



Cisco Nexus 1000V Troubleshooting Guide, Release 4.2(1)SV2(2.2)

April 13, 2015

Cisco Systems, Inc.

www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Text Part Number: OL-31383-01

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

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

Cisco Nexus 1000V Troubleshooting Guide, Release 4.2(1)SV2(2.2)
© 2014 Cisco Systems, Inc. All rights reserved.



New and Changed Information **xiii**

Preface **xv**

Audience **xv**

Document Conventions **xv**

Related Documentation **xvi**

xviii

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

xviii

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**

Cisco Support Communities **1-7**

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**

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
Resolving Extension Key Conflicts	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
Unregister the Extension Key in the vCenter Server	3-12
Problems with the Nexus 1000V Installation Management Center	3-13
Licenses	4-1
Information About Licenses	4-1
Contents of the License File	4-2
Prerequisites to License Troubleshooting	4-2
Problems with Licenses	4-3
License Troubleshooting Commands	4-4

Upgrade 5-1

- Information about Upgrades 5-1
- Problems with the In Service Software Upgrade 5-1
- Problems with the VEM Upgrade 5-5
- Problems with the GUI Upgrade 5-6
 - Recovering a Secondary VSM with Active Primary 5-7
 - Stopping a VSM Upgrade 5-8
 - Changing Boot Variables 5-9
 - Powering On the VSM 5-10
 - Changing the HA Role 5-11
 - Recovering a Primary VSM with Active Secondary 5-12
 - Disconnecting the Port Groups 5-12
 - Powering Off the VSM 5-14
 - Connecting the Port Groups 5-14
- Problems with VSM-VEM Layer 2 to 3 Conversion Tool 5-16
- Upgrade Troubleshooting Commands 5-16

High Availability 6-1

- Information About High Availability 6-1
 - System-Level High Availability 6-2
 - Network-Level High Availability 6-2
- Problems with High Availability 6-2
- High Availability Troubleshooting Commands 6-5

VSM and VEM Modules 7-1

- Information About Modules 7-1
- Troubleshooting a Module Not Coming Up on the VSM 7-1
 - Guidelines for Troubleshooting Modules 7-2
 - Flow Chart for Troubleshooting Modules 7-3
- Problems with the VSM 7-4
 - Verifying the VSM Is Connected to the vCenter Server 7-6
 - Verifying the VSM Is Configured Correctly 7-7
 - Verifying the Domain Configuration 7-7
 - Verifying the System Port Profile Configuration 7-8
 - Verifying the Control and Packet VLAN Configuration 7-8
 - Checking the vCenter Server Configuration 7-10
 - Checking Network Connectivity Between the VSM and the VEM 7-10
 - Recovering Management and Control Connectivity of a Host when a VSM is Running on a VEM 7-12
 - Using the VEM Connect Script 7-12

Checking the VEM Configuration	7-14
Collecting Logs	7-17
VSM and VEM Troubleshooting Commands	7-18
Ports	8-1
Information About Ports	8-1
Information About Interface Characteristics	8-1
Information About Interface Counters	8-2
Information About Link Flapping	8-2
Information About Port Security	8-2
Port Diagnostic Checklist	8-3
Problems with Ports	8-3
Cannot Enable an Interface	8-4
Port Link Failure or Port Not Connected	8-4
Link Flapping	8-5
Port ErrDisabled	8-5
VM Cannot Ping a Secured Port	8-6
Port Security Violations	8-7
Port Troubleshooting Commands	8-8
Port Profiles	9-1
Information About Port Profiles	9-1
Problems with Port Profiles	9-2
Recovering a Quarantined Offline Interface	9-4
Port Profile Logs	9-5
Port Profile Troubleshooting Commands	9-5
Port Channels and Trunking	10-1
Overview	10-1
Port Channel Overview	10-1
Port Channel Restriction	10-2
Trunking Overview	10-2
Initial Troubleshooting Checklist	10-2
Troubleshooting Asymmetric Port Channels	10-3
Cannot Create Port Channel	10-4
Newly Added Interface Does Not Come Online In a Port Channel	10-4
Forcing Port Channel Characteristics onto an Interface	10-4
Verifying a Port Channel Configuration	10-5
VLAN Traffic Does Not Traverse Trunk	10-5

