 18: Deploy Template Wizard Window



- 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.

The Choosing a Host Window opens.

Figure 19: Choosing a Host Window



Step 9 Choose a host.


Step 10

Example:

Click Next.

The Choosing a Datastore window opens.


Figure 20: Choosing a Datastore Window



- 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 window opens. Make sure that the Power on this virtual machine after creation check box is not checked.

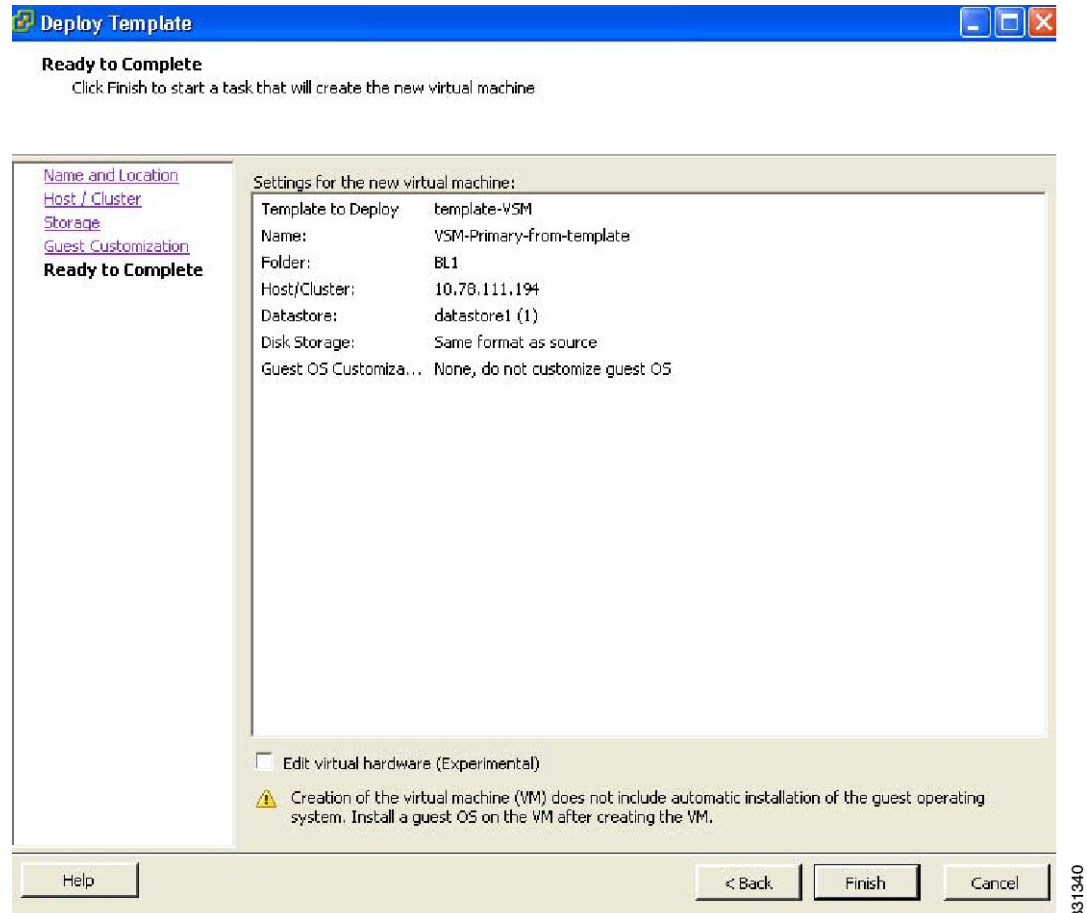
Figure 21: Guest Customization Window



Step 14 Click Next.

The Deploy Template - Ready to Complete window opens.


Figure 22: Guest Customization Window



- Step 15** Confirm the settings for the new virtual machine and click Finish. 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.

Figure 23: Guest Customization Window



Step 18 In the Hardware / Summary pane, choose Network adapter 1.

Step 19 Uncheck the Connect at power on check box.

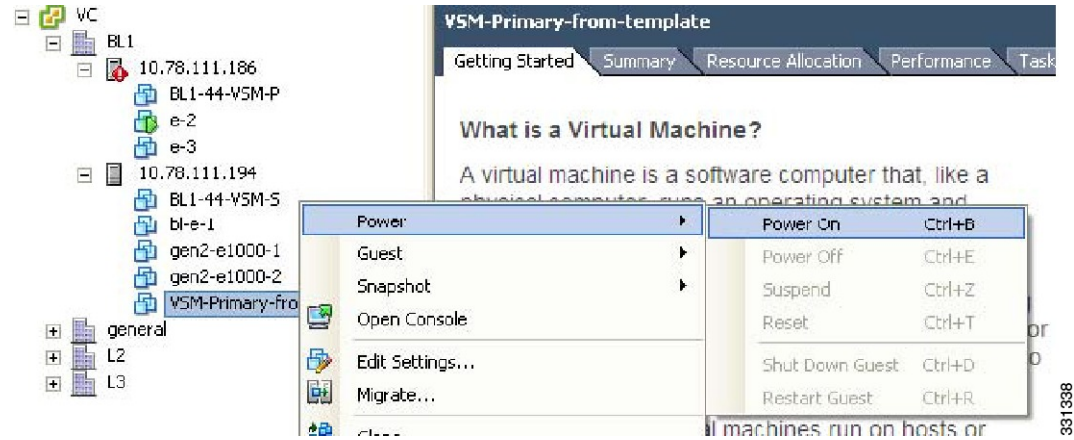
Step 20 Choose Network adapter 2.

Step 21 In the Device Status area, uncheck the Connect at power on check box.

Step 22 Click OK.

The Power On window opens.

