



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



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

April 15, 2011

Americas Headquarters

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

Text Part Number: OL-22825-01

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

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

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

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

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

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Internet Protocol (IP) addresses that are used in the examples, command display output, and figures within this document are for illustration only. If an actual IP address appears in this document, it is coincidental.

Cisco Nexus 1000V Troubleshooting Guide, Release 4.0(4)SV1(3a)
© 2010 Cisco Systems, Inc. All rights reserved.



New and Changed Information

This section describes the information in this document that is either new or has changed with each release.

To find additional information about new features or command changes, see the following:

- [Release Notes](#).
- [Command Reference](#).

New or changed content	Description	Changed in release	Where Documented
vemcmd set port-mode command	Restore lost connectivity due to incorrect port setting.	4.0(4)SV1(3a)	“Checking the VEM Configuration” section on page 6-13
VEM health check	The VEM health check shows the cause of a connectivity problem and recommends next steps in troubleshooting.	4.0(4)SV1(3)	“Verifying Network Connectivity Between the VSM and the VEM” section on page 6-11
VSM connection failure after ESX reboot	The section describes how to prevent loss of connectivity related to an MTU mismatch following a reboot of the ESX.	4.0(4)SV1(3)	“Connection Failure After ESX Reboot” section on page 18-6

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



CONTENTS

New and Changed Information iii

Preface xiii

Audience	xiii
Document Organization	xiii
Document Conventions	xiv
Available Documents	xv
Obtaining Documentation and Submitting a Service Request	xvi

Overview of Troubleshooting 1-1

Overview of the Troubleshooting Process	1-1
Overview of Best Practices	1-1
Troubleshooting Basics	1-2
Troubleshooting Guidelines	1-2
Gathering Information	1-2
Verifying Ports	1-3
Verifying Layer 2 Connectivity	1-3
Verifying Layer 3 Connectivity	1-3
Overview of Symptoms	1-4
System Messages	1-4
System Message Text	1-4
Syslog Server Implementation	1-5
Troubleshooting with Logs	1-6
Viewing Logs	1-6
Contacting Cisco or VMware Customer Support	1-7

Tools Used in Troubleshooting 2-1

Commands	2-1
Ping	2-1
Traceroute	2-2
Monitoring Processes and CPUs	2-2
Identifying the Processes Running and their States	2-2

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

Displaying CPU Utilization	2-3
Displaying CPU and Memory Information	2-4
RADIUS	2-4
Syslog	2-5
Logging Levels	2-5
Enabling Logging for Telnet or SSH	2-6
Installation	3-1
Isolating Installation Problems	3-1
Verifying Your VMware License Version	3-1
Host is Not Visible from Distributed Virtual Switch	3-2
Refreshing the vCenter Server Connection	3-3
Improving Performance	3-4
Verifying the Domain Configuration	3-4
Verifying the Port Group Assignments for a VSM VM Virtual Interface	3-4
Verifying VSM and vCenter Server Connectivity	3-5
Troubleshooting Connections to a vCenter Server	3-5
Recovering the Network Administrator Password	3-6
Managing Extension Keys	3-6
Known Extension Problems and Resolutions	3-7
Resolving a Plug-In Conflict	3-7
Finding the Extension Key on the Cisco Nexus 1000V	3-7
Finding the Extension Key Tied to a Specific DVS	3-8
Verifying Extension Keys	3-9
Recreating the Cisco Nexus 1000V Installation	3-10
Removing the Hosts from the Cisco Nexus 1000V DVS	3-11
Removing the Cisco Nexus 1000V From the vCenter Server	3-11
Unregistering the Extension Key in the vCenter Server	3-12
Upgrade	4-1
Identifying Upgrade Solutions	4-1
Verifying Bundle IDs	4-3
Verifying that the correct VEM Software Version is Installed	4-3
Verifying that the VEM is Running on the Host	4-3
Verifying the VMware NICs on a DVS	4-3
Verifying the VMware Patches Installed on a Host	4-4
Verifying Upgrade Status	4-4

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

Licensing 5-1

- Licensing Overview 5-1
- Troubleshooting Unlicensed Modules 5-2
- Troubleshooting License Installation Issues 5-3
 - License Troubleshooting Checklist 5-3
 - Contents of the License File 5-3
 - Removing an Evaluation License File 5-3
- Determining License Usage 5-3
- Installed License Information 5-4
- Troubleshooting Post License Installation Problems 5-4
- Troubleshooting the Removal of a License 5-5

VSM and VEM Modules 6-1

- Information About Modules 6-1
- Troubleshooting a Module Not Coming Up on the VSM 6-1
 - Guidelines for Troubleshooting Modules 6-2
 - Flow Chart for Troubleshooting Modules 6-3
- Problems with the VSM 6-4
 - Verifying the VSM Is Connected to the vCenter Server 6-6
 - Verifying the VSM Is Configured Correctly 6-7
 - Verifying the Domain Configuration 6-7
 - Verifying the System Port Profile Configuration 6-8
 - Verifying the Control and Packet VLAN Configuration 6-9
 - Checking the vCenter Server Configuration 6-10
 - Verifying Network Connectivity Between the VSM and the VEM 6-11
 - Checking the VEM Configuration 6-13
 - Collecting Logs 6-15
- Commands 6-16

Ports and Port Profiles 7-1

- Overview 7-1
- Guidelines for Configuring a Port Interface 7-2
 - Verifying the Module State 7-2
- Diagnostic Checklist 7-2
- Viewing the Port State 7-3
- Using Port Counters 7-4
- Port Interface Symptoms and Solutions 7-5

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

Cannot Enable an Interface	7-5
Port Remains in a Link Failure or Not Connected State	7-6
Link Flapping	7-6
About the Link Flapping Cycle	7-6
Troubleshooting Prerequisites	7-6
Symptoms, Causes, and Solutions	7-7
Port State Is ErrDisabled	7-7
About the ErrDisabled Port State	7-7
Verifying the ErrDisable State	7-7
MTU Configuration Failure on ESX	7-8
Port Security	7-8
Troubleshooting Port Security Problems	7-9
Cannot Ping from a VM with Port Security Enabled	7-9
Port Enabled with Port Security is Error Disabled	7-11
Port Security Restrictions and Limitations	7-12
Collecting Debugging Output for Port Security	7-12
Symptoms, Causes, and Solutions	7-13
Port Profiles	7-13
Troubleshooting Commands for Port Profiles	7-14
System Port Profiles	7-18
Port Profiles Symptoms and Solutions	7-19
Transferring Port Profiles from the VSM to the vCenter Server	7-19
Port Channels and Trunking	8-1
Overview	8-1
Port Channel Overview	8-1
Port Channel Restriction	8-2
Trunking Overview	8-2
Initial Troubleshooting Checklist	8-2
Troubleshooting Asymmetric Port Channels	8-3
Cannot Create Port Channel	8-4
Newly Added Interface Does Not Come Online In a Port Channel	8-4
Forcing Port Channel Characteristics onto an Interface	8-4
Verifying a Port Channel Configuration	8-6
VLAN Traffic Does Not Traverse Trunk	8-6

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

Layer 2 Switching	9-1
Information About Layer 2 Ethernet Switching	9-1
Port Model	9-1
Viewing Ports from the VEM	9-2
Viewing Ports from the VSM	9-3
Port Types	9-4
Layer 2 Switching Problems	9-4
Verifying a Connection Between VEM Ports	9-4
Verifying a Connection Between VEMs	9-5
Isolating Traffic Interruptions	9-6
Verifying Layer 2 Switching	9-7
ACLs	10-1
About Access Control Lists (ACLs)	10-1
ACL Configuration Limits	10-1
ACL Restrictions	10-2
Troubleshooting ACLs	10-2
Displaying ACL Policies on the VEM	10-2
Debugging Policy Verification Issues	10-3
Quality of Service	11-1
Information About Quality of Service	11-1
QoS Configuration Limits	11-1
QoS Troubleshooting Commands	11-2
Troubleshooting the VEM	11-2
Debugging Policing Verification Errors	11-3
NetFlow	12-1
Information About NetFlow	12-1
NetFlow Troubleshooting Commands	12-2
Common NetFlow Problems	12-3
Debugging a Policy Verification Error	12-3
Debugging Statistics Export	12-3
VLANs	13-1
Information About VLANs	13-1
Initial Troubleshooting Checklist	13-2
Cannot Create a VLAN	13-3

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

Private VLANs 14-1

- Information About Private VLANs 14-1
 - Private VLAN Domain 14-1
 - Spanning Multiple Switches 14-1
 - Private VLAN Ports 14-2
- Troubleshooting Guidelines 14-2
- Private VLAN Troubleshooting Commands 14-2

Multicast IGMP 15-1

- Information About Multicast 15-1
 - Multicast IGMP Snooping 15-1
- Problems with Multicast IGMP Snooping 15-2
 - Troubleshooting Guidelines 15-2
 - Troubleshooting Commands 15-2
 - Symptoms, Causes, and Solutions 15-4

SPAN 16-1

- Information About SPAN 16-1
 - SPAN Sources 16-1
 - Source Ports 16-2
 - SPAN Destinations 16-2
 - Destination Ports 16-2
 - ERSPAN Destinations 16-2
 - SPAN Sessions 16-2
- Troubleshooting SPAN Problems 16-3
 - Local SPAN Session Problems 16-3
 - Troubleshooting Commands 16-3
 - Problems and Solutions 16-4
 - Examples 16-4

High Availability 17-1

- Information About High Availability 17-1
 - System-Level High Availability 17-2
 - Single or Dual Supervisors 17-2
 - Network-Level High Availability 17-2
- Problems with High Availability 17-3
- High Availability Troubleshooting Commands 17-6
 - Displaying Cores 17-6

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

Displaying Logs of Processes	17-6
Checking Redundancy	17-7
Checking System Redundancy Status	17-7
Checking System Internal Redundancy Status	17-7
Checking the System Manager State	17-8
Reloading a Module	17-9
Attaching to the Standby VSM Console	17-10

System 18-1

Information About the System	18-1
General Restrictions for vCenter Server	18-2
Extension Key	18-2
Recovering a DVS	18-2
Recovering a DVS With a Saved Copy of the VSM	18-3
Recovering a DVS Without a Saved Copy of the VSM	18-4
Problems Related to VSM and vCenter Server Connectivity	18-5
Connection Failure After ESX Reboot	18-6
Information About Changing the MTU	18-6
Verifying that an MTU Mismatch Caused a Connection Failure	18-7
Restoring Network Connectivity to the VEM	18-8
Displaying VC Event History	18-8
VSM Creation	18-9
Port Profiles	18-9
Problems with Port Profiles	18-9
Problems with Hosts	18-10
Problems with VM Traffic	18-10
VEM Troubleshooting Commands	18-10
VEM Log Commands	18-11
Error Messages	18-12

Before Contacting Technical Support 19-1

Gathering Information for Technical Support	19-1
Obtaining a File of Core Memory Information	19-2
Copying Files	19-3

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



Preface

The Troubleshooting document provides information about how to recognize a problem, determine its cause, and find possible solutions.

This preface describes the following aspects of this document:

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

Audience

This publication is for experienced network administrators who configure and maintain a Cisco Nexus 1000V.

Document Organization

This document is organized into the following chapters:

Title	Description
Chapter 1, “Overview of Troubleshooting”	Describes basic troubleshooting information.
Chapter 2, “Tools Used in Troubleshooting”	Describes the available troubleshooting tools.
Chapter 3, “Installation”	Describes how to troubleshoot installation problems.
Chapter 5, “Licensing”	Describes how to identify and resolve problems related to licensing for the Cisco Nexus 1000V.
Chapter 6, “VSM and VEM Modules”	Describes how to identify and resolve problems with modules.
Chapter 7, “Ports and Port Profiles”	Describes how to identify and resolve problems with ports.
Chapter 8, “Port Channels and Trunking”	Describes how to identify and resolve problems with port channels and trunks.

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

Title	Description
Chapter 9, “Layer 2 Switching”	Describes how to identify and resolve problems with Layer 2 switching.
Chapter 10, “ACLs”	Describes how to identify and resolve problems with ACLs.
Chapter 11, “Quality of Service”	Describes how to identify and resolve problems related to Quality of Service (QoS).
Chapter 12, “NetFlow”	Describes how to identify and resolve problems with NetFlow.
Chapter 13, “VLANs”	Describes how to identify and resolve problems with VLANs.
Chapter 14, “Private VLANs”	Describes how to identify and resolve problems related to private VLANs.
Chapter 15, “Multicast IGMP”	Describes how to identify and resolve problems with multicast.
Chapter 16, “SPAN”	Describes how to identify and resolve problems with SPAN.
Chapter 17, “High Availability”	describes how to identify and resolve problems related to High Availability (HA).
Chapter 18, “System”	Describes how to identify and resolve problems with VMware.
Chapter 19, “Before Contacting Technical Support”	Describes the steps to take before requesting technical support.

Document Conventions

Command descriptions use these conventions:

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

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information that the switch displays are in screen font.
<code>boldface screen font</code>	Information that you must enter is in boldface screen font.
<i><code>italic screen font</code></i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.

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

[]	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.

Available Documents

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

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

General Information

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

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

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

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

Install and Upgrade

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

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

Cisco Nexus 1000V VEM Software Installation and Upgrade Guide, Release 4.0(4)SV1(3a)

Cisco Nexus 1010 Virtual Services Appliance Installation Guide

Configuration Guides

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

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

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

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

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

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

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

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

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

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

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

Programming Guide

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

Reference Guides

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

Cisco Nexus 1000V MIB Quick Reference

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

Troubleshooting and Alerts

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

Cisco Nexus 1000V Password Recovery Guide

Cisco NX-OS System Messages Reference

Network Analysis Module Documentation

Cisco Network Analysis Module Software Documentation Guide, 4.2

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

Network Analysis Module Command Reference Guide, 4.2

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

Cisco Network Analysis Module Software Release Notes, 4.2

Obtaining Documentation and Submitting a Service Request

For information 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

Overview of Troubleshooting

This chapter introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that may occur when configuring and using Cisco Nexus 1000V.

This chapter includes the following sections:

- [Overview of the Troubleshooting Process, page 1-1](#)
- [Overview of Best Practices, page 1-1](#)
- [Troubleshooting Basics, page 1-2](#)
- [Overview of Symptoms, page 1-4](#)
- [Overview of Symptoms, page 1-4](#)
- [System Messages, page 1-4](#)
- [Troubleshooting with Logs, page 1-6](#)
- [Contacting Cisco or VMware Customer Support, page 1-7](#)

Overview of the Troubleshooting Process

To troubleshoot your network, follow these general steps:

-
- | | |
|---------------|--|
| Step 1 | Gather information that defines the specific symptoms. |
| Step 2 | Identify all potential problems that could be causing the symptoms. |
| Step 3 | Systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear. |
-

Overview of Best Practices

Best practices are the recommended steps you should take to ensure the proper operation of your network. We recommend the following general best practices for most networks:

- Maintain a consistent Cisco Nexus 1000V release across all network devices.
- Refer to the release notes for your Cisco Nexus 1000V release for the latest features, limitations, and caveats.

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

- Enable system message logging. See the “Overview of Symptoms” section on page 1-4.
- Verify and troubleshoot any new configuration changes after implementing the change.

Troubleshooting Basics

This section introduces questions to ask when troubleshooting a problem with Cisco Nexus 1000V or connected devices. Use the answers to these questions to identify the scope of the problem and to plan a course of action.

This section includes the following topics:

- [Troubleshooting Guidelines, page 1-2](#)
- [Gathering Information, page 1-2](#)
- [Verifying Ports, page 1-3](#)
- [Verifying Layer 2 Connectivity, page 1-3](#)
- [Verifying Layer 3 Connectivity, page 1-3](#)

Troubleshooting Guidelines

By answering the questions in the following subsections, you can determine the paths you need to follow and the components that you should investigate further.

Answer the following questions to determine the status of your installation:

- Is this a newly installed system or an existing installation? (It could be a new host, switch, or VLAN).
- Has the host ever been able to see the network?
- Are you trying to solve an existing application problem (too slow, too high latency, excessively long response time) or did the problem show up recently?
- What changed in the configuration or in the overall infrastructure immediately before the applications started to have problems?

To discover a network problem, use the following general network troubleshooting steps:

-
- Step 1** Gather information on problems in your system. See the “[Gathering Information](#)” section on page 1-2.
 - Step 2** Verify the layer 2 connectivity. See the “[Verifying Layer 2 Connectivity](#)” section on page 1-3.
 - Step 3** Verify the configuration for your end devices (storage subsystems and servers).
 - Step 4** Verify end-to-end connectivity. See the “[Verifying Layer 3 Connectivity](#)” section on page 1-3.
-

Gathering Information

This section highlights the tools that are commonly used to troubleshoot problems within your network. These tools are a subset of what you may use to troubleshoot your specific problem.

Each chapter in this guide may include additional tools and commands specific to the symptoms and possible problems covered in that chapter.

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

You should also have an accurate topology of your network to help isolate problem areas.

Issue the following commands and examine the outputs:

- **show module**
- **show version**
- **show running-config**
- **show logging log**
- **show interfaces brief**
- **show vlan**
- **show accounting log**
- **show tech support svcs**

**Note**

To issue commands with the **internal** keyword, you must log in with a network-admin role.

Verifying Ports

Answer the following questions to verify ports:

- Are you using the correct media copper or optical; fiber type.
- Is the media broken or damaged?
- Are you checking a virtual Ethernet port? If so, then use the **show interface brief** command. The status should be up.
- Are you checking a physical Ethernet port? If so, you need to check it by looking at the server, or by looking at an upstream switch.
- Check if the network adapters of the VSM VM are assigned the right port groups and if all of them are connected from the vSphere Client.

Verifying Layer 2 Connectivity

Answer the following questions to verify layer 2 connectivity:

- Are the necessary interfaces in the same VLANs?
- Are all ports in a port channel configured the same for speed, duplex, trunk mode?

Use the **show vlan brief** command. The status should be up.

Use the **show port-profile** command to check a port profile configuration?

Use the **show interface-brief** command to check the status of a virtual Ethernet port or a physical Ethernet port.

Verifying Layer 3 Connectivity

Answer the following questions to verify layer 3 connectivity:

- Have you configured a gateway of last resort?

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

- Are any IP access lists, filters, or route maps blocking route updates?

Use the **ping** or **trace** commands to verify connectivity. See the following for more information:

- [“Ping” section on page 2-1](#)
- [“Traceroute” section on page 2-2](#)

Overview of Symptoms

The symptom-based troubleshooting approach provides multiple ways to diagnose and resolve problems. By using multiple entry points with links to solutions, this guide best serves users who may have identical problems that are perceived by different indicators. Search this guide in PDF form, use the index, or rely on the symptoms and diagnostics listed in each chapter as entry points to access necessary information in an efficient manner.

Using a given a set of observable symptoms on a network, it is important to be able to diagnose and correct software configuration issues and inoperable hardware components so that the problems are resolved with minimal disruption to the network. Those problems and corrective actions include the following:

- Identify key Cisco Nexus 1000V troubleshooting tools.
- Obtain and analyze protocol traces using SPAN or Ethalyzer on the CLI.
- Identify or rule out physical port issues.
- Identify or rule out switch module issues.
- Diagnose and correct layer 2 issues.
- Diagnose and correct layer 3 issues.
- Obtain core dumps and other diagnostic data for use by the TAC.
- Recover from switch upgrade failures.

System Messages

The system software sends the syslog (system) messages to the console (and, optionally, to a logging server on another system) during operation. Not all messages indicate a problem with your system. Some messages are purely informational, while others might help diagnose problems with links, internal hardware, or the system software.

This section contains the following topics:

- [System Message Text, page 1-4](#)
- [Syslog Server Implementation, page 1-5](#)

System Message Text

Message-text is a text string that describes the condition. This portion of the message might contain detailed information about the event, including terminal port numbers, network addresses, or addresses that correspond to locations in the system memory address space. Because the information in these variable fields changes from message to message, it is represented here by short strings enclosed in square brackets ([]). A decimal number, for example, is represented as [dec].

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

```
2009 Apr 29 12:35:51 n1000v %KERN-1-SYSTEM_MSG: stun_set_domain_id : Setting domain ID
(1024) - kernel
```

Use this string to find the matching system message in the *Cisco NX-OS System Messages Reference System Messages Reference*.

Each system message is followed by an explanation and recommended action. The action may be as simple as “No action required.” It may involve a fix or a recommendation to contact technical support as shown in the following example:

```
Error Message 2009 Apr 29 14:57:23 n1000v %MODULE-5-MOD_OK: Module 3 is online
(serial: )
```

Explanation VEM module inserted successfully on slot 3.

Recommended Action None. This is an information message. Use "show module" to verify the module in slot 3.

Syslog Server Implementation

The syslog facility allows the Cisco Nexus 1000V device to send a copy of the message log to a host for more permanent storage. This can be useful if the logs need to be examined over a long period of time or when the Cisco Nexus 1000V device is not accessible.

This example demonstrates how to configure a Cisco Nexus 1000V device to use the syslog facility on a Solaris platform. Although a Solaris host is being used, syslog configuration on all UNIX and Linux systems is very similar.

Syslog uses the concept of a facility to determine how it should be handled on the syslog server (the Solaris system in this example), and the message severity. Therefore, different message severities can be handled differently by the syslog server. They could be logged to different files or e-mailed to a particular user. Specifying a severity determines that all messages of that level and greater severity (lower number) will be acted upon.



Note

The Cisco Nexus 1000V messages should be logged to a different file from the standard syslog file so that they cannot be confused with other non-Cisco syslog messages. The logfile should not be located on the / file system, to prevent log messages from filling up the / file system.

Syslog Client: switch1

Syslog Server: 172.22.36.211 (Solaris)

Syslog facility: local1

Syslog severity: notifications (level 5, the default)

File to log Cisco Nexus 1000V messages to: /var/adm/nxos_logs

To configure a syslog server, follow these steps:

Step 1 Configure the Cisco Nexus 1000V:

```
n1000v# config terminal
```

Enter configuration commands, one per line. End with CNTL/Z.

```
n1000v(config)# logging server 192.0.2.1 6 facility local1
```

To display the configuration:

```
n1000v# show logging server
```

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

```
Logging server: enabled
{192.0.2.1}
  server severity: notifications
  server facility: local1
```

Step 2 Configure the syslog server:

- a. Modify `/etc/syslog.conf` to handle `local1` messages. For Solaris, there needs to be at least one tab between the `facility.severity` and the action (`/var/adm/nxos_logs`).

```
#Below is for the NX-OS logging
local1.notice /var/adm/nxos_logs
```

- b. Create the log file.

```
#touch /var/adm/nxos_logs
```

- c. Restart syslog.

```
# /etc/init.d/syslog stop
# /etc/init.d/syslog start
syslog service starting.
```

- d. Verify syslog started.

```
# ps -ef |grep syslogd
root 23508 1 0 11:01:41 ? 0:00 /usr/sbin/syslogd
```

Step 3 Test the syslog server by creating an event in Cisco Nexus 1000V. In this case, port e1/2 was bounced and the following was listed on the syslog server. Notice that the IP address of the switch is listed in brackets.

```
# tail -f /var/adm/nxos_logs
Sep 17 11:07:41 [172.22.36.142.2.2] : 2004 Sep 17 11:17:29 pacific:
%PORT-5-IF_DOWN_INITIALIZING: %$VLAN 1%$ Interface e 1/2 is down (Initializing)
Sep 17 11:07:49 [172.22.36.142.2.2] : 2004 Sep 17 11:17:36 pacific: %PORT-5-IF_UP:
%$VLAN 1%$ Interface e 1/2 is up in mode access
Sep 17 11:07:51 [172.22.36.142.2.2] : 2004 Sep 17 11:17:39 pacific:
%VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(dhcp-171-71-49-125.cisco.com)
```

Troubleshooting with Logs

Cisco Nexus 1000V generates many types of system messages on the switch and sends them to a syslog server. These messages can be viewed to determine what events may have led up to the current problem condition you are facing.

Viewing Logs

Use the following commands to access and view logs in Cisco Nexus 1000V:

```
n1000v# show logging ?
```

```
console      Show console logging configuration
info         Show logging configuration
internal     syslog syslog internal information
last         Show last few lines of logfile
level        Show facility logging configuration
logfile      Show contents of logfile
```

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

```

loopback      Show logging loopback configuration
module        Show module logging configuration
monitor       Show monitor logging configuration
nvr           Show NVRAM log
pending       server address pending configuration
pending-diff  server address pending configuration diff
server        Show server logging configuration
session       Show logging session status
status        Show logging status
timestamp     Show logging timestamp configuration
|             Pipe command output to filter

```

Example 1-1 shows an example of the **show logging** command output.

Example 1-1 show logging Command

```

n1000v# show logging server
Logging server: enabled
{192.0.1.1}
server severity: critical
server facility: user

```

Contacting Cisco or VMware Customer Support

If you are unable to solve a problem after using the troubleshooting suggestions in this guide, contact a customer service representative for assistance and further instructions. Before you call, have the following information ready to help your service provider assist you as quickly as possible:

- Version of the software that you are running
- Version of the ESX and vCenter Server software that you are running
- Contact phone number.
- Brief description of the problem
- Brief explanation of the steps you have already taken to isolate and resolve the problem

If you purchased the product and support contract from Cisco, contact Cisco for support. Cisco provides L1, L2, and L3 support.

If you purchased the Cisco Nexus 1000V and an SNS through VMware, call VMware for support. VMware provides L1 and L2 support. Cisco provides L3 support.

After you have collected this information, see the [“Obtaining Documentation and Submitting a Service Request”](#) section on page -xvi.

For more information on steps to take before calling Technical Support, see the [“Gathering Information for Technical Support”](#) section on page 19-1.

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



CHAPTER 2

Tools Used in Troubleshooting

This chapter describes the troubleshooting tools available for the Cisco Nexus 1000V and includes the following topics:

- [Commands, page 2-1](#)
- [Ping, page 2-1](#)
- [Traceroute, page 2-2](#)
- [Monitoring Processes and CPUs, page 2-2](#)
- [RADIUS, page 2-4](#)
- [Syslog, page 2-5](#)

Commands

You use the CLI from a local console or remotely using a Telnet or SSH session. The CLI provides a command structure similar to Cisco NX-OS software, with context-sensitive help, **show** commands, multi-user support, and role-based access control.

Each feature has show commands that provide information about the feature configuration, status, and performance. Additionally, you can use the following commands for more information:

- **show system**—Provides information on system-level components, including cores, errors, and exceptions. Use the show system error-id command to find details on error codes:

```
n1000v# copy running-config startup-config
[#####] 100%
2008 Jan 16 09:59:29 zoom %$ VDC-1 %$ %BOOTVAR-2-AUTOCOPY_FAILED: Autocopy of file
/bootflash/n1000-s1-dk9.4.0.0.837.bin.S8 to standby failed, error=0x401e0008

n1000v# show system error-id 0x401e0008
Error Facility: sysmgr
Error Description: request was aborted, standby disk may be full
```

Ping

The ping utility generates a series of *echo* packets to a destination across a TCP/IP internetwork. When the echo packets arrive at the destination, they are rerouted and sent back to the source. Using ping, you can verify connectivity and latency to a particular destination across an IP routed network.

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

The ping allows you to ping a port or end device. By specifying the IPv4 address, you can send a series of frames to a target destination. Once these frames reach the target, they are looped back to the source and a time-stamp is taken. Ping helps you to verify the connectivity and latency to destination.

Traceroute

Use traceroute to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

Traceroute identifies the path taken on a hop-by-hop basis and includes a timestamp at each hop in both directions. You can use traceroute to test the connectivity of ports along the path between the generating switch and the switch closest to the destination.

Use the **traceroute** CLI command to access this feature.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of failure.

Monitoring Processes and CPUs

There are features in the CLI for monitoring switch processes and CPU status and utilization.

This section contains the following topics:

- [Identifying the Processes Running and their States, page 2-2](#)
- [Displaying CPU Utilization, page 2-3](#)
- [Displaying CPU and Memory Information, page 2-4](#)

Identifying the Processes Running and their States

Use the **show processes command** to identify the processes that are running and the status of each process. (See [Example 2-1](#).) The command output includes:

- PID = process ID.
- State = process state.
- PC = current program counter in hex format.
- Start_cnt = how many times a process has been started (or restarted).
- TTY = terminal that controls the process. A “-” usually means a daemon not running on any particular TTY.
- Process = name of the process.

Process states are:

- D = uninterruptible sleep (usually I/O).
- R = runnable (on run queue).
- S = sleeping.
- T = traced or stopped.

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

- Z = defunct (“zombie”) process.
- NR = not-running.
- ER = should be running but currently not-running.



Note

The ER state typically designates a process that has been restarted too many times, causing the system to classify it as faulty and disable it.

Example 2-1 show processes Command

```
n1000v# show processes ?
cpu      Show processes CPU Info
log      Show information about process logs
memory   Show processes Memory Info

n1000v# show processes

PID      State  PC          Start_cnt  TTY  Process
-----  -----  -----  -----  ---  -----
   1      S     b7f9e468    1          -    init
   2      S           0          1          -    migration/0
   3      S           0          1          -    ksoftirqd/0
   4      S           0          1          -    desched/0
   5      S           0          1          -    migration/1
   6      S           0          1          -    ksoftirqd/1
   7      S           0          1          -    desched/1
   8      S           0          1          -    events/0
   9      S           0          1          -    events/1
  10      S           0          1          -    khelper
  15      S           0          1          -    kthread
  24      S           0          1          -    kacpid
 101      S           0          1          -    kblockd/0
 102      S           0          1          -    kblockd/1
 115      S           0          1          -    khubd
 191      S           0          1          -    pdflush
 192      S           0          1          -    pdflushn
...

```

Displaying CPU Utilization

Use the **show processes cpu** command to display CPU utilization. The command output includes:

- Runtime(ms) = CPU time the process has used, expressed in milliseconds.
- Invoked = number of times the process has been invoked.
- uSecs = microseconds of CPU time in average for each process invocation.
- 1Sec = CPU utilization in percentage for the last one second.

Example 2-2 show processes cpu Command

```
n1000v# show processes cpu

PID      Runtime(ms)  Invoked  uSecs  1Sec  Process
-----  -----  -----  -----  ---  -----
   1           922  4294967295    0      0    init
   2           580   377810      1      0  migration/0

```

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

```

3          889   3156260   0   0  ksoftirqd/0
4         1648   532020   3   0  desched/0
5          400   150060   2   0  migration/1
6         1929   2882820   0   0  ksoftirqd/1
7         1269   183010   6   0  desched/1
8         2520  47589180   0   0  events/0
9         1730   2874470   0   0  events/1
10         64    158960   0   0  khelper
15         0     106970   0   0  kthread
24         0     12870   0   0  kacpid
101        62   3737520   0   0  kblockd/0
102        82   3806840   0   0  kblockd/1
115        0     67290   0   0  khubd
191        0     5810   0   0  pdflush
192       983   4141020   0   0  pdflush
194        0     5700   0   0  aio/0
193        0     8890   0   0  kswapd0
195        0     5750   0   0  aio/1
...

```

Displaying CPU and Memory Information

Use the **show system resources** command to display system-related CPU and memory statistics. The output includes the following:

- Load is defined as number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states shows the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage provides the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the used memory statistics.

Example 2-3 *show system resources* Command

```

n1000v# show system resources
Load average: 1 minute: 0.30 5 minutes: 0.34 15 minutes: 0.28
Processes : 606 total, 2 running
CPU states : 0.0% user, 0.0% kernel, 100.0% idle
Memory usage: 2063268K total, 1725944K used, 337324K free
              2420K buffers, 857644K cache

```

RADIUS

RADIUS is a protocol used for the exchange of attributes or credentials between a head-end RADIUS server and a client device. These attributes relate to three classes of services:

- Authentication
- Authorization
- Accounting

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

Authentication refers to the authentication of users for access to a specific device. You can use RADIUS to manage user accounts for access to an Cisco Nexus 1000V device. When you try to log into a device, Cisco Nexus 1000V validates you with information from a central RADIUS server.