Layer 2 Switching	11-1
Information About Layer 2 Ethernet Switching	11-1
Port Model	11-1
Viewing Ports from the VEM	11-2
Viewing Ports from the VSM	11-3
Port Types	11-4
Layer 2 Switching Problems	11-4
Verifying a Connection Between VEM Ports	11-4
Verifying a Connection Between VEMs	11-5
Isolating Traffic Interruptions	11-6
Verifying Layer 2 Switching	11-7
Troubleshooting Microsoft NLB Unicast Mode	11-12
Limitations and Restrictions	11-12
Disabling Automatic Static MAC Learning on vEthernet	11-12
Checking Status on a VSM	11-12
Checking Status on a VEM	11-13
Configuring MS NLB for Multiple VM NICs in the Same Subnet	11-14
Enabling UUFB	11-14
Disabling UUFB for VMs that use Dynamic MAC Addresses	11-14
VLANs	12-1
Information About VLANs	12-1
Initial Troubleshooting Checklist	12-2
Cannot Create a VLAN	12-3
Private VLANs	13-1
Information About Private VLANs	13-1
Private VLAN Domain	13-1
Spanning Multiple Switches	13-1
Private VLAN Ports	13-2
Troubleshooting Guidelines	13-2
Private VLAN Troubleshooting Commands	13-2
NetFlow	14-1
Information About NetFlow	14-1
NetFlow Troubleshooting Commands	14-2
Common NetFlow Problems	14-3
Debugging a Policy Verification Error	14-3
Debugging Statistics Export	14-3

ACLs	15-1
About Access Control Lists (ACLs)	15-1
ACL Configuration Limits	15-1
ACL Restrictions	15-2
Troubleshooting ACLs	15-2
Displaying ACL Policies on the VEM	15-2
Debugging Policy Verification Issues	15-3
Troubleshooting ACL Logging	15-3
Using the CLI to Troubleshoot ACL Logging on a VEM	15-4
Viewing Current Flows	15-4
Viewing Active Flows	15-4
Flushing All ACL Flows	15-4
Showing Flow Debug Statistics	15-5
ACL Logging Troubleshooting Scenarios	15-5
Troubleshooting a Syslog Server Configuration	15-5
Troubleshooting an ACL Rule That Does Not Have a Log Keyword	15-6
Troubleshooting a Maximum Flow Limit Value That is Too Low	15-6
Troubleshooting a Mismatched Configuration between a VSM and a VEM	15-7
Quality of Service	16-1
Information About Quality of Service	16-1
QoS Configuration Limits	16-1
QoS Troubleshooting Commands	16-2
Troubleshooting the VEM	16-2
Debugging Policing Verification Errors	16-3
SPAN	17-1
Information About SPAN	17-1
SPAN Session Guidelines	17-1
Problems with SPAN	17-2
SPAN Troubleshooting Commands	17-3
Multicast IGMP	18-1
Information About Multicast	18-1
Multicast IGMP Snooping	18-1
Problems with Multicast IGMP Snooping	18-2
Troubleshooting Guidelines	18-2
Troubleshooting Commands	18-2
Symptoms, Causes, and Solutions	18-4

DHCP, DAI, and IPSG	19-1
Information About DHCP Snooping	19-1
Information About Dynamic ARP Inspection	19-2
Information About IP Source Guard	19-2
Guidelines and Limitations for Troubleshooting	19-2
Problems with DHCP Snooping	19-3
Troubleshooting Dropped ARP Responses	19-4
Problems with IP Source Guard	19-5
Collecting and Evaluating Logs	19-5
VSM Logging	19-5
Host Logging	19-6
DHCP, DAI, and IPSG Troubleshooting Commands	19-6
Virtual Service Domain	20-1
Information about Virtual Service Domain	20-1
Problems with Virtual Service Domain	20-1
Collecting and Evaluating Logs	20-2
Virtual Service Domain Troubleshooting Commands	20-3
System	21-1
Information About the System	21-1
General Restrictions for vCenter Server	21-2
Extension Key	21-2
Recovering a DVS	21-2
Recovering a DVS With a Saved Copy of the VSM	21-3
Recovering a DVS Without a Saved Copy of the VSM	21-4
Problems Related to VSM and vCenter Server Connectivity	21-5
Connection Failure After ESX Reboot	21-6
Setting the System MTU	21-7
Recovering Lost Connectivity Due to MTU Mismatch	21-8
VSM Creation	21-9
Port Profiles	21-9
Problems with Port Profiles	21-10
Problems with Hosts	21-10
Problems with VM Traffic	21-10
VEM Troubleshooting Commands	21-11
VEM Log Commands	21-12
Error Messages	21-12