Figure 24: Guest Customization 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:
- Example:**
- ```
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:

**Example:**

```
switch# copy running-config startup-config
scp: sftp: startup-config
[#####] 100%
switch#
```

- Step 4** Erase the startup configuration by entering the following command:

**Example:**

```
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:

**Example:**

```
switch# reload
This command will reboot the system. (y/n)? [n] y
```

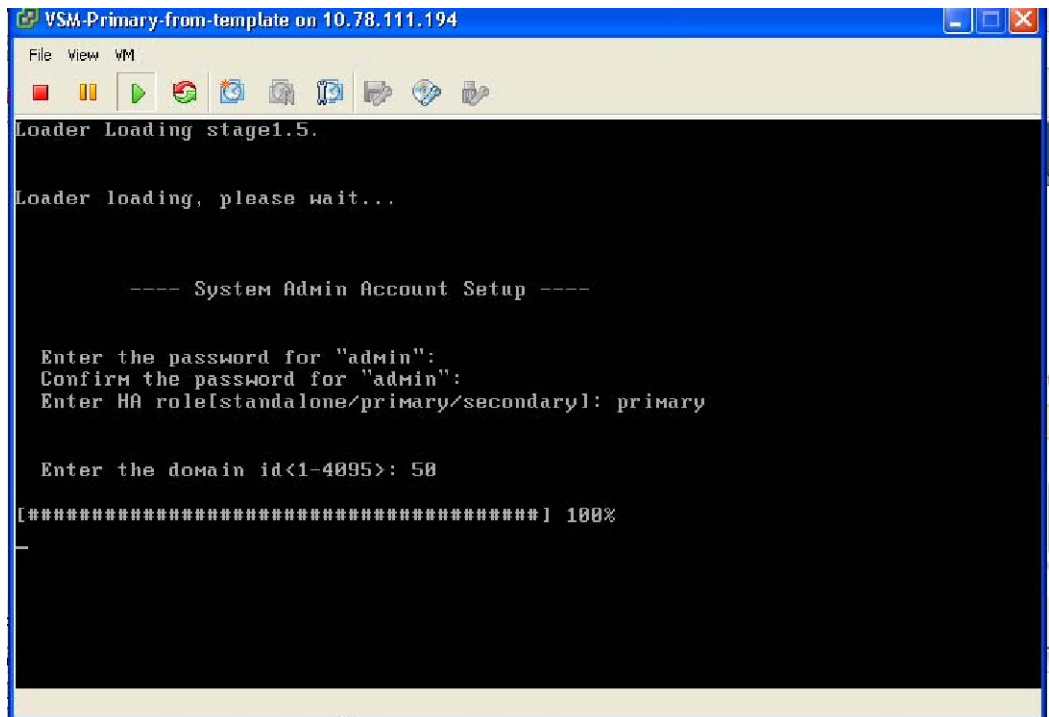
## Restoring the Backup Configuration on the VSM

This section describes how to restore the backup configuration on the VSM.

### Procedure

**Step 1** When the VSM reboots, the System Admin Account Setup window opens.

*Figure 25: System Admin Account Setup Window*



**Step 2** Enter and confirm the Administrator password.

**Example:**

```
---- System Admin Account Setup ----
Enter the password for "admin":
Confirm the password for "admin":
```

**Step 3** Enter the domain ID.

**Example:**

```
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.

**Example:**

```
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
to skip the remaining dialogs.
```

**Step 5** Enter yes when you are prompted to enter the basic configuration dialog.

**Example:**

```
Would you like to enter the basic configuration dialog (yes/no): yes
```

**Step 6** Enter no when asked to create another Login account.

**Example:**

```
Create another login account (yes/no) [n]: no
```

**Step 7** Enter no when asked to configure a read-only SNMP community string.

**Example:**

```
Configure read-only SNMP community string (yes/no) [n]: no
```

**Step 8** Enter no when asked to configure a read-write SNMP community string.

**Example:**

```
Configure read-write SNMP community string (yes/no) [n]: no
```

**Step 9** Enter a name for the switch.

**Example:**

```
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.

**Example:**

```
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 no when asked to configure the default gateway.



**Example:**

```
Configure the default-gateway: (yes/no) [y]: no
```

```
IPv4 address of the default gateway : 172.23.233.1
```

**Step 12** Enter yes when asked to enable the Telnet service.

**Example:**

```
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*.

**Example:**

```
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.

**Example:**

```
Enable the http-server? (yes/no) yes
```

**Step 15** Enter no when asked to configure the NTP server

**Example:**

```
Configure NTP server? (yes/no) [n]: no
```

**Step 16** Enter no when asked to configure the VEM feature level.

**Example:**

```
Vem feature level will be set to 4.2(1)SV1(4a).
```

```
Do you want to reconfigure? (yes/no) [n] no
```

The system now summarizes the complete configuration and prompts you to edit it.

**Example:**

```
The following configuration will be applied:
```

```
interface Mgmt0
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
svs mode L2
control vlan 1
packet vlan 1
domain id 1
```

**Step 17** Enter no when asked if you would like to edit the configuration.

**Example:**

```
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 18** Enter yes when asked to use and save this configuration.

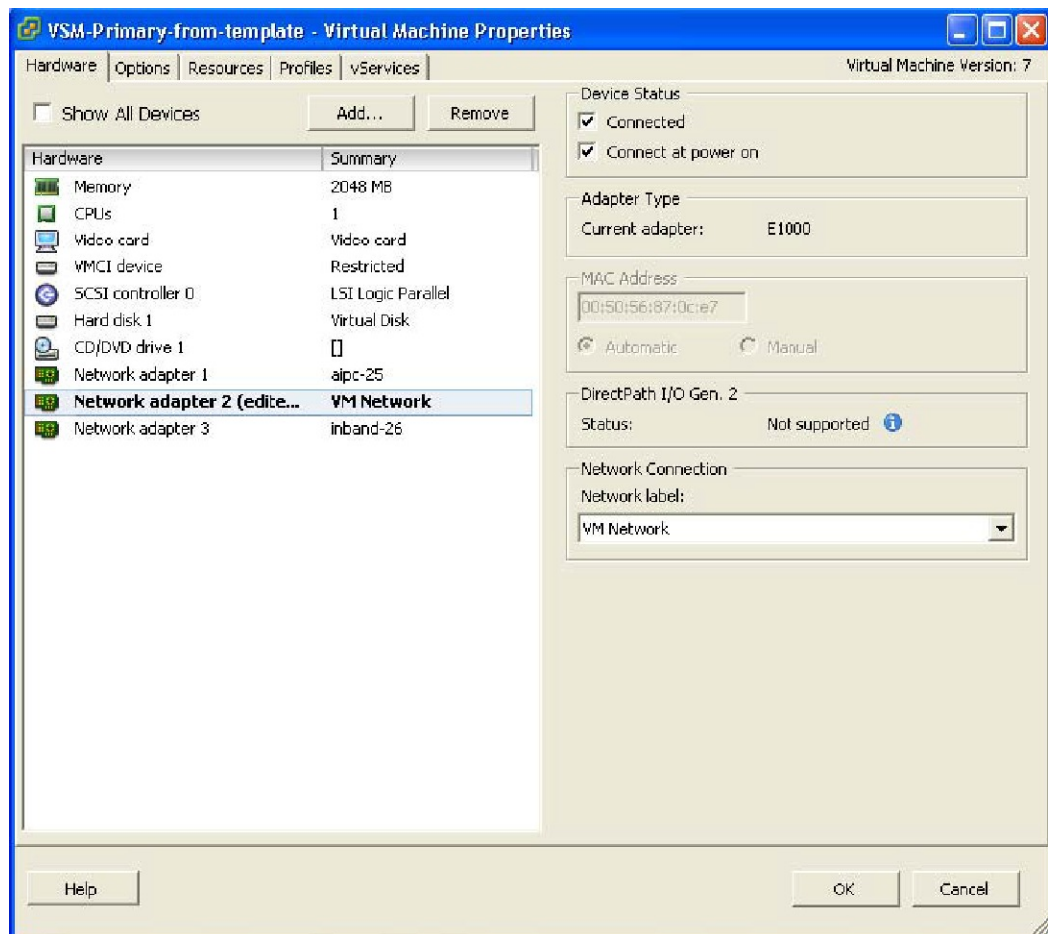
**Example:**

```
Use this configuration and save it? (yes/no) [y]: yes
[#####] 100%
```

If you do not save the configuration now, then 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 also automatically configured.