Authorization refers to the scope of access that you have once you have been authenticated. Assigned roles for users can be stored in a RADIUS server along with a list of actual devices that the user should have access to. Once the user has been authenticated, then switch can then refer to the RADIUS server to determine the extent of access the user will have within the switch network.

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

The following is an example of an accounting log entries.

```
n1000v# show accounting log
Sun Dec 15 04:02:27 2002:start:/dev/pts/0_1039924947:admin
Sun Dec 15 04:02:28 2002:stop:/dev/pts/0_1039924947:admin:vsh exited normally
Sun Dec 15 04:02:33 2002:start:/dev/pts/0_1039924953:admin
Sun Dec 15 04:02:34 2002:stop:/dev/pts/0_1039924953:admin:vsh exited normally
Sun Dec 15 05:02:08 2002:start:snmp_1039928528_172.22.95.167:public
Sun Dec 15 05:02:08 2002:update:snmp_1039928528_172.22.95.167:public:Switchname
```

**Note**

The accounting log only shows the beginning and ending (start and stop) for each session.

Syslog

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides the following capabilities:

- Logging information for monitoring and troubleshooting.
- Selection of the types of logging information to be captured.
- Selection of the destination of the captured logging information.

Syslog lets you store a chronological log of system messages locally or sent to a central Syslog server. Syslog messages can also be sent to the console for immediate use. These messages can vary in detail depending on the configuration that you choose.

Syslog messages are categorized into 7 severity levels from *debug* to *critical* events. You can limit the severity levels that are reported for specific services within the switch.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) can be logged to a local file or server.

Logging Levels

Cisco Nexus 1000V supports the following logging levels:

- 0-emergency
- 1-alert
- 2-critical
- 3-error
- 4-warning

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

- 5-notification
- 6-informational
- 7-debugging

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. Users can specify which system messages should be saved based on the type of facility and the severity level. Messages are time-stamped to enhance real-time debugging and management.

Enabling Logging for Telnet or SSH

System logging messages are sent to the console based on the default or configured logging facility and severity values.

Users can disable logging to the console or enable logging to a given Telnet or SSH session.

- To disable console logging, use the **no logging console** command in CONFIG mode.
- To enable logging for telnet or SSH, use the **terminal monitor** command in EXEC mode.



Note

Note: When logging to a console session is disabled or enabled, that state is applied to all future console sessions. If a user exits and logs in again to a new session, the state is preserved. However, when logging to a Telnet or SSH session is enabled or disabled, that state is applied only to that session. The state is not preserved after the user exits the session.

The **no logging console** command shown in [Example 2-4](#):

- Disables console logging
- Enabled by default

Example 2-4 *no logging console Command*

```
n1000v(config)# no logging console
```

The **terminal monitor** command shown in [Example 2-5](#):

- Enables logging for telnet or SSH
- Disabled by default

Example 2-5 *terminal monitor Command*

```
n1000v# terminal monitor
```

For more information about configuring syslog, see the *Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SVI(3)*.



CHAPTER 3

Installation

This chapter describes how to identify and resolve installation problems, and includes the following topics:

- [Isolating Installation Problems, page 3-1](#)
- [Improving Performance, page 3-4](#)
- [Verifying the Domain Configuration, page 3-4](#)
- [Verifying the Port Group Assignments for a VSM VM Virtual Interface, page 3-4](#)
- [Verifying VSM and vCenter Server Connectivity, page 3-5](#)
- [Troubleshooting Connections to a vCenter Server, page 3-5](#)
- [Recovering the Network Administrator Password, page 3-6](#)
- [Managing Extension Keys, page 3-6](#)
- [Recreating the Cisco Nexus 1000V Installation, page 3-10](#)

Isolating Installation Problems

Use this section to isolate a problem with the installation, including the following.

- [Verifying Your VMware License Version, page 3-1](#)
- [Host is Not Visible from Distributed Virtual Switch, page 3-2](#)
- [Refreshing the vCenter Server Connection, page 3-3](#)

Verifying Your VMware License Version

Use this procedure, before beginning to troubleshoot any installation issues, to verify that your ESX server has the VMware Enterprise Plus license which includes the Distributed Virtual Switch feature.

BEFORE YOU BEGIN

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

- You are logged in to the vSphere client on the ESX server.
- You are logged in to the Cisco Nexus 1000V CLI in EXEC mode.

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

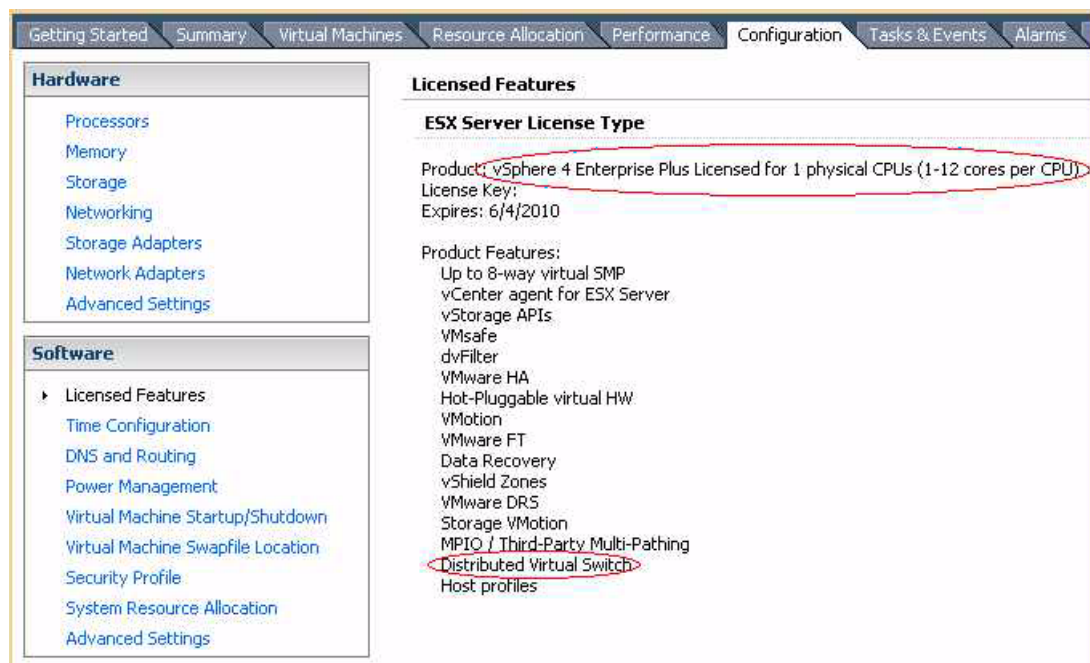
- This procedure verifies that your ESX server uses the VMware Enterprise Plus license. This license includes the feature, Distributed Virtual Switch, which allows visibility to the Cisco Nexus 1000V.
- If your vSphere ESX server does not have the Enterprise Plus license, then you must upgrade your license.

DETAILED STEPS

Step 1 From the vSphere client, select the host whose Enterprise Plus license you want to check.

Step 2 Click the **Configuration** tab and select **Licensed Features**.

The Enterprise Plus licensed features are displayed.



Step 3 Verify that the following are included in the Licensed Features:

- Enterprise Plus license
- Distributed Virtual Switch feature

Step 4 Do one of the following:

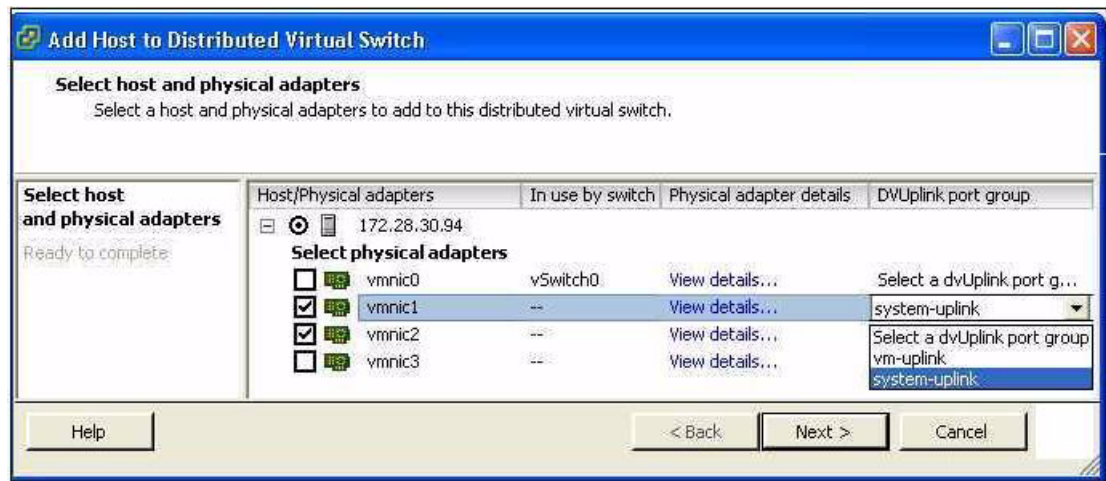
- If your ESX server has an Enterprise Plus license, then you have the correct license and visibility to the Cisco Nexus 1000V.
- If your ESX server does not have an Enterprise Plus license, then you must upgrade your VMware License to an Enterprise Plus license in order to have visibility to the Cisco Nexus 1000V.

Host is Not Visible from Distributed Virtual Switch

If you have added hosts and adapters with your VSM, then you must also add them in the vCenter Client Add Host to Distributed Virtual Switch dialog box shown in [Figure 3-1](#).

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

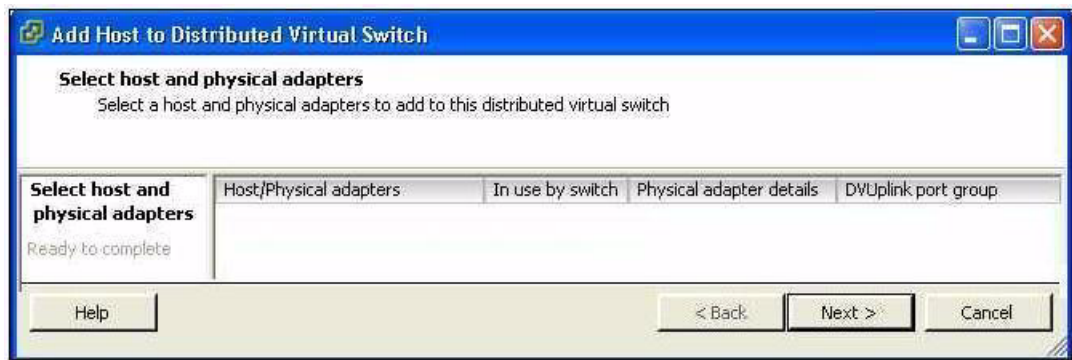
Figure 3-1 Host is Visible from the Distributed Virtual Switch



If the hosts and adapters do not appear in this dialog box, as shown in Figure 3-2, then you may have the incorrect VMware license installed on your ESX server.

Use the “[Verifying Your VMware License Version](#)” procedure on page 3-1 to confirm.

Figure 3-2 Host is Not Visible from the Distributed Virtual Switch



Refreshing the vCenter Server Connection

Use this procedure to refresh the connection between the Cisco Nexus 1000V and vCenter Server.

- Step 1** From the Cisco Nexus 1000V Connection Configuration mode on the VSM, enter the following command sequence:

```

Example:
n1000v# config t
n1000v(config)# svcs connection s1
n1000v(config-svs-conn)# no connect
n1000v(config-svs-conn)# connect
  
```

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

Step 2 You have completed this procedure.

Improving Performance

Use the following pointers to improve performance on the ESX host and the VMs.

- Install VMware Tools on the vCenter Server VM, with Hardware Acceleration enabled to the full.
- Use the command line interface in the VMs instead of the graphical interface where possible.

Verifying the Domain Configuration

The Virtual Supervisor Module (VSM) and Virtual Ethernet Module (VEM) are separated within a Layer 2 domain. To allow VSM-VEM pairs to communicate within the same Layer 2 domain, each pair must have a unique identifier. The domain ID serves as the unique identifier that allows multiple VSM-VEM pairs to communicate inside the same Layer 2 domain.

Following the installation of the Cisco Nexus 1000V, make certain that you configure a domain ID. Without a domain ID, the VSM will not be able to connect to the vCenter Server. Follow these guidelines:

- The domain ID should be a value within the range of 1 to 4095.
- All the control traffic between the VSM and the VEM is carried over the configured control VLAN.
- All the data traffic between the VSM and the VEM is carried over the configured packet VLAN.
- Make sure the control VLAN and the packet VLAN are allowed on the port in the upstream switch to which the physical NIC of the host hosting the VSM and VEM VM are connected.

Verifying the Port Group Assignments for a VSM VM Virtual Interface

Use this procedure to verify that two port groups are created on the ESX hosting the VSM VM through the vCenter Server. The following port groups (PG) should be created:

- Control PG (Vlan = Control VLAN)
- Packet PG (Vlan = Packet VLAN)
- Management PG (Vlan = Management VLAN)

Make sure the port groups are assigned to the three virtual interfaces of the VSM VM in the following order:

Virtual Interface Number	Port Group
Network Adapter 1	Control PG
Network Adapter 2	MGMT PG
Network Adapter 3	Packet PG

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

To verify if the VSM VM network adapter 1, network adapter 2, and network adapter 3 are carrying the control VLAN, management VLAN, and the packet VLAN, follow these steps:

-
- Step 1** Enter the **show mac address-table dynamic interface vlan** *control-vlan* command on the upstream switch.
- Expected Output: the network adapter1 MAC address of the VSM VM.
- Step 2** Enter the **show mac address-table dynamic interface vlan** *mgmt-vlan* command on the upstream switch.
- Expected Output: the network adapter2 MAC address of the VSM VM.
- Step 3** Enter the **show mac address-table dynamic interface vlan** *packet-vlan* command on the upstream switch.
- Expected Output: the network adapter3 MAC address of the VSM VM.
-

Verifying VSM and vCenter Server Connectivity

When troubleshooting connectivity between the VSM and vCenter Server, follow these guidelines:

- Make sure that domain parameters are configured correctly.
- Make sure the Windows VM machine hosting the vCenter Server has the following ports open.
 - Port 80
 - Port 443
- Try reloading the VSM if after verifying the preceding steps, the connect still fails.
- Check if the VSM extension is created by the vCenter Server by pointing your web browser to <https://your-virtual-center/mob/>, and then clicking Content > Extension Manager.

Use this procedure to troubleshoot connectivity between a VSM and a vCenter Server:

-
- Step 1** Ensure that Nexus N1000V VSM VM network adapters are configured properly.
- Step 2** Make sure the Windows VM machine hosting the vCenter Server has the following ports open.
- Port 80
 - Port 443
- Step 3** Ping the vCenter Server from the VSM.
- Step 4** Ensure the VMware VirtualCenter Server service is running.
-

Troubleshooting Connections to a vCenter Server

Use this procedure to troubleshoot connections between a Cisco Nexus 1000V VSM and a vCenter Server:

-
- Step 1** In a web browser, enter the path: <http://<VSM-IP>>

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

- Step 2** Download the `cisco_nexus_1000v_extension.xml` file to your desktop.
- Step 3** From the vCenter Server menu, choose **Plugins** → **Manage Plugins**. Right click an empty area and select the plugin in Step2 as the New Extension.
-

If these steps fail, then you may be using an out-of-date .xml file.

Use this procedure to confirm that the extension is available:

- Step 1** In a web browser, enter the path: `http://<vCenter-Server-IP>/mob`.
- Step 2** Click **Content**.
- Step 3** Click **extensionManager**.
- Step 4** If `extensionList[Cisco_Nexus_1000v_584325821]` is displayed in the value column, then proceed to connect to the VSM.
-

**Note**

The actual value of “Cisco_Nexus_1000V_584325821” will vary. It should match the extension key from the `cisco_nexus_1000v_extension.xml` file.

Recovering the Network Administrator Password

For information about recovering the network administrator password, see the *Cisco Nexus 1000V Password Recovery Guide*.

Managing Extension Keys

This section includes the following topics:

- [Known Extension Problems and Resolutions, page 3-7](#)
- [Resolving a Plug-In Conflict, page 3-7](#)
- [Finding the Extension Key on the Cisco Nexus 1000V, page 3-7](#)
- [Finding the Extension Key Tied to a Specific DVS, page 3-8](#)
- [Verifying Extension Keys, page 3-9](#)

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

Known Extension Problems and Resolutions

Use [Table 3-1](#) to troubleshoot and resolve known problems with plug-ins and extensions.

Table 3-1 Known Extension Problems and Resolutions

Problem	Resolution
The extension does not show up immediately in the plugin.	Close the VI client and then open the VI client again.
You cannot delete the extension from the VI client.	If you delete the extension using MOB, then the VI client screen may not refresh and indicate that the extension was deleted. In this case, close the VI client and then open the VI client again.
If you click the download and install link for the extension, you see an error of invalid URI.	None. You do not need to click download and install . If you do, it has no effect on the installation or connectivity. The plug-in only needs to be registered with the vCenter.

Resolving a Plug-In Conflict

If you see the error, “The specified parameter was not correct,” when Creating a Plug-In on the vCenter Server, then you have tried to register a plugin that is already registered.

Use the following procedure to resolve this problem.

-
- Step 1** Make sure that you are using the correct `cisco_nexus1000v_extension.xml` file.
 - Step 2** Make sure that you have refreshed your browser since it caches this file and unless refreshed it might cache obsolete content with the same file name.
 - Step 3** Follow the steps described in the [“Verifying Extension Keys”](#) section on page 3-9 to compare the extension key installed on the VSM with the plug-in installed on the vCenter Server.
-

Finding the Extension Key on the Cisco Nexus 1000V

You can use this procedure to find the extension key on the Cisco Nexus 1000V.

BEFORE YOU BEGIN

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

- You are logged in to the the Cisco Nexus 1000V VSM CLI in EXEC mode.
- You can use the extension key found in this procedure in the [“Unregistering the Extension Key in the vCenter Server”](#) procedure on page 3-12.

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

DETAILED STEPS

- Step 1** From the Cisco Nexus 1000V for the VSM whose extension key you want to view, enter the following command:

```
show vmware vc extension-key
```

Example:

```
n1000v# show vmware vc extension-key
Extension ID: Cisco_Nexus_1000V_1935882621
n1000v#
```

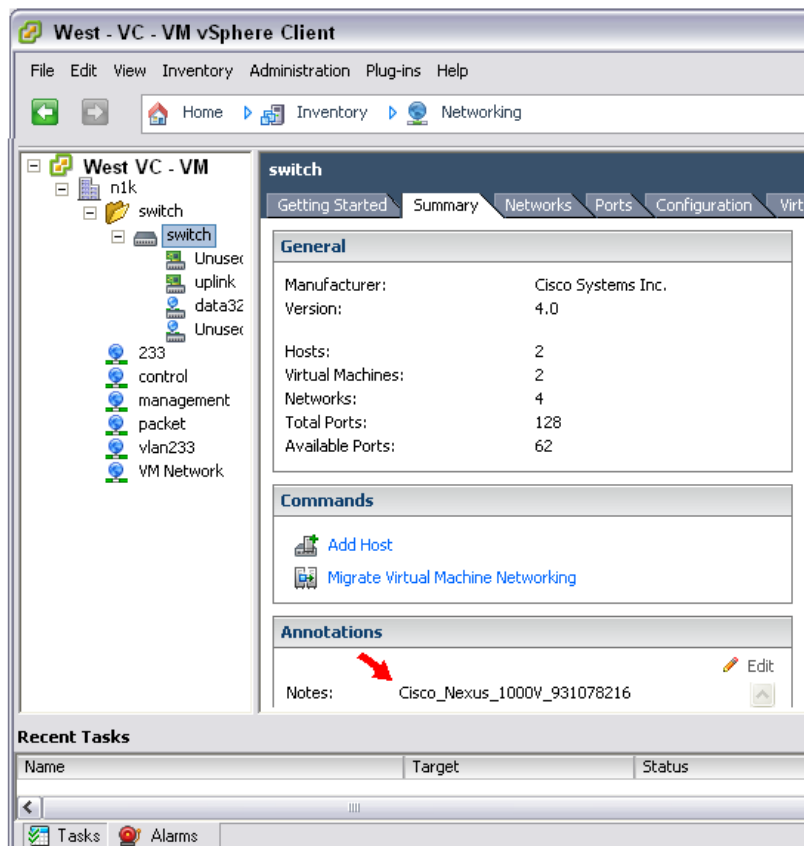
Finding the Extension Key Tied to a Specific DVS

You can use this procedure to find the extension key tied to a specific DVS.

- Step 1** From the vSphere client, select the DVS whose extension key you want to find.

- Step 2** Click the **Summary** tab.

The Summary tab opens with the extension key displayed in the Notes section of the Annotations block.



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

Verifying Extension Keys

You can use this procedure to verify that the Cisco Nexus 1000V and vCenter Server are using the same extension key.

DETAILED STEPS

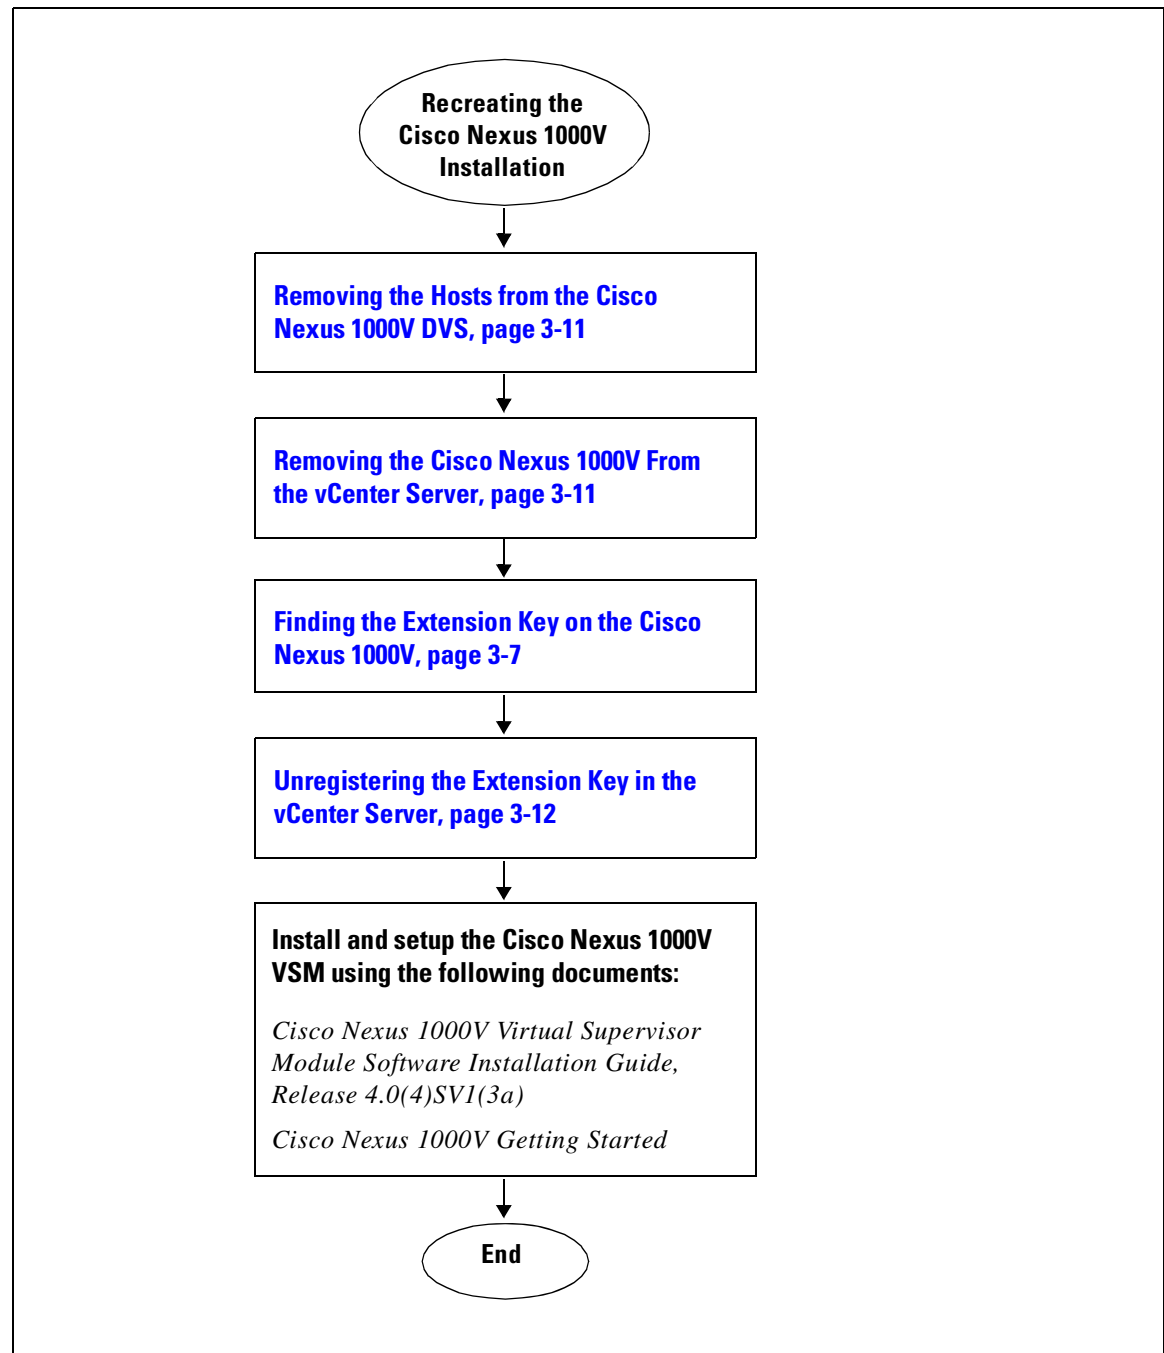
-
- Step 1** Find the extension key used on the Cisco Nexus 1000V using the [Finding the Extension Key on the Cisco Nexus 1000V, page 3-7](#).
 - Step 2** Find the extension key used on the vCenter Server using the [Finding the Extension Key Tied to a Specific DVS, page 3-8](#).
 - Step 3** Verify that the two extension keys (the one found in [Step 1](#) with that in [Step 2](#)) are the same.
-

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

Recreating the Cisco Nexus 1000V Installation

Use this section to recreate the complete Cisco Nexus 1000V configuration in the event of a persistent problem that cannot be resolved using any other workaround.

FlowChart: Recreating the Cisco Nexus 1000V Installation



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

Removing the Hosts from the Cisco Nexus 1000V DVS

Use this procedure to remove the hosts from the Cisco Nexus 1000V DVS.

BEFORE YOU BEGIN

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

- You are logged in to the vSphere Client.
- You know the name of the Cisco Nexus 1000V DVS to remove from vCenter Server.

DETAILED STEPS

-
- Step 1** From the vSphere Client, choose **Inventory** → **Networking**.
- Step 2** Select the DVS for the Cisco Nexus 1000V and click the **Hosts** tab.
The Host tab opens.
- Step 3** Right-click each host, and choose **Remove from Distributed Virtual Switch**.
The hosts are now removed from the DVS.
-

Removing the Cisco Nexus 1000V From the vCenter Server

You can use this procedure to remove the Cisco Nexus 1000V DVS from vCenter Server.

BEFORE YOU BEGIN

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

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

DETAILED STEPS

-
- Step 1** From the Cisco Nexus 1000V VSM, use the following commands to remove the DVS from the vCenter Server.
- ```
config t
svs connection vc
no vmware dvs
```
- Example:**
- ```
n1000v# conf t  
n1000v(config)# svs connection vc  
n1000v(config-svs-conn)# no vmware dvs  
n1000v(config-svs-conn)#
```
- The DVS is removed from the vCenter Server.
- Step 2** You have completed this procedure.
Return to [FlowChart: Recreating the Cisco Nexus 1000V Installation, page 3-10](#).
-

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

Unregistering the Extension Key in the vCenter Server

You can use this procedure to unregister the Cisco Nexus 1000V extension key in vCenter Server. After the extension key is unregistered

BEFORE YOU BEGIN

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

- You have a browser window open.
- This procedure requires you to paste the extension key name into the vCenter Server MOB. You should already have the extension key found in the [“Finding the Extension Key on the Cisco Nexus 1000V” procedure on page 3-7.](#)
- After using this procedure to unregister the extension key in vCenter Server, you can start a fresh installation of the Cisco Nexus 1000V VSM software.

DETAILED STEPS

Step 1 Point your browser to the following url:

<https://<vc-ip>/mob/?moid=ExtensionManager>

The Extension Manager opens in your Manager Object Browser (MOB).

Home

Managed Object Type:
ManagedObjectReference:ExtensionManager

Managed Object ID: ExtensionManager

Properties

NAME	TYPE	VALUE
extensionList	Extension []	<ul style="list-style-type: none"> • extensionList["Cisco Nexus 1000V 1265583024"] • extensionList["Cisco Nexus 1000V 1410054174"] • extensionList["Cisco Nexus 1000V 1596939501"] • extensionList["Cisco Nexus 1000V 2018829329"] • extensionList["Cisco Nexus 1000V 2095452616"] • extensionList["Cisco Nexus 1000V 413176078"] • extensionList["Cisco Nexus 1000V 597460431"] • extensionList["Cisco Nexus 1000V 41882082"]

Methods

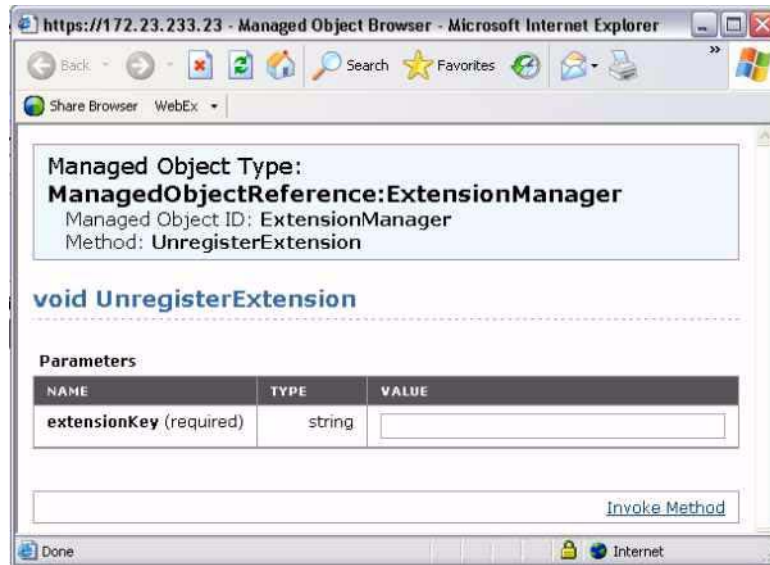
RETURN TYPE	NAME
Extension	FindExtension
string	GetPublicKey
void	RegisterExtension
void	SetExtensionCertificate
void	SetPublicKey
void	UnregisterExtension

Step 2 Click **Unregister Extension**.

<https://<vc-ip>/mob/?moid=ExtensionManager&method=unregisterExtension>

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

A dialog box opens for unregistering the extension.



Step 3 In the value field, paste the extension key you found in the [“Finding the Extension Key on the Cisco Nexus 1000V” procedure on page 3-7](#), and then click **Invoke Method**.

The extension key is unregistered in vCenter Server so that you can start a new installation of the Cisco Nexus 1000V VSM software or add a new extension key.

Step 4 You have completed this procedure.

Do one of the following:

Return to [FlowChart: Recreating the Cisco Nexus 1000V Installation, page 3-10](#).

Return to the

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



CHAPTER 4

Upgrade

This chapter describes how to identify and resolve upgrade problems.