Network Segmentation Manager 22-1

- Information About Network Segmentation Manager 22-1
- Problems with Network Segmentation Manager 22-2
- Network Segmentation Manager Troubleshooting Commands 22-7

VXLANs 23-1

- Information About VXLANs 23-1
 - Overview 23-1
 - VXLAN Tunnel EndPoint 23-2
 - VXLAN Gateway 23-2
 - VXLAN Trunks 23-2
 - Multi-MAC Capability 23-3
 - Fragmentation 23-3
 - Scalability 23-3
 - Maximum Number of VXLANs 23-3
 - Supported Features 23-3
 - Jumbo Frames 23-3
 - Disabling the VXLAN Feature Globally 23-4
- VXLAN Troubleshooting Commands 23-4
 - VSM Commands 23-4
 - VXLAN Gateway Commands 23-5
 - VEM Commands 23-7
 - VEM Packet Path Debugging 23-9
 - VEM Multicast Debugging 23-10
 - VXLAN Datapath Debugging 23-11
 - Vemlog Debugging 23-11
 - Vempkt 23-12
 - Statistics 23-12
 - Show Commands 23-13

VSI Discovery and Configuration Protocol 24-1

- Information About VDP 24-1
- Problems with VDP 24-2
- VDP Troubleshooting Commands 24-2
 - VSM Commands 24-2
 - VEM Commands 24-4

Cisco TrustSec 25-1

- Information About Cisco TrustSec 25-1

Guidelines and Limitations for Troubleshooting Cisco TrustSec	25-1
Cisco TrustSec Troubleshooting Commands	25-2
Debugging Commands	25-2
Host Logging Commands	25-3
Example	25-3
Show Commands	25-4
Problems with Cisco TrustSec	25-4
vCenter Plug-in	26-1
vCenter Plug-in Overview	26-1
Requirements for VMware vSphere Web Client	26-1
Generating a Log Bundle	26-2
Ethalyzer	27-1
Using Ethalyzer	27-1
Before Contacting Technical Support	28-1
Cisco Support Communities	28-1
Gathering Information for Technical Support	28-1
Obtaining a File of Core Memory Information	28-2
Copying Files	28-3



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](#).

Feature	Description	Changed in release	Where Documented
VSI Discovery and Configuration Protocol	Added new section for troubleshooting commands for VSI Discovery and Configuration Protocol (VDP).	4.2(1)SV2(2.2)	“VSI Discovery and Configuration Protocol”
VXLAN Gateway	Added section for troubleshooting commands for VXLAN Gateway.	4.2(1)SV2(2.1)	“VXLANS”
Upgrade	Added section for problems with VSM-VEM Layer 2 to 3 Conversion Tool.	4.2(1)SV2(1.1)	Upgrade
Ethalyzer	Added Ethalyzer as a Nexus 1000V protocol analyzer tool content.	4.2(1)SV2(1.1)	Ethalyzer
DHCP Enhancements	Added the troubleshooting commands for DHCP.	4.2(1)SV2(1.1)	DHCP, DAI, and IPSG Troubleshooting Commands
High Availability	Updated the high availability section. Added command output for the new command show system internal active-active remote accounting logs and updated the output for the show system redundancy status command.	4.2(1)SV2(1.1)	High Availability
Licensing	Added the svs license transfer src-vem <vem no> license_pool command to troubleshoot the issues with checking out the licenses or returning them to the license pool.	4.2(1)SV2(1.1)	License Troubleshooting Commands
Nexus 1000V VC Plugin Installation	Added new section to troubleshoot the Nexus 1000V VC Plugin Installation.	4.2(1)SV2(1.1)	vCenter Plug-in