- Step 19** In the vSphere Client, right-click the VSM and choose Edit Settings. The VSM Virtual Machine Properties window opens.

**Figure 26: VSM Virtual Machine Properties Window**



- Step 20** In the Hardware/Summary pane, choose Network adapter 2.
- Step 21** Check the Connect at power on check box.
- Step 22** Log in to the VSM.
- Step 23** Copy the backup configuration to the VSM bootflash by entering the following command:

**Example:**

```

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 24** Copy the backup configuration to the running configuration by entering the following command:

**Example:**

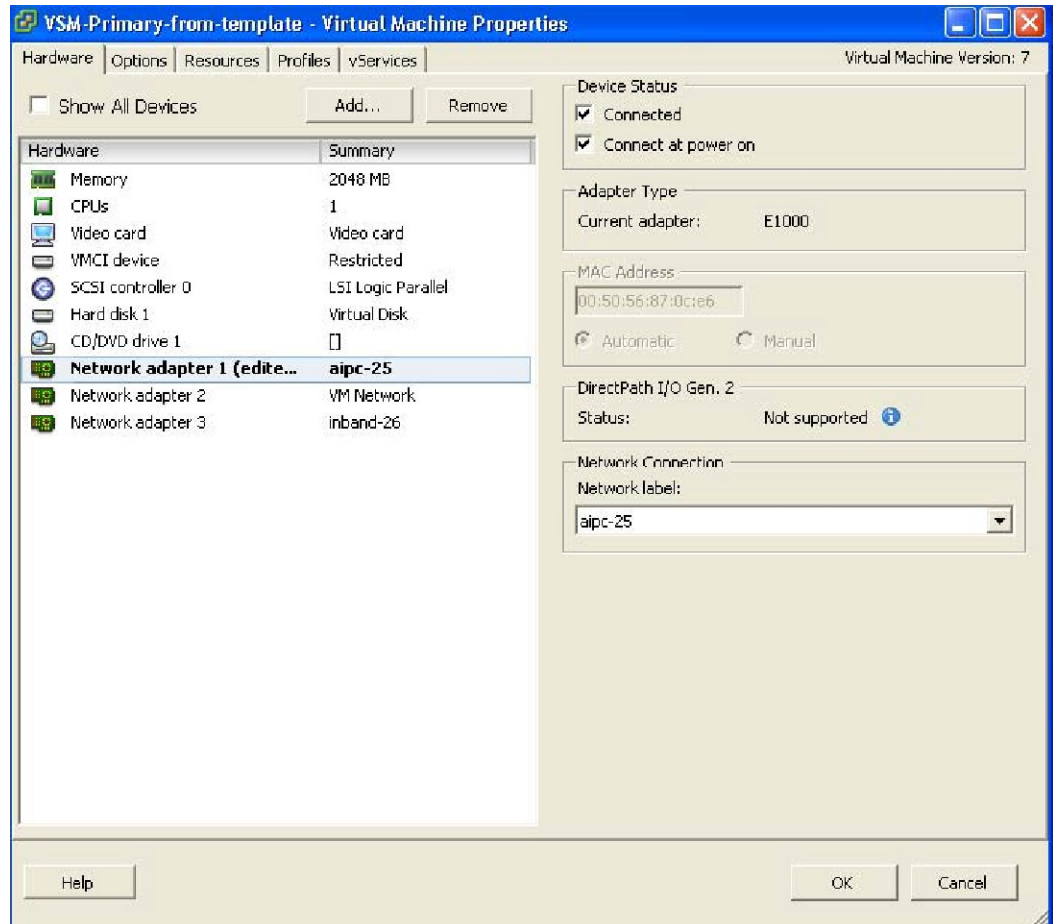
```

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
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!
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

```

The Virtual Machine Properties window displays.

**Figure 27: Virtual Machine Properties Window**



**Step 25** In the Hardware / Summary pane, choose Network adapter 1.

**Step 26** In the Device Status area, check the Connect at power on check box.

**Step 27** Confirm that the VEMs are attached to the VSM by entering the following command:

**Example:**

```
switch# show module
```

```
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 28** Copy the backup configuration to the running configuration by entering the following command:

**Example:**

```

switch# 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 %ETHERPORT-5-SPEED: Interface port-channel1,
operational speed changed to 1 Gbps
2011 Apr 26 12:21:35 switch %ETHERPORT-5-IF_DUPLEX: Interface port-channel1,
operational duplex mode changed to Full
2011 Apr 26 12:21:35 switch %ETHERPORT-5-IF_RX_FLOW_CONTROL: Interface portchannel1,
operational Receive Flow Control state changed to on
2011 Apr 26 12:21:35 switch %ETHERPORT-5-IF_TX_FLOW_CONTROL: Interface portchannel1,
operational Transmit Flow Control state changed to on
VSM backup and Recovery Procedure EDCS-1017832Cisco 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 %ETHERPORT-5-IF_UP: Interface Ethernet3/2 is up in
mode trunk
2011 Apr 26 12:21:35 switch %ETHERPORT-5-IF_UP: Interface port-channel1 is up in
mode trunk
switch#

```

This step is necessary if features are configured directly through the interface configuration mode for Ethernet interfaces and for features like ERSPAN/NFM.

**Step 29** Copy the running-configuration to the startup-configuration by entering the following command:

**Example:**

```
switch# copy running-config startup-config
[#####] 100%
switch#
```

**Step 30** 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 Installation and Upgrade Guide*.

- 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.