Identifying Upgrade Solutions

Table 4-1 provides symptoms related to upgrade, their possible causes, and recommended solutions.

Table 4-1 Identifying Upgrade Solutions

Symptom	Possible Cause	Solution
VUM is downgrading my VIB package	<p>The VC does not have the correct VSM software bundle ID.</p> <p>If the VSM and VC bundle IDs do not match, VUM installs the VIB corresponding to the older VSM version on the host.</p>	<p>Upgrade the VSM bundle ID on the VC.</p> <ol style="list-style-type: none"> 1. “Verifying Bundle IDs” procedure on page 4-3 2. <i>Upgrading the VSMs</i> procedure in the following document: <ul style="list-style-type: none"> – <i>Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(3a)</i>.
Unable to add the host to VSM	The VEM is not installed on the host	<p>Verify that VEM software is installed and running on the host. If not, install VEM software.</p> <ol style="list-style-type: none"> 1. “Verifying that the correct VEM Software Version is Installed” procedure on page 4-3 2. “Verifying that the VEM is Running on the Host” procedure on page 4-3 3. If not installed, install VEM. <p>For more information, see the <i>Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(3a)</i>.</p> <ol style="list-style-type: none"> 4. “Verifying that the correct VEM Software Version is Installed” procedure on page 4-3

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

Table 4-1 Identifying Upgrade Solutions (continued)

Symptom	Possible Cause	Solution
Unable to add a host	The host is part of another DVS.	<p>Verify whether the host is part of another DVS. If so, then remove it and add it to the correct DVS.</p> <ol style="list-style-type: none"> “Verifying the VMware NICs on a DVS” procedure on page 4-3 Remove the host from the DVS where it is currently installed. Add the host to the correct DVS. <p>For more information, see the <i>Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(3a)</i></p>
The upgrade fails	<p>VMware patch level is 219382 or higher. If so, you cannot install VEM using VUM.</p>	<p>Verify the patch level. If required, upgrade the VUM and install VEM software.</p> <ol style="list-style-type: none"> “Verifying that the correct VEM Software Version is Installed” procedure on page 4-3 Do one of the following: <ul style="list-style-type: none"> Upgrade your VUM to the latest version. Install the VEM software. <p>For more information, see the <i>Cisco Nexus 1000V Virtual Ethernet Module Software Installation Guide, Release 4.0(4)SV1(3a)</i>.</p>
	<p>The upgrade commands were not used in the correct sequence, causing a mismatch between the bundle ID on the VSM and VC.</p> <p>If the VSM and VC bundle IDs do not match, VUM installs the VIB corresponding to the older VSM version on the host.</p> <p>The following is the correct sequence for entering upgrade commands:</p> <ol style="list-style-type: none"> vmware vem upgrade notify vmware vem upgrade proceed vmware vem upgrade complete 	<p>Identify the problem host and repeat the upgrade.</p> <ol style="list-style-type: none"> “Verifying Upgrade Status” procedure on page 4-4 Identify which host failed to upgrade. Verify that the VUM repository has the VIB that must be installed on that host.
	<p>Your vCenter VUM build is older than Update 1, and your proxy is enabled.</p>	<p>Disable the proxy on the VUM and retry the upgrade.</p> <p>Upgrade your VC to Update 1 and VUM to VMware vCenter Update Manager 4.0 Update 1 Patch 1.</p>

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

Verifying Bundle IDs

Use the following command to verify that the software bundle IDs on the VSM and VC match.

show vmware vem upgrade status

Example:

```
switch # show vmware vem upgrade status
Upgrade VIBs: System VEM Image
Upgrade Status:
Upgrade Notification Sent Time:
Upgrade Status Time(vCenter):
Upgrade Start Time:
Upgrade End Time(vCenter):
Upgrade Error:
Upgrade Bundle ID:
VSM: VEM400-201003101-BG
DVS: VEM400-200912267-BG
```

Verifying that the correct VEM Software Version is Installed

Use the following command to verify that the correct VEM software version is installed.

vemcmd show version

Example:

```
# vemcmd show version
VEM Version: 4.2.1.1.0.1.0-1.9.7
VSM Version: 4.2(1)SV1(1) [build 4.2(1)SV1(0.1)]
System Version: VMware ESXi 4.0.0 Releasebuild-208167
```

Verifying that the VEM is Running on the Host

Use the following command to verify that the VEM is running on the host.

vem status

Example:

```
~ # vem status
VEM modules are loaded
Switch Name Num Ports Used Ports Configured Ports MTU Uplinks
vSwitch0 64 3 64 1500 vmnic0
DVS Name Num Ports Used Ports Configured Ports Uplinks
n1000v 256 9 256 vmnic1 VEM Agent is running
```

Verifying the VMware NICs on a DVS

Use the following command to verify the VMware NICs associated with a DVS.

esxcfg-vswitch -1

Example:

```
~ # esxcfg-vswitch -1
Switch Name      Num Ports  Used Ports  Configured Ports  MTU  Uplinks
vSwitch0        128        3           128              1500  vmnic0
  PortGroup Name  VLAN ID   Used Ports  Uplinks
  VM Network      0         0           vmnic0
```

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

Management Network	0	1	vmnic0			
DVS Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks	
bldvs3	256	52	256	1500	vmnic2,vmnic1	
DVPort ID	In Use	Client				
352	1	vmnic1				
353	1	vmnic2				

Verifying the VMware Patches Installed on a Host

Use the following command to verify the VMware patches installed on a host.

esxupdate -vib-view query

Example:

Verifying Upgrade Status

Use the following command to verify the status of the VEM upgrade on the VSM and whether any host failed to upgrade.

show vmware vem upgrade status

Example:

```
switch# show vmware vem upgrade status
2010 Mar 17 15:19:41 pvk-vsm %VMS-5-DVS_UPGRADE_INFO: VEM Upgrade Notification has been
accepted by vCenter Admin.
Upgrade VIBs: System VEM Image
Upgrade Status: Upgrade Accepted by vCenter Admin
Upgrade Notification Sent Time: Wed Mar 17 15:19:05 2010
Upgrade Status Time(vCenter): Wed Mar 17 17:28:46 2010
Upgrade Start Time:
Upgrade End Time(vCenter):
Upgrade Error:
Upgrade Bundle ID:
VSM: VEM400-201003101-BG
DVS: VEM400-200912267-BG
```




CHAPTER 5

Licensing

This chapter describes how to identify and resolve problems related to licensing for the Cisco Nexus 1000V.

This chapter includes the following sections.

- [Licensing Overview, page 5-1](#)
- [Troubleshooting Unlicensed Modules, page 5-2](#)
- [Troubleshooting License Installation Issues, page 5-3](#)
- [Determining License Usage, page 5-3](#)
- [Installed License Information, page 5-4](#)
- [Troubleshooting Post License Installation Problems, page 5-4](#)
- [Troubleshooting the Removal of a License, page 5-5](#)

Licensing Overview

The name for the Cisco Nexus 1000V license package is NEXUS1000V_LAN_SERVICES_PKG.

The licensing model for Cisco Nexus 1000V is based on the number of CPU sockets of the ESX servers attached as VEMs to the VSM.

A module is licensed or unlicensed according to the following definitions:

- Licensed module—A VEM is considered to be licensed if it is able to acquire licenses for all of its CPU sockets.
- Unlicensed module—A VEM is considered to be unlicensed if it is not able to acquire licenses for any, or a subset of, its CPU sockets.

In the case that a VEM is unlicensed, all the virtual Ethernet ports on the VEM corresponding to the virtual machines (VMs) are kept down, with a reason code indicating that the VEM is unlicensed.

The VSM does not contain any licenses by default.

For additional information about licensing, see the *Cisco Nexus 1000V License Configuration Guide, Release 4.0(4)SV1(3)*.

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

Troubleshooting Unlicensed Modules

By default, the VSM does not contain any licenses. As a result, if you add a VEM to the VSM, the VEM comes up unlicensed.

To identify an unlicensed module, enter the **show module** command on the VSM.

```
n1000v# show module
Mod  Ports  Module-Type                Model                Status
---  ---  -
1    0      Virtual Supervisor Module  Nexus1000V          active *
5    248    Virtual Ethernet Module    NA                   unlicensed
Mod  Sw                Hw
---  ---  -
1    4.0(4)SV1(1)      0.0
5    4.0(4)SV1(1)      0.4
Mod  MAC-Address(es)      Serial-Num
---  ---  -
1    00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8  NA
5    02-00-0c-00-05-00 to 02-00-0c-00-05-80  NA
Mod  Server-IP          Server-UUID                Server-Name
---  ---  -
1    172.23.232.140     NA                          NA
5    172.23.233.100     33393935-3234-5553-4539-30364e345630  172.23.233.100
```

As shown, the status field for VEM 5 is unlicensed.

To gather information about why vEthernet interfaces are in VEM unlicensed state, enter the **show interface veth1** command.

```
n1000v# show int veth1
Vethernet1 is down (VEM Unlicensed)
  Port description is VM-Pri, Network Adapter 1
  Hardware is Virtual, address is 0050.56b7.1c7b
  Owner is VM "VM-Pri", adapter is Network Adapter 1
  Active on module 5
  VMware DVS port 32
  Port-Profile is dhcp-profile
  Port mode is access
  Rx
  5002 Input Packets 4008 Unicast Packets
  85 Multicast Packets 909 Broadcast Packets
  846478 Bytes
  Tx
  608046 Output Packets 17129 Unicast Packets
  502543 Multicast Packets 88374 Broadcast Packets 0 Flood Packets
  38144480 Bytes
  20 Input Packet Drops 0 Output Packet Drops
```

If you power on a virtual machine with ports on a Cisco Nexus 1000V port group set, the interfaces are kept down with the status as VEM Unlicensed if the VEM is unlicensed.



Note

The server administrator has no information on whether the VEMs are licensed or unlicensed. Therefore, the license state of the VEMs must be communicated to the server administrators so that they are aware that the vEthernet interfaces on unlicensed modules will not be able to pass traffic.

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

Troubleshooting License Installation Issues

This section assumes that you have a valid Cisco Nexus 1000V license file.

For information on how to purchase or install a license file, see the *Cisco Nexus 1000V License Configuration Guide, Release 4.0(4)SV1(3)*.

License Troubleshooting Checklist

Before you start the troubleshooting process, follow these requirements:

- Make sure the name of the license file is less than 32 characters.
- Make sure no other license file with the same name is installed on the VSM. If there is a license file with the same name, rename your new license file to something else.
- Do not edit the contents of the license file. If you have already done so, please contact your Cisco Customer Support Account Team.
- Make sure the Host-ID in the license file is the same as that on the switch.

Contents of the License File

The Cisco Nexus 1000V license file looks as follows:

```
SERVER this_host ANY
VENDOR cisco
INCREMENT NEXUS1000V_LAN_SERVICES_PKG cisco 1.0 26-jun-2009 16 \

HOSTID=VDH=4724514071229227089 \
NOTICE="<LicFileID>20090427142506674</LicFileID><LicLineID>1</LicLineID> \
<PAK12345></PAK12345>" SIGN=E0AF5428C434

Host-ID of the VSM
n1000v#show license host-id
License hostid: VDH=4724514071229227089
```

Notice that both instances of the host-id match and are equal to VDH=4724514071229227089.

Removing an Evaluation License File

If an evaluation license file is already installed on the VSM, then it must be removed from the VSM before installing a permanent license file. For more information, see the *Cisco Nexus 1000V License Configuration Guide, Release 4.0(4)SV1(3)*.

Determining License Usage

To view the license state of the VEMs on your VSM and the number of CPU sockets per VEM, use the following command:

```
show module vem internal license-info
```

Example:

```
n1000v# show module vem internal license-info
```

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

```
License Sync Initiator      : VEM 5
License Sync Stages        : Complete
Num of Def Licenses in Use : 0
Num of Sync participants   : 1
License Host-ID           : 4724514071229227089
-----VEM License Info -----
Vem      Current License Operation      License Status  License Flags
-----
5                None                    licensed      None
-----VEM Socket License Info -----
Vem      Sync      License Usage      Sockets License Version
-----
5        Yes                2                2                1.0
```

In this output, VEM 5 is Licensed. It has 2 CPU sockets and it currently uses 2 licenses.

Installed License Information

Use the **show license usage** command to view the installed license count.

```
n1000v#show license usage
Feature                               Ins Lic  Status Expiry Date Comments
                               Count
-----
NEXUS1000V_LAN_SERVICES_PKG    Yes  16   In use 26 Jun 2009  -
-----
```

The output shows that 16 licenses have been installed and they will expire on June 26, 2009.

Troubleshooting Post License Installation Problems

After you install a license, you might see syslog messages like the following, which indicate a problem:

PLATFORM-2-PFM-VEM_UNLICENSED Syslog

Error Message 2008 Dec 19 22:28:30 N1KV %PLATFORM-2-PFM_VEM_UNLICENSED: License for VEM 5 could not be obtained. Please contact your Cisco account team or partner to purchase Licenses. To activate your purchased licenses, click on www.cisco.com/go/license.

Explanation It means that enough licenses were not installed to license the CPU Sockets of all the VEMs connected to the VSM.

Recommended Action Install additional licenses for the VEMs that have not been able to acquire licenses.



Note

To determine how many CPU sockets a particular VEM has, see the [“Determining License Usage” procedure on page 5-3](#).

If the license file is about to expire, then you might see the following syslog message:

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

PLATFORM-2-PFM_LIC_WARN_EXP Syslog

Error Message 2008 Dec 19 22:28:30 N1KV %PLATFORM-2-PFM_LIC_WARN_EXP: WARNING License for VEMs is about to expire in 1 days! The VEMs' VNICS will be brought down if license is allowed to expire. Please contact your Cisco account team or partner to purchase Licenses. To activate your purchased licenses, click on www.cisco.com/go/license.

Explanation This is a warning message saying that the license file is going to expire within a certain period of time.

Recommended Action Contact your Cisco Account team to purchase a license file.

Troubleshooting the Removal of a License

You cannot clear a license file that is currently being used.

To see how many licenses have currently been checked out, use the **show module vem internal license-info** command.

```
n1000v#show module vem internal license-info
License Sync Initiator      : VEM 5
License Sync Stages        : Complete
Num of Def Licenses in Use : 0
Num of Sync participants   : 1
License Host-ID            : 4724514071229227089
-----VEM License Info -----
Vem      Current License Operation      License Status  License Flags
--      -
5                None                licensed      None
6                None                licensed      None
-----VEM Socket License Info -----
Vem      Sync      License Usage      Sockets License Version
--      -
5        Yes        2                2                1.0
6        Yes        2                2                1.0
```

To clear a license file, you need to manually release the licenses bound to each of the VEMs, and then clear the license file using the **clear license** command.

```
n1000v#clear license n1kv_license.lic
Clearing license failed: License is in use
n1000v#svs license transfer src-vem 5 license_pool
n1000v#svs license transfer src-vem 6 license_pool
```

```
n1000v#show module 5
Mod  Ports  Module-Type          Model          Status
--  -
5    248    Virtual Ethernet Module  NA            unlicensed
Mod  Sw      Hw
-----
5    NA     NA
Mod  MAC-Address(es)      Serial-Num
-----
5    02-00-0c-00-05-00 to 02-00-0c-00-05-80  NA
Mod  Server-IP      Server-UUID      Server-Name
-----
```

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

```

5      172.23.233.100      33393935-3234-5553-4539-30364e345630      172.23.233.100

n1000v#show module 6
Mod  Ports  Module-Type                Model                Status
---  ---
6    248    Virtual Ethernet Module    NA                    unlicensed
Mod  Sw      Hw
--  --
6    NA     NA
Mod  MAC-Address(es)          Serial-Num
---  ---
6    02-00-0c-00-05-00 to 02-00-0c-00-05-80  NA
Mod  Server-IP                Server-UUID          Server-Name
---  ---
6    172.23.233.101          12354635-3192-7653-3690-12375a5345721  172.23.233.101

n1000v#clear license n1kv_license.lic
Clearing license .....done

```



CHAPTER 6

VSM and VEM Modules

This chapter describes how to identify and resolve problems that relate to modules and includes the following sections:

- [Information About Modules, page 6-1](#)
- [Troubleshooting a Module Not Coming Up on the VSM, page 6-1](#)
- [Problems with the VSM, page 6-4](#)
- [Commands, page 6-16](#)

Information About Modules

Cisco Nexus 1000V manages a data center defined by a VirtualCenter. Each server in the data center is represented as a module and can be managed as if it were a module in a physical Cisco switch.

The Cisco Nexus 1000V implementation has 2 parts:

- Virtual supervisor module (VSM) – This is the control software of the Cisco Nexus 1000V distributed virtual switch. It runs on a virtual machine (VM) and is based on Cisco NX-OS software.
- Virtual Ethernet module (VEM) – This is the part of Cisco Nexus 1000V that actually switches data traffic. It runs on a VMware ESX 4.0 host. Several VEMs are controlled by one VSM. All the VEMs that form a switch domain should be in the same virtual Data Center as defined by VMware VirtualCenter.

Troubleshooting a Module Not Coming Up on the VSM

This section describes the process you can use when a module does not come up on the VSM. This section includes the following topics:

- [Guidelines for Troubleshooting Modules, page 6-2](#)
- [Flow Chart for Troubleshooting Modules, page 6-3](#)
- [Verifying the VSM Is Connected to the vCenter Server, page 6-6](#)
- [Verifying the VSM Is Configured Correctly, page 6-7](#)
- [Checking the vCenter Server Configuration, page 6-10](#)
- [Verifying Network Connectivity Between the VSM and the VEM, page 6-11](#)

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

- [Checking the VEM Configuration, page 6-13](#)
- [Collecting Logs, page 6-15](#)

Guidelines for Troubleshooting Modules

Follow these guidelines when troubleshooting a module controlled by the VSM.

- You must have a VSM VM and a VEM up and running.
- Make sure you are running compatible versions of vCenter Server and VSM.

For more information, see the *Cisco Nexus 1000V Compatibility Information, Release 4.0(4)SV1(3a)*.

- To verify network connectivity between the VSM and vCenter Server, ping the IP address of the vCenter Server. If you are using a domain name service (DNS) name, use the DNS name in the ping. If a ping to the vCenter Server fails, check to see if you can ping the gateway. Otherwise, check the mgmt0 interface configuration settings.
- Make sure the firewall settings are OFF on the vCenter Server. If you want the firewall settings, then check to see if these ports are open.
 - Port 80
 - Port 443
- If you see the following error, verify that the VSM extension was created from vCenter Server.
 - ERROR: [VMware vCenter Server 4.0.0 build-150489]
Extension key was not registered before its use

To verify that the extension or plugin was created, see the [“Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8](#).

For more information about extension keys or plugins, see the [“Managing Extension Keys” section on page 3-6](#).

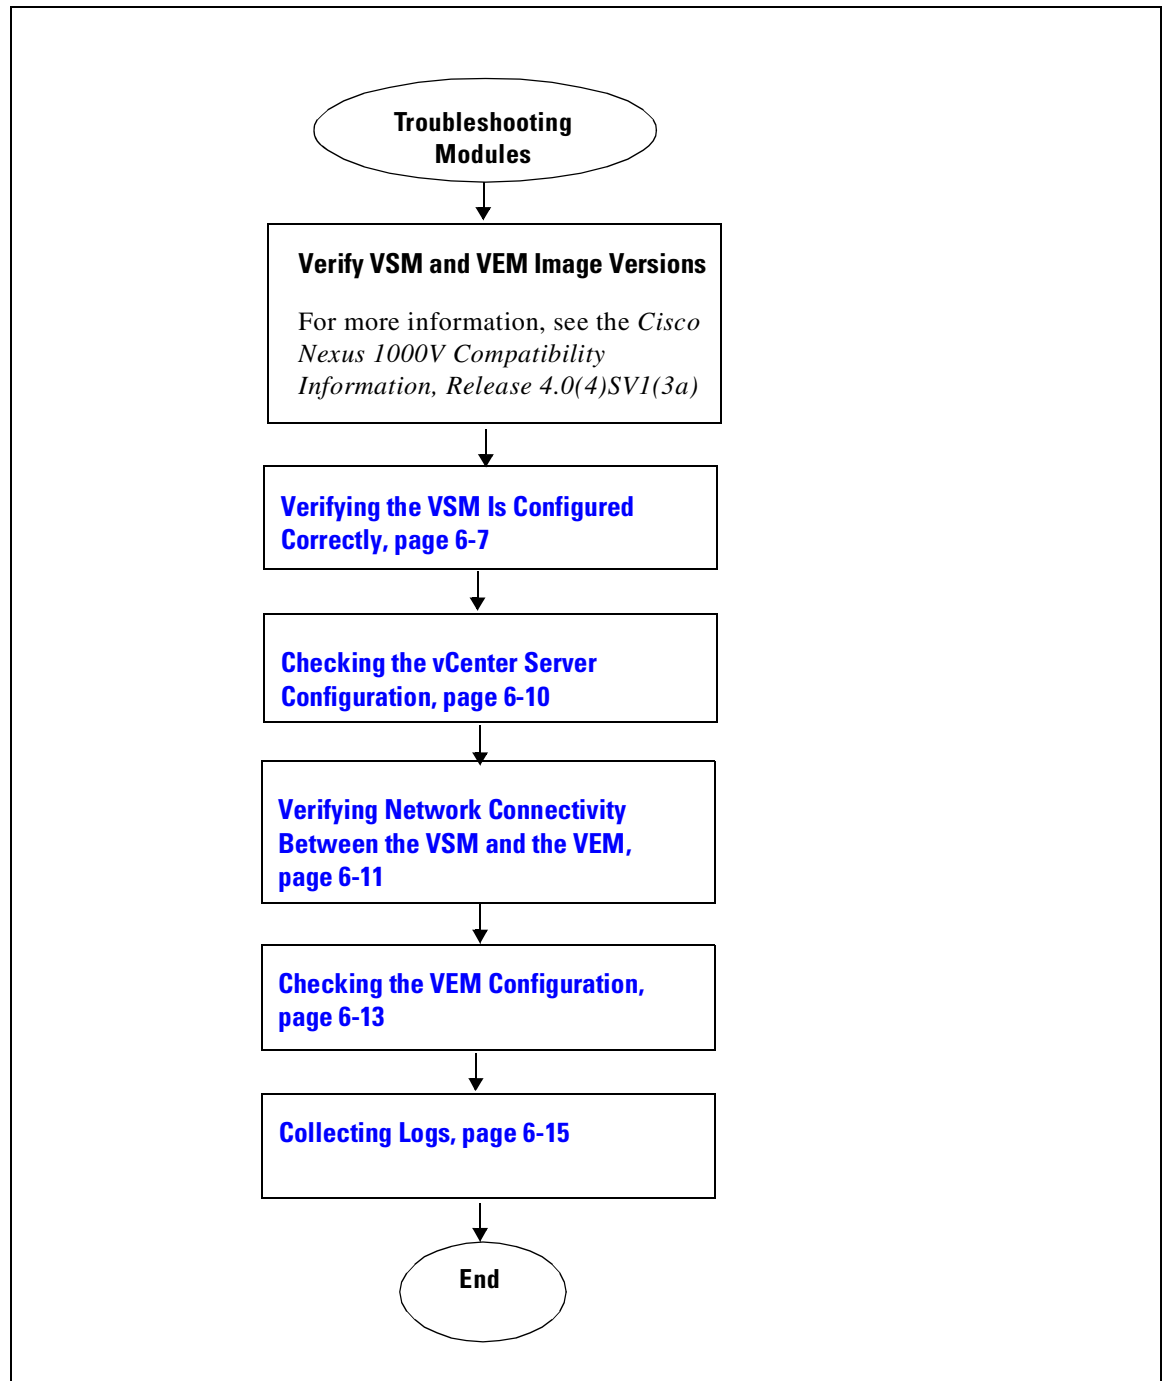
- If you see the following error, see the [“Checking the vCenter Server Configuration” procedure on page 6-10](#).
 - ERROR: Datacenter not found
- For a list of terms used with Cisco Nexus 1000V, see the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

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

Flow Chart for Troubleshooting Modules

Use the following flowchart to troubleshoot modules.

Flowchart: Troubleshooting Modules



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

Problems with the VSM

The following are symptoms, possible causes, and solutions for problems with the VSM.

Table 6-1 Problems with the VSM

Symptom	Possible Causes	Solution
<p>You see the following error on the VSM:</p> <pre>ERROR: [VMware vCenter Server 4.0.0 build-150489] Extension key was not registered before its use</pre>	An extension or plug-in was not created for the VSM.	<ol style="list-style-type: none"> 1. Verify the extension or plug-in was created. “Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8 2. If the plug-in is not found, then create one using the following procedure in the <i>Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)</i>: Creating a Cisco Nexus 1000V Plug-In on the vCenter Server
Following a reboot of the VSM, the system stops functioning in one of the following states and does not recover on its own. Attempts to debug fail.		
After boot, VSM stopped at “Loader Loading.”	Corrupt boot menu file.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk unformatted and bringup new image. 3. Do one of the following: <ul style="list-style-type: none"> • If you have a saved backup copy of your configuration file, restore the configuration on the VSM. copy source filesystem: filename system:running-config • If not, reconfigure the VSM using the following section in the <i>Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)</i>: Setting Up the Software
If using VSM serial port, after boot, VSM in loader prompt.	Corrupt VSM kickstart image.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM with the same software version currently installed. 2. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk unformatted and bringup new image. Follow the VSM installation procedure. 3. Verify the boot variables are configured properly: show boot 4. If not, set the boot variables used to boot the VSM: boot system bootflash:system-boot-variable-name boot kickstart bootflash:kickstart-boot-variable-name 5. Save the configuration. copy running-config startup-config 6. Reload. reload

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

Table 6-1 Problems with the VSM (continued)

Symptom	Possible Causes	Solution
If using VSM serial port, after boot, VSM in loader prompt. (continued)	Boot variables are not set.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk unformatted and bringup new image. 3. Verify the boot variables are configured properly: show boot 4. If not, set the boot variables used to boot the VSM: boot system bootflash:system-boot-variable-name boot kickstart bootflash:kickstart-boot-variable-name 5. Save the configuration. copy running-config startup-config
After boot, VSM in boot prompt.	Corrupt VSM system image.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 1, Install Nexus1000v and bring up new image. 3. Follow the VSM installation procedure.
After boot, VSM re-configured.	Startup configuration is deleted.	<p>Do one of the following:</p> <ul style="list-style-type: none"> • If you have a saved backup copy of your configuration file, restore the configuration on the VSM. copy source filesystem: filename system:running-config • If not, reconfigure the VSM using the following section in the <i>Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)</i>: Setting Up the Software
Standby VSM reloads unexpectedly and it does not come back up as HA Standby.	Control VLAN or control interface down	<ol style="list-style-type: none"> 1. Verify the standby VSM has booted up. show module 2. Verify active communication between VSMs. show system internal redundancy trace 3. If the output indicates either of the following, then troubleshoot the communication over the control VLAN. DEGRADED_MODE PRI_HB_STOPPED <p>Check for the following:</p> <ul style="list-style-type: none"> – Links down between VSMs. – Broadcast traffic not forwarded between VSMs. – Control VLAN not properly configured. – Port profiles not defined properly, or wrong port profiles chosen for VSM or host where they reside.

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

Table 6-1 Problems with the VSM (continued)

Symptom	Possible Causes	Solution
After boot, secondary VSM reboots continuously.	Active and standby VSMs fail to synchronize, remote storage is used, and the active has re-mounted its file system as read only.	<ol style="list-style-type: none"> From active VSM, check for a system message indicating the file system was remounted. show logging If you see the following message in the output, then reload the active VSM. %KERN-2-SYSTEM_MSG: Remounting filesystem read-only - kernel
	Active and standby VSMs fail to synchronize, and a process crashes at the standby.	<ol style="list-style-type: none"> From active VSM, look for a process crash. show cores If the output indicates a process crash, then see the “Before Contacting Technical Support” section on page 19-1 to open a support case.

Verifying the VSM Is Connected to the vCenter Server

You can use the following procedure to verify that the VSM is connected to the vCenter Server.

Step 1 Verify the connection between the VSM and vCenter Server.

show svcs connections

The output should indicate that the operational status is **Connected**.

Example:

```
n1000v# show svcs connections
connection vc:
  ip address: 172.23.231.223
  protocol: vmware-vim https
  certificate: user-installed
  datacenter name: hamilton-dc
  DVS uuid: 92 7a 14 50 05 11 15 9c-1a b0 f2 d4 8a d7 6e 6c
  config status: Disabled
  operational status: Disconnected
```

Step 2 Do one of the following:

- If the status is **Connected**, then return to the [Flowchart: Troubleshooting Modules, page 6-3](#).
- If not, then continue with the next step.

Step 3 Connect to the vCenter Server.

config t

svcs connection *datacenter_name*

connect

Example:

```
n1000v# conf t
n1000v(config)# svcs connection HamiltonDC
n1000v(config-svs-conn)# connect
```

Example:

```
n1000v# conf t
n1000v(config)# svcs connection HamiltonDC
```

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

```
n1000v(config-svs-conn)# connect
ERROR: [VMWARE-VIM] Extension key was not registered before its use.
```

Step 4 Do one of the following:

- If you see an error message about the Extension key, continue with the next step [Table 6-1](#).
- If not, go to [Step 6](#).

Step 5 Do the following and then go to [Step 6](#).

- Unregister the extension key using the “[Unregistering the Extension Key in the vCenter Server](#)” procedure on page 3-12.
- Install a new extension key using the following procedure in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.
 - [Creating a Cisco Nexus 1000V Plug-In on the vCenter Server](#)

Step 6 Verify the connection between the VSM and vCenter Server.

show svcs connections

The output should indicate that the operational status is **Connected**.

Example:

```
n1000v# show svcs connections
connection vc:
  ip address: 172.23.231.223
  protocol: vmware-vim https
  certificate: user-installed
  datacenter name: hamilton-dc
  DVS uuid: 92 7a 14 50 05 11 15 9c-1a b0 f2 d4 8a d7 6e 6c
  config status: Disabled
  operational status: Disconnected
```

Step 7 Do one of the following:

- If the status is **Connected**, then you have completed this procedure.
- If not, then return to the [Flowchart: Troubleshooting Modules, page 6-3](#).

Verifying the VSM Is Configured Correctly

This section includes the following procedures to verify the VSM configuration.

- [Verifying the Domain Configuration, page 6-7](#)
- [Verifying the System Port Profile Configuration, page 6-8](#)
- [Verifying the Control and Packet VLAN Configuration, page 6-9](#)

Verifying the Domain Configuration

You can use the following procedure to verify the domain configuration.

BEFORE YOU BEGIN

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

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

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

- The output of the show **svs domain** command should indicate the following:
 - The presence of a control VLAN and a packet VLAN.
 - The domain configuration was successfully pushed to VC.

Step 1 On the VSM, verify the domain configuration.

show svs domain

Example:

```
n1000v# show svs domain
SVS domain config:
  Domain id:      682
  Control vlan:   3002
  Packet vlan:    3003
  L2/L3 Control VLAN mode: L2
  L2/L3 Control VLAN interface: mgmt0
  Status: Config push to VC successful
```

Verifying the System Port Profile Configuration

You can use this procedure to verify that there is at least one port profile configured as follows:

- The port profile type is Ethernet.
- The port profile has a system VLAN.
- The port profile is configured in trunk mode.
- The list of allowed VLANs is either a superset of or the same as the list of system VLANs.

For more information about system port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)*.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You know the name of your system port profile.

DETAILED STEPS

Step 1 On the VSM, verify the system port profile configuration.

show port-profile name *system-port-profile-name*

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

```

Example:
n1000v# show port-profile name system-uplink
port-profile system-uplink
  description: "System profile for critical ports"
  type: ethernet
  status: enabled
  capability l3control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: 260
  port-group: system-uplink
  max ports: -
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 260
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 260
    no shutdown
  assigned interfaces:
n1000v(config-port-prof)#

```

Verifying the Control and Packet VLAN Configuration

You can use the following procedure to verify that the control and packet VLANs are configured on the VSM.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- You know the VLAN IDs of the control and packet VLANs

The control and packet VLANs can be the same VLAN or separate VLANs.