Feature	Description	Changed in release	Where Documented
Nexus 1000V Installation Management Center	Added new section to troubleshoot the Nexus 1000V Installation Management Center.	4.2(1)SV1(5.1)	“Problems with the Nexus 1000V Installation Management Center”
Recovering Management and Control Connectivity of a Host	Added new section to recover management and control connectivity of a host when a VSM is running on a VEM	4.2(1)SV1(5.1)	“Recovering Management and Control Connectivity of a Host when a VSM is Running on a VEM”
ACL Logging	Added new section to troubleshoot ACL Logging.	4.2(1)SV1(5.1)	“Troubleshooting ACL Logging”
NSM	Added new chapter to troubleshoot the Network Segmentation Manager (NSM).	4.2(1)SV1(5.1)	“Network Segmentation Manager”
VXLAN	Added new chapter to troubleshoot the Virtual Extensible Local Area Network (VXLAN).	4.2(1)SV1(5.1)	“VXLANs”
Microsoft NLBUncast Mode	Added new section for troubleshooting Microsoft Network Load Balancing (NLB) Unicast mode	4.2(1)SV1(5.1)	“Layer 2 Switching”
In service software upgrade (ISSU)	Added new section for troubleshooting ISSU.	4.2(1)SV1(4a)	“Upgrade”
VEM software upgrade	Added new section for troubleshooting VEM software upgrade.	4.2(1)SV1(4a)	“Upgrade”
DHCP, DAI, IPSG	Added new section for troubleshooting DHCP, Dynamic ARP Inspection, and IP Source Guard.	4.2(1)SV1(4)	“DHCP, DAI, and IPSG”
Virtual Service Domain	Added new section for troubleshooting Virtual Service Domain.	4.2(1)SV1(4)	“Virtual Service Domain”
Port profiles	Added section for port profiles, new information about quarantined port profiles.	4.2(1)SV1(4)	“Port Profiles”
Upgrade	Added new section for troubleshooting upgrade problems.	4.2(1)SV1(4)	“Upgrade”
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)	“Checking Network Connectivity Between the VSM and the VEM”
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”



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 xv](#)
- [Document Conventions, page xv](#)
- [Related Documentation, page xvi](#)
- [Obtaining Documentation and Submitting a Service Request, page xviii](#)

Audience

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

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:

screen font	Terminal sessions and information that the switch displays are in screen font.
boldface screen font	Information that you must enter is in boldface screen font.

<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

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



Caution

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

Related Documentation

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

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

General Information

Cisco Nexus 1000V Documentation Roadmap

Cisco Nexus 1000V Release Notes

Cisco Nexus 1000V Compatibility Information

Install and Upgrade

Cisco Nexus 1000V Installation and Upgrade Guide

Configuration Guides

Cisco Nexus 1000V High Availability and Redundancy Configuration Guide

Cisco Nexus 1000V Interface Configuration Guide

Cisco Nexus 1000V Layer 2 Switching Configuration Guide

Cisco Nexus 1000V License Configuration Guide

Cisco Nexus 1000V Network Segmentation Manager Configuration Guide

Cisco Nexus 1000V Port Profile Configuration Guide

Cisco Nexus 1000V Quality of Service Configuration Guide

Cisco Nexus 1000V REST API Plug-in Configuration Guide

Cisco Nexus 1000V Security Configuration Guide

Cisco Nexus 1000V System Management Configuration Guide
Cisco Nexus 1000V vCenter Plugin Configuration Guide
Cisco Nexus 1000V VXLAN Configuration Guide
Cisco Nexus 1000V VDP Configuration Guide
Cisco Nexus 1000V DFA Configuration Guide
Cisco Nexus 1000V vCenter Plugin Configuration Guide

Programming Guide

Cisco Nexus 1000V XML API User Guide

Reference Guides

Cisco Nexus 1000V Command Reference
Cisco Nexus 1000V Resource Availability Reference

Troubleshooting, Password Recovery, System Messages Guides

Cisco Nexus 1000V Troubleshooting Guide
Cisco Nexus 1000V Password Recovery Guide
Cisco NX-OS System Messages Reference