---

## 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. |



# CHAPTER 16

## Enabling vTracker

---


This chapter contains the following sections:

- [Information About vTracker, page 200](#)
- [Guidelines and Limitations, page 201](#)
- [Default Settings for vTracker Parameters, page 201](#)
- [Enabling vTracker Globally, page 201](#)
- [Upstream View, page 203](#)
- [Virtual Machine \(VM\) View, page 205](#)
- [Module pNIC View, page 211](#)
- [VLAN View, page 212](#)
- [VMotion View, page 214](#)
- [Feature History for vTracker, page 216](#)

## Information About vTracker

The following illustration displays the vTracker setup diagram:

**Figure 28: vTracker Setup Diagram in the Cisco Nexus 1000V Environment**



The vTracker feature on the Cisco Nexus 1000V switch provides information about the virtual network environment. Once you enable vTracker, it becomes aware of all the modules and interfaces that are connected with the switch. vTracker provides various views that are based on the data sourced from the vCenter, the Cisco Discovery Protocol (CDP), and other related systems connected with the virtual switch. You can use vTracker to troubleshoot, monitor, and maintain the systems. Using vTracker show commands, you can access consolidated network information across the following views:

- Upstream View—Provides information on all the virtual ports connected to an upstream physical switch. The view is from top of the network to the bottom.
- VM View—Supports two sets of data:
  - VM vNIC View—Provides information about the virtual machines (VMs) that are managed by the Cisco Nexus 1000V switch. The vNIC view is from the bottom to the top of the network.
  - VM Info View—VM Info View—Provides information about all the VMs that run on each server module.
- Module pNIC View—Provides information about the physical network interface cards (pNIC) that are connected to each Virtual Ethernet Module (VEM).
- VLAN View—Provides information about all the VMs that are connected to specific VLANs.
- vMotion View—Provides information about all the ongoing and previous VM migration events.



**Note**

vTracker is available with both Essential and Advanced edition of Cisco Nexus 1000V.

## Guidelines and Limitations

vTracker has the following configuration guidelines and limitations:

- For VM and VMotion views, you should connect the Virtual Supervisor Module (VSM) with the vCenter for the vTracker **show** commands to work.
- vTracker is disabled by default.
- While the Cisco Nexus 1000V switch information is validated, the information sourced by vTracker from the vCenter is not verifiable.
- All vTracker views are valid for a given time only, because the virtual environment is dynamic and constantly changing.
- In a scaled-up environment, vTracker can experience delays in retrieving real-time information, which is distributed across VEMs and vCenter, among other components.

## Default Settings for vTracker Parameters

### Default vTracker Parameters

| Parameters       | Default           |
|------------------|-------------------|
| feature vtracker | Disabled globally |

## Enabling vTracker Globally

- vTracker can be configured only globally, not on individual interfaces.
- By default, vTracker is disabled.

### Before You Begin

- You are logged in to the VSM CLI in EXEC mode or the configuration mode of any node.
- vTracker does not change any VSM configuration settings or behavior. Rather, it only tracks and displays the current configuration views.

**Procedure**

|               | <b>Command or Action</b>                                  | <b>Purpose</b>                                                                                                                              |
|---------------|-----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | switch# <b>configure terminal</b>                         | Enters global configuration mode.                                                                                                           |
| <b>Step 2</b> | switch(config)# <b>[no] feature vtracker</b>              | Enables the vTracker feature.<br>Use the <b>no</b> form of this command to disable this feature.                                            |
| <b>Step 3</b> | switch(config)# <b>copy running-config startup-config</b> | (Optional)<br>Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration. |

The following example enables vTracker:


```
switch# configure terminal
switch(config)# feature vtracker
switch(config)# copy running-config startup-config
```

# Upstream View

## Upstream View Overview

The upstream view provides end-to-end network information from the VM to the physical switch. The following is the upstream view set-up diagram:

*Figure 29: Upstream View Setup Diagram in the Cisco Nexus 1000V Environment*



### Note

Cisco Discovery Protocol (CDP) neighbor information must be accessible to generate the required upstream view output. You can enter the **show cdp neighbors** command on the VSM, if CDP is enabled globally and all the interfaces.

## Displaying Upstream View

To display the upstream view, follow the given step.

### Procedure

**show vtracker upstream-view** [**device-id** *name* | **device-ip** *IP address*]

The following examples show the vTracker upstream view in a VSM:

#### Example:

```
switch(config)# show vtracker upstream-view
```

| Device-Name<br>Device-IP      | Device-Port<br>Local-Port | Server-Name<br>Adapter Status | PC-Type<br>PO-Intf | Veth-interfaces |
|-------------------------------|---------------------------|-------------------------------|--------------------|-----------------|
| Upstream-SW-A<br>203.0.113.66 | Gig2/7<br>Eth3/3          | 203.0.113.118<br>vmnic2 up    | MacPinn<br>Po1     | 10-11           |
| Upstream-SW-B<br>203.0.113.54 | Gig3/10<br>Eth3/4         | 203.0.113.117<br>vmnic3 up    | MacPinn<br>Po1     | 9               |
|                               | Gig3/8<br>Eth4/3          | 203.0.113.99<br>vmnic2 up     | Default<br>Po2     | 1-2             |
|                               | Gig3/9<br>Eth4/4          | 203.0.113.99<br>vmnic3 up     | Default<br>Po2     | 1-2             |

#### Example:

```
switch(config)# show vtracker upstream-view device-id Upstream-SW-A
```

| Device-Name<br>Device-IP      | Device-Port<br>Local-Port | Server-Name<br>Adapter Status | PC-Type<br>PO-Intf | Veth-interfaces |
|-------------------------------|---------------------------|-------------------------------|--------------------|-----------------|
| Upstream-SW-A<br>203.0.113.66 | Gig2/7<br>Eth3/3          | 203.0.113.118<br>vmnic2 up    | MacPinn<br>Po1     | 10-11           |

## Upstream View Field Description

The column headings in the upstream view examples above is described in the following table:

| Column      | Description                                                                                    |
|-------------|------------------------------------------------------------------------------------------------|
| Device-Name | Name of the neighboring device.                                                                |
| Device-IP   | IP address of the device.                                                                      |
| Device-Port | Port interface of the device that is connected to the Cisco Nexus 1000V Ethernet (local) port. |
| Local-Port  | Local port interface, which is connected to the neighboring device port.                       |
| Server-Name | Name or IP address of the server module to which the local port is connected.                  |

| Column          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adapter         | Local port name as known by the hypervisor. For VMWare ESX or ESXi, it is known as VMNic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Status          | Local port's operational status.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| PC-Type         | Port-channel type of the local port. Each PC-Type has a corresponding channel-group configuration in the port profile or the interface. Supported values are as follows: <ul style="list-style-type: none"> <li>• Default—channel-group auto or channel-group auto mode on</li> <li>• MacPinn—channel-group auto mode on mac-pinning</li> <li>• MacPinnRel—channel-group auto mode on mac-pinning relative</li> <li>• SubGrpCdp—channel-group auto mode on sub-group cdp</li> <li>• SubGrpMan—channel-group auto mode on sub-group manual</li> <li>• LACP-A—channel-group auto mode active</li> <li>• LACP-P—channel-group auto mode passive</li> </ul> |
| PO-Intf         | Port channel interface of the local port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| veth-interfaces | Available virtual Ethernet interfaces for which traffic can flow through the upstream switch. <p><b>Note</b> You can get similar information by entering the <b>show int virtual pinning</b> command at the VSM prompt.</p>                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Virtual Machine (VM) View

### Virtual Machine (VM) View Overview

The VM view provides you with comprehensive information about the VMs that are connected with the Cisco Nexus 1000V switch. The VM perspective is divided into two views:

- VM vNIC View—Provides information about all the vNICs (virtual network interface cards) adapters that are managed by the Cisco Nexus 1000V switch.