If you do not know the VLAN IDs, use the **show running-config** command.

- The control and packet VLANs are also system VLANs in the system port profile.

Step 1 On the VSM, verify the control and packet VLANs.

show running-config vlan *vlan-id*

```

Example:
n1000v# show running-config vlan 260-261
version 4.0(4)SV1(3)
vlan 260
  name cp_control
vlan 261
  name cp_packet
n1000v#

```

```

Example:
n1000v# show running-config
version 4.0(4)SV1(3)
. . .

```

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

```
vlan 260
  name cp_control
vlan 261
  name cp_packet
. . .
n1000v#
```

Step 2 On the VSM, verify that the control and packet VLANs are also system VLANs in the system port profile.

show port-profile name *system-port-profile-name*

Example:

```
n1000v# show port-profile name system-uplink
port-profile system-uplink
  description: "System profile for critical ports"
  type: ethernet
  status: enabled
  capability l3control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: 260
  port-group: system-uplink
  max ports: -
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 260
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 260
    no shutdown
  assigned interfaces:
n1000v(config-port-prof)#
```

Step 3 Do one of the following:

- If the control and packet VLANs are present and are system VLANs in the system port profile, you have completed this procedure.
- If not, configure them using the following procedure in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

Configuring the System Port Profile for VSM-VEM Communication

Checking the vCenter Server Configuration

You can use the following procedure from vSphere client to verify the configuration on the vCenter Server.

Step 1 Confirm that the host is added to the data center and the **n1000V** DVS in that data center.

Step 2 Confirm that at least one pnic of the host is added to the DVS, and that pnic is assigned to the **system-uplink** profile.

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

- Step 3** Confirm that the three VSM vnics are assigned to the port groups containing the control VLAN, packet VLAN, and management network.

Verifying Network Connectivity Between the VSM and the VEM

You can use the following procedure to verify Layer 2 network connectivity between the VSM and VEM. you are logged in to the ESX host.
one vem per esx host

BEFORE YOU BEGIN

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

- You are logged in to the VSM CLI in EXEC mode.
- You are logged in to the VEM.
- You are logged in to the CLI of the upstream switch.

- Step 1** On the VSM, find its MAC address.

show svcs neighbors

The VSM MAC address displays as the AIPC Interface MAC.

The user VEM Agent MAC address of the host displays as the Src MAC.

Example:

```
n1000v# show svcs neighbors
```

```
Active Domain ID: 1030
```

```
AIPC Interface MAC: 0050-568e-58b7
```

```
Inband Interface MAC: 0050-568e-2a39
```

Src MAC	Type	Domain-id	Node-id	Last learnt (Sec. ago)
0002-3d44-0602	VEM	1030	0302	261058.59

- Step 2** Do one of the following:

- If the output of the **show svcs neighbors** command in [Step 1](#) does not display the VEM MAC address, then there is a problem with connectivity between the server hosting the VSM and the upstream switch. Recheck the VSM configuration and vCenter Server configuration.
- Otherwise, continue with the next step.

- Step 3** On the VEM, run the vem-health script using the VSM MAC address you found in [Step 1](#).



Note If the vem-health script is not in the PATH, you can find it under `/usr/lib/ext/cisco/nexus/vem*/sbin/`.

vem-health check `vsm_mac_address`

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

The vem-health script output shows the cause of the connectivity problem and recommends next steps in troubleshooting.

Example:

```
~ # vem-health check 00:50:56:a3:36:90
VSM Control MAC address: 00:50:56:a3:36:90
Control VLAN: 90
DPA MAC: 00:02:3d:40:5a:03
```

```
VSM heartbeats are not reaching the VEM.
Your uplink configuration is correct.
Recommended action:
Check if the VEM's upstream switch has learned the VSM's Control MAC.
```

Step 4 Do one of the following:

- If the VEM health check in [Step 3](#) indicates a problem with connectivity to the upstream switch, continue with the next step.
- Otherwise, go to [Step 7](#).

Step 5 On the upstream switch, display the MAC address table to verify the network configuration.

Example:

```
switch# show mac address-table interface Gi3/1 vlan 3002
Legend: * - primary entry
         age - seconds since last seen
         n/a - not available
```

vlan	mac address	type	learn	age	ports
Active Supervisor:					
* 3002	0050.56be.7ca7	dynamic	Yes	0	Gi3/1

```
switch# show mac address-table interface Gi3/2 vlan 3002
Legend: * - primary entry
         age - seconds since last seen
         n/a - not available
```

vlan	mac address	type	learn	age	ports
Active Supervisor:					
* 3002	00:02:3d:40:0b:0c	dynamic	Yes	0	Gi3/2

Step 6 Do one of the following:

- If the output from [Step 5](#) does not display the MAC address of the VSM, then there is a problem with connectivity between the server hosting the VSM and the upstream switch. Recheck the VSM configuration and vCenter Server configuration.
- Otherwise, continue with the next step.

Step 7 On the VEM, enter the following commands to verify that the VSM MAC appears in the control and packet VLANs.

config t

module vem module_number execute vemcmd show l2 control_vlan_id

module vem module_number execute vemcmd show l2 packet_vlan_id

The VSM eth0 and eth1 MAC addresses should display in the host control and packet VLANs.

Example:

```
n100v# config t
```

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

```
n1000v(config)# module vem 3 execute vemcmd show 12 3002
Bridge domain 3002 brtmax 100, brtcnt 3, timeout 120
  Dynamic MAC 00:50:56:be:7c:a7 LTL 16 pvlan 0 timeout 110
  Dynamic MAC 00:02:3d:40:0b:0c LTL 10 pvlan 0 timeout 110

n1000v(config)# module vem 3 execute vemcmd show 12 3003
Bridge domain 3002 brtmax 100, brtcnt 3, timeout 120
  Dynamic MAC 00:50:56:be:7c:a7 LTL 16 pvlan 0 timeout 110
  Dynamic MAC 00:02:3d:20:0b:0c LTL 10 pvlan 0 timeout 110
```

Step 8 Do one of the following:

- If the MAC address of the VSM does not appear in the output of [Step 7](#), then check the VEM configuration as explained in “[Checking the VEM Configuration](#)” section on page 6-13.
- Otherwise, you have completed this procedure.

Checking the VEM Configuration

You can use the following procedure to verify that the ESX host received the VEM configuration and setup.

BEFORE YOU BEGIN

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

- You are logged in to the VEM.
- The port LTL number is the port index number, for example, 48.

Step 1 On the ESX host, use the following command to confirm that the VEM Agent is running, and that the correct host uplinks are added to the DVS.

vem status

Example:

```
~ # vem status
VEM modules are loaded
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	3	64	1500	vmnic0
DVS Name	Num Ports	Used Ports	Configured Ports	Uplinks	
n1000v	256	9	256		vmnic1 VEM Agent is running

Step 2 Use the following commands to restore connectivity that is lost due to an incorrect MTU value on an uplink:

vemcmd show port *port-LTL-number*

vemcmd set mtu *size ltl port-LTL-number*

The MTU is reset on the ESX causing the connectivity to be restored.

Example:

```
~ # vemcmd show port 48
LTL   IfIndex  Vlan   Bndl  SG_ID  Pinned_SGID  Type  Admin State  CBL Mode
Name  . . .
17   1a030100   1 T    304    1          32  PHYS      UP    UP    1  Trunk vmnic1
```

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

```
~# vemcmd set mtu 9000 ltl 17
```



Note Use these **vemcmds** only as a recovery measure and then update the MTU value in the port profile configuration for system uplinks or in the interface configuration for non-system uplinks.

Step 3 Use the following command to verify that the domain ID, control VLANs, and packet VLANs are configured correctly on the host.

vemcmd show card

Example:

```
~ # vemcmd show card
Card UUID type 2: 58f8afd7-e1e3-3c51-85e2-6e6f2819a7b8
Card name: sfish-srvr-1
Switch name: n1000v
Switch alias: DvsPortset-0
Switch uuid: 56 e0 36 50 91 1c 32 7a-e9 9f 31 59 88 0c 7f 76
Card domain: 1024
Card slot: 4
VEM Control (Control VLAN) MAC: 00:02:3d:14:00:03
VEM Packet (Inband) MAC: 00:02:3d:24:00:03
VEM Control Agent (DPA) MAC: 00:02:3d:44:00:03
VEM SPAN MAC: 00:02:3d:34:00:03
Management IP address: 172.23.232.102
Max physical ports: 32
Max virtual ports: 216
Card control VLAN: 3002
Card packet VLAN: 3003
    Processors: 4
    Processor Cores: 4
    Processor Sockets: 2
    Physical Memory: 4290351104
```

Step 4 Use the following command to verify that the ports of the host added to the DVS are listed, and that the ports are correctly configured as access or trunk on the host.

vemcmd show port

Example:

```
~ # vemcmd show port
LTL    IfIndex  Vlan    Bndl  SG_ID Pinned_SGID  Type  Admin State  CBL Mode  Name
8      0        3969    0      2      2  VIRT  UP    UP    4 Access 120
9      0        3969    0      2      2  VIRT  UP    UP    4 Access 121
10     0        3002    0      2      2  VIRT  UP    UP    4 Access 122
11     0        3968    0      2      2  VIRT  UP    UP    4 Access 123
12     0        3003    0      2      2  VIRT  UP    UP    4 Access 124
13     0         1      0      2      2  VIRT  UP    UP    0 Access 125
14     0        3967    0      2      2  VIRT  UP    UP    4 Access 126
16    1a030100  1 T     0      2      2  PHYS  UP    UP    4 Trunk vmnic1
```

The last line of output indicates that vmnic1 should be in Trunk mode, with the CBL value of 4. The CBL value of the native VLAN does not have to be 4. It will be 0 if it is not allowed, or 1 if it is VLAN 1 and not allowed. This is not an issue unless the native VLAN is the Control VLAN VLAN. The Admin state and Port state should be UP.

Step 5 Use the following commands to verify in the broadcast domain that the vmnic port that is supposed to carry the control VLAN and packet VLAN is present.

```
vemcmd show bd control_vlan
vemcmd show bd packet_vlan
```

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

```

Example:
~ # vemcmd show bd 3002
BD 3002, vdc 1, vlan 3002, 2 ports
Portlist:
  10 122
  16 vmn1c1
~ # vemcmd show bd 3003
BD 3003, vdc 1, vlan 3003, 2 ports
Portlist:
  12 124
  16 vmn1c1

```

Step 6 Use the **vemcmd show trunk** command to verify the following:

- The control and packet VLANs are shown in the command output, indicating that the DV port groups are successfully pushed from the vCenter Server to the host.
- The correct physical trunk port vmnic is used.

```

Example:
~ # vemcmd show trunk
Trunk port 16 native_vlan 1 CBL 4vlan(1) cbl 4, vlan(3002) cbl 4, vlan(3003) cbl 4,

```

At least one physical uplink must be carrying the control and packet VLANs. If more than one uplink is carrying the control and packet VLANs, the uplinks must be in a port channel profile. The port channel itself would not be visible because the VEM is not yet added to the VSM.

Step 7 Restore connectivity that is lost due to incorrect port and system VLAN settings.

```
vemcmd show port port-ltl-number
```

```
vemcmd set system-vlan vlan_id ltl port-ltl-number
```

```
vemcmd set port-mode [trunk | access] [native-vlan vlan_id] ltl port-ltl-number
```



Note When using the **set port-mode** command, if you specify access port mode, you must set a native VLAN ID. If you specify trunk port mode, the native VLAN ID is optional. If no native VLAN ID is specified, the native VLAN ID is set to non-system VLAN 1.

```

Example:
~ # vemcmd show port 48
LTL   IfIndex  Vlan   Bndl  SG_ID  Pinned_SGID  Type  Admin State  CBL Mode
Name  . . .
48    1b030000  1      0     32     1  VIRT    UP DOWN DOWN  0 Access vnk1
~ # vemcmd set system-vlan 99 ltl 48
~ # vemcmd set port-mode trunk ltl 48

```



Note Use these **vemcmds** only as a recovery measure and then update the port profile configuration with correct system VLANs.

Collecting Logs

After you have verified network connectivity between the VEM and the VSM, you can use the following procedure to collect log files to help identify the problem.

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

Step 1 On the VEM, use the following command to verify its UUID.

vemcmd show card info

Example:

```
~ # module vem 3 vemcmd show card info
Card UUID type 0: 4908a717-7d86-d28b-7d69-001a64635d18
Card name: sfish-srvr-7
Switch name: N1000v
Switch uuid: 50 84 06 50 81 36 4c 22-9b 4e c5 3e 1f 67 e5 ff
Card domain: 11
Card slot: 12
Control VLAN MAC: 00:02:3d:10:0b:0c
Inband MAC: 00:02:3d:20:0b:0c
SPAN MAC: 00:02:3d:30:0b:0c
USER DPA MAC: 00:02:3d:40:0b:0c
Management IP address: 172.28.30.56
Max physical ports: 16
Max virtual ports: 32
Card control VLAN: 3002
Card packet VLAN: 3003
```

Step 2 On the VSM, use the following command to verify the module number to which the corresponding UUID entry is mapped.

show server_info

Example:

```
n1000v# show server_info
Mod      Status          UUID
---      -
13       absent          4908a717-7d86-d28b-7d69-001a64635d18
```

Step 3 Using the module number from [Step 2](#), collect the output of the following commands:

- **show platform internal event-history module 13**
- **show module internal event-history module 13**
- **show system internal im event-history module 13**
- **show system internal vmm event-history module 13**
- **show system internal ethpm event-history module 13**



Note

If you need to contact Cisco TAC for assistance in resolving an issue, you will need the output of the commands listed in [Step 3](#).

Commands

Use the following commands to troubleshoot bringing the VSM module into service:

- **show svcs neighbors** – to show all svcs neighbors
- **show platform internal event-history module** – to display platform manager module state machines

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

Use the following commands to troubleshoot problems with the VSM:

- **show platform internal event-history module** *module-number*
- **show module internal event-history module** *module-number*
- **show system internal im event-history module** *module-number*
- **show system internal vmm event-history module** *module-number*
- **show system internal ethpm event-history module** *module-number*
- **show system internal ethpm event-history int** *type slot*

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



CHAPTER 7

Ports and Port Profiles

This chapter describes how to identify and resolve problems with ports and includes the following topics:

- [Overview, page 7-1](#)
- [Guidelines for Configuring a Port Interface, page 7-2](#)
- [Diagnostic Checklist, page 7-2](#)
- [Viewing the Port State, page 7-3](#)
- [Using Port Counters, page 7-4](#)
- [Port Interface Symptoms and Solutions, page 7-5](#)
- [Port Security, page 7-8](#)
- [Port Profiles, page 7-13](#)
- [Transferring Port Profiles from the VSM to the vCenter Server, page 7-19](#)

Overview

Before a switch can relay frames from one data link to another, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Ethernet (physical) interfaces, virtual Ethernet interfaces, and the management interface (mgmt0),.

Each interface has the following:

- **Administrative Configuration**
The administrative configuration does not change unless you modify it. This configuration has attributes that you can configure in administrative mode.
- **Operational state**
The operational state of a specified attribute, such as the interface speed. This state cannot be changed and is read-only. Some values may not be valid when the interface is down (such as the operation speed).

For a complete description of port modes, administrative states, and operational states, see the *Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(3)*.

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

Guidelines for Configuring a Port Interface

Use the following guidelines when configuring a port interface.

- Using the procedure, [Verifying the Module State, page 7-2](#), make sure that the module is active.

Verifying the Module State

Use this procedure to verify the state of a module.

BEFORE YOU BEGIN

- The output of this command should indicate that the module is OK (active)

DETAILED STEPS

Step 1 From EXEC mode, enter the following command:

```
show module module-number
```

Example:

```
n1000v# show mod 3
Mod  Ports  Module-Type                Model                Status
---  ---
3    248    Virtual Ethernet Module
                                     ok

Mod  Sw          Hw
---  ---
3    NA         0.0

Mod  MAC-Address(es)          Serial-Num
---  ---
3    02-00-0c-00-03-00 to 02-00-0c-00-03-80  NA

Mod  Server-IP          Server-UUID          Server-Name
---  ---
3    192.168.48.20     496e48fa-ee6c-d952-af5b-001517136344  frodo
```

Diagnostic Checklist

Use the following checklist to begin diagnosing port interface activity.

Checklist	
Verify that the VSM is connected to the vCenter Server by using the show sys connections command.	✓
Verify that appropriate port profiles are assigned to the physical NICS and the virtual NICS by verifying the same on the vSphere Client connected to vCenter Server.	

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

Checklist (continued)	✓
Verify that the ports have been created using the show interface brief command.	
Using the procedure, Viewing the Port State, page 7-3 , verify the state of the interface. For more information about port states, see the <i>Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(3)</i> .	

Use the following commands to troubleshoot ports:

- **show interface status**
- **show interfaces capabilities**
- **show system internal ethpm errors**
- **show system internal ethpm event-history**
- **show system internal ethpm info**
- **show system internal ethpm mem-stats**
- **show system internal ethpm msgs**
- **show system internal vim errors**
- **show system internal vim event-history**
- **show system internal vim info**
- **show system internal vim mem-stats**
- **show system internal vim msgs**

Viewing the Port State

Use this procedure to view the port state.

BEFORE YOU BEGIN

- The output of this command includes the following:
 - Administrative state
 - Speed
 - Trunk VLAN status
 - Number of frames sent and received
 - Transmission errors, including discards, errors, CRCs, and invalid frames

DETAILED STEPS

Step 1 From EXEC mode, enter the following command:

```
show interface ethernet slot-number
```

Example:

```
n1000v# show int eth3/2
```

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

```

Ethernet3/2 is up
  Hardware: Ethernet, address: 0050.5653.6345 (bia 0050.5653.6345)
  MTU 1500 bytes, BW -598629368 Kbit, DLY 10 usec,
    reliability 0/255, txload 0/255, rxload 0/255
  Encapsulation ARPA
  Port mode is trunk
  full-duplex, 1000 Mb/s
  Beacon is turned off
  Auto-Negotiation is turned off
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned on
  Switchport monitor is off
    Rx
    18775 Input Packets 10910 Unicast Packets
    862 Multicast Packets 7003 Broadcast Packets
    2165184 Bytes
    Tx
    6411 Output Packets 6188 Unicast Packets
    216 Multicast Packets 7 Broadcast Packets 58 Flood Packets
    1081277 Bytes
    1000 Input Packet Drops 0 Output Packet Drops
    1 interface resets
n1000v#

```

Using Port Counters

Counters can identify synchronization problems by showing a significant disparity between received and transmitted frames.

BEFORE YOU BEGIN

- Create a baseline first by clearing the counters.
The values stored in the counters can be meaningless for a port that has been active for an extended period. Clearing the counters provides a better idea of the actual link behavior at this time.

DETAILED STEPS

Step 1 From EXEC mode, enter the following command to zero out the counters for the interface:

```
clear counters interface ethernet slot-number
```

Example:

```
n1000v# clear counters interface eth 2/45
n1000v#
```

Step 2 Enter the following command to view the port counters:

```
show interface ethernet slot number counters
```

Example:

```
n1000v# show interface eth3/2 counters
```

```

-----
Port                InOctets      InUcastPkts    InMcastPkts    InBcastPkts
-----
Eth3/2              2224326       11226          885             7191

```

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

Port	OutOctets	OutUcastPkts	OutMcastPkts	OutBcastPkts
Eth3/2	1112171	6368	220	7

Port Interface Symptoms and Solutions

This section includes possible causes and solutions for the following symptoms:

- [Cannot Enable an Interface, page 7-5](#)
- [Port Remains in a Link Failure or Not Connected State, page 7-6](#)
- [Link Flapping, page 7-6](#)
- [Port State Is ErrDisabled, page 7-7](#)

Cannot Enable an Interface

Symptom	Possible Cause	Solution
Cannot enable an interface.	Layer 2 port is not associated with an access VLAN or the VLAN is suspended.	Use the show interface brief CLI command to see if the interface is configured in a VLAN. Use the show vlan brief CLI command to determine the status of the VLAN. Use the state active CLI command in VLAN configuration mode to configure the VLAN as active.

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

Port Remains in a Link Failure or Not Connected State

Symptom	Possible Cause	Solution
Port remains in a link-failure state.	Port connection is bad.	Use the show system internal ethpm info CLI command to verify the port status is in link-failure. Use the shut command followed by the no shut command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.
	Link is stuck in initialization state or the link is in a point-to-point state.	Use the show logging CLI command to check for a Link Failure, Not Connected system message. Use the shut CLI command followed by the no shut command to disable and enable the port. If this does not clear the problem, try moving the connection to a different port on the same or another module.
	—	If these steps are inconclusive on the VSM, use the vss-support command to collect the ESX side NIC configuration.

Link Flapping

This section includes the following topics:

- [About the Link Flapping Cycle, page 7-6](#)
- [Troubleshooting Prerequisites, page 7-6](#)
- [Symptoms, Causes, and Solutions, page 7-7](#)

About the Link Flapping Cycle

When a port is flapping, it cycles through the following states, in this order, and then starts over again:

1. Initializing - The link is initializing.
2. Offline - The port is offline.
3. Link failure or not connected - The physical layer is not operational and there is no active device connection.

Troubleshooting Prerequisites

When troubleshooting unexpected link flapping, it is important to know the following information:

- Who initiated the link flap.
- The actual reason for the link being down.

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

Symptoms, Causes, and Solutions

Symptom	Possible Cause	Solution
Unexpected link flapping occurs.	The bit rate exceeds the threshold and puts the port into an error disabled state.	Right-click the port in Device Manager and select disable and then enable , or use the shut CLI command followed by the no shut command to return the port to the normal state.
	Some problem in the switch triggers the link flap action by the end device. Some of the causes are: <ul style="list-style-type: none"> • Packet drop in the switch, because of either a hardware failure or an intermittent hardware error. • Packet drop resulting from a software error. • A control frame is erroneously sent to the device. 	Determine link flap reason as indicated by the MAC driver. Use the debug facilities on the end device to troubleshoot the problem. An external device may choose to initialize the link again when encountering the error. If so, the exact method of link initialization varies by device.
	The link flapping can be caused by ESX errors, or link flapping on the upstream switch.	

Port State Is ErrDisabled

This section includes the following topics:

- [About the ErrDisabled Port State, page 7-7](#)
- [Verifying the ErrDisable State, page 7-7](#)
- [Verifying the ErrDisable State, page 7-7](#)

About the ErrDisabled Port State

The ErrDisabled state indicates that the switch detected a problem with the port and disabled the port. This state could be caused by a flapping port or a high amount of bad frames (CRC errors), potentially indicating something wrong with the media.

Verifying the ErrDisable State

To resolve the ErrDisable state using the CLI, follow these steps:

-
- Step 1** Use the **show interface** command to verify that the switch detected a problem and disabled the port. Check cables.
- ```
n1000v# show interface e1/14
e1/7 is down (errDisabled)
```
- Step 2** Use the **show port internal event-history interface** command to view information about the internal state transitions of the port. In this example, porte1/7 entered the ErrDisabled state because of a capability mismatch, or “CAP MISMATCH.” You might not know how to interpret this event, but you can look for more information with other commands.
- ```
n1000v# show port internal event-history interface e1/7
```

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

```
>>>>FSM: <e1/7> has 86 logged transitions<<<<<
1) FSM:<e1/7> Transition at 647054 usecs after Tue Jan  1 22:44..
   Previous state: [PI_FSM_ST_IF_NOT_INIT]
   Triggered event: [PI_FSM_EV_MODULE_INIT_DONE]
   Next state: [PI_FSM_ST_IF_INIT_EVAL]
2) FSM:<e1/7> Transition at 647114 usecs after Tue Jan  1 22:43..
   Previous state: [PI_FSM_ST_IF_INIT_EVAL]
   Triggered event: [PI_FSM_EV_IE_ERR_DISABLED_CAP_MISMATCH]
   Next state: [PI_FSM_ST_IF_DOWN_STATE]
```

Step 3 Use the **show logging logfile** command to display the switch log file and view a list of port state changes. In this example, an error was recorded when someone attempted to add port e1/7 to port channel 7. The port was not configured identically to port channel 7, so the attempt failed.

```
n1000v# show logging logfile
. . .
Jan  4 06:54:04 switch %PORT_CHANNEL-5-CREATED: port-channel 7 created
Jan  4 06:54:24 switch %PORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel 7
is down (No operational members)
Jan  4 06:54:40 switch %PORT_CHANNEL-5-PORT_ADDED: e1/8 added to port-channel 7
Jan  4 06:54:56 switch %PORT-5-IF_DOWN_ADMIN_DOWN: Interface e1/7 is down (Administratively
down)
Jan  4 06:54:59 switch %PORT_CHANNEL-3-COMPAT_CHECK_FAILURE: speed is not compatible
Jan  4 06:55:56 switch%PORT_CHANNEL-5-PORT_ADDED: e1/7 added to port-channel 7
```

MTU Configuration Failure on ESX

An MTU configuration fails if you attempt to configure a value that is not supported by the physical NIC on the ESX host and VEM. A system message similar to the following is generated as a warning:

```
2010 Nov 15 04:35:27 my-n1k %VEM_MGR-SLOT3-1-VEM_SYSLOG_ALERT: vssnet :
sf_platform_set_mtu: Failed setting MTU for VMW port wiht portID 33554475.
```

Port Security

The port security feature allows you to secure a port by limiting and identifying the MAC addresses that can access the port. Secure MACs can be manually configured or dynamically learned.

There are two type of security violations:

- Addr-Count-Exceed Violation
- MAC Move Violation

The following port types support port security:

- VEthernet access ports
- VEthernet trunk ports

VEthernet SPAN destination ports do not support port security. In addition, port security is not supported on standalone Ethernet interfaces or on members of a Port Channel.

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

Troubleshooting Port Security Problems

This section describes how to troubleshoot the following connectivity issues when you have port security enabled on an interface:

- Cannot Ping from a VM with Port Security Enabled
- Port Enabled with Port Security is Error Disabled

Cannot Ping from a VM with Port Security Enabled

If you cannot send a ping from a VM with port security enabled, follow these steps:

- Step 1** Enter the **module vem 3 execute vemcmd show portsec stats** command to view the actual port security configuration applied on the port.

Syntax: **module vem vem number execute vemcmd show portsec stats**

```
n1000V#module vem 3 execute vemcmd show portsec stats
LTL  if_index  cp-cnt  Max      Aging  Aging  DSM  Sticky  VM
      Secure   Addresses
      Time    Type    Bit   Enabled Name
47   1b020000  0       1       0     Absolute Clr    No   VM-Pri.eth1
```

The output shows that port security is enabled on the interface with LTL 47 connected to the Network Adapter 1 of the VM-Pri Virtual Machine

In addition, it shows other port security configuration attributes: Maximum No of Secured Addresses is 1, Aging Type is Absolute, Aging Time is 0 seconds (which means aging is disabled), and Sticky MAC is disabled.



Caution

If Drop on Source Miss (DSM) is set, it means that no new MAC addresses can be learned by this port.

To clear the DSM bit, enter the **no port-security stop learning** command on the VSM:

```
n1000V# no port-security stop learning
```

If the DSM bit is not set, proceed to step 2.

- Step 2** Log in to the ESX Host containing the VM and enter the **module vem 3 execute vemcmd show portsec macs all** command to view all secure MACs on that VEM.

```
~ #module vem 3 execute vemcmd show portsec macs all
VLAN 65's Secure MAC list:
  cp MAC 08:66:5c:99:72:f2 LTL 48 timeout 960
```

cp means currently being processed, which means that the packet is not yet acknowledged by the port security process running on the VSM.

This verification notification is sent over the inband channel.

Because the verification notification is sent through the inband channel, the inband VLAN must be on one of the uplink ports on the VEM as well as the corresponding ports on the upstream switch.

- Step 3** Use the **show svcs domain** command to find out the packet VLAN (inband VLAN)

```
n1000v(config-port-prof)# show svcs domain
SVS domain config:
  Domain id: 559
```

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

```
Control vlan: 3002
Packet vlan: 3003
L2/L3 Aipc mode: L2
L2/L3 Aipc interface: mgmt0
Status: Config push to VC successful.
```

In this output, the packet VLAN is 69

Step 4 Verify that the packet VLAN is allowed on any of the uplink ports of the VEM.

Assume there is one uplink and it is bound to a port-profile uplink-profile. Enter the **show port-profile na uplink-all** command:

```
n1000v# show port-profile na uplink-all
port-profile uplink-all
description:
type: vethernet
status: enabled
capability l3control: no
pinning control-vlan: -
pinning packet-vlan: -
system vlans: 69-69
port-group: uplink-all
max ports:
inherit: port-profile xyz
config attributes:
  switchport mode trunk
  switchport access vlan 1
  switchport trunk allowed vlan 1, 68-69,231-233
  channel-group auto mode on sub-group cdp
  no shutdown
evaluated config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 1,68-69,231-233
  channel-group auto mode on sub-group cdp
  no shutdown
assigned interfaces:
  Ethernet3/2
```

As shown in the output, the uplink profile is assigned to Ethernet 3/2 and the inband VLAN (69) is allowed on the port. If it is not, add the packet VLAN (69) to the allowed VLAN list.

Step 5 Enter the **show cdp neighbors** command to find out the upstream neighbors connected to Ethernet interface 3/2.

```
n1000V#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	Eth3/2	149	R S I	WS-C6506-E	Gig1/38

The output shows that Ethernet interface 3/2 is connected to the switch n1000v-6k-2 on Gigabit interface 1/38.

Log in to n1000v-6k-2 and verify that the packet VLAN is allowed on the port.

```
n1000v-6k-2#show running-config interface gigabitEthernet 1/38
Building configuration...
```

```
Current configuration : 161 bytes
!
```

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

```
interface GigabitEthernet1/38
description sfish-srvr-100:vmnic1
switchport
switchport trunk allowed vlan 1,60-69,231-233
switchport mode trunk
end
```

The output shows that the packet VLAN 69 is allowed on the port. If it is not, add the packet VLAN to the allowed VLAN list.

Port Enabled with Port Security is Error Disabled

The ErrDisabled state of a port indicates that the VSM detected a problem with the port and disabled the port. Port security could be responsible for error disabling the port for the following reasons:

- Address Count Exceed Violation
- MAC Move Violation

Address Count Exceed Violation

This issues occurs when more than the configured maximum number of secured addresses are seen on the port. The default violation action is to error disable the port. One way to discover this is to use a **grep** command for the search pattern PORT-SECURITY-2- on the output of a **show logging** command.

```
n1000v#show port-security address interface vethernet 1
Total Secured Mac Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 8192
```

```
-----
Secure Mac Address Table
-----
Vlan    Mac Address           Type           Ports           Remaining age
              (mins)
-----
 65     0050.56B7.7DE2       DYNAMIC       Vethernet1      0
=====
```

The output shows that MAC 0050.56B7.7DE2 is secured on veth1.

```
n1000V#show port-security
Total Secured Mac Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 8192
```

```
-----
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
              (Count)         (Count)         (Count)
-----
Vethernet1   1                0                0                Shutdown
=====
```

The Max Secured Address is 1.

Another MAC E276.DECF.7DE2 appears on VEthernet 1. Now the port is error disabled.

```
n1000v# show logging | inc "PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLATION_MAX_MAC_VLAN"
```

```
2008 Dec 20 21:33:44 N1KV %PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLATION_MAX_MAC_VLAN:
Port Vethernet1 moved to SHUTDOWN state as host E276.DECF.7DE2 is trying to access the
port in vlan 65
```

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

MAC Move Violation

A MAC Move Violation occurs when a MAC that is already secured on one port, such as port A, is seen on another secure port, such as port B.

```
n1000v#show port-security address interface vethernet 1
Total Secured Mac Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 8192
```

```
-----
                        Secure Mac Address Table
-----
Vlan    Mac Address                Type                Ports                Remaining age
-----  -
65      0050.56B7.7DE2            DYNAMIC            Vethernet1          0
=====
```

The output shows that MAC 0050.56B7.7DE2 is secured on veth1

```
n1000v#show port-security
Total Secured Mac Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 8192
```

```
-----
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
              (Count)        (Count)      (Count)
-----
Vethernet1   1              0            0                  Shutdown
=====
```

The output shows the Max Secured Address is 1.