- VM Info View—Provides information about all the VMs that run on each server module. The data is fetched from the attributes of each VM via the vCenter.

**Note**

The VSM must be connected with the vCenter in order to generate the required VM view output. You can enter the **show vsx connections** command on the VSM to verify the connection.

## Displaying the VM vNIC View

To display the VM vNIC view, follow the given step.

### Procedure

**show vtracker vm-view vnic** [**module number** | **vm name**]

**Note** The timeout for this command is 180 seconds.

The following examples show the vTracker VM vNIC view in a VSM:

#### Example:

```
switch(config)# show vtracker vm-view vnic
* Network: For Access interface - Access vlan, Trunk interface - Native vlan,
 VXLAN interface - Segment Id.
```

| Mod | VM-Name<br>HypvPort | VethPort<br>Adapter | Drv Type<br>Mode  | Mac-Addr<br>IP-Addr             | State | Network | Pinning |
|-----|---------------------|---------------------|-------------------|---------------------------------|-------|---------|---------|
| 3   | gentoo-2<br>1025    | Veth3<br>Adapter 3  | Vmxnet3<br>access | 0050.56b5.37de<br>n/a           | up    | 339     | Eth3/8  |
| 3   | gentoo-2<br>1026    | Veth4<br>Adapter 4  | E1000<br>access   | 0050.56b5.37df<br>n/a           | up    | 339     | Eth3/8  |
| 3   | gentoo-2<br>1024    | Veth5<br>Adapter 2  | Vmxnet2<br>access | 0050.56b5.37dd<br>n/a           | up    | 339     | Eth3/8  |
| 4   | Fedora-VM1<br>4258  | Veth7<br>Adapter 2  | E1000<br>pvlan    | 0050.56bb.4fc1<br>10.104.249.49 | up    | 406     | Eth4/3  |
| 5   | Fedora-VM2<br>100   | Veth1<br>Adapter 1  | E1000<br>trunk    | 0050.56b5.098b<br>n/a           | up    | 1       | Po9     |
| 5   | Fedora-VM2<br>3232  | Veth2<br>Adapter 3  | E1000<br>pvlan    | 0050.56b5.098d<br>10.104.249.60 | up    | 405     | Po9     |

#### Example:

```
switch(config)# show vtracker vm-view vnic module 4
* Network: For Access interface - Access vlan, Trunk interface - Native vlan,
 VXLAN interface - Segment Id.
```

| Mod | VM-Name<br>HypvPort | VethPort<br>Adapter | Drv Type<br>Mode | Mac-Addr<br>IP-Addr             | State | Network | Pinning |
|-----|---------------------|---------------------|------------------|---------------------------------|-------|---------|---------|
| 4   | Fedora-VM1<br>4258  | Veth7<br>Adapter 2  | E1000<br>pvlan   | 0050.56bb.4fc1<br>10.104.249.49 | up    | 406     | Eth4/3  |

## VM vNIC View Field Description

The column headings in the VM vNIC view examples above is described in the following table:

| Column   | Description                                                                                                                                                                                                                                                                       |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mod      | Module number on which the VM resides.                                                                                                                                                                                                                                            |
| VM-Name  | VM name.                                                                                                                                                                                                                                                                          |
| HypvPort | Generated port ID in the hypervisor. For VMware hypervisor, it is called the dvPort ID.                                                                                                                                                                                           |
| VethPort | vEthernet interface number in the Cisco Nexus 1000V switch.                                                                                                                                                                                                                       |
| Adapter  | Network adapter number of the vEthernet interface.                                                                                                                                                                                                                                |
| Drv Type | Driver type of the network adapter. Supported values are as follows: <ul style="list-style-type: none"> <li>• E1000</li> <li>• E1000e</li> <li>• PCNet32</li> <li>• Vmxnet2</li> <li>• Vmxnet3</li> </ul>                                                                         |
| Mode     | Interface modes. Supported values are as follows: <ul style="list-style-type: none"> <li>• access—Access port/Virtual Extensible Local Area Network (VXLAN) port</li> <li>• trunk—Trunk port</li> <li>• pvlan—Private VLAN (PVLAN) host mode or pvlan promiscuous mode</li> </ul> |
| Mac-Addr | MAC address of the network adapter.                                                                                                                                                                                                                                               |
| IP-Addr  | IPv4 address of the network adapter, if the VMware tools are installed on the OS.                                                                                                                                                                                                 |
| State    | Operational status of the network adapter.                                                                                                                                                                                                                                        |

| Column  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network | <p>Network interface ID. Supported values are as follows:</p> <ul style="list-style-type: none"> <li>• access vlan—Access interface</li> <li>• trunk interface—Native VLAN</li> <li>• vxlan interface—Segment ID</li> <li>• pvlan interface—Promiscuous - primary VLAN; Isolated - secondary VLAN; Community-secondary VLAN</li> </ul> <p><b>Note</b> To know the interface type, refer the Mode value.</p>                                                                                                                                                                                                  |
| Pinning | <ul style="list-style-type: none"> <li>• For LACP or static port-channels, pinning columns only display the port-channel number. The link the VM traffic travels depends upon the hashing algorithm the port-channel is using.</li> <li>• For a vPC CDP/Manual/MAC Pinning port-channel, each vEthernet port is pinned to a sub-group of the port-channel. The sub-group corresponds to an Ethernet or its uplink interface. This column shows the Ethernet port members of the sub-group.</li> <li>• If the Ethernet ports are not part of the port channel in any module, this column is blank.</li> </ul> |

## Displaying the VM Info View

To display the VM Info view, follow the given step.

### Procedure

**show vtracker vm-view info** [**module number** | **vm name**]

**Note** The timeout for this command is 180 seconds.

The following examples show the vTracker VM Info view in a VSM:

### Example:

```
switch(config)# show vtracker vm-view info
Module 4:
 VM Name: Fedora-VM1
 Guest Os: Other Linux (32-bit)
 Power State: Powered On
 VM Uuid: 421871bd-425e-c484-d868-1f65f4f1bc50
```



```

Virtual CPU Allocated: 1
CPU Usage: 1 %
Memory Allocated: 256 MB
Memory Usage: 1 %
VM FT State: Unknown
Tools Running status: Not Running
Tools Version status: not installed
Data Store: NFS1_4
VM Uptime: 1 day 29 minutes 46 seconds

VM Name: Fedora-VM2
Guest Os: Other Linux (32-bit)
Power State: Powered On
VM Uuid: 4218ab37-d56d-63e4-3b00-77849401071e
Virtual CPU Allocated: 1
CPU Usage: 1 %
Memory Allocated: 256 MB
Memory Usage: 1 %
VM FT State: Unknown
Tools Running status: Not Running
Tools Version status: not installed
Data Store: NFS1_4
VM Uptime: 58 minutes 30 seconds

```

```

Module 5:
VM Name: gentoo-cluster2-1
Guest Os: Other (64-bit)
Power State: Powered Off
VM Uuid: 4235edf5-1553-650f-ade8-39565ee3cd57
Virtual CPU Allocated: 1
CPU Usage: 0 %
Memory Allocated: 512 MB
Memory Usage: 0 %
VM FT State: Unknown
Tools Running status: Not Running
Tools Version status: not installed
Data Store: datastore1 (2)
VM Uptime: n/a

```

**Example:**

```

switch(config)# show vtracker vm-view info vm Fedora-VM1
Module 4:
VM Name: Fedora-VM1
Guest Os: Other Linux (32-bit)
Power State: Powered On
VM Uuid: 421871bd-425e-c484-d868-1f65f4f1bc50
Virtual CPU Allocated: 1
CPU Usage: 1 %
Memory Allocated: 256 MB
Memory Usage: 1 %
VM FT State: Unknown
Tools Running status: Not Running
Tools Version status: not installed
Data Store: NFS1_4
VM Uptime: 1 day 29 minutes 46 seconds

```

## VM Info View Field Description

The column headings in the VM Info view examples above is described in the following table:

| Column  | Description                            |
|---------|----------------------------------------|
| Module  | Module number on which the VM resides. |
| VM Name | VM name.                               |

| Column                | Description                                                                                                                                                                                                                 |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Guest OS              | Guest operating system name, which is running on the VM.                                                                                                                                                                    |
| Power State           | Operational state of the VM. Supported status values are as follows: <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Powered On</li> <li>• Powered Off</li> <li>• Suspended</li> <li>• Non Available</li> </ul> |
| VM Uuid               | UUID of the VM.                                                                                                                                                                                                             |
| Virtual CPU Allocated | Number of the virtual CPUs allocated for the VM.                                                                                                                                                                            |
| CPU Usage             | VM usage in percentage.                                                                                                                                                                                                     |
| Memory Allocated      | Memory allocated to the VM in megabytes.                                                                                                                                                                                    |
| Memory Usage          | VM memory usage in percentage.                                                                                                                                                                                              |
| VM FT State           | Fault tolerance state of the VM. Supported values are as follows: <ul style="list-style-type: none"> <li>• Unknown</li> <li>• FT Primary</li> <li>• FT Secondary</li> <li>• Not Available</li> </ul>                        |
| Tools Running status  | VMware tools running status. Supported values are as follows: <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Starting</li> <li>• Running</li> <li>• Not Running</li> <li>• Not Available</li> </ul>            |

| Column               | Description                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tools Version status | VMware tools that display the version status. Supported values are as follows: <ul style="list-style-type: none"> <li>• Unknown</li> <li>• Current</li> <li>• Need Upgrade</li> <li>• Not Installed</li> <li>• Unmanaged</li> <li>• Blacklisted</li> <li>• Supported New</li> <li>• Supported Old</li> <li>• Too New</li> <li>• Too Old</li> <li>• Not Available</li> </ul> |
| Data Store           | Data store name on which the VM resides.                                                                                                                                                                                                                                                                                                                                    |
| VM Uptime            | How long the VM has been running.                                                                                                                                                                                                                                                                                                                                           |

## Module pNIC View

### Module pNIC View Overview

The Module pNIC View provides information about the physical network interface cards (pNICs) that are connected to each of the VEM server module in the network.

### Displaying the Module pNIC View

To display the Module pNIC view, follow the given step.

#### Procedure

```
show vtracker module-view pnic [module number]
```

The following examples show the vTracker Module pNIC view in a VSM:

#### Example:

```
switch(config)# show vtracker module-view pnic
```

-----

```

Mod EthIf Adapter Mac-Address Driver DriverVer FwVer

3 Eth3/8 vmnic7 0050.5652.f935 igb 2.1.11.1 1.4-3
 Intel Corporation 82576 Gigabit Network Connection

4 Eth4/3 vmnic2 0050.565e.df74 e1000 8.0.3.2-1vmw-NAPI N/A
 Intel Corporation 82546GB Gigabit Ethernet Controller

4 Eth4/4 vmnic3 0050.565e.df75 e1000 8.0.3.2-1vmw-NAPI N/A
 Intel Corporation 82546GB Gigabit Ethernet Controller

```

**Example:**

```
switch(config)# show vtracker module-view pnic module 3
```

```

Mod EthIf Adapter Mac-Address Driver DriverVer FwVer

3 Eth3/8 vmnic7 0050.5652.f935 igb 2.1.11.1 1.4-3
 Intel Corporation 82576 Gigabit Network Connection

```

## Module pNIC View Field Description

The column headings in the Module pNIC view examples above is described in the following table:

| Column      | Description                                      |
|-------------|--------------------------------------------------|
| Mod         | Server module name on which the VM resides.      |
| EthIf       | Ethernet interface ID of the server module.      |
| Adapter     | Ethernet adapter name as seen by the Hypervisor. |
| Description | Manufacturer name of the above adapter.          |
| Mac-Address | MAC address of the Ethernet interface.           |
| Driver      | Driver type of the interface.                    |
| DriverVer   | Driver version of the interface.                 |
| FwVer       | Firmware version of the interface.               |

## VLAN View

### VLAN View Overview

The VLAN view provides information about all the VMs that are connected to a specific VLAN or a range of VLANs. It is a view from the VLAN perspective.

## Displaying the VLAN View

To display the VLAN view, follow the given step.