MAC E276.DECF.7DE2 appears on VEthernet 1. Now the port is error disabled.

```
n1000v# show logging | inc "PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLATION_MAX_MAC_VLAN"

2008 Dec 20 21:33:44 N1KV
%PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLATION_MAX_MAC_VLAN: Port
Vethernet1 moved to SHUTDOWN state as host E276.DECF.7DE2 is trying to access the port in vlan 65
```

Port Security Restrictions and Limitations

When troubleshooting port security issues, make sure you follow these guidelines:

- Dynamic secure MACs cannot be cleared using the **clear mac address-table** command. Use the **clear port-security** command instead.
- Port security cannot be enabled on a Veth on a VLAN if there are static MACs configured on the same VLAN. You need to delete any static MACs that are present on the VLAN on any interface to enable port security on a Veth on that VLAN.
- Restrict Violation Action is not supported. Only Shutdown and Protect Violation Modes can be configured as a Port Security Violation Action.

Collecting Debugging Output for Port Security

Use the following commands to troubleshoot port security:

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

- **show port-security**
- **show port-security interface veth**
- **show port -security address**

On the VSM, use the following commands to collect information and troubleshoot port security:

- **show system internal port-security msgs**
- **show system internal port-security errors**
- **show system internal l2fm msgs**
- **show system internal l2fm errors**
- **show system internal l2fm info detail**
- **show system internal pktmgr interface brief**
- **show system internal pktmgr client detail**

Symptoms, Causes, and Solutions

Symptom	Possible Causes	Solution
A ping from the VM fails on an interface that has Port Security enabled on it.	—	<p>Verify that the first packet from the VM has been sent to the VSM.</p> <p>Ensure that the uplink port on the ESX host and the port on the uplink switch is carrying the inband VLAN.</p> <p>Ensure that the uplink port on the ESX port (and the corresponding port on the uplink switch) hosting the CPVA is carrying the inband VLAN.</p> <p>Check that the Veth interface state is up in Packet Manager. If it is not, enter a shutdown command followed by a no shutdown command on the Veth interface.</p>

Port Profiles

Port profiles are used to configure interfaces. A port profile can be assigned to multiple interfaces giving them all the same configuration. Changes to the port profile will be propagated automatically to the configuration of any interface assigned to it.

In the VMware vCenter Server, a port profile is represented as a port group. The VEthernet or Ethernet interfaces are assigned in vCenter Server to a port profile for:

- Defining port configuration by policy.
- Applying a single policy across a large number of ports.
- Supporting both VEthernet and Ethernet ports.

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

Port profiles that are configured as uplinks, can be assigned by the server administrator to physical ports (a vmnic or a pn1c). Port profiles that are not configured as uplinks can be assigned to a VM virtual port.



Note

While manual interface configuration overrides that of the port profile, it is not recommended. Manual interface configuration is only used, for example, to quickly test a change or allow a port to be disabled without having to change the inherited port profile.

For more information about assigning port profiles, see your VMware documentation.

To verify that the profiles are assigned as expected, use the following show commands:

- **show port-profile usage**
- **show running-config interface *interface-id***

The output of the **show running-config interface *interface-id*** command shows a config line such as, `inherit port-profile MyProfile`, indicating the inherited port profile.



Note

Inherited port profiles cannot be changed or removed from an interface using the CLI. This can only be done through the vCenter Server.



Note

Inherited port profiles are automatically configured when the ports are attached on the hosts. This is done by matching up the VMware port group assigned by the system administrator with the port profile that created it.

For detailed information about port profiles, see the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)*.

Troubleshooting Commands for Port Profiles

To collect detailed logs for port profiles, execute the following commands that enable debug logs:

- **debug port-profile trace**
- **debug port-profile error**
- **debug port-profile all**

After enabling the debug log, re-execute a port-profile operation and capture the output in a log file.

Use the following commands to troubleshoot port profiles:

- **show port-profile**

```
n1000v# show port-profile
port-profile UpLinkProfile1
description:
type: vethernet
status: disabled
capability l3control: no
pinning control-vlan: -
pinning packet-vlan: -
system vlans: none
port-group:
max ports: 32
```

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

```

inherit:
config attributes:
  channel-group auto mode on mac-pinning
evaluated config attributes:
  channel-group auto mode on mac-pinning
assigned interfaces:
port-profile UpLinkProfile2
description:
type: vethernet
status: disabled
capability l3control: no
pinning control-vlan: -
pinning packet-vlan: -
system vlans: none
port-group:
max ports: 32
inherit:
config attributes:
  channel-group auto mode on sub-group cdp
evaluated config attributes:
  channel-group auto mode on sub-group cdp
assigned interfaces:
port-profile UpLinkProfile3
description:
type: vethernet
status: disabled
capability l3control: no
pinning control-vlan: -
pinning packet-vlan: -
system vlans: none
port-group:
max ports: 32
inherit:
config attributes:
  channel-group auto mode on sub-group manual
evaluated config attributes:
  channel-group auto mode on sub-group manual
assigned interfaces:n1000v#

```

- **show port-profile expand-interface**

```

n1000v# show port-profile expand-interface

port-profile uplink1
Ethernet3/2
  switchport mode trunk
  switchport trunk allowed vlan 1,110-119
  no shutdown
Ethernet4/2
  switchport mode trunk
  switchport trunk allowed vlan 1,110-119
  no shutdown

port-profile data
Vethernet1
  switchport mode access
  switchport access vlan 118
  no shutdown
n1000v#

```

- **show port-profile usage**

```

n1000v# show port-profile usage

```

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

```
-----
Port Profile          Port          Adapter      Owner
-----
uplink1              Eth3/2        vmn1c1       172.23.232.57
                    Eth4/2        vmn1c1       172.23.232.58
data                 Veth1         Net Adapter 1 ubuntu-2
n1000v#
```

- **show port-profile internal info**

```
n1000v# show port-profile internal info
port-profile Unused_Or_Quarantine_Uplink
  ppid: 00000001
  flags: 00000000
  fsm_state: PPM_PROFILE_FSM_ST_CREATED
  state: enabled
  capability: 00000002
  description: "Port-group created for Nexus1000V internal usage. Do not use."
  alias_id: Unused_Or_Quarantine_Uplink (type=1)
  num_aliases: 1
  alias (type=2):
    name: dvportgroup-1060
    flags: 00000000
  alias name: dvportgroup-1060 type: 2 (pss)
  parent port-profile: none
  num_child_profiles: 0
  num_active_ifs: 0
port-profile Unused_Or_Quarantine_Veth
  ppid: 00000002
  flags: 00000000
  fsm_state: PPM_PROFILE_FSM_ST_CREATED
  state: enabled
  capability: 00000000
  description: "Port-group created for Nexus1000V internal usage. Do not use."
  alias_id: Unused_Or_Quarantine_Veth (type=1)
  num_aliases: 1
  alias (type=2):
    name: dvportgroup-1061
    flags: 00000000
  alias name: dvportgroup-1061 type: 2 (pss)
  parent port-profile: none
  num_child_profiles: 0
  num_active_ifs: 0
port-profile uplink1
  ppid: 00000003
  flags: 00000000
  fsm_state: PPM_PROFILE_FSM_ST_CREATED
  state: enabled
  capability: 00000003
  description: ""
  alias_id: uplink1 (type=1)
  num_aliases: 1
  alias (type=2):
    name: dvportgroup-1062
    flags: 00000000
  alias name: dvportgroup-1062 type: 2 (pss)
  parent port-profile: none
  num_child_profiles: 0
  num_active_ifs: 1
  Ethernet3/2:
    flags: 00000000
    is_active: true
    is_user_configured: false
    bind_count: 1
```


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

```

    is_bound_by_eth_attach: 1
port-profile data
  ppid: 00000005
  flags: 00000000
  fsm_state: PPM_PROFILE_FSM_ST_CREATED
  state: enabled
  capability: 00000000
  description: ""
  alias_id: data (type=1)
  num_aliases: 1
  alias (type=2):
    name: dvportgroup-1064
    flags: 00000000
  alias name: dvportgroup-1064 type: 2 (pss)
  parent port-profile: none
  num_child_profiles: 0
  num_active_ifs: 0
vms info flag: 00000001
n1000v#

```

- **show port-profile internal event-history msgs**

```

n1000v# show port-profile internal event-history msgs
1) Event:E_MTS_RX, length:60, at 553112 usecs after Thu May 14 00:28:52 2009
  [REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X0028B018, Ret:SUCCESS
  Src:0x00000101/3929, Dst:0x00000101/429, Flags:None
  HA_SEQNO:0X00000000, RRtoken:0x0028B018, Sync:NONE, Payloadsize:212
  Payload:
  0x0000: 01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 31 37

2) Event:E_MTS_RX, length:60, at 472402 usecs after Thu May 14 00:28:48 2009
  [REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X0028AF64, Ret:SUCCESS
  Src:0x00000101/3928, Dst:0x00000101/429, Flags:None
  HA_SEQNO:0X00000000, RRtoken:0x0028AF64, Sync:NONE, Payloadsize:212
  Payload:
  0x0000: 01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 31 37

3) Event:E_MTS_RX, length:60, at 897349 usecs after Thu May 14 00:24:59 2009
  [REQ] Opc:MTS_OPC_VSH_CMD_TLV(7679), Id:0X00289DB3, Ret:SUCCESS
  Src:0x00000101/3899, Dst:0x00000101/429, Flags:None
  HA_SEQNO:0X00000000, RRtoken:0x00289DB3, Sync:NONE, Payloadsize:228
  Payload:
  0x0000: 04 03 02 01 e4 00 00 00 00 00 00 00 00 00 00 00

4) Event:E_MTS_RX, length:60, at 171002 usecs after Thu May 14 00:19:27 2009
  [REQ] Opc:MTS_OPC_VSH_CMD_TLV(7679), Id:0X00288A62, Ret:SUCCESS
  Src:0x00000101/3899, Dst:0x00000101/429, Flags:None
  HA_SEQNO:0X00000000, RRtoken:0x00288A62, Sync:NONE, Payloadsize:220
  Payload:
  0x0000: 04 03 02 01 dc 00 00 00 00 00 00 00 00 00 00 00

```

- **show port-profile internal event-history port-profile *profile-name***

```

n1000v# show port-profile internal event-history port-profile data

>>>>FSM: <port-profile/5> has 6 logged transitions<<<<<

1) FSM:<port-profile/5> Transition at 212488 usecs after Mon May 11 19:45:02 2009
  Previous state: [PPM_PROFILE_FSM_ST_NOT_EXISTENT]
  Triggered event: [PPM_PROFILE_FSM_EV_INIT]
  Next state: [PPM_PROFILE_FSM_ST_CREATED]

2) FSM:<port-profile/5> Transition at 212494 usecs after Mon May 11 19:45:02 2009
  Previous state: [PPM_PROFILE_FSM_ST_CREATED]

```

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

```

Triggered event: [PPM_PROFILE_FSM_EV_CFG_CHANGED]
Next state: [PPM_PROFILE_FSM_ST_UPDATING_EVAL_CFG]

3) FSM:<port-profile/5> Transition at 212516 usecs after Mon May 11 19:45:02 2009
Previous state: [PPM_PROFILE_FSM_ST_UPDATING_EVAL_CFG]
Triggered event: [PPM_PROFILE_FSM_EV_EVAL_CFG_CHANGED]
Next state: [PPM_PROFILE_FSM_ST_MSP_HANDSHAKE_CFG_CHANGE]

4) FSM:<port-profile/5> Transition at 212535 usecs after Mon May 11 19:45:02 2009
Previous state: [PPM_PROFILE_FSM_ST_MSP_HANDSHAKE_CFG_CHANGE]
Triggered event: [PPM_PROFILE_FSM_EV_MSP_HANDSHAKE_FAIL]
Next state: [PPM_PROFILE_FSM_ST_UPDATING_CLIENTS]

5) FSM:<port-profile/5> Transition at 212542 usecs after Mon May 11 19:45:02 2009
Previous state: [PPM_PROFILE_FSM_ST_UPDATING_CLIENTS]
Triggered event: [PPM_PROFILE_FSM_EV_UPDATE_DONE]
Next state: [PPM_PROFILE_FSM_ST_WAIT_FOR_CHILD]

6) FSM:<port-profile/5> Transition at 213668 usecs after Mon May 11 19:45:02 2009
Previous state: [PPM_PROFILE_FSM_ST_WAIT_FOR_CHILD]
Triggered event: [PPM_PROFILE_FSM_EV_CHILD_PROFILE_DONE]
Next state: [PPM_PROFILE_FSM_ST_CREATED]

```

System Port Profiles

System port profiles are special port profiles that must be configured before the VSM and the VEM can communicate with each other. System port profiles are used to convey the control and packet VLAN IDs from the VSM to the VEM via the vCenter Server.

When configuring system port profiles, follow these guidelines:

- For trunk ports, the system VLAN list must be a subset of the allowed VLAN list.
- For access ports, there must be one system VLAN, and it must be the same as the access VLAN.
- Issue the **no system vlan** command only when no interface is using the profile.
- Once a system profile is in use by at least one interface, you can only add to the list of system VLANs, but not delete any VLANs from the list.
- For a profile with system VLANs, the **no port-profile** command, the **no vmware port-group** command, and the **no state enabled** command can be issued only when no interface is using the profile.
- The maximum number of port profiles is 128.

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

Port Profiles Symptoms and Solutions

Symptom	Possible Causes	Solution
You do not see a port group on a vCenter Server or see the message Warning: Operation succeeded locally but update failed on vCenter server. Please check if you are connected to vCenter Server.	—	Issue the show svcs connections command to verify that the connection to the vCenter Server is active. The switch output should display Enabled and Connected. Issue the show svcs domain command and check the status for success. Also verify that the following commands have been specified for the port profile: <ul style="list-style-type: none"> • vmware port-group • state enabled
A port configuration is not applied to an interface.	—	Issue the show port-profile usage command to show the interface. Use the show run command and the show port-profile expand-interface command to verify that the interface level configuration did not overwrite the port profile configuration.
An Ethernet interface or Veth interface is administratively down.	The interface is inheriting one of the quarantine port profiles. Use the show port-profile usage command to verify this situation.	Reassign the vnic or pnic to a non-quarantine port group to enable the Veth or Ethernet interface to be up and able to forward traffic. This action requires changing the port group on the vCenter Server.

Transferring Port Profiles from the VSM to the vCenter Server

When transferring a Port Profile from the VSM to the vCenter Server, follow these guidelines:

- Make sure that an Uplink Port Profile (UPP) has the following essential attributes:
 - Uplink capability.
 - System VLANs configured if it is a system port profile.



Note For a privileged profile, make sure you explicitly allow VLANs in the profile if you are configuring trunk mode. Enter the **switchport trunk allowed vlan *your-vlan -list*** command for this type of configuration.

- VMware port group.
- Switchport mode trunk or access.
- No shutdown.

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

- State enabled.
- Make sure you explicitly create any VLANs which you configure in the Port Profiles.



CHAPTER 8

Port Channels and Trunking

Use this chapter to troubleshoot port channels and trunking.

This chapter includes the following topics:

- [Overview, page 8-1](#)
- [Initial Troubleshooting Checklist, page 8-2](#)
- [Troubleshooting Asymmetric Port Channels, page 8-3](#)
- [Cannot Create Port Channel, page 8-4](#)
- [Newly Added Interface Does Not Come Online In a Port Channel, page 8-4](#)
- [VLAN Traffic Does Not Traverse Trunk, page 8-6](#)

Overview

This section includes the following topics:

- [Port Channel Overview, page 8-1](#)
- [Trunking Overview, page 8-2](#)

Port Channel Overview

Port channels aggregate multiple physical interfaces into one logical interface to provide higher bandwidth, load balancing, and link redundancy.

A port channel performs the following functions:

- Increases the aggregate bandwidth on a link by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth usage.
- Provides high availability. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a port channel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The MAC address tables are not affected by link failure.

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

Port Channel Restriction

The following are port channel restrictions.

- Port channels do not support ACLs.
- Port channels do not support NetFlow.

Trunking Overview


Trunking, also known as VLAN trunking, enables interconnected ports to transmit and receive frames in more than one VLAN, over the same physical link.

Trunking and port channels function as follows:

- Port channels enable several physical links to be combined into one aggregated logical link.
- Trunking enables a link to carry (trunk) multiple VLAN traffic.

Initial Troubleshooting Checklist

Use the following checklist to begin troubleshooting port channel and trunking issues:

Checklist	√
Use the show port-channel compatibility-parameters CLI command to determine port channel requirements.	
Ensure that all interfaces in the port channel have the same destination device for LACP channels. By using Asymmetric Port Channel (APC), ports in a ON mode channel can be connected to two different destination devices.	
 Note APC is supported only on mode channels. It is not supported for LACP channels.	
Verify that either side of a port channel is connected to the same number of interfaces.	
Verify that each interface is connected to the same type of interface on the other side.	
Verify that all required VLANs on a trunk port are in the allowed VLAN list.	
Verify that all the members trying to form a port channel are on the same module.	
Verify that the port channel configuration is present in the profile used by the physical ports.	
Configure APC if the ports are connected to different upstream switches.	
If the upstream switch does not support port channels, make sure to configure APC in the profile. In addition, make sure that there are two ports at most in the APC.	

The following commands help troubleshoot port channels and trunking:

- **show port-channel summary**
- **show port-channel internal event-history interface port-channel *channel-number***

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

- **show port-channel internal event-history interface ethernet *slot-number***
- **show system internal ethpm event-history interface port-channel *channel-number***
- **show system internal ethpm event-history interface ethernet *slot-number***
- **show vlan internal trunk interface ethernet *slot-number***
- **show vlan internal trunk interface port-channel *channel-number***
- **debug port-channel error**
- **module vem *module-number* execute vemcmd show port**
- **module vem *module-number* execute vemcmd show pc**
- **module vem *module-number* execute vemcmd show trunk**

Example 8-1 shows output of the **show port-channel summary** command.

Example 8-1 show port-channel summary Command

```
n1000v# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual   H - Hot-standby (LACP only)
        s - Suspended    r - Module-removed
        S - Switched     R - Routed
        U - Up (port-channel)

-----
Group Port-      Type      Protocol  Member Ports
      Channel
-----
1     Po1 (SU)    Eth       NONE      Eth3/4 (P)
2     Po2 (SU)    Eth       NONE      Eth3/2 (P)  Eth3/6 (P)
```

Troubleshooting Asymmetric Port Channels

When troubleshooting asymmetric port channels, follow these guidelines:

- Use APC when you want to configure a port channel whose members are connected to two different upstream switches.
- APC depends on Cisco Discovery Protocol (CDP). Make sure CDP is enabled on VSM and upstream switches.
- Physical ports within an APC get assigned subgroup IDs based on the CDP information received from upstream switches.
- A user can manually configure subgroup IDs in interface configuration submode.
- Make sure that you configured sub-group CDP either with a port profile or on the port channel interface.
- Ports in APC will come up only when they are assigned subgroup IDs manually or through CDP.
- Issue the **show cdp neighbors** command on the VSM and check the output.
- Once the ports came up, check that ports are put in the correct sub-groups by issuing the **module vem *module-number* execute vemcmd show pc** command on the VEM.
- Use the **debug port-channel trace** command to collect information.

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

Cannot Create Port Channel

Symptom	Possible Cause	Solution
Cannot create a port channel.	Maximum number of port channels reached for system.	Use the command, show port-channel summary , to verify the number of port-channels already configured. You can have a maximum of 256 port channels on the Cisco Nexus 1000V.

Newly Added Interface Does Not Come Online In a Port Channel

Symptom	Possible Cause	Solution
Newly added interface does not come online in a port channel.	Port channel mode is on.	<ol style="list-style-type: none"> 1. Make sure you have the port channel configuration in the port profile (port group) used by that interface. 2. Check if there is a port channel already present on the module that is using the same port profile. If there is, check the running configuration on the port channel and the newly added interface. The interface will not come up if the port channel configurations are different. 3. If the port channel configuration is different, apply the difference on the newly added interface. Remove the port, and then add it back.
	Interface parameters are not compatible with those of the existing port.	Use the procedure, Forcing Port Channel Characteristics onto an Interface , page 8-4, to force the physical interface to take on the parameters of the port channel. Use this procedure only if you want to configure the port channel manually and not through the port profile.

Forcing Port Channel Characteristics onto an Interface

Use this procedure to force the physical interface to take on the characteristics of the port channel. Use this procedure only if you want to configure the port channel manually and not through the port profile.

BEFORE YOUR BEGIN

- You are logged in to the CLI in configuration mode.
- The forced interface must have the same speed, duplex, and flow control settings as the channel group.

DETAILED STEPS

Step 1 From CLI configuration mode, enter the following command.

```
interface ethernet slot/port
```

You are placed into interface configuration mode.

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

Example:

```
switch(config)# interface ethernet 1/4  
switch(config-if)
```

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

Step 2 Enter the following command:

channel-group *channel-number* **force**

The physical interface with an incompatible configuration is forced to join the channel group.

Example:

```
switch(config-if)# channel-group 5 force
switch(config-if)
```

Verifying a Port Channel Configuration

Use this procedure to debug port channels configured through a port profile.

BEFORE YOUR BEGIN

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

DETAILED STEPS

- Step 1** Issue the **show port-profile name** *profile-name* command to verify that you have configured a port channel in the profile.
- Step 2** Issue the **show port-channel summary** command.
- Step 3** Issue the **debug port-channel trace** command.

VLAN Traffic Does Not Traverse Trunk

Symptom	Possible Cause	Solution
VLAN traffic does not traverse trunk.	VLAN not in allowed VLAN list.	Add the VLAN to allowed VLAN list. Use the switchport trunk allowed vlan add <i>vlan-id</i> command in the profile used by the interface.



CHAPTER 9

Layer 2 Switching

This chapter describes how to identify and resolve problems that relate to Layer 2 switching.

This chapter includes the following sections:

- [Information About Layer 2 Ethernet Switching, page 9-1](#)
- [Port Model, page 9-1](#)
- [Layer 2 Switching Problems, page 9-4](#)
- [Verifying Layer 2 Switching, page 9-7](#)

Information About Layer 2 Ethernet Switching

Nexus1000V provides a distributed, layer 2 virtual switch that extends across many virtualized hosts.

It consists of two components:

- Virtual Supervisor Module (VSM), which is also known as the Control Plane (CP), acts as the Supervisor and contains the Cisco CLI, configuration, and high-level features.
- Virtual Ethernet Module (VEM), which is also known as the Data Plane (DP), acts as a line card and runs in each virtualized server to handle packet forwarding and other localized functions.

Port Model

This section describes the following port perspectives:


- [Viewing Ports from the VEM, page 9-2](#)
- [Viewing Ports from the VSM, page 9-3.](#)

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

Viewing Ports from the VEM

The Nexus1000V differentiates between virtual and physical ports on each of the VEMs. [Figure 9-1](#) shows how ports on the Nexus1000V switch are bound to physical and virtual VMware ports within a VEM.

Figure 9-1 VEM View of Ports



On the virtual side of the switch, there are three layers of ports that are mapped together:

- **Virtual NICs:** There are three types of Virtual NICs in VMware. The virtual NIC (vnic) is part of the VM, and represents the physical port of the host which is plugged into the switch. The virtual kernel NIC (vmknic) is used by the hypervisor for management, VMotion, iSCSI, NFS and other network access needed by the kernel. This interface would carry the IP address of the hypervisor itself, and is also bound to a virtual Ethernet port. The vswif (not shown) appears only in COS-based systems, and is used as the VMware management port. Each of these types maps to a veth port within Nexus1000V.
- **Virtual Ethernet Ports (VEth):** A VEth port is a port on the Cisco Nexus 1000V Distributed Virtual Switch. Cisco Nexus 1000V has a flat space of VEth ports 0..N. The virtual cable plugs into these VEth ports that are moved to the host running the VM.
VEth ports are assigned to port groups.
- **Local Virtual Ethernet Ports (lveth):** Each host has a number of local VEth ports. These ports are dynamically selected for VEth ports that are needed on the host.
These local ports do not move, and are addressable by the module-port number method.

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

On the physical side of the switch, from bottom to top:

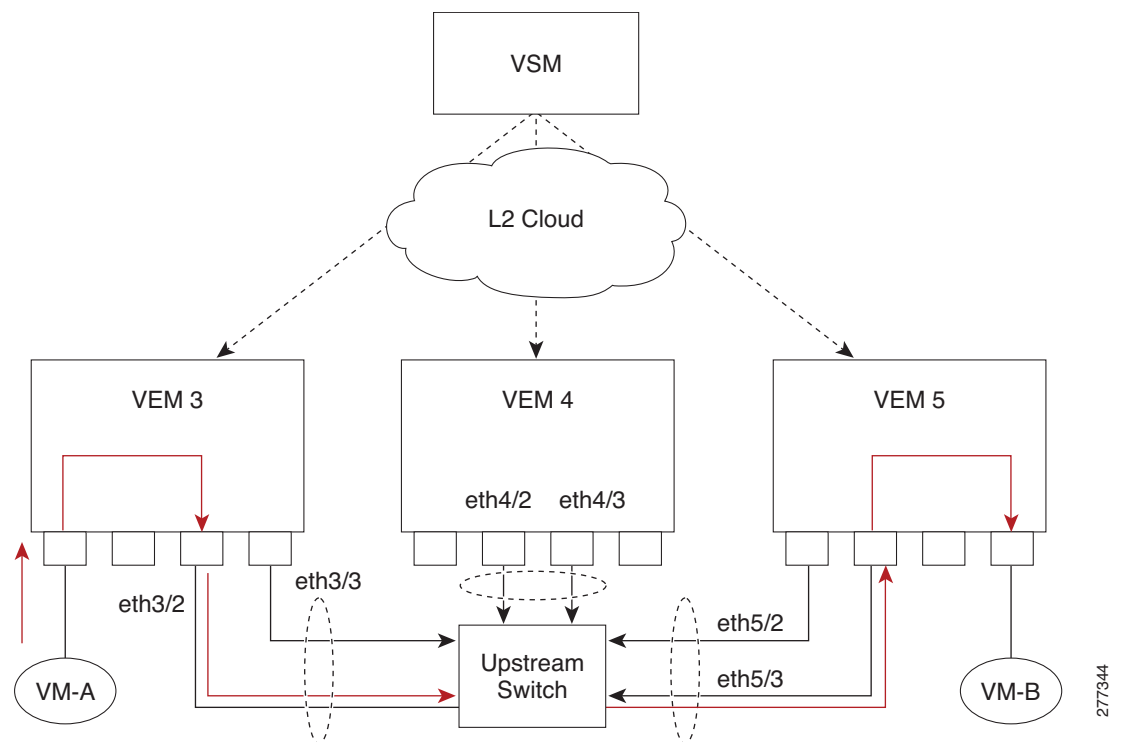
- Each physical NIC in VMware is represented by an interface called a vmnic. The vmnic number is allocated during VMware installation, or when a new physical NIC is installed, and remains the same for the life of the host.
- Each uplink port on the host represents a physical interface. It acts a lot like an lvech port, but because physical ports do not move between hosts, the mapping is 1:1 between an uplink port and a vmnic.
- Each physical port added to Nexus1000V switch appears as a physical Ethernet port, just as it would on a hardware-based switch.

The uplink port concept is handled entirely by VMware, and is used to associate port configuration with vmnics. There is no fixed relationship between the uplink # and vmnic #, and these can be different on different hosts, and can change throughout the life of the host. On the VSM, the Ethernet interface number, such as ethernet 2/4, is derived from the vmnic number, not the uplink number.

Viewing Ports from the VSM

Figure 9-2 shows the VSM view ports.

Figure 9-2 VSM View of Ports



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

Port Types

The following types of ports are available:

- Veths (Virtual Ethernet Interfaces) can be associated with any one of the following:
 - VNICs of a Virtual Machine on the ESX Host.
 - VMKNICs of the ESX Host
 - VSWIFs of an ESX COS Host.
- Eths (Physical Ethernet Interfaces) – correspond to the Physical NICs on the ESX Host.
- Po (Port Channel Interfaces) – The physical NICs of an ESX Host can be bundled into a logical interface. This logical bundle is referred to as a port channel interface.

For more information about Layer 2 switching, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(3)*.

Layer 2 Switching Problems

This section describes how to troubleshoot Layer 2 problems and lists troubleshooting commands. This section includes the following topics:

- [Verifying a Connection Between VEM Ports, page 9-4](#)
- [Verifying a Connection Between VEMs, page 9-5](#)
- [Isolating Traffic Interruptions, page 9-6](#)
- [Verifying Layer 2 Switching, page 9-7](#)

Verifying a Connection Between VEM Ports

To verify a connection between two veth ports on a VEM, follow these steps:

-
- Step 1** On the VSM, enter the **show vlan** command to view the state of the VLANs associated with the port. If the VLAN associated with a port is not active, then the port may be down. In this case, you must create the VLAN and activate it.
- Step 2** To see the state of the port on the VSM, enter a **show interface brief** command.
- Step 3** Enter the **module vem module-number execute vemcmd show port** command to display the ports that are present on the VEM, their local interface indices, VLAN, type (physical or virtual), CBL state, port mode, and port name.

The key things to look for in the output are:

- State of the port.
- CBL.
- Mode.
- Attached device name.
- The LTL of the port you are trying to troubleshoot. It will help you identify the interface quickly in other VEM commands where the interface name is not displayed.

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

- Make sure the state of the port is up. If not, verify the configuration of the port on the VSM.
- Step 4** To view the VLANs and their port lists on a particular VEM, use the **module vem module-number execute vemcmd show bd** command:

```
n1000v# module vem 5 execute vemcmd show bd
```

If you are trying to verify that a port belongs to a particular VLAN, make sure you see the port name or LTL in the port list of that VLAN.

Verifying a Connection Between VEMs

To verify a connection between veth ports on two separate VEMs, follow these steps:

- Step 1** Issue the **show vlan** command to check if the VLAN associated with the port is created on the VSM.
- Step 2** Issue the **show interface brief** command to check if the ports are up in the VSM.
- Step 3** On the VEM, issue the **module vem 3 execute vemcmd show port** command to check if the CBL state of the two ports is set to the value of 4 for forwarding.
- Step 4** On the VEM, issue the **module vem 3 execute vemcmd show bd** command to check if the two veth ports are listed in the flood list of the VLAN to which they are trying to communicate.
- Step 5** Verify that the uplink switch to which the VEMs are connected is carrying the VLAN to which the ports belong.
- Step 6** Find out the port on the upstream switch to which the pnic (that is supposed to be carrying the VLAN) on the VEM is connected to.

```
n1000v# 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 Interface	Hldtme	Capability	Platform	Port ID
swordfish-6k-2	Eth5/2	168	R S I	WS-C6506-E	Gig1/38

The PNIC (Eth 5/2) is connected to swordfish-6k-2 on port Gig1/38.

- Step 7** Log in to the upstream switch and make sure the port is configured to allow the VLAN you are looking for.

```
n1000v#show running-config interface gigabitEthernet 1/38
Building configuration...
```

```
Current configuration : 161 bytes
!
interface GigabitEthernet1/38
 description Srvr-100:vmnic1
 switchport
 switchport trunk allowed vlan 1,60-69,231-233
 switchport mode trunk
end
```

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

As this output shows, VLANs 1,60-69, 231-233 are allowed on the port. If a particular VLAN is not in the allowed VLAN list, make sure to add it to the allowed VLAN list of the port.

Isolating Traffic Interruptions

Use the following steps to isolate the cause for no traffic passing across VMs on different VEMs.

- Step 1** In output of the **show port-profile name** command, verify the following information:
- The control and packet VLANs that you configured are present (in the example, these are 3002 and 3003)
 - If the physical NIC in your configuration carries the VLAN for VM, then that VLAN is also present in the allowed VLAN list.