### Procedure

**show vtracker vlan-view vnic [vlan number/range]**

The following examples show the vTracker VLAN view in a VSM:

#### Example:

```
switch(config)# show vtracker vlan-view
* R = Regular Vlan, P = Primary Vlan, C = Community Vlan
 I = Isolated Vlan, U = Invalid
```

| VLAN | Type | VethPort | VM Name    | Adapter Name  | Mod |
|------|------|----------|------------|---------------|-----|
| 1    | R    | -        | -          | -             | -   |
| 233  | R    | -        | -          | -             | -   |
| 335  | R    | -        | -          | -             | -   |
| 336  | R    | -        | -          | -             | -   |
| 337  | R    | -        | -          | -             | -   |
| 338  | R    | -        | -          | -             | -   |
| 339  | R    | Veth3    | gentoo-2   | Net Adapter 3 | 3   |
|      |      | Veth4    | gentoo-2   | Net Adapter 4 | 3   |
|      |      | Veth5    | gentoo-2   | Net Adapter 2 | 3   |
| 340  | R    | -        | -          | -             | -   |
| 341  | R    | -        | -          | -             | -   |
| 400  | R    | Veth1    | Fedora-VM2 | Net Adapter 1 | 5   |
| 401  | R    | Veth1    | Fedora-VM2 | Net Adapter 1 | 5   |
| 402  | R    | Veth1    | Fedora-VM2 | Net Adapter 1 | 5   |
| 403  | R    | -        | -          | -             | -   |
| 404  | P    | Veth6    | Fedora-VM1 | Net Adapter 1 | 4   |
| 405  | C    | Veth2    | Fedora-VM2 | Net Adapter 3 | 5   |
| 406  | I    | Veth7    | Fedora-VM1 | Net Adapter 2 | 4   |

#### Example:

```
switch(config)# show vtracker vlan-view vlan 233-340
* R = Regular Vlan, P = Primary Vlan, C = Community Vlan
 I = Isolated Vlan, U = Invalid
```

| VLAN | Type | VethPort | VM Name  | Adapter Name  | Mod |
|------|------|----------|----------|---------------|-----|
| 233  | R    | -        | -        | -             | -   |
| 335  | R    | -        | -        | -             | -   |
| 336  | R    | -        | -        | -             | -   |
| 337  | R    | -        | -        | -             | -   |
| 338  | R    | -        | -        | -             | -   |
| 339  | R    | Veth3    | gentoo-2 | Net Adapter 3 | 3   |
|      |      | Veth4    | gentoo-2 | Net Adapter 4 | 3   |
|      |      | Veth5    | gentoo-2 | Net Adapter 2 | 3   |
| 340  | R    | -        | -        | -             | -   |

## VLAN View Field Description

The column headings in the VLAN view examples above is described in the following table:

| Column       | Description                                                                                                                                                                                                              |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VLAN         | VLAN ID on which the VM resides.                                                                                                                                                                                         |
| Type         | VLAN type. Supported types are as follows: <ul style="list-style-type: none"> <li>• R—Regular VLAN</li> <li>• P—Primary VLAN</li> <li>• C—Community VLAN</li> <li>• I—Isolated VLAN</li> <li>• U—Invalid VLAN</li> </ul> |
| VethPort     | vEthernet interface port number used by the VLAN.                                                                                                                                                                        |
| VM Name      | VM name of the interface.                                                                                                                                                                                                |
| Adapter Name | Adapter name of the interface.                                                                                                                                                                                           |
| Mod          | Module number on which the interface resides.                                                                                                                                                                            |

## VMotion View

### VMotion View Overview

The vMotion view provides information about all the ongoing (if any) as well as previous VM migration events. However, only VMs that are currently being managed by the Cisco Nexus 1000V switch are displayed in the output.



#### Note

The VSM must be connected with the vCenter in order to generate the required VMotion view output. You can enter the **show svcs connections** command on the VSM to verify the connection.

### Displaying the VMotion View

To display the VMotion view, follow the given step.

#### Procedure

**show vtracker vmotion-view** [**now** | **last number 1-100**]

**Note** The timeout for this command is 180 seconds.

The following examples show the vTracker VMotion view in a VSM:

**Example:**

```
switch(config)# show vtracker vmotion-view last 20
Note: Command execution is in progress...
```

Note: VM Migration events are shown only for VMs currently managed by Nexus 1000v.

\* '-' = Module is offline or no longer attached to Nexus1000v DVS

```

VM-Name Src Dst Start-Time Completion-Time
 Mod Mod

rk-ubt-1-0046 6 4 Mon Sep 3 10:42:27 2012 OnGoing
rk-ubt-1-0045 6 4 Mon Sep 3 10:42:27 2012 OnGoing
rk-ubt-1-0031 6 4 Mon Sep 3 10:42:27 2012 Mon Sep 3 10:44:10 2012
rk-ubt-1-0021 6 4 Mon Sep 3 10:42:27 2012 Mon Sep 3 10:43:42 2012
rk-ubt-1-0023 6 3 Thu Aug 16 14:25:26 2012 Thu Aug 16 14:27:55 2012
rk-ubt-1-0029 6 3 Thu Aug 16 14:25:26 2012 Thu Aug 16 14:27:50 2012
rk-ubt-1-0024 6 3 Thu Aug 16 14:25:26 2012 Thu Aug 16 14:26:13 2012
rk-ubt-1-0025 6 3 Thu Aug 16 14:25:26 2012 Thu Aug 16 14:26:12 2012
rk-ubt-1-0026 6 3 Thu Aug 16 14:25:26 2012 Thu Aug 16 14:26:09 2012
RHEL-Tool-VmServer - 3 Wed Aug 8 12:57:48 2012 Wed Aug 8 12:58:37 2012

```

**Example:**

```
switch(config)# show vtracker vmotion-view now
Note: Command execution is in progress...
```

\*Note: VM Migration events are shown only for VMs currently managed by Nexus 1000v.

\* '-' = Module is offline or no longer attached to Nexus1000v DVS

```

VM-Name Src Dst Start-Time Completion-Time
 Mod Mod

rk-ubt-1-0046 6 4 Mon Sep 3 10:42:27 2012 OnGoing
rk-ubt-1-0045 6 4 Mon Sep 3 10:42:27 2012 OnGoing

```

## VMotion View Field Description

The column headings in the VMotion view examples above is described in the following table:

| Column  | Description                                 |
|---------|---------------------------------------------|
| VM-Name | VM name.                                    |
| Src Mod | Source module number of the migration.      |
| Dst Mod | Destination module number of the migration. |

| Column          | Description                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------|
| Start-Time      | Migration start time per the time zone defined in the Virtual Supervisor Module (VSM).                |
| Completion-Time | Migration completion time in VSM time zone. For migration in progress, the status shows as “OnGoing.” |

## Feature History for vTracker

| Feature Name   | Releases     | Feature Information          |
|----------------|--------------|------------------------------|
| vTracker Views | 4.2(1)SV2(1) | This feature was introduced. |





# CHAPTER 17

## Configuring Virtualized Workload Mobility

---

This chapter contains the following sections:

- [Information About Virtualized Workload Mobility \(DC to DC vMotion\)](#), page 217
- [Prerequisites for Virtualized Workload Mobility \(DC to DC vMotion\)](#), page 218
- [Guidelines and Limitations](#), page 218
- [Migrating a VSM](#), page 219
- [Verifying and Monitoring the Virtualized Workload Mobility \(DC to DC vMotion\) Configuration](#), page 220
- [Feature History for Virtualized Workload Mobility \(DC to DC vMotion\)](#), page 221

### 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
- Split Cluster

#### 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.

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.

## 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.

## 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.

**Note**

---

If a VEM loses the connection to its VSM, the Vmotions to that particular VEM are blocked. The VEM shows up in the VCenter Server as having a degraded (yellow) status.

---

## 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 **showshow** 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).

## Migrating a VSM Hosted on an ESX

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:

**Note**

For information on vMotion or storage vMotion, see the VMware documentation.

### 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.

### Procedure

- 
- Step 1** Migrate the standby VSM to the backup site.
- Step 2** Perform a storage vMotion for the standby VSM storage.
- Step 3** `switch#system switchover`  
Initiates a 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 and Monitoring the Virtualized Workload Mobility (DC to DC vMotion) Configuration

Refer to the following section for verifying and monitoring the Virtualized Workload Mobility (DC to DC vMotion) configuration:

### Procedure

`switch#show module`

Displays the virtualized workload mobility (DC to DC vMotion) configuration.

## 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. |





## INDEX

### A

- accessing netflow data [119](#)
- adding hardware NICs to DVS [167](#)
- advertising CDP version [9](#)

### B

- binding VMkernel NICs [165](#)

### C

- CDP [8, 15](#)
  - default settings [8](#)
  - feature history [15](#)
  - guidelines [8](#)
- CDP configuration [15](#)
- CDP global configuration [9](#)
- CDP interface configuration [12](#)
- changed information [1](#)
  - description [1](#)
- changing access vlan [169](#)
- changing VMkernel NIC access vlan [168](#)
- characteristics of ERSPAN destinations [78](#)
- characteristics of local SPAN destinations [78](#)
- characteristics of SPAN sources [78](#)
- clearing [14](#)
  - CDP statistics [14](#)
- clearing NTP session [74](#)
- Configuration Management [43](#)
- configuration to enable SPAN monitoring [97](#)
- configuring [19, 82, 85](#)
  - domain [19](#)
  - ERSPAN port profile [85](#)
  - local SPAN session [82](#)
- configuring CDP [9](#)
- configuring domain [18](#)
  - guideline [18](#)
- configuring ERSPAN flow IDs [94](#)
- configuring ERSPAN session [87](#)

- configuring iSCSI multipath [158](#)
- configuring SNMP [103](#)
- configuring SPAN [82](#)
- configuring SPAN session [96](#)
- configuring system message logging [143](#)
- configuring VSM backup and recovery [176](#)
- converting to hardware iSCSI configuration [166](#)
- creating [19, 24, 26, 27](#)
  - control vlan [26](#)
  - domain [19](#)
  - packet vlan [27](#)
  - port profile for layer 3 control [24](#)
- creating VMkernel NICs [161](#)
  - attaching port profile [161](#)

### D

- default settings [19, 201](#)
  - vTracker [201](#)
- displaying [204, 206, 208, 211, 213, 214](#)
  - upstream view [204](#)
- displaying log files [148](#)

### E

- enabling [201](#)
  - globally [201](#)
  - vTracker [201](#)
- enabling CDP globally [9](#)
- enabling CDP on interface [12](#)
- encapsulated remote SPAN [79](#)
- example for ERSPAN session [95](#)

### F

- feature history [28](#)
  - VSM domain [28](#)
- field description [204, 207, 209, 212, 213, 215](#)

Flow [115](#)  
 about [115](#)  
 flow record [117](#)

## H

high availability [8](#)  
 history [216](#)

## I

identifying iSCSI adapters [164](#)  
 on host server [164](#)  
 on vSphere client [164](#)  
 information [200](#)  
 vTracker [200](#)  
 information about CDP [7](#)  
 information about domain [17](#)  
 information about files [53](#)  
 information about iSCSI multipath [153](#)  
 information about NTP [71](#)  
 information about SPAN and ERSPAN [77](#)  
 iSCSI [157](#)  
 guidelines [157](#)  
 iSCSI multipath [154, 158, 174](#)  
 default settings [158](#)  
 feature history [174](#)  
 overview [154](#)  
 prerequisites [158](#)  
 related documents [174](#)  
 iSCSI multipath setup [155](#)

## L

layer 2 transport [23](#)  
 layer 3 control [17](#)  
 layer 3 transport [21](#)  
 local SPAN [79](#)  
 log files [148](#)  
 displaying [148](#)

## M

managing [172](#)  
 storage loss detection [172](#)  
 manually pinning the NICs [162](#)  
 module pnic view [211, 212](#)  
 monitoring CDP [14](#)

## N

NetFlow [115](#)  
 about [115](#)  
 network analysis module [80](#)  
 new and changed [1](#)  
 new information [1](#)  
 description [1](#)  
 NTP [73](#)  
 guidelines [73](#)  
 NTP peers [72](#)

## O

overview [203, 205, 211, 212, 214](#)

## P

port profile for VMkernel NIC [159](#)  
 process [168](#)  
 changing access vlan [168](#)

## R

removing binding [167](#)  
 software iSCSI adapter [167](#)  
 resuming SPAN session [92, 93](#)  
 global configuration mode [92](#)

## S

server connections [31](#)  
 shutting down SPAN session [90, 91](#)  
 monitor configuration [91](#)  
 SPAN [81, 82](#)  
 default settings [82](#)  
 guidelines [81](#)  
 SPAN and ERSPAN [97](#)  
 feature history [97](#)  
 SPAN destinations [78](#)  
 SPAN sessions [80](#)  
 SPAN sources [77](#)  
 supported iSCSI adapters [154](#)

## U

uplink pinning and storage binding [159](#)  
 upstream view [203, 204](#)



user management [69](#)

## V

verifying [95](#), [171](#)

    iSCSI multipath configuration [171](#)

    SPAN configuration [95](#)

verifying CDP configuration [14](#)

virtual machine view [205](#)

virtualized workload mobility [217](#)

vlan view [212](#), [213](#)

vm info view [208](#), [209](#)

vm vnic view [206](#), [207](#)

vmotion view [214](#), [215](#)

vtracker [216](#)

vTracker [201](#)

    guidelines [201](#)

    limitations [201](#)