```
n1000v#show port-profile name alluplink
port-profile alluplink
  description:
  status: enabled
  system vlans: 3002,3003
  port-group: alluplink
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,610,620,630-650
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 1,80,3002,3003,610,620,630-650
    no shutdown
  assigned interfaces:
    Ethernet2/2
```

- Step 2** Inside the VM, use the following command to verify that the Ethernet interface is up.

ifconfig -a

If not, consider deleting that NIC from the VM, and adding another NIC.

- Step 3** Using any sniffer tool, verify that ARP requests and responses are received on the VM interface.

- Step 4** On the upstream switch, use the following commands to look for the association between the IP and MAC address:

debug arp

show arp

Example:

```
n1000v_CAT6K# debug arp
ARP packet debugging is on
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
11w4d: RARP: Rcvd RARP req for 0050.56b7.4d35
11w4d: RARP: Rcvd RARP req for 0050.56b7.52f4
11w4d: IP ARP: rcvd req src 10.78.1.123 0050.564f.3586, dst 10.78.1.24 Vlan3002
11w4d: RARP: Rcvd RARP req for 0050.56b7.3031
n1000v_CAT6K#
```

Example:

```
n1000v_CAT6K# sh arp
```


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

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	10.78.1.72	-	001a.6464.2008	ARPA	
Internet	7.114.1.100	-	0011.bcac.6c00	ARPA	Vlan140
Internet	41.0.0.1	-	0011.bcac.6c00	ARPA	Vlan410
Internet	7.61.5.1	-	0011.bcac.6c00	ARPA	Vlan1161
Internet	10.78.1.5	-	0011.bcac.6c00	ARPA	Vlan3002
Internet	7.70.1.1	-	0011.bcac.6c00	ARPA	Vlan700
Internet	7.70.3.1	-	0011.bcac.6c00	ARPA	Vlan703
Internet	7.70.4.1	-	0011.bcac.6c00	ARPA	Vlan704
Internet	10.78.1.1	0	0011.bc7c.9c0a	ARPA	Vlan3002
Internet	10.78.1.15	0	0050.56b7.52f4	ARPA	Vlan3002
Internet	10.78.1.123	0	0050.564f.3586	ARPA	Vlan3002

Step 5 You have completed this procedure.

Verifying Layer 2 Switching

Use the following commands to display and verify the Layer 2 MAC address configuration.

Command	Purpose
show mac address-table	Displays the MAC address table to verify all MAC addresses on all VEMs controlled by the VSM. See Example 9-1 on page 9-8
show mac address-table module <i>module-number</i>	Displays all the MAC addresses on the specified VEM.
show mac address-table static <i>HHHH.WWWW.HHHH</i>	Displays the MAC address table static entries. See Example 9-2 on page 9-9
show mac address-table address <i>HHHH.WWWW.HHHH</i>	Displays the interface on which the MAC address specified is learned or configured. <ul style="list-style-type: none"> For dynamic MACs, if the same MAC appears on multiple interfaces, then each of them is displayed separately. For static MACs, if the same MAC appears on multiple interfaces, then only the entry on the configured interface is displayed.
show running-config vlan <vlan-id>	Displays VLAN information in the running configuration.
show vlan [all-ports brief id <vlan-id> name <name> dot1q tag native]	Displays VLAN information as specified. See Example 9-3 on page 9-9 .
show vlan summary	Displays a summary of VLAN information.
show interface brief	Displays a table of interface states. See Example 9-4 on page 9-10 .

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

Command	Purpose
module vem <i>module-number</i> execute vemcmd show port	On the VEM, displays the port state on a particular VEM. This command can only be used from the VEM. See Example 9-5 on page 9-10 .
module vem <i>module-number</i> execute vemcmd show bd command	For the specified VEM, displays its VLANs and their port lists. See Example 9-6 on page 9-11 .
module vem <i>module-number</i> execute vemcmd show trunk	For the specified VEM, displays the VLAN state on a trunk port. <ul style="list-style-type: none"> • If a VLAN is active on a port, then its CBL state should be 4. • If a VLAN is blocked, then its CBL state is 1. See Example 9-7 on page 9-11 .
module vem <i>module-number</i> execute vemcmd show l2 <i>vlan-id</i>	For the specified VEM, displays the VLAN forwarding table for a specified VLAN. See Example 9-8 on page 9-11 .
show interface <i>interface_id</i> mac	Displays the MAC addresses and the burn-in MAC address for an interface.

Example 9-1 show mac address-table Command



Note The Cisco Nexus 1000V MAC address table does not display multicast MAC addresses.



Tip VEM indicates on which VEM this MAC is seen.

N1KV Internal Port refers to an internal port created on the VEM. This port is used for control and management of the VEM and is not used for forwarding packets.

```
n1000v# show mac address-table
Legend:
      * - primary entry, G - Gateway MAC, (R) - Routed MAC
      age - seconds since last seen
      VEM      VLAN      MAC Address      Type      age      Ports
-----+-----+-----+-----+-----+-----+-----
*   3         1         0002.3d22.e300   static    -        N1KV Internal Port
*   3         1         0002.3d22.e302   static    -        N1KV Internal Port
*   4         1         0002.3d22.e303   static    -        N1KV Internal Port
*   3         1         0002.3d32.e300   static    -        N1KV Internal Port
*   3         1         0002.3d32.e302   static    -        N1KV Internal Port
*   4         1         0002.3d32.e303   static    -        N1KV Internal Port
*   3         1         0002.3d62.e300   static    -        N1KV Internal Port
*   3         1         0002.3d62.e302   static    -        N1KV Internal Port
*   4         1         0002.3d62.e303   static    -        N1KV Internal Port
      4         1         0023.7d34.f4e2   dynamic   23        Eth4/2
      3         115        0002.3d42.e302   dynamic    0        N1KV Internal Port
      4         115        0002.3d42.e303   dynamic    0        N1KV Internal Port
```

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

```

4      115      0050.56bb.49d9   dynamic  0      Eth4/2
3      115      0050.56bb.49d9   dynamic  0      Eth3/4
3      116      0002.3d22.e302   dynamic  1      N1KV Internal Port
4      116      0002.3d22.e302   dynamic  1      Eth4/2
4      116      0002.3d22.e303   dynamic  1      N1KV Internal Port
3      116      0002.3d22.e303   dynamic  1      Eth3/4

```

Example 9-2 show mac address-table address Command



Tip This command shows all interfaces on which a MAC is learned dynamically. In this example, the same MAC appears on Eth4/2 and Eth3/4.

```

n1000v# show mac address-table address 0050.56bb.49d9
Legend:
      * - primary entry, G - Gateway MAC, (R) - Routed MAC
      age - seconds since last seen

```

VEM	VLAN	MAC Address	Type	age	Ports
4	115	0050.56bb.49d9	dynamic	0	Eth4/2
3	115	0050.56bb.49d9	dynamic	0	Eth3/4

Example 9-3 show vlan Command



Tip This command shows the state of each VLAN created on the VSM.

```

n1000v# show vlan

```

VLAN	Name	Status	Ports
1	default	active	Eth3/3, Eth3/4, Eth4/2, Eth4/3
110	VLAN0110	active	
111	VLAN0111	active	
112	VLAN0112	active	
113	VLAN0113	active	
114	VLAN0114	active	
115	VLAN0115	active	
116	VLAN0116	active	
117	VLAN0117	active	
118	VLAN0118	active	
119	VLAN0119	active	
800	VLAN0800	active	
801	VLAN0801	active	
802	VLAN0802	active	
803	VLAN0803	active	
804	VLAN0804	active	
805	VLAN0805	active	
806	VLAN0806	active	
807	VLAN0807	active	
808	VLAN0808	active	
809	VLAN0809	active	
810	VLAN0810	active	
811	VLAN0811	active	
812	VLAN0812	active	
813	VLAN0813	active	
814	VLAN0814	active	

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

```

815 VLAN0815          active
816 VLAN0816          active
817 VLAN0817          active
818 VLAN0818          active
819 VLAN0819          active
820 VLAN0820          active
VLAN Name            Status      Ports
-----
-----

Remote SPAN VLANs
-----

Primary  Secondary  Type          Ports
-----
-----

```

Example 9-4 show interface brief Command

```

n1000v# show int brief

-----
Port      VRF      Status IP Address      Speed  MTU
-----
mgmt0    --      up      172.23.232.143  1000  1500
-----

Ethernet  VLAN  Type Mode  Status Reason      Speed  Port
Interface
-----
Eth3/4    1     eth trunk up    none      1000(D) --
Eth4/2    1     eth trunk up    none      1000(D) --
Eth4/3    1     eth trunk up    none      1000(D) --

```

Example 9-5 module vem module-number execute vemcmd show port Command



Tip Look for the state of the port.

```

~ # module vem 3 execute vemcmd show port
LTL   IfIndex  Vlan   Bndl  SG_ID  Pinned_SGID  Type  Admin  State  CBL  Mode  Name
  8     0     3969   0     2      2           VIRT  UP     UP    4   Access 120
  9     0     3969   0     2      2           VIRT  UP     UP    4   Access 121
 10    0     115    0     2      0           VIRT  UP     UP    4   Access 122
 11    0     3968   0     2      2           VIRT  UP     UP    4   Access 123
 12    0     116    0     2      0           VIRT  UP     UP    4   Access 124
 13    0     1      0     2      2           VIRT  UP     UP    0   Access 125
 14    0     3967   0     2      2           VIRT  UP     UP    4   Access 126
 16   1a030100   1 T     0     0      2           PHYS  UP     UP    4   Trunk
vmnic1
 17   1a030200   1 T     0     2      2           PHYS  UP     UP    4   Trunk
vmnic2

```

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

Example 9-6 *module vem module-number execute vemcmd show bd Command*



Tip If a port belongs to a particular VLAN, the port name or LTL should be in the port list for the VLAN.

```

~ # module vem 5 execute vemcmd show bd
Number of valid BDS: 8
BD 1, vdc 1, vlan 1, 2 ports
Portlist:
16 vmnic1
17 vmnic2
BD 100, vdc 1, vlan 100, 0 ports
Portlist:
BD 110, vdc 1, vlan 110, 1 ports
Portlist:
16 vmnic1
BD 111, vdc 1, vlan 111, 1 ports
Portlist:
16 vmnic1
BD 112, vdc 1, vlan 112, 1 ports
Portlist:
16 vmnic1
BD 113, vdc 1, vlan 113, 1 ports
Portlist:
16 vmnic1
BD 114, vdc 1, vlan 114, 1 ports
Portlist:
16 vmnic1
BD 115, vdc 1, vlan 115, 2 ports
Portlist:
10 l22
16 vmnic1

```

Example 9-7 *module vem module-number execute vemcmd show trunk Command*



Tip If a VLAN is active on a port, then its CBL state should be 4.
If a VLAN is blocked, then its CBL state is 1.

```

~ # module vem 5 execute vemcmd show trunk
Trunk port 16 native_vlan 1 CBL 4
vlan(1) cbl 4, vlan(110) cbl 4, vlan(111) cbl 4, vlan(112) cbl 4, vlan(113) cbl 4,
vlan(114) cbl 4, vlan(115) cbl 4, vlan(116) cbl 4, vlan(117) cbl 4, vlan(118) cbl 4,
vlan(119) cbl 4,
Trunk port 17 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(117) cbl 4,
~ #

```

Example 9-8 *module vem module-number execute vemcmd show l2 Command*

```

Bridge domain 115 brtmax 1024, brtcnt 2, timeout 300
Dynamic MAC 00:50:56:bb:49:d9 LTL 16 timeout 0
Dynamic MAC 00:02:3d:42:e3:03 LTL 10 timeout 0

```

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



CHAPTER 10

ACLs

This chapter describes how to identify and resolve problems that relate to Access Control Lists (ACLs).

This chapter includes the following sections:

- [About Access Control Lists \(ACLs\), page 10-1](#)
- [ACL Configuration Limits, page 10-1](#)
- [ACL Restrictions, page 10-2](#)
- [Troubleshooting ACLs, page 10-2](#)
- [Displaying ACL Policies on the VEM, page 10-2](#)
- [Debugging Policy Verification Issues, page 10-3](#)

About Access Control Lists (ACLs)

An ACL is an ordered set of rules for filtering traffic. When the device determines that an ACL applies to a packet, it tests the packet against the rules. The first matching rule determines whether the packet is permitted or denied. If there is no match, the device applies a default rule. The device processes packets that are permitted and drops packets that are denied.

ACLs protect networks and specific hosts from unnecessary or unwanted traffic. For example, ACLs are used to disallow HTTP traffic from a high-security network to the Internet. ACLs also allow HTTP traffic but only to specific sites, using the IP address of the site to identify it in an IP ACL.

The following types of ACLs are supported for filtering traffic:

- IP ACLs—The device applies IP ACLs only to IP traffic.
- MAC ACLs—The device applies MAC ACLs only to non-IP traffic.

For detailed information about how ACL rules are used to configure network traffic, see the *Cisco Nexus 1000V Security Configuration Guide, Release 4.0(4)SV1(3)*.

ACL Configuration Limits

The following configuration limits apply to ACLs:

- You cannot have more than 128 rules in an ACL.
- You cannot have more than 10,000 ACLs (spread across all the ACLs) in one VEM.

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

ACL Restrictions

The following restrictions apply to ACLs:

- You cannot apply more than one IP ACL and one MAC ACL in each direction on an interface.
- A MAC ACL applies only to Layer 2 packets.
- VLAN ACLs are not supported.
- IP fragments are not supported in ACL rules.
- Non initial fragments are not subject to ACL lookup.
- The established option to specify TCP flags is not supported.
- You cannot have two not-equal-to (neq) operators in the same rule.
- ACL is not supported in port channels.

Troubleshooting ACLs

The commands listed in this section can be used on the VSM to see the policies that are configured and applied on the interfaces.

Use the following command to display configured ACLs:

- **show access-list summary**

Use following commands on the VSM to see run-time information of the ACLMGR and ACLCOMP during configuration errors, and to collect ACLMGR process run-time information configuration errors:

- **show system internal aclmgr event-history errors**
- **show system internal aclmgr event-history msgs**
- **show system internal aclmgr ppf**
- **show system internal aclmgr mem-stats (to debug memory usage and leaks)**
- **show system internal aclmgr status**
- **show system internal aclmgr dictionary**

Use the following commands to collect ACLCOMP process run-time information configuration errors:

- **show system internal aclcomp event-history errors**
- **show system internal aclcomp event-history msgs**
- **show system internal aclcomp pdl detailed**
- **show system internal aclcomp mem-stats (to debug memory usage and leaks)**

Displaying ACL Policies on the VEM

The commands listed in this section can be used to display configured ACL policies on the VEM.

Use the following command to list the ACLs installed on that server

```
~ # module vem 3 execute vemcmd show acl
Acl-id Ref-cnt Type Numrules Stats
      1      1  IPv4      1  disabled
```


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

The Acl-id is the local ACLID for this VEM. Ref-cnt refers to the number of instances of this ACL in this VEM.

Use the following command to list the interfaces on which ACLs have been installed

```
~ # module vem 3 execute vemcmd show acl pinst
LTL  Acl-id  Dir
 16      1  ingress
```

Debugging Policy Verification Issues

To debug a policy verification failure, follow these steps:

Step 1 On the VSM, enter the **debug logfile *filename*** command to redirect the output to a file in bootflash.

Step 2 Enter the **debug aclmgr all** command.

Step 3 Enter the **debug aclcomp all** command.

For the VEMs where the policy exists, or is being applied, enter the following these steps from the VSM. The output goes to the console.

Step 4 Enter the **module vem *module-number* execute vemdpalog debug sfaclagent all** command.

Step 5 Enter the **module vem *module-number* execute vemdpalog debug sfpdlagent all** command.

Step 6 Enter the **module vem *module-number* execute vemlog debug sfacl all** command.

Step 7 Enter the **module vem *module-number* execute vemlog start** command.

Step 8 Enter the **module vem *module-number* execute vemlog start** command.

Step 9 Configure the policy that was causing the verify error.

Step 10 Enter the **module vem *module-number* execute vemdpalog show all** command.

Step 11 Enter **module vem *module-number* execute vemlog show all** command.

Save the Telnet or SSH session buffer to a file. Copy the logfile created in bootflash.

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



CHAPTER 11

Quality of Service

This chapter describes how to identify and resolve problems related to Quality of Service (QoS).

This chapter includes the following sections:

- [Information About Quality of Service, page 11-1](#)
- [QoS Configuration Limits, page 11-1](#)
- [QoS Troubleshooting Commands, page 11-2](#)
- [Troubleshooting the VEM, page 11-2](#)
- [Debugging Policing Verification Errors, page 11-3](#)

Information About Quality of Service

QoS lets you classify network traffic so that it can be policed and prioritized in a way that prevents congestion. Traffic is processed based on how you classify it and the QoS policies that you put in place. Classification, marking, and policing are the three main features of QoS.

- **Traffic Classification**—Groups network traffic based on defined criteria.
- **Traffic Marking**—Modifies traffic attributes such as DSCP, COS, and Precedence by class.
- **Policing**—Monitors data rates and burst sizes for a particular class of traffic. QoS policing on a network determines whether network traffic is within a specified profile (contract).

For detailed information about QoS, refer to the *Cisco Nexus 1000V Quality of Service Configuration Guide, Release 4.0(4)SV1(3)*.

QoS Configuration Limits

[Table 11-1](#) and [Table 11-2](#) list the configuration limits for QoS.

Table 11-1 QoS Configuration Limits

Item	DVS Limit	Per Server Limit
Class map	1000	64 (with policies)
Policy map	128	16
Service policy	—	128

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

Table 11-2 QoS Configuration Limits

Item	Limit
Match criteria per class map	32
Class maps per policy map	64

QoS Troubleshooting Commands

The commands listed in this section can be used on the VSM to see the policies that are configured and applied on the interfaces.

Use the following commands to display configured policies and class-maps:

- **Show policy-map [policy-map-name]**
- **Show class-map [class-map-name]**

Use the following command to display installed policies:

- **Show policy-map interface brief**

Use following commands on the VSM to see run-time information of the QOSMGR and ACLCOMP during configuration errors.

The commands to collect QOSMGR process run-time information configuration errors are as follows:

- **show system internal ipqos event-history errors**
- **show system internal ipqos event-history msgs**
- **show system internal ipqos port-node**
- **show system internal ipqos mem-stats** (to debug memory usage and leaks)
- **show system internal ipqos status**
- **show system internal ipqos log** (to show aborted plan information)
- **show system internal ipqos**

The commands to collect ACLCOMP process run-time information configuration errors are as follows:

- **show system internal aclcomp event-history errors**
- **show system internal aclcomp event-history msgs**
- **show system internal aclcomp pdl detailed**
- **show system internal aclcomp mem-stats** (to debug memory usage and leaks)

Troubleshooting the VEM

The commands listed in this section can be used to display configured QoS policies on the VEM.

Use the following command to list all class maps and polices in use on the server:

- **module vem module-number execute vemcmd show qos node**

```
~ # module vem 3 execute vemcmd show qos node
nodeid  type      details
-----  -
```

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

```

0  policer
    cir:50 pir:50
    bc:200000 be:200000
    cir/pir units 1 bc/be units 3 flags 2
1  class  op_AND
    DSCP
2  class op_DEFAULT

```

Use the following command to list all the installed policy maps in use on the server:

- **module vem *module-number* execute vemcmd show qos policy**

```

~ # module vem 3 execute vemcmd show qos policy
policyid classid policerid set_type value
-----
0          1          -1          dscp          5
          2          0          dscp          0

```

Use the following command to list all service policies installed on the server:

- **module vem *module-number* execute vemcmd show qos pinst**

```

~ # module vem 3 execute vemcmd show qos pinst

id      type
-----
17 Ingress
    class          bytes matched          pkts matched
-----
    1              0              0
    2              85529             572
    0
    policer stats: conforming (85529, 572)
    policer stats: exceeding (0, 0)
    policer stats: violating (0, 0)

```

Debugging Policing Verification Errors

To debug a policy verification failure caused by processing on the VSM, follow these steps:

-
- Step 1** Enter the **debug aclmgr all** command if the policy references an ACL.
 - Step 2** Enter **debug ipqos all** command.
 - Step 3** Enter the **debug aclcomp all** command.
 - Step 4** Enter the **service-policy** command which will execute the command once again with debug traces output to the console. This command allows you to collect logs for all operations.
 - Step 5** Save the Telnet SSH session buffer to a file.
-

If you are debugging a policy on a port profile, it may be easier to first install it directly on an interface.

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

To debug a policy verification failure on the VEM, follow these steps:

-
- Step 1** Enter the **module vem *module-number* execute vemdpallog clear** command.
 - Step 2** Enter the **module vem *module-number* execute vemdpallog sfqosagent all** command.
 - Step 3** Enter **module vem *module-number* execute vemdpallog start** command.
 - Step 4** Enter the **service-policy** command which will execute the command once again with the DPA debug traces output to vemdpallog.
 - Step 5** Enter **module vem *module-number* execute vemdpallog stop** command.
 - Step 6** Enter the **module vem *module-number* execute vemdpallog show all** command to see the logs on console.

The output will look similar to the following:

```
calling add policy 81610ac len 220 classmaps 3- --> Session actions
...
Adding classmap 1 (108) with op 1 and 2 filters
...
Adding classmap 2 (116) with op 2 and 2 filters
...
Adding classmap 3 (56) with op 0 and 0 filters
...
init pinst lt1 11 policy id 0 if_index 1a020200 --> Service-policy being applied
installing pinst type 0 17 for policy 0
dpa_sf_qos_verify returned 0
...
Session commit complete and successful --> Session ending
```



CHAPTER 12

NetFlow

This chapter describes how to identify and resolve problems that relate to Netflow.

This chapter includes the following sections:

- [Information About NetFlow, page 12-1](#)
- [NetFlow Troubleshooting Commands, page 12-2](#)
- [Common NetFlow Problems, page 12-3](#)

Information About NetFlow

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

A flow is a one-directional stream of packets that arrives on a source interface (or sub-interface), 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. 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.

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 flow record.

For detailed information about configuring NetFlow, see the *Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(3)*.

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

NetFlow Troubleshooting Commands

Use the commands listed in this section to troubleshoot NetFlow problems.

- **debug logfile *filename***—Use this command to redirect the output of the following debug commands to a file stored in bootflash.
 - **debug nfm all**
 - **debug sf_nf_srv all**
- **vemdebug netflow dump policy**— Use this command to verify if the correct policy is installed on an interface on a VEM. The output of this command goes to the vemlog internal buffer. Make sure the output shows the cache type as normal, and shows the correct cache size and cache timer values.

```
Apr 14 12:20:51.504410      19  2    2  16  Debug Port 49 has 1
monitors for dir INPUT traffic IPV4
Apr 14 12:20:51.504412      20  2    2  16  Debug Flow Monitor fml:
Apr 14 12:20:51.504413      21  2    2  16  Debug Description:
fml
Apr 14 12:20:51.504413      22  2    2  16  Debug Monitor ID:
3
Apr 14 12:20:51.504413      23  2    2  16  Debug Cache:
Apr 14 12:20:51.504414      24  2    2  16  Debug Type:
normal
Apr 14 12:20:51.504414      25  2    2  16  Debug Status:
allocated
Apr 14 12:20:51.504415      26  2    2  16  Debug Size:
256 entries
Apr 14 12:20:51.504415      27  2    2  16  Debug Inactive
Timeout: 15 secs
Apr 14 12:20:51.504416      28  2    2  16  Debug Active
Timeout: 1800 secs
```

- **vemdebug netflow dump pakstore**—Use this command to dump pakstore usage for a policy on an interface. The output goes to a vemlog internal buffer. Make sure the output shows the correct monitor name and interface.

```
Apr 14 12:25:30. 29787      260  0    2  16  Debug Pak Store for
Client: fml
Apr 14 12:25:30. 29793      266  0    2  16  Debug Pak Store for
Client: LTL49
```

- **vemlog debug sfnetflow_cache all**
- **vemlog debug sfnetflow_config all**
- **vemlog debug sfnetflow_flowapi all**

Use these command to enable NetFlow debugging for policy installation on the VEM. Debug messages are printed for every PDL session open, verify, and commit requests coming from the DPA.

- **vemlog debug sfnetflow all**

Use this command to enable packet path debugging for Netflow policies on the VEM. Debug messages are printed for every packet that hits a NetFlow policy. Use this command with caution. High traffic could result in lot of debug messages.

Use the following commands to collect information about NFM process run-time configuration errors:

- **show flow internal event-history errors**
- **show flow internal event-history msgs**

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

- **show flow internal ddb b**
- **show flow internal mem-stats** (to debug memory usage and leaks)

Use the following commands to collect sf_nf_srv process run-time information:

- **show system internal sf_nf_srv event-history errors**
- **show system internal sf_nf_srv event-history msgs**
- **show system internal sf_nf_srv pdl detailed**
- **show system internal sf_nf_srv mem-stats** (to debug memory usage and leaks)

Common NetFlow Problems

Common NetFlow configuration problems on the VSM can occur if you attempt to do the following:

- Use undefined records, exporters, samplers, or monitors
- Use invalid records, exporters, samplers, or monitors
- Modify records, exporters, samplers, or monitors after they are applied to an interface
- Configure a monitor on an interface which causes the VEM to run out of memory and results in a verification error
- Use NetFlow in a port channel. NetFlow is not supported in port channels.

In addition, a configuration error can occur if there is a mismatch between the UDP port configured on the exporter and the port NetFlow Collector has listening turned on.

Debugging a Policy Verification Error

To debug a policy verification failure due to some processing on the VSM, follow these steps:

-
- Step 1** Issue the **debug nfm all** command.
 - Step 2** Issue the **debug sf_nf_srv_all** command.
 - Step 3** Save the Telnet SSH session buffer to a file.
 - Step 4** Issue the **ip flow mon monitor name direction** command.

The command will execute once again and the debug traces will be output to the console.

You can also use the policy verification procedure to collect logs for operations such as defining a flow record or tracing exporter functionality.

Debugging Statistics Export

When debugging a NetFlow statistics export problem, follow these guidelines:

- Ensure the destination IP address is reachable from the VSM.
- Ensure the UDP port configured on the exporter matches that used by the NetFlow Collector.

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

- Issue the **show flow exporter** command to view statistics for the exporter and identify any drops on the VSM.



CHAPTER 13

VLANs

This chapter describes how to identify and resolve problems that might occur when implementing VLANs.

This chapter includes the following sections:

- [Information About VLANs, page 13-1](#)
- [Initial Troubleshooting Checklist, page 13-2](#)
- [Cannot Create a VLAN, page 13-3](#)

Information About VLANs

VLANs can isolate devices that are physically connected to the same network, but are logically considered to be part of different LANs that do not need to be aware of one another.

We recommend using only following characters in a VLAN name:

- a-z or A-Z
- 0 - 9
- - (hyphen)
- _ (underscore)

Consider the following guidelines for VLANs:

- Keep user traffic off the management VLAN; keep the management VLAN separate from user data.



Note

We recommend that you enable sticky Address Resolution Protocol (ARP) when you configure private VLANs. ARP entries learned on Layer 3 private VLAN interfaces that are sticky ARP entries. For security reasons, private VLAN port sticky ARP entries do not age out.

- IGMP only runs on the primary VLAN and uses the configuration of the primary VLAN for all secondary VLANs.
- Any IGMP join request in the secondary VLAN is treated as if it is received in the primary VLAN.
- Private VLANs support these Switched Port Analyzer (SPAN) features:
 - You can configure a private VLAN port as a SPAN source port.
 - You can use VLAN-based SPAN (VSPAN) on primary, isolated, and community VLANs or use SPAN on only one VLAN to separately monitor egress or ingress traffic.

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

- A private VLAN host or promiscuous port cannot be a SPAN destination port. If you configure a SPAN destination port as a private VLAN port, the port becomes inactive.
- A destination SPAN port cannot be an isolated port. (However, a source SPAN port can be an isolated port.) V
- SPAN could be configured to span both primary and secondary VLANs or, alternatively, to span either one if the user is interested only in ingress or egress traffic.
- A MAC address learned in a secondary VLAN is placed in the shared table of the primary VLAN. When the secondary VLAN is associated to the primary VLAN, their MAC address tables are merged into one, shared MAC table.

Initial Troubleshooting Checklist

Troubleshooting a VLAN problem involves gathering information about the configuration and connectivity of individual devices and the entire network. In the case of VLANs, begin your troubleshooting activity as follows:

Checklist	✓
Verify the physical connectivity for any problem ports or VLANs.	
Verify that both end devices are in the same VLAN.	

The following CLI commands are used to display VLAN information:

- **show system internal private-vlan info**
- **show system internal private-vlan event-history errors**
- **show system internal private-vlan event-history traces**
- **show vlan id *vlan-id***
- **show vlan private-vlan**
- **show vlan all-ports**
- **show vlan private-vlan type**
- **show vlan internal bd-info vlan-to-bd 1**
- **show vlan internal errors**
- **show vlan internal info**
- **show vlan internal event-history errors**

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

Cannot Create a VLAN

Symptom	Possible Cause	Solution
Cannot create a VLAN.	Using a reserved VLAN ID	VLANs 3968 to 4047 and 4094 are reserved for internal use and cannot be changed.

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



CHAPTER 14

Private VLANs

This chapter describes how to identify and resolve problems related to private VLANs.

This chapter includes the following sections:

- [Information About Private VLANs, page 14-1](#)
- [Troubleshooting Guidelines, page 14-2](#)
- [Private VLAN Troubleshooting Commands, page 14-2](#)

Information About Private VLANs

Private VLANs (PVLANS) are used to segregate Layer 2 ISP traffic and convey it to a single router interface. PVLANS achieve device isolation by applying Layer 2 forwarding constraints that allow end devices to share the same IP subnet while being Layer 2 isolated. In turn, the use of larger subnets reduces address management overhead. Three separate port designations are used, each having its own unique set of rules regulating each connected endpoint's ability to communicate with other connected endpoints within the same private VLAN domain.

Private VLAN Domain

A private VLAN domain consists of one or more pairs of VLANs. The primary VLAN makes up the domain; and each VLAN pair makes up a subdomain. The VLANs in a pair are called the primary VLAN and the secondary VLAN. All VLAN pairs within a private VLAN have the same primary VLAN. The secondary VLAN ID is what differentiates one subdomain from another.

Spanning Multiple Switches

Private VLANs can span multiple switches, just like regular VLANs. Inter-switch link ports need not be aware of the special VLAN type and carry frames tagged with these VLANs just like they do any other frames. Private VLANs ensure that traffic from an isolated port in one switch does not reach another isolated or community port in a different switch even after traversing an inter-switch link. By embedding the isolation information at the VLAN level and by transporting it along with the packet, it is possible to maintain consistent behavior throughout the network. Therefore, the mechanism which restricts Layer 2 communication between two isolated ports in the same switch, also restricts Layer 2 communication between two isolated ports in two different switches.

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

Private VLAN Ports

Within a private VLAN domain, there are three separate port designations. Each port designation has its own unique set of rules which regulate the ability of one endpoint to communicate with other connected endpoints within the same private VLAN domain. The following are the three port designations:

- promiscuous
- isolated
- community

For additional information about private VLANs, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0(4)SV1(3)*.

Troubleshooting Guidelines

Follow these guidelines when troubleshooting private VLAN issues:

- Use the **show vlan private-vlan** command to verify that a private VLAN is configured correctly.
- Use the **show interface slot-port** command to verify the interface is up.
- Use the **module vem module-number execute vemcmd show port** command to verify the VEM is configured correctly.

Private VLAN Troubleshooting Commands

Use the commands listed in this section to troubleshoot problems related to private VLANs.

To verify that a private VLAN is configured correctly, use the following command:

- **show vlan private-vlan**

```
n1000V# show vlan private-vlan
Primary Secondary Type Ports
-----
152      157      community
152      158      isolated
156      153      community
156      154      community
156      155      isolated
```

To verify if a physical Ethernet interface in a private VLAN trunk promiscuous mode is up, use the following command:

- **show interface**

```
n1000V# show int eth3/4
Ethernet3/4 is up
Hardware: Ethernet, address: 0050.565a.ca50 (bia 0050.565a.ca50)
MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 0/255, txload 0/255, rxload 0/255
Encapsulation ARPA
Port mode is Private-vlan trunk promiscuous
full-duplex, 1000 Mb/s
Beacon is turned off
Auto-Negotiation is turned off
Input flow-control is off, output flow-control is off
Auto-mdix is turned on
```


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

```
Switchport monitor is off
Rx
158776 Input Packets 75724 Unicast Packets
76 Multicast Packets 82976 Broadcast Packets
13861581 Bytes
Tx
75763 Output Packets 75709 Unicast Packets
3 Multicast Packets 51 Broadcast Packets 0 Flood Packets
7424670 Bytes
5507 Input Packet Drops 0 Output Packet Drops
2 interface resets
```

To verify if a virtual Ethernet interface in private VLAN host mode is up, use the following command:

- **show interface**

```
n1000V# show int v3
Vethernet3 is up
Hardware is Virtual, address is 0050.56bb.6330
Owner is VM "fedora9", adapter is Network Adapter 1
Active on module 3
VMware DVS port 10
Port-Profile is pvlancomm153
Port mode is Private-vlan host
Rx
14802 Input Packets 14539 Unicast Packets
122 Multicast Packets 141 Broadcast Packets
1446568 Bytes
Tx
15755 Output Packets 14492 Unicast Packets
0 Multicast Packets 1263 Broadcast Packets 0 Flood Packets
1494886 Bytes
45 Input Packet Drops 0 Output Packet Drops
```

To verify if a VEM is configured correctly, use the following command:

- **module vem module-number execute vemcmd show port**

```
n1000V# module vem 3 execute vemcmd show port
  LTL    IfIndex  Vlan    Bndl  SG_ID  Pinned_SGID  Type  Admin  State  CBL  Mode  Name
    8      0    3969     0     2      2          VIRT  UP    UP    4  Access  120
    9      0    3969     0     2      2          VIRT  UP    UP    4  Access  121
   10      0     150     0     2      2          VIRT  UP    UP    4  Access  122
   11      0    3968     0     2      2          VIRT  UP    UP    4  Access  123
   12      0     151     0     2      2          VIRT  UP    UP    4  Access  124
   13      0      1      0     2      2          VIRT  UP    UP    0  Access  125
   14      0    3967     0     2      2          VIRT  UP    UP    4  Access  126
   16  1a020100     1 T     0     2      2          PHYS  UP    UP    4  Trunk
vmnic1
   18  1a020300     1 T     0     2      2          PHYS  UP    UP    4  Trunk
vmnic3
  pvlan promiscuous trunk port
    153 --> 156
    154 --> 156
    155 --> 156
    157 --> 152
    158 --> 152
   19  1a020400     1 T     0     2      2          PHYS  UP    UP    4  Trunk
vmnic4
  pvlan promiscuous trunk port
    153 --> 156
    154 --> 156
    155 --> 156
    157 --> 152
    158 --> 152
```

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

```
47 1b020000 154 0 2 0 VIRT UP UP 4 Access
fedora9.eth0
pvlan community 156 153
```

If additional information is required for Cisco Technical Support to troubleshoot a private VLAN issue, use the following commands:

- **show system internal private-vlan info**
- **show system internal private-vlan event-history traces**



CHAPTER 15

Multicast IGMP

This chapter describes how to identify and resolve problems that relate to multicast Internet Group Management Protocol (IGMP) snooping.

This chapter includes the following sections:

- [Information About Multicast, page 15-1](#)
- [Multicast IGMP Snooping, page 15-1](#)
- [Problems with Multicast IGMP Snooping, page 15-2](#)

Information About Multicast

IP multicast is a method of forwarding the same set of IP packets to a number of hosts within a network. You can use multicast in both IPv4 and IPv6 networks to provide efficient delivery of data to multiple destinations.

Multicast involves both a method of delivery and discovery of senders and receivers of multicast data, which is transmitted on IP multicast addresses called groups. A multicast address that includes a group and source IP address is often referred to as a channel.

Multicast IGMP Snooping

IGMP snooping software examines Layer 2 IP multicast traffic within a VLAN to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help the routers forward IGMP membership reports. The IGMP snooping software responds to topology change notifications.

In general, IGMP snooping works as follows:

- Ethernet switches, like Cisco Catalyst 6000 switches, parse and intercept all IGMP packets and forward them to a CPU, such as a Supervisor module, for protocol processing.
- Router ports are learned using IGMP queries. The switch returns IGMP queries, it remembers which port the query comes from, and marks the port as a router port.
- IGMP membership is learned using IGMP reports. The switch parses IGMP report packets, and updates its multicast forwarding table to keep track of IGMP membership.
- When the switch receives multicast traffic, it check its multicast table, and forwards the traffic only to those ports interested in the traffic.

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

- IGMP queries are flooded to the whole VLAN.
- IGMP reports are forwarded to the uplink port (the router ports).
- Multicast data traffic is forwarded to uplink ports (the router ports).

Problems with Multicast IGMP Snooping

The operation of multicast IGMP snooping depends on the correct configuration of the upstream switch. Because the IGMP process needs to know which upstream port connects to the router that supports IGMP routing, you must turn on IP multicast routing on the upstream switch by issuing the **ip multicast-routing** command.

The following example shows how to turn on global multicast-routing, configure an SVI interface, and turn on the PIM routing protocol:

```
switch#conf terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#ip multicast-routing
switch(config)#end

switch#conf terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#int vlan159
switch(config-if)#ip pim dense-mode
switch(config-if)#end
```

Troubleshooting Guidelines

Follow these guidelines when troubleshooting multicast IGMP issues:

- Use the **show ip igmp snooping** command to verify that IGMP snooping is enabled.
- Make sure the upstream switch has IGMP configured.
- Use the **show ip igmp snooping groups** command to verify that the switch is configured correctly and is ready to forward multicast traffic. In the command output, look for the letter R under the port heading. The R indicates that the VSM has learned the uplink router port from the IGMP query that was sent by the upstream switch, and is ready to forward multicast traffic.

Troubleshooting Commands

To troubleshoot issues with multicast IGMP snooping, use the following commands:

- **show cdp neighbor**

You can use the **show cdp neighbor** command because IGMP uses the packet VLAN to forward IGMP packets to the VSM, which is the same mechanism that CDP uses. However, if you have disabled the CDP protocol on the upstream switch using the **no cdp enable** command, then the **show cdp neighbor** command will not display any information.

Example 15-1 *show cdp neighbor Command*

```
n1000V# show cdp neighbor
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
```

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

V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device ID	Local Intrfce	Hldtme	Capability	Platform	Port ID
n1000V	Eth3/2	179	R S I	WS-C6506-E	Gig5/16
n1000V	Eth3/4	179	R S I	WS-C6506-E	Gig5/23

- **show ip igmp groups**

Use the show ip igmp groups command to make sure IGMP snooping is enabled on the VLAN.

Example 15-2 show ip igmp snooping vlan Command

```
n1000V# show ip igmp snooping vlan 159
IGMP Snooping information for vlan 159
  IGMP snooping enabled      <-- IGMP SNOOPING is enabled for vlan 159
  IGMP querier none
  Switch-querier disabled
  IGMPv3 Explicit tracking enabled (initializing, time-left: 00:03:20)
  IGMPv2 Fast leave disabled
  IGMPv1/v2 Report suppression enabled
  IGMPv3 Report suppression disabled
  Router port detection using PIM Hellos, IGMP Queries
  Number of router-ports: 0
  Number of groups: 0show ip igmp snooping
```

- **show ip igmp snooping groups**
- **debug ip igmp snooping vlan**

Example 15-3 debug ip igmp snooping vlan Command

```
n1000V(config)# debug ip igmp snooping vlan
2008 Sep  2 13:29:36.125661 igmp: SNOOP: <vlan 159> Process a valid IGMP packet
2008 Sep  2 13:29:36.126005 igmp: SNOOP: <vlan 159> Received v2 report: group
224.0.0.251 fro 7.159.159.54 on Vethernet3
2008 Sep  2 13:29:36.126086 igmp: SNOOP: <vlan 159> Added oif Vethernet3 for (*,
224.0.0.251) entry
2008 Sep  2 13:29:36.126157 igmp: SNOOP: <vlan 159> Forwarding report for (*,
224.0.0.251) came on Vethernet3
2008 Sep  2 13:29:36.126225 igmp: SNOOP: <vlan 159> Forwarding the packet to
router-ports
2008 Sep  2 13:29:36.126323 igmp: SNOOP: <vlan 159> Forwarding packet to router-port
Ethernet3/6 (iod 42)
```

On the VSM, use the following command:

- **module vem module-number execute vemcmd show vlan**

In [Example 15-4](#), the output shows that LTL 18 corresponds to vmnic3, and LTL 47 corresponds to VM fedora8, interface eth0.

The multicast group table for 224.1.2.3, shows the interfaces the VEM will forward to when it receives multicast traffic for group 224.1.2.3. If fedora8 has multicast group 224.1.2.3 on its eth0 interface, then LTL 47 should be in the multicast group table for 224.1.2.3.

LTL 18 is also in multicast group 224.1.2.3, which means it is a VM and generates multicast traffic to 224.1.2.3. The traffic will be forwarded to vmnic3, which is the uplink to the upstream switch.

The multicast group table entry for 0.0.0.0 serves as a default route. If any multicast group traffic does not match any of the multilcast group, the address will use the default route, which means, in this case, that the traffic will be forwarded to an upstream switch through vmnic3.

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

Example 15-4 *module vem module-number execute vemcmd show vlan Command*

```
n1000V# module vem 3 execute vemcmd show vlan 159
BD 159, vdc 1, vlan 159, 3 ports
Portlist:
    18  vmmnic3
    47  fedora8.eth0

Multicast Group Table:
Group 224.1.2.3 RID 1 Multicast LTL 4408
    47
    18
Group 0.0.0.0 RID 2 Multicast LTL 4407
    18
```

Symptoms, Causes, and Solutions

Symptom	Possible Causes	Solution
A VM which is interested in multicast traffic, but is not receiving the multicast traffic.	—	Use the debug ip igmp snooping vlan command to determine if IGMP snooping is working as expected. Examine the output to see if the port is receiving the IGMP report and if the interface has been added to the multicast traffic interface list for the VM.
	—	Use module vem module-number execute vemcmd show vlan command to verify that the multicast distribution table in the VEM has the correct information in it.
	—	Use the module vem module-number execute vemcmd show port command to see the port table. Make sure the table has the correct information in it. Make sure that the state of the trunk port and the access port is UP/UP.



CHAPTER 16

SPAN

This chapter describes how to identify and resolve problems that relate to SPAN.

This chapter includes the following sections

- [Information About SPAN, page 16-1](#)
- [Troubleshooting SPAN Problems, page 16-3](#)

Information About SPAN

The Switched Port Analyzer (SPAN) feature (sometimes called port mirroring or port monitoring) selects network traffic for analysis by a network analyzer. The network analyzer can be a Cisco SwitchProbe or other Remote Monitoring (RMON) probe.

Cisco Nexus 1000V supports two types of SPAN:

- SPAN (local SPAN) that can monitor sources within a host or VEM
- Encapsulated remote SPAN (ERSPAN) that can send monitored traffic to an IP destination

For detailed information about how to configure SPAN, see the *Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(3)*.

SPAN Sources

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

- 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.

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

Source Ports

Cisco Nexus 1000V supports multiple source ports and multiple source VLANs. A source port has these characteristics:

- Can be port type Ethernet, virtual Ethernet, port-channel, or VLAN.
- Cannot be a destination port.
- 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.
- Must be on the same host (linecard) as the destination port.

SPAN Destinations

Ethernet and virtual Ethernet interfaces are supported as SPAN destinations.

Destination Ports

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 port type Ethernet, virtual Ethernet, or a port channel.
- Cannot be a source port.
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session.
- Receives copies of transmitted and received traffic for all monitored source ports. 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 be on the same host (linecard) as the source port.

ERSPAN Destinations

ERSPAN destinations refer to an IP address to which the monitored traffic is sent. The destination IP can belong to an IP of a sniffer device, ERSPAN capable switch (such as a Catalyst 6000 series switch), or a PC running a sniffer application. The only limitation is that the destination IP should be reachable through the configured ERSPAN enabled VMKnic on the host. For detailed information about how to configure ERSPAN, see the *Cisco Nexus 1000V System Management Configuration Guide, Release 4.0(4)SV1(3)*.

SPAN Sessions

You can create up to a total of 64 SPAN and ERSPAN sessions to define sources and destinations on the local device. You can also create a SPAN session to monitor multiple VLAN sources and choose only VLANs of interest to transmit on multiple destination ports. For example, you can configure SPAN on a trunk port and monitor traffic from different VLANs on different destination ports.

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

Troubleshooting SPAN Problems

When troubleshooting issues with SPAN, make sure you have followed these configuration guidelines and limitations:

- A maximum total of 64SPAN and ERSPAN sessions can be configured per VSM.
- You can configure a particular destination port in only one SPAN session.
- You cannot configure a port as both a source and destination port.
- When a SPAN session contains multiple transmit source ports, packets that these ports receive may be replicated even though they are not transmitted on the ports. Some examples of this behavior on source ports are as follows:
 - Traffic that results from flooding
 - Broadcast and multicast traffic
- For VLAN SPAN sessions with both receive and transmit configured, two packets (one from receive and one from transmit) are forwarded from the destination port if the packets get switched on the same VLAN.
- After VMotion:
 - A session is stopped if the source and destination ports are separated
 - A session resumes if the source and destination ports end up on the same host

Local SPAN Session Problems

A running SPAN session must meet these requirements:

- The limit of 64 SPAN sessions has not been exceeded.
- At least one operational source has been configured.
- At least one operational destination has been configured.
- The configured source and destination are on the same host.
- The session has been enabled with the **no shut** command.

A session is stopped if the follow events occur:

- All the source ports go down or are removed.
- All the destination ports go down or are removed.
- All the source and destination ports are separated by a VMotion.
- The session is disabled by a **shut** command.

Troubleshooting Commands

Uses the **show monitor session** command to troubleshoot a SPAN session. The output of this command shows the current state of the session and the reason it is down.

To collect additional information, use the following commands:

- **show monitor internal errors**
- **show monitor internal event-history msgs**
- **show monitor internal info global-info**

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

- `show monitor internal mem-stats`
- `module vem module-number execute vemcmd show span`

Problems and Solutions

Symptom	Possible Causes	Solution
You observe issues with VM traffic after configuring a session with Eth destinations.	—	Ensure that the Eth destination is not connected to the same uplink switch. The SPAN packets might cause problems with the IP tables, the MAC tables, or both on the uplink switch, which can cause problems with the regular traffic.
The session state is up and the packets are not received at the destination ports.	—	Check if the correct VLANs are allowed on the trunk destination ports.
The session displays an error.	—	Make sure that VEM connectivity is working correctly. Enter a shut command followed by a no shut command for the session to force reprogramming of the session on the VEM.
The ERSPAN session is up, but does not see packets at the destination.	The erspan-id is not configured.	Make sure that the correct erspan-id that matches with the destination session is configured.
	An ERSPAN enabled VMKNic is not configured on the host or VEM.	Make sure you use create a VMKNic on the host using an erspan-capable port profile.
	The ERSPAN enabled VMKNic is not configured with a proper IP, gateway, or both.	Make sure the ERSPAN IP destination is reachable from the host VMKNic. To test this, issue the vmkping dest-id command on the command line of the host.

Examples

The following example shows the output of the **show monitor session** command.

```
n1000v(config)# show monitor session 1
  session 1
  -----
type           : erspan-source
state          : up
source intf    :
  rx           : Eth3/3
  tx           : Eth3/3
  both         : Eth3/3
source VLANs   :
  rx           :
  tx           :
```

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

```
both :
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
```

The following example shows the output of the **module vem *module-number* execute vemcmd show span** command.

```
n1000v# module vem 3 execute vemcmd show span
VEM SOURCE IP: 10.54.54.10
HW SSN ID DST LTL/IP ERSPAN ID
0 10.54.54.1 999
1 48 local
```

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



CHAPTER 17

High Availability

This chapter describes how to identify and resolve problems related to High Availability.

This chapter includes the following sections:

- [Information About High Availability, page 17-1](#)
- [Problems with High Availability, page 17-3](#)
- [High Availability Troubleshooting Commands, page 17-6](#)

Information About High Availability

The purpose of High Availability (HA) is to limit the impact of failures—both hardware and software—within a system. The Cisco NX-OS operating system is designed for high availability at the network, system, and service levels.

For detailed information about High Availability, see the *Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)*.

The following HA features minimize or prevent traffic disruption in the event of a failure:

- Redundancy— redundancy at every aspect of the software architecture.
- Isolation of processes— isolation between software components to prevent a failure within one process disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- Supervisor stateful switchover— Active/standby dual supervisor configuration. State and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide seamless and stateful switchover in the event of a VSM failure.

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

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

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

System-Level High Availability

The Cisco Nexus 1000V supports redundant VSM virtual machines — a primary and a secondary — running as an HA pair. Dual VSMs operate in an active/standby capacity in which only one of the VSMs is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two VSMs to provide a stateful switchover if the active VSM fails.

Single or Dual Supervisors

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

- Virtual Ethernet Modules (VEMs) running within virtualization servers (these are represented as modules within the VSM)
- A remote management component, for example, VMware vCenter Server.
- One or two Virtual Supervisor Modules (VSMs) running within Virtual Machines (VMs)

Single VSM Operation	Dual VSM Operation
<ul style="list-style-type: none"> • Stateless—Service restarts from the startup configuration • Stateful—Service resumes from previous state. 	<ul style="list-style-type: none"> • One active VSM and one standby VSM. • The active VSM runs all the system applications and controls the system. • On the standby VSM, the applications are started and initialized in standby mode. They are also synchronized and kept up to date with the active VSM in order to maintain the runtime context of “ready to run.” • On a switchover, the standby VSM takes over for the active VSM.

Network-Level High Availability

The Cisco Nexus 1000V HA at the network level includes port channels and Link Aggregation Control Protocol (LACP). A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to eight physical links. If a member port within a port channel fails, the traffic previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, LACP lets you configure up to 16 interfaces into a port channel. A maximum of eight interfaces can be active, and a maximum of eight interfaces can be placed in a standby state.

For additional information about port channels and LACP, see the *Cisco Nexus 1000V Layer 2 Switching Configuration Guide, Release 4.0*.

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

Problems with High Availability

Table 17-1 provides symptoms related to high availability, their possible causes, and recommended solutions.

Table 17-1 ***Problems with High Availability***

Symptom	Possible Cause	Solution
The active VSM does not see the standby VSM.	Roles are not configured properly. <ul style="list-style-type: none"> primary secondary 	Verify roles, update an incorrect role, and save the configuration. <ol style="list-style-type: none"> Verify the role of each VSM. show system redundancy status Update an incorrect role. system redundancy role Save the configuration. copy run start
	Network connectivity problems between the VSM and the upstream and virtual switches. Problem could be in the control or management VLAN.	Restore connectivity. <ol style="list-style-type: none"> From the vSphere client, shut down the VSM, which should be in standby mode. From the vSphere client, bring up the standby VSM after network connectivity is restored.
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs.	Verify that the VSMs are using the same software version. If not, then reinstall the image. <ol style="list-style-type: none"> Verify software version on both VSMs. show version Reinstall the secondary VSM with the same version used in the primary.
	Fatal errors during gsync process. <ul style="list-style-type: none"> Check the gsyncctrl log using the show system internal log sysmgr gsyncctrl command and look for fatal errors. 	Reload the standby VSM. reload module standby_module_number See the “Reloading a Module” procedure on page 17-9.

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

Table 17-1 Problems with High Availability (continued)

Symptom	Possible Cause	Solution
The standby VSM reboots periodically.	The VSM has connectivity only through the management interface. When a VSM is able to communicate through the management interface, but not through the control interface, the active VSM resets the standby to prevent the two VSMs from being in HA mode and out of sync.	Check control VLAN connectivity between the primary and secondary VSM. show system internal redundancy info In the output, degraded_mode flag = true . If there is no connectivity, restore it through the control interface.
	VSMs have different versions. Enter the debug system internal sysmgr all command and look for the active_verctrl entry that indicates a version mismatch, as the following output shows: <pre>2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.</pre>	Isolate the standby VSM and boot it. Use the show version command to check the software version in both VSMs. Install the image matching the Active VSM on the standby. For more information, see the <i>Cisco Nexus 1000V High Availability and Redundancy Configuration Guide, Release 4.0(4)SV1(3)</i> .
Both VSMs are in active mode.	Network connectivity problems. <ul style="list-style-type: none"> Check for control and management VLAN connectivity between the VSM at the upstream and virtual switches. When the VSM cannot communicate through any of these two interfaces, they will both try to become active. 	If network problems exist: <ol style="list-style-type: none"> From the vSphere client, shut down the VSM, which should be in standby mode. From the vSphere client, bring up the standby VSM after network connectivity is restored.
	The VSMs have different domain IDs	Verify the domain IDs in each VSM and, if needed, change the incorrect domain ID. <ol style="list-style-type: none"> Verify domain IDs in each VSM. show system internal redundancy info Isolate the VSM with the incorrect domain ID so that it cannot communicate with the other VSM. Change the domain ID in the isolated VSM, save configuration, and power off the VSM. Reconnect and power on the isolated VSM.

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

Table 17-1 Problems with High Availability (continued)

Symptom	Possible Cause	Solution
Active and standby VSMs are not synchronized	<p>Incompatible versions</p> <p>The boot variables for active and standby VSMs are set to different image names, or if image names are the same, the files are not the correct files.</p> <p>When active and standby VSMs are running different versions that are not HA compatible, they are unable to synchronize.</p>	<p>Update the software version or the boot variables.</p> <ol style="list-style-type: none"> From each VSM (active and standby), verify the software version. <p>show version</p> <ol style="list-style-type: none"> Reload the standby VSM with the version that is running in the active by doing one of the following: <ul style="list-style-type: none"> correcting the boot variable names replacing the incorrect software files <p>See the “Reloading a Module” procedure on page 17-9.</p>
	<p>Broadcast traffic problem:</p> <p>Broadcast traffic from standby to active VSM may prevent the VSMs from synchronizing. Standby VSM tries to contact the active VSM periodically, but if broadcast traffic problems persist for over a minute when the standby is booting up, the system cannot synchronize.</p>	<p>Fix the traffic problem and reload the standby VSM.</p> <ol style="list-style-type: none"> From the standby VSM, verify the broadcast traffic problem. <p>show system internal log sysmgr verctrl</p> <p>If so, the following message will appear in the output:</p> <pre>standby_verctrl: no response from the active System Manager</pre> <ol style="list-style-type: none"> Fix network connectivity. Reload standby VSM. <p>reload module standby_module_number</p> <p>See the “Reloading a Module” procedure on page 17-9.</p>
	<p>False standby removal</p> <p>The active VSM falsely detects a disconnect with the standby. The standby is removed and reinserted and synchronization does not occur.</p>	<p>Verify redundancy states and reload the standby VSM.</p> <ol style="list-style-type: none"> Verify active VSM redundancy. <p>show system internal redundancy status</p> <p>Output = RDN_DRV_ST_AC_NP</p> <ol style="list-style-type: none"> Verify standby VSM redundancy. <p>show system internal redundancy status</p> <p>Output = RDN_DRV_ST_SB_AC</p> <ol style="list-style-type: none"> Reload the standby VSM. <p>reload module standby_module_number</p> <p>See the “Reloading a Module” procedure on page 17-9.</p>

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

High Availability Troubleshooting Commands

This section lists commands that can be used troubleshoot problems related to High Availability and includes the following topics:

- [Displaying Cores](#), page 17-6
- [Checking Redundancy](#), page 17-7
- [Checking the System Manager State](#), page 17-8
- [Reloading a Module](#), page 17-9
- [Attaching to the Standby VSM Console](#), page 17-10

Displaying Cores

To list cores, use the following command:

show cores

Example:

```
n1000V# show cores
VDC No Module-num      Process-name      PID      Core-create-time
-----
1      1      private-vlan      3207     Apr 28 13:29
```

Displaying Logs of Processes

To display logs of processes, use the following command and the pid from the output of the **show cores** command:

show processes log [pid pid]

```
n1000V# show processes log
VDC Process      PID      Normal-exit  Stack  Core  Log-create-time
---
1 private-vlan    3207     N      Y      N      Tue Apr 28 13:29:48 2009
```

Example:

```
n1000V# show processes log pid 3207
```

```
=====
Service: private-vlan
Description: Private VLAN

Started at Wed Apr 22 18:41:25 2009 (235489 us)
Stopped at Tue Apr 28 13:29:48 2009 (309243 us)
Uptime: 5 days 18 hours 48 minutes 23 seconds

Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2) <-- Reason for the process abort
Last heartbeat 46.88 secs ago
System image name: nexus-1000v-mzg.4.0.4.SV1.1.bin
System image version: 4.0(4)SV1(1) S25

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

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

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

Checking Redundancy

This section includes the following topics:

- [Checking System Redundancy Status, page 17-7](#)
- [Checking System Internal Redundancy Status, page 17-7](#)

Checking System Redundancy Status

To check system redundancy status, use the following command:

show system redundancy status

```
N1000V# show system redundancy status
Redundancy role
-----
      administrative:  primary <-- Configured redundancy role
      operational:    primary <-- Current operational redundancy role

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

This supervisor (sup-1)
-----
      Redundancy state:  Active <-- Redundancy state of this VSM
      Supervisor state:  Active
      Internal state:    Active with HA standby

Other supervisor (sup-2)
-----
      Redundancy state:  Standby <-- Redundancy state of the other VSM
      Supervisor state:  HA standby
      Internal state:    HA standby <-- The standby VSM is in HA mode and in sync
```

Checking System Internal Redundancy Status

To check the system internal redundancy status, use the following command:

show system internal redundancy info

```
n1000V# show system internal redundancy info
My CP:
  slot: 0
  domain: 184 <-- Domain id used by this VSM
  role: primary <-- Redundancy role of this VSM
  status: RDN_ST_AC <-- Indicates redundancy state (RDN_ST) of the this VSM is Active (AC)
  state: RDN_DRV_ST_AC_SB
  intr: enabled
  power_off_reqs: 0
  reset_reqs: 0
Other CP:
  slot: 1
```

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

```

status: RDN_ST_SB <-- Indicates redundancy state (RDN_ST) of the other VSM is Standby
(SB)
active: true
ver_rcvd: true
degraded_mode: false <-- When true, it indicates that communication through the control
interface is faulty
Redun Device 0: <-- This device maps to the control interface
name: ha0
pdev: ad7b6c60
alarm: false
mac: 00:50:56:b7:4b:59
tx_set_ver_req_pkts: 11590
tx_set_ver_rsp_pkts: 4
tx_heartbeat_req_pkts: 442571
tx_heartbeat_rsp_pkts: 6
rx_set_ver_req_pkts: 4
rx_set_ver_rsp_pkts: 1
rx_heartbeat_req_pkts: 6
rx_heartbeat_rsp_pkts: 442546 <-- Counter should be increasing, as this indicates that
communication between VSM is working properly.
rx_drops_wrong_domain: 0
rx_drops_wrong_slot: 0
rx_drops_short_pkt: 0
rx_drops_queue_full: 0
rx_drops_inactive_cp: 0
rx_drops_bad_src: 0
rx_drops_not_ready: 0
rx_unknown_pkts: 0
Redun Device 1: <-- This device maps to the mgmt interface
name: hal
pdev: ad7b6860
alarm: true
mac: ff:ff:ff:ff:ff:ff
tx_set_ver_req_pkts: 11589
tx_set_ver_rsp_pkts: 0
tx_heartbeat_req_pkts: 12
tx_heartbeat_rsp_pkts: 0
rx_set_ver_req_pkts: 0
rx_set_ver_rsp_pkts: 0
rx_heartbeat_req_pkts: 0
rx_heartbeat_rsp_pkts: 0 <-- When communication between VSM through the control
interface is interrupted but continues through the mgmt interface, the
rx_heartbeat_rsp_pkts will increase.
rx_drops_wrong_domain: 0
rx_drops_wrong_slot: 0
rx_drops_short_pkt: 0
rx_drops_queue_full: 0
rx_drops_inactive_cp: 0
rx_drops_bad_src: 0
rx_drops_not_ready: 0
rx_unknown_pkts: 0

```

Checking the System Manager State

To check the system internal sysmgr state, use the following command:

```
show system internal sysmgr state
```

```
N1000V# show system internal sysmgr state
```

```
The master System Manager has PID 1988 and UUID 0x1.
```

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

```
Last time System Manager was gracefully shutdown.
The state is SRV_STATE_MASTER_ACTIVE_HOTSTDBY entered at time Tue Apr 28 13:09:13 2009.

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

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

Hap-reset is currently enabled.

Watchdog checking is currently disabled.

Watchdog kgdb setting is currently enabled.
```

Debugging info:

```
The trace mask is 0x00000000, the syslog priority enabled is 3.
The '-d' option is currently disabled.
The statistics generation is currently enabled.
```

HA info:

```
slotid = 1    supid = 0
cardstate = SYSMGR_CARDSTATE_ACTIVE .
cardstate = SYSMGR_CARDSTATE_ACTIVE (hot switchover is configured enabled).
Configured to use the real platform manager.
Configured to use the real redundancy driver.
Redundancy register: this_sup = RDN_ST_AC, other_sup = RDN_ST_SB.
EOBC device name: eth0.
Remote addresses:  MTS - 0x00000201/3      IP - 127.1.1.2
MSYNC done.
Remote MSYNC not done.
Module online notification received.
Local super-state is: SYSMGR_SUPERSTATE_STABLE
Standby super-state is: SYSMGR_SUPERSTATE_STABLE
Swover Reason : SYSMGR_SUP_REMOVED_SWOVER <-- Reason for the last switchover
Total number of Switchovers: 0 <-- Number of switchovers
>> Duration of the switchover would be listed, if any.
```

Statistics:

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

Reloading a Module

To reload a module, use the following command:

```
reload module
```



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

Example:

```
n1000V# reload module 2
```

In this example, the secondary VSM is reloaded.

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

Attaching to the Standby VSM Console

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

To attach to the standby VSM console, use the following command.

attach module

Example:

```
n1000V# attach module 2
```

This example shows how to attach to the console of the secondary VSM.



CHAPTER 18

System

This chapter describes the Cisco Nexus 1000V system and how to identify and correct problems related to the system.

This chapter includes the following sections:

- [Information About the System, page 18-1](#)
- [General Restrictions for vCenter Server, page 18-2](#)
- [Recovering a DVS, page 18-2](#)
- [Problems Related to VSM and vCenter Server Connectivity, page 18-5](#)
- [VSM Creation, page 18-9](#)
- [Port Profiles, page 18-9](#)
- [Problems with Hosts, page 18-10](#)
- [Problems with VM Traffic, page 18-10](#)
- [Connection Failure After ESX Reboot, page 18-6](#)
- [VEM Troubleshooting Commands, page 18-10](#)
- [VEM Log Commands, page 18-11](#)
- [Error Messages, page 18-12](#)

Information About the System

Cisco Nexus 1000V provides Layer 2 switching functions in a virtualized server environment. you can configure and monitor the virtual switch in an ESX server using the Cisco NX-OS command line interface. You also have visibility into the networking components of the ESX servers and access to the virtual switches within the network.

The Cisco Nexus 1000V manages a data center defined by the vCenter Server. Each server in the Datacenter is represented as a linecard and can be managed as if it were a line card in a physical Cisco switch. The Cisco Nexus 1000V has two components:

- Virtual supervisor module (VSM) – This is the control software of the Cisco Nexus 1000V distributed virtual switch. It runs on a virtual machine (VM) and is based on Cisco NX-OS.
- Virtual Ethernet module (VEM) – This is the part of Cisco Nexus 1000V that actually switches data traffic. It runs on a VMware ESX 4.0 host. Several VEMs are controlled by one VSM. All the VEMs that form a switch domain should be in the same virtual Datacenter as defined by VMware vCenter Server.

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

For more information, see the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

General Restrictions for vCenter Server

When you are troubleshooting issues related to vCenter Server, make sure that you observe the following restrictions:

- The name of a distributed virtual switch (DVS) name must be unique across Datacenters
- You create a DVS in a network folder
- A Datacenter cannot be removed unless the DVS folder or the underlying DVS is deleted.
- A DVS can be deleted only with the help of VSM using the **no vmware dvs** command in config-svs-conn mode.
- The no vmware dvs command can succeed only if there are no VMs using the DVS port-groups.
- A port group on vCenter Server can be deleted only if there are no interfaces associated with it.
- A sync operation performed in conjunction with the **connect** command helps VSM keep in sync with vCenter Server.
- Each VSM uses a unique extension key to communicate with vCenter Server and perform operations on a DVS.

Extension Key

The VSM uses the extension key when communicating with the vCenter Server. Each VSM has its own unique extension key, such as Cisco_Nexus_1000V_32943215

Use the **show vmware vc extension-key** command to find the extension key of the VSM. It is also listed in the .xml file.

The extension key registered on the vCenter Server can be found through the MOB. For more information, see the [“Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8](#).

The same extension key cannot be used to create more than one DVS on the vCenter Server.

Recovering a DVS

You can use this procedure to recover a DVS if the VSM VM that was used to create it is lost or needs to be replaced. This section includes the following procedures:

- [Recovering a DVS With a Saved Copy of the VSM, page 18-3](#)
- [Recovering a DVS Without a Saved Copy of the VSM, page 18-4](#)

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

Recovering a DVS With a Saved Copy of the VSM

You can use this procedure to recover a DVS when you have previously saved a back up copy of the VSM configuration file.

BEFORE YOU BEGIN

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

- Use this procedure if you have previously saved a back up copy of the VSM configuration file. If you have not previously saved a back up copy, see the [“Recovering a DVS Without a Saved Copy of the VSM” procedure on page 18-4](#).
- Make sure that the VSM VM switchname is the same as the DVS switchname on the vCenter Server. This allows the VSM configuration to synchronize with the correct DVS on the vCenter Server.

To change the VSM switchname use the **switchname newname** command.

-
- Step 1** From the MOB, find the DVS extension key.
For more information, see the [“Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8](#).
- Step 2** On the VSM, add the DVS extension key found in [Step 1](#).
The extension key allows the VSM to log in to the vCenter server.
Example:
n1000v# **config t**
n1000v(config)# **vmware vc extension-key Cisco_Nexus_1000V_32943215**
- Step 3** From the MOB, unregister the extension key found in [Step 1](#).
For more information, see the [“Unregistering the Extension Key in the vCenter Server” procedure on page 3-12](#).
- Step 4** From the VC client, register the extension (plug-in) for the VSM.
For more information see the following procedure in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.
 - [Creating a Cisco Nexus 1000V Plug-In on the vCenter Server](#)
- Step 5** On the VSM, restore the configuration using a previously saved copy of the VSM configuration file.
copy path/filename running-config
Example:
n1000v# **copy sftp://user1@172.22.36.10/backup/hamilton_cfg running-config**
- Step 6** Do one of the following:
 - If the vCenter server connection is not part of the previously saved configuration, continue with the next step.
 - Otherwise, go to [Step 8](#).
- Step 7** On the VSM, restore the configuration for the vCenter server connection.
Example:
n1000v# **config t**
n1000v (config)# **svs connection VC**
n1000v(config-svs-conn#) **protocol vmware-vim**
n1000v(config-svs-conn#) **remote ip address 192.168.0.1**
n1000v(config-svs-conn#) **vmware dvs datacenter-name Hamilton-DC**

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

Step 8 Connect to vCenter Server.

Example:

```
n1000v(config-svs-conn#) connect
```

You can now use the old DVS or remove it.

Recovering a DVS Without a Saved Copy of the VSM

You can use this procedure to recover a DVS when you have not previously saved a back up copy of the VSM configuration file.

BEFORE YOU BEGIN

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

- The folder in which the VSM resides must be:
 - At the root-level of the Data Center in which it resides. It cannot be embedded in another folder.
 - Of the same name as the VSM.

If the folder does not meet the above criteria, the connection to vCenter server fails with the error, *the VSM already exists*.

- Use this procedure if you have not previously saved a back up copy of the VSM configuration file. If you have previously saved a back up copy, then see the [“Recovering a DVS With a Saved Copy of the VSM” procedure on page 18-3](#).
- If you have not previously saved a back up copy of the VSM configuration file, then you may try recreating the old port profiles before connecting to the VC. This procedure has a step for recreating port profiles. If you do not recreate these before connecting to VC, then all the port groups present on the VC are removed and all ports in use are moved to the quarantine port groups.
- Make sure that the VSM VM switchname is the same as the DVS switchname on the vCenter Server. This allows the VSM configuration to synchronize with the correct DVS on the vCenter Server.

To change the VSM switchname use the **switchname** *newname* command.

Step 1 From the MOB, find the DVS extension key.
For more information, see the [“Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8](#).

Step 2 On the VSM, add the DVS extension key found in [Step 1](#).
The extension key allows the VSM to log in to the vCenter server.

Example:

```
n1000v# config t
n1000v(config)# vmware vc extension-key Cisco_Nexus_1000V_32943215
```

Step 3 From the MOB, unregister the extension key found in [Step 1](#).
For more information, see the [“Unregistering the Extension Key in the vCenter Server” procedure on page 3-12](#).

Step 4 From the VC client, register the extension (plug-in) for the VSM.

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

For more information see the following procedure in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

- Creating a Cisco Nexus 1000V Plug-In on the vCenter Server

Step 5 Manually recreate the old port profiles from your previous configuration.

For more information, see the following procedures in the *Cisco Nexus 1000V Getting Started Guide, Release 4.0(4)SV1(3)*.

- Configuring the System Port Profile for VSM-VEM Communication
- Configuring the Uplink Port Profile for VM Traffic
- Configuring the Data Port Profile for VM Traffic



Note If you do not manually recreate the port profiles, then all port groups on the vCenter Server are removed when the VSM connects.

Step 6 On the VSM, restore the configuration for the vCenter server connection.

Example:

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

Step 7 Connect to vCenter Server.

Example:

```
n1000v(config-svs-conn#) connect
```

You can now use the old DVS or remove it.

Problems Related to VSM and vCenter Server Connectivity

Symptom	Solution
The vCenter Server connection seems to succeed, but does not.	Make sure that the domain ID is configured correctly.
The svcs connection command fails.	<ol style="list-style-type: none"> 1. Make sure you have configured all parameters for the svcs connection command. 2. Make sure you can ping the vCenter Server IP address. 3. Make sure that the proxy.xml file is correct for both the IP address and length. 4. Restart the vCenter Server.
The connection fails after an ESX reboot	“Connection Failure After ESX Reboot” procedure on page 18-6

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

Symptom	Solution
The host does not show up in the Add host to DVS screen.	Make sure that the Host is installed with VMware Enterprise plus license containing the Distributed Virtual Switch feature.
Add host to DVS returns an error.	Confirm that the VEM software is installed on the ESX server,
The server name column of the show module command output shows the IP address.	The server name shows the host-name or IP address, whichever was used to add the host to the DVS on the vCenter Server.

Connection Failure After ESX Reboot

If you lose connectivity between the VSM and VEM following a reboot of the ESX server host, it may be caused by a reset of the MTU to its default on the host.

This section includes the following topics:

- [Information About Changing the MTU, page 18-6](#)
- [Verifying that an MTU Mismatch Caused a Connection Failure, page 18-7](#)
- [Restoring Network Connectivity to the VEM, page 18-8](#)

Information About Changing the MTU

If you use an MTU other than 1500 (the default) for a physical NIC attached to the Cisco Nexus 1000V, then reboots of the ESX can revert this setting to the default setting of 1500. This results in a mismatch with the MTU setting on one of the following and a potential connection failure of the VSM and VEM.

- VMware kernel NIC
- Upstream switchport of the VEM uplink

For example, in networks where you use jumbo frames, you may manually configure an MTU of other than 1500. If the ESX reboots during a power cycle, the MTU of the physical NIC reverts to the default of 1500 but the MTU setting on the VMware kernel NIC or upstream switchport does not.

To verify that an MTU setting mismatch has caused a connection failure, see the [“Verifying that an MTU Mismatch Caused a Connection Failure” procedure on page 18-7](#).

To restore connectivity that was lost due to an MTU mismatch, see the [“Restoring Network Connectivity to the VEM” procedure on page 18-8](#).

To prevent a loss of connectivity between the VSM and VEM, and preserve a physical NIC MTU setting of other than 1500 (the default) across reboots of the ESX, you must configure a system MTU in the system uplink port profile.

To configure the **system mtu** on the system uplink port profile, see the following procedure in the *Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)*

- *Creating a System Port Profile*

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

Verifying that an MTU Mismatch Caused a Connection Failure

You can use this procedure to verify that the cause of a VEM connection failure is an MTU mismatch.

BEFORE YOU BEGIN

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

- You are logged in to the VEM.
- You are logged in to the CLI of the upstream switch.

Step 1 On the upstream switch, use the following command to verify the MTU on the interface connected to the VEM uplink.

show interface

show port

Example:

```
Console> show port jumbo
Jumbo frames MTU size is 9216 bytes.
Jumbo frames enabled on port(s) 6/1-2,7/1-8.
Console>
```

Example:

```
switch# show interface vlan 10
Vlan10 is up, line protocol is up
  Hardware is EtherSVI, address is 0005.9b78.6e7c
  MTU 8000 bytes, BW 1000000 Kbit, DLY 10 usec,
switch#
```

Step 2 On the ESX, verify the MTU setting for the physical NIC attached to the Cisco Nexus 1000V.

vem status

```
~ # vem status
VEM modules are loaded
```

Switch Name	Num Ports	Used Ports	Configured Ports	MTU	Uplinks
vSwitch0	64	3	64	1500	vmnic0
DVS Name	Num Ports	Used Ports	Configured Ports	Uplinks	
n1000v	256	9	256		vmnic1 VEM Agent is running

Step 3 On the ESX, identify the VMware kernel NICs that are attached to the Cisco Nexus 1000V.

vemcmd show port

Step 4 On the ESX, verify the MTU settings for the VMware kernel NICs identified in [Step 3](#),

esxcfg-vmknic -l

Step 5 Do one of the following:

- If the MTU settings do not match, go to the [“Restoring Network Connectivity to the VEM” procedure on page 18-8](#).
- If the MTU settings match, then this is not the cause of the connection failure.

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

Restoring Network Connectivity to the VEM

You can use this procedure to restore network connectivity with a VEM that was lost due to an MTU mismatch between the physical NIC and the VMware kernel NIC.

BEFORE YOU BEGIN

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

- You are logged in to the ESX host.

Step 1 Use the following commands to reset the MTU and restore connectivity that was lost due to an MTU mismatch with the physical NIC attached to the Cisco Nexus 1000V:

```
vemcmd show port port-LTL-number
```

```
vemcmd set mtu size ltl port-LTL-number
```

The MTU is reset on the ESX and connectivity with the VEM is restored.

Example:

```
~ # vemcmd show port 48
LTL    IfIndex  Vlan   Bndl  SG_ID Pinned_SGID  Type  Admin State  CBL Mode
Name   . . .
17    1a030100    1 T   304    1          32  PHYS    UP    UP    1  Trunk vmnic1
~# vemcmd set mtu 9000 ltl 17
```

Step 2 On the VSM, update the MTU setting using one of the following documents:

- To update the system port profile for system uplinks:
Cisco Nexus 1000V Port Profile Configuration Guide, Release 4.0(4)SV1(3)
 - To update the interface configuration for non-system uplinks:
Cisco Nexus 1000V Interface Configuration Guide, Release 4.0(4)SV1(3)
-

Displaying VC Event History

Example 18-1 shows the **show vms internal event-history errors** command that is useful for examining VC errors in detail. It shows whether an error is caused by a VSM (client) or the server.

Example 18-1 show vms internal event-history error Command

```
n1000v# show vms internal event-history errors

Event:E_DEBUG, length:239, at 758116 usecs after Tue Feb  3 18:21:58 2009
  [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: "":ServerFaultCode [VMWARE-VIM]
A DVS n1000v with spec.name as n1000v already exists, cannot create DVS n1000v. A
specified parameter was not correct.spec.name

Event:E_DEBUG, length:142, at 824006 usecs after Tue Feb  3 18:18:30 2009
  [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: SOAP-ENV:Client [VMWARE-VIM]
Operation could not be completed due to connection failure.

Event:E_DEBUG, length:134, at 468208 usecs after Tue Feb  3 18:15:37 2009
  [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: "":ServerFaultCode [VMWARE-VIM]
Extension key was not registered before its use.
```

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

VSM Creation

Symptom	Possible Causes	Solution
The VSM VM is stuck at the boot prompt.	—	Make sure that you have three e1000 NICs.
The VSM VM cannot ping itself.	—	Configure the mgmt0 interface.
The VSM VM can ping itself, but not the gateway.	—	Make sure the NIC order is correct: control, mgmt, inband.
The VSM VM can ping the gateway, but not the outside subnet.	—	Configure vrf context management.

Port Profiles

When creating a port profile, use the following commands to create the corresponding port groups on the vCenter Server:

- **vmware port-group**
- **state enabled**

Profiles that have the system VLAN configuration allow the VEM to communicate with the VSM.

Make sure that the system port-profile is defined with the right system VLANs.

Use the **show port-profile** and **show port-profile usage** commands to collect basic required information.

Problems with Port Profiles

Symptom	Possible Causes	Solution
You receive an error message “Possible failure in communication with vCenter Server.”	The VSM is not connected to the vCenter Server.	Issue the svs connection vc command to connect to the vCenter Server.
	The port group name is not unique.	Port group names must be unique within a vCenter Server Datacenter.
Port profile or port groups do not appear on the vCenter Server.	—	Make sure you have issued the vmware port-group command and state enable command.

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

Problems with Hosts

Symptom	Solution
You receive an error message, DVS Operation failed for one or more members.”	<p>Issue the vem status -v command to verify if the VEM is running on the host.</p> <p>Issue the vem unload command to unload the VEM.</p> <p>In the vSphere Client, remove the stale DVS:</p> <ol style="list-style-type: none"> Go to the Host tab Networking->Configuration->Distributed Virtual Switch Click Remove.
The host is visible on the vCenter Server, but not the VSM.	<p>Issue the vemcmd show trunk command to verify that there is an uplink carrying the control VLAN. The profile applied to the uplink must be a system profile with a control VLAN as a system VLAN.</p> <p>Verify the control VLAN in the upstream switch port and the path to the VSM VM. Make sure that one uplink at most carries the control VLAN, or that all uplinks and upstream ports carrying the control VLAN are in port channels.</p>
A module flap occurs.	The VSM may be overloaded. Make sure that you have 1 GB of memory and CPU shares for the VSM VM on the vCenter Server.

Problems with VM Traffic

When troubleshooting problems with intra-host VM traffic, follow these guidelines:

- Make sure that at least one of the VMware virtual NICs is on the correct DVS port group and is connected.
- If the VMware virtual NIC is down, determine if there is a conflict between the MAC address configured in the OS and the MAC address assigned by VMware. You can see the assigned MAC addresses in the vmx file.

When troubleshooting problems with inter-host VM traffic, follow these guidelines:

- Determine if there is exactly one uplink sharing a VLAN with the VMware virtual NIC. If there is more than one, they must be in a port channel.
- Ping a SVI on the upstream switch using the **show intX counters** command.

VEM Troubleshooting Commands

Use the following commands to display VEM information:

- **vemlog** – displays and controls VEM kernel logs

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

- **vemcmd** – displays configuration and status information
- **vem-support all** – collects support information
- **vem status**– collects status information
- **vem version**– collects version information
- **vemlog show last *number-of-entries*** – displays the circular buffer

Example 18-2 vemlog show last Command

```
[root@esx-cos1 ~]# vemlog show last 5
Timestamp                Entry CPU Mod Lv      Message
Oct 13 13:15:52.615416   1095  1    1  4 Warning vssnet_port_pg_data_ ...
Oct 13 13:15:52.620028   1096  1    1  4 Warning vssnet_port_pg_data_ ...
Oct 13 13:15:52.630377   1097  1    1  4 Warning sv_sswitch_state ...
Oct 13 13:15:52.633201   1098  1    1  8 Info    vssnet new switch ...
Oct 13 13:16:24.990236   1099  1    0  0      Suspending log
```

- **vemlog show info** – displays information about entries in the log

Example 18-3 vemcmd show info Command

```
[root@esx-cos1 ~]# vemlog show info
Enabled: Yes
Total Entries: 1092
Wrapped Entries: 0
Lost Entries: 0
Skipped Entries: 0
Available Entries: 6898
Stop After Entry: Not Specified
```

- **vemcmd help** – displays the type of information you can display

Example 18-4 vemcmd help Command

```
[root@esx-cos1 ~]# vemcmd help
show card                Show the card's global info
show vlan [vlan]         Show the VLAN/BD table
show bd [bd]             Show the VLAN/BD table
show l2 <bd-number>     Show the L2 table for a given BD/VLAN
show l2 all              Show the L2 table
show port [priv|vsm]     Show the port table
show pc                  Show the port channel table
show portmac             Show the port table MAC entries
show trunk [priv|vsm]   Show the trunk ports in the port table
show stats               Show port stats
```

VEM Log Commands

Use the following commands to control the vemlog:

- **vemlog stop** – stops the log
- **vemlog clear** – clear s the log
- **vemlog start *number-of-entries*** – starts the log and stops it after the specified number of entries
- **vemlog stop *number-of-entries*** – stops the log after the next specified number of entries

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

- **vemlog resume** – starts the log, but does not clear the stop value

Error Messages

On the vSphere Client, you can see error messages under the recent tasks tab. You can find detailed description of the error under the Tasks and Events tab. The same messages are also propagated to the VSM.

Table 18-1 lists error messages that you might see on the VSM.

Table 18-1 Error Messages on the VSM

Error	Description
ERROR: [VMWARE-VIM] Extension key was not registered before its use	This error indicates that VSM extension key is not registered.
ERROR: [VMWARE-VIM] A DVS n1000v with spec.name as n1000v already exists, cannot create DVS n1000v. A specified parameter was not correct. spec.name.	This error is displayed after you enter the first connect command, and indicates that a DVS already exists with the same name.
ERROR: [VMWARE-VIM] A DVS n1000v with spec.extensionKey as Cisco_Nexus_1000V_2055343757 already exists, cannot create DVS new-n1000v. A specified parameter was not correct. spec.extensionKey	This error is displayed when the VSM tries to create a different DVS after changing the switch name.
ERROR: [VMWARE-VIM] A DVS n1000v with name as n1000v already exists, cannot reconfigure DVS test. A specified parameter was not correct. Spec.name	This error indicates that a DVS with the same name already exists.
Warning: Operation succeeded locally but update failed on vCenter server.[VMWARE-VIM] DVPortgroup test port 0 is in use. The resource vim.dvs.DistributedVirtualPort 0 is in use.	This warning is displayed when the VSM tries to delete the port profile if the VSM is not aware of the nics attached to the port groups.



CHAPTER 19

Before Contacting Technical Support

This chapter describes the steps to take before calling for technical support and includes the following sections:

- [Gathering Information for Technical Support, page 19-1](#)
- [Obtaining a File of Core Memory Information, page 19-2](#)
- [Copying Files, page 19-3](#)



Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtm>

Gathering Information for Technical Support

Use this procedure to gather information about your network that is needed by your customer support representative or Cisco TAC.



Note Required logs and counters are part of volatile storage and will not persist through a reload. Do not reload the module or the switch until you have completed this procedure.

DETAILED STEPS

-
- Step 1** Configure your Telnet or SSH application to log screen output to a text file.
 - Step 2** Set the number of lines that appear on the screen so that pausing is disabled.
terminal length 0
 - Step 3** Display the configuration information needed to troubleshoot your network.
show tech-support svcs
 - Step 4** Capture the error codes that appear in your message logs.
 - **show logging logfile**
Displays the contents of the logfile.
 - **show logging last *number***
Displays the last few lines of the logfile.

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

- Step 5** Gather your answers to the following questions:
- On which switch or port is the problem occurring?
 - Cisco Nexus 1000V software, driver versions, operating systems versions and storage device firmware are in your fabric?
 - ESX and vCenter Server software that you are running?
 - What is the network topology?
 - Were any changes being made to the environment (VLANs, adding modules, upgrades) prior to or at the time of this event?
 - Are there other similarly configured devices that could have this problem, but do not?
 - Where was this problematic device connected (which switch and interface)?
 - When did this problem first occur?
 - When did this problem last occur?
 - How often does this problem occur?
 - How many devices have this problem?
 - Were any traces or debug output captured during the problem time? What troubleshooting steps have you attempted? Which, if any, of the following tools were used?
 - Ethalyzer, local or remote SPAN
 - CLI debug commands
 - traceroute, ping
- Step 6** Is your problem related to a software upgrade attempt?
- What was the original Cisco Nexus 1000V version?
 - What is the new Cisco Nexus 1000V version?
-

Obtaining a File of Core Memory Information

Cisco customer support engineers often use files from your system for analysis. One of these is a file containing memory information, and is referred to as a core dump. The file is sent to a TFTP server or to a Flash card in slot0: of the local switch. You should set up your switch to generate this file under the instruction of your customer support representative, and send it to a TFTP server so that it can be e-mailed to them.

To generate a file of core memory information, or a core dump, use the command in the following example.

```
n1000v# system cores tftp://10.91.51.200/jsmith_cores
n1000v# show system cores
Cores are transferred to tftp://10.91.51.200/jsmith_cores
```



Note

The file name (indicated by `jsmith_cores`) must exist in the TFTP server directory.

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

Copying Files

It may be required to move files to or from the switch. These files may include log, configuration, or firmware files.

Cisco Nexus 1000V always acts as a client, such that an ftp/scp/tftp session will always originate from the switch and either push files to an external system or pull files from an external system.

```
File Server: 172.22.36.10
File to be copied to the switch: /etc/hosts
```

The **copy** CLI command supports four transfer protocols and 12 different sources for files.

```
n1000v# copy ?
bootflash: Select source filesystem
core: Select source filesystem
debug: Select source filesystem
ftp: Select source filesystem
licenses Backup license files
log: Select source filesystem
modflash: Select source filesystem
nvram: Select source filesystem
running-config Copy running configuration to destination
scp: Select source filesystem
sftp: Select source filesystem
slot0: Select source filesystem
startup-config Copy startup configuration to destination
system: Select source filesystem
tftp: Select source filesystem
volatile: Select source filesystem
```

Use the following syntax to use secure copy (scp) as the transfer mechanism:

```
"scp://[username@]server[/path]"
```

To copy `/etc/hosts` from 172.22.36.10 using the user `user1`, where the destination would be `hosts.txt`, use the following command:

```
n1000v# copy scp://user1@172.22.36.10/etc/hosts bootflash:hosts.txt
user1@172.22.36.10's password:
hosts 100% |*****| 2035 00:00
```

To back up the startup-configuration to a sftp server, use the following command:

```
n1000v# copy startup-config sftp://user1@172.22.36.10/test/startup-configuration.bak1
Connecting to 172.22.36.10...
User1@172.22.36.10's password:
n1000v#
```



Tip

Backing up the startup-configuration to a server should be done on a daily basis and prior to any changes. A short script could be written to be run on Cisco Nexus 1000V to perform a save and then backup of the configuration. The script only needs to contain two commands: **copy running-configuration startup-configuration** and then **copy startup-configuration tftp://server/name**. To execute the script use: **run-script filename**.

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



INDEX

A

Access Control Lists. See ACLs.

ACLs

- commands for troubleshooting [10-2](#)
- configuration limits [10-1](#)
- debugging policy verification [10-3](#)
- description [10-1](#)
- displaying ACL policies on the VEM [10-2](#)
- restrictions [10-2](#)

B

bundle, software

- ID match [4-3, 4-4](#)

C

Cisco Nexus 1000V

- system overview [18-1](#)

Cisco Nexus 1000V switch

- copying files to or from [19-3](#)

CLI

- ping command [2-1](#)
- traceroute command [2-2](#)

connectivity

- vCenter Server [3-5](#)
- verifying [1-3](#)
- verifying between VSM and vCenter Server [3-5](#)
- verifying between VSM and VEM [6-11](#)
- VSM and vCenter Server problem symptoms and solutions [18-5](#)

core dumps [19-2](#)

CPU status

- monitoring [2-2](#)

customer support

- collecting information [19-1](#)
- contacting Cisco or VMware [1-7](#)

D

documentation

- additional publications [i-xv, i-xvi](#)
- conventions [i-xiv](#)

domain parameters [3-4](#)

DVS

- find extension key [3-8](#)
- recovering [18-2](#)

E

error messages

- vSphere Client [18-12](#)

extension key

- finding for specific DVS [3-8](#)
- unregistering in vCenter Server [3-12](#)

H

HA

- commands to troubleshoot [17-6](#)
- description [17-1](#)
- network level support [17-2](#)
- problem symptoms and solutions [17-3](#)
- system level support [17-2](#)

High Availability. See HA

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

hosts

problem symptoms and solutions [18-9, 18-10](#)

I

IGMP snooping

commands for troubleshooting [15-2](#)

description [15-1](#)

problem symptoms and solutions [15-4](#)

troubleshooting guidelines [15-2](#)

installation

recreating the installation flowchart [3-3, 3-10](#)

Internet Group Management Protocol snooping. See IGMP snooping

J

jumbo frames and MTU [18-6](#)

K

key, extension [3-8](#)

L

Layer 2 switching

inter-VEM ping [9-5](#)

intra-VEMping [9-4](#)

overview [9-1](#)

problem symptoms and solutions [9-7](#)

traffic interruptions [9-6](#)

license

Cisco Nexus N1000V license package [5-1](#)

contents of Cisco Nexus N1000V license file [5-3](#)

installation issues [5-3](#)

post-installation issues [5-4](#)

removal [5-5](#)

troubleshooting checklist [5-3](#)

usage [5-3](#)

VMware Enterprise Plus [3-1](#)

licensed module [5-1](#)

Link Aggregation Control Protocol (LACP) [17-2](#)

logging levels [2-5](#)

logs [1-6](#)

M

MAC address tables

verifying [9-7](#)

module

licensed [5-1](#)

not coming up on the VSM [6-1](#)

unlicensed [5-1](#)

verifying module state [7-2](#)

virtual Ethernet module (VEM) [6-1](#)

virtual supervisor module (VSM) [6-1](#)

MTU settings, with jumbo frames [18-6](#)

multicast

description [15-1](#)

N

NetFlow

commands for troubleshooting [12-2](#)

configuration problems [12-3](#)

description [12-1](#)

network adapter [3-5](#)

Nexus 1000V switch

plug-in [3-7](#)

P

port channels

asymmetric [8-3](#)

cannot create port channel [8-4](#)

commands to troubleshoot [8-2](#)

description [8-1](#)

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

forcing port channel characteristics onto an interface [8-4](#)

initial checklist [8-2](#)

interface does not come online [8-4](#)

troubleshooting checklist [8-2](#)

verifying a port channel configuration [8-6](#)

port groups

assigning to VSM VM [3-4](#)

virtual interfaces [3-4](#)

port profiles

commands to troubleshoot [7-14](#)

creating corresponding port groups on vCenter Server [18-9](#)

debug logs [7-14](#)

description [7-13](#)

problem symptoms and solutions [7-19](#)

system port profiles [7-18](#)

transferring from the VSM to vCenter Server [7-19](#)

ports

error disabled [7-7](#)

flapping [7-6](#)

interface configuration [7-2](#)

interface description [7-1](#)

link flapping [7-6](#)

overview [7-1](#)

ping from a VM with port security enabled [7-9](#)

port counters [7-4](#)

port enabled and port security is ErrDisabled [7-11](#)

port interface cannot be enabled [7-5](#)

port security address count exceed violation [7-11](#)

port security MAC move violation [7-12](#)

port security problems [7-8](#)

port security restrictions and limitations [7-12](#)

port state is ErrDisabled [7-7](#)

port state is link failure or not connected [7-6](#)

port types [9-4](#)

troubleshooting checklist [7-2](#)

troubleshooting with CLI [7-3](#)

verifying [1-3](#)

viewing port state [7-3](#)

private VLANs

commands to troubleshoot [14-2](#)

description [14-1](#)

troubleshooting guidelines [14-2](#)

Q

QoS

commands to troubleshoot [11-2](#)

configuration limits [11-1](#)

debugging policy verification errors [11-3](#)

description [11-1](#)

troubleshooting QoS policies on the VEM [11-2](#)

Quality of Service. See QoS

R

RADIUS

accounting logs [2-5](#)

related documents [i-xv, i-xvi](#)

S

software

core dumps [19-2](#)

software bundle ID, match [4-3, 4-4](#)

SPAN

commands to troubleshoot [16-3](#)

configuration guidelines [16-3](#)

description [16-1](#)

destinations [16-2](#)

ERSPAN destinations [16-2](#)

problem symptoms and solutions [16-4](#)

session requirements [16-3](#)

sessions [16-2](#)

sources [16-1](#)

Switched Port Analyzer. See SPAN

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

synchronization problems [7-4](#)

syslog

See system messages

system messages

explanation and recommended action [1-5](#)

logging levels [2-5](#)

overview [1-4, 2-5](#)

syslog server [1-5](#)

using CLI [1-5](#)

system port profiles

description [7-18](#)

system processes

monitoring [2-2](#)

T

troubleshooting process

best practices [1-1](#)

common CLI commands [1-3](#)

general process steps [1-2](#)

guidelines [1-2](#)

overview [1-1](#)

trunking

initial checklist [8-2](#)

overview [8-2](#)

U

unlicensed module [5-1](#)

unregistering an extension key [3-12](#)

V

vCenter Server

refreshing connection [3-3](#)

removing the VSM [3-11](#)

restrictions [18-2](#)

unregistering the VSM [3-12](#)

verifying connection to VSM [6-6](#)

verifying correct configuration [6-10](#)

VEM

commands for vemlog [18-11](#)

commands to troubleshoot [18-10](#)

domain parameters [3-4](#)

physical ports [9-2](#)

unlicensed [5-2](#)

verifying correct configuration [6-13](#)

view of ports [9-2](#)

virtual ports [9-2](#)

verifying

MAC address tables [9-7](#)

Virtual Ethernet Module. See VEM.

virtual Ethernet port (veth) [9-2](#)

virtual NIC [9-2](#)

Virtual Supervisor Module. See VSM.

VLAN

cannot create [13-3](#)

traffic does not traverse trunk [8-6](#)

troubleshooting checklist [13-2](#)

VLANs

description [13-1](#)

VM

improving performance [3-4](#)

traffic problems [18-10](#)

vmnic

number allocation [9-3](#)

VSM

commands to troubleshoot [6-16](#)

creating [18-9](#)

domain parameters [3-4](#)

identifying extension key [3-12](#)

removing hosts from [3-11](#)

verifying correct configuration [6-7](#)

view of ports [9-3](#)

vSphere Client

error messages [18-12](#)

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

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