Virtual Services Appliance Documentation

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

Virtual Security Gateway Documentation

The Cisco Virtual Security Gateway documentation is available at
http://www.cisco.com/en/US/products/ps13095/tsd_products_support_series_home.html

Virtual Network Management Center

The Cisco Virtual Network Management Center documentation is available at
http://www.cisco.com/en/US/products/ps11213/tsd_products_support_series_home.html

Virtual Wide Area Application Services (vWAAS)

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

ASA 1000V Cloud Firewall

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

Obtaining Documentation and Submitting a Service Request

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

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



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](#)
- [System Messages, page 1-4](#)
- [Troubleshooting with Logs, page 1-6](#)
- [Cisco Support Communities, page 1-7](#)
- [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.

- 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.

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?

- 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].

```
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
```

```
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
```



```

loopback      Show logging loopback configuration
module        Show module logging configuration
monitor       Show monitor logging configuration
nvram         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

```

Cisco Support Communities

For additional information, visit one of the following support communities:

- [Cisco Support Community for Server Networking](#)
- [Cisco Communities: Nexus 1000V](#)

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 Nexus 1000V 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 Cisco Nexus 1000V and support contract from Cisco, contact Cisco for Nexus 1000V support. Cisco provides L1, L2, and L3 support.

If you purchased the Cisco Nexus 1000V and an SNS through VMware, you should call VMware for Nexus 1000V 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 -xviii.

For more information on steps to take before calling Technical Support, see the [“Gathering Information for Technical Support”](#) section on page 28-1.



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 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.

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.

- 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

```

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

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

- 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*.



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](#)
- [Problems with the Nexus 1000V Installation Management Center, page 3-13](#)

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.


- 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](#).

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
```

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

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 Nexus 1000V 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>>

- 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](#)

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

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#
```

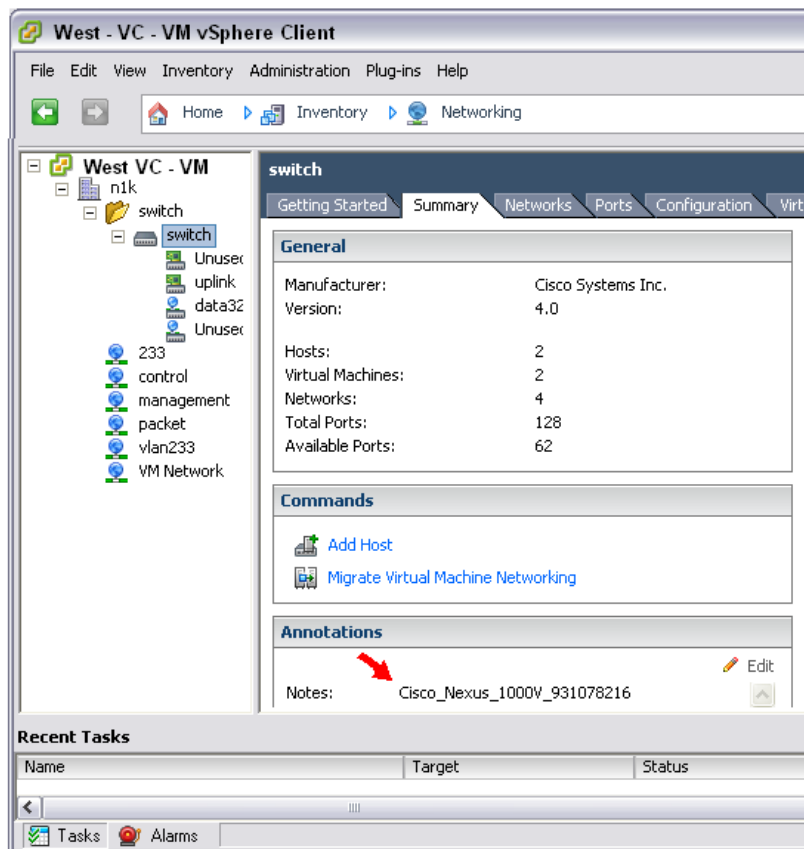
Finding the Extension Key Tied to a Specific DVS

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.



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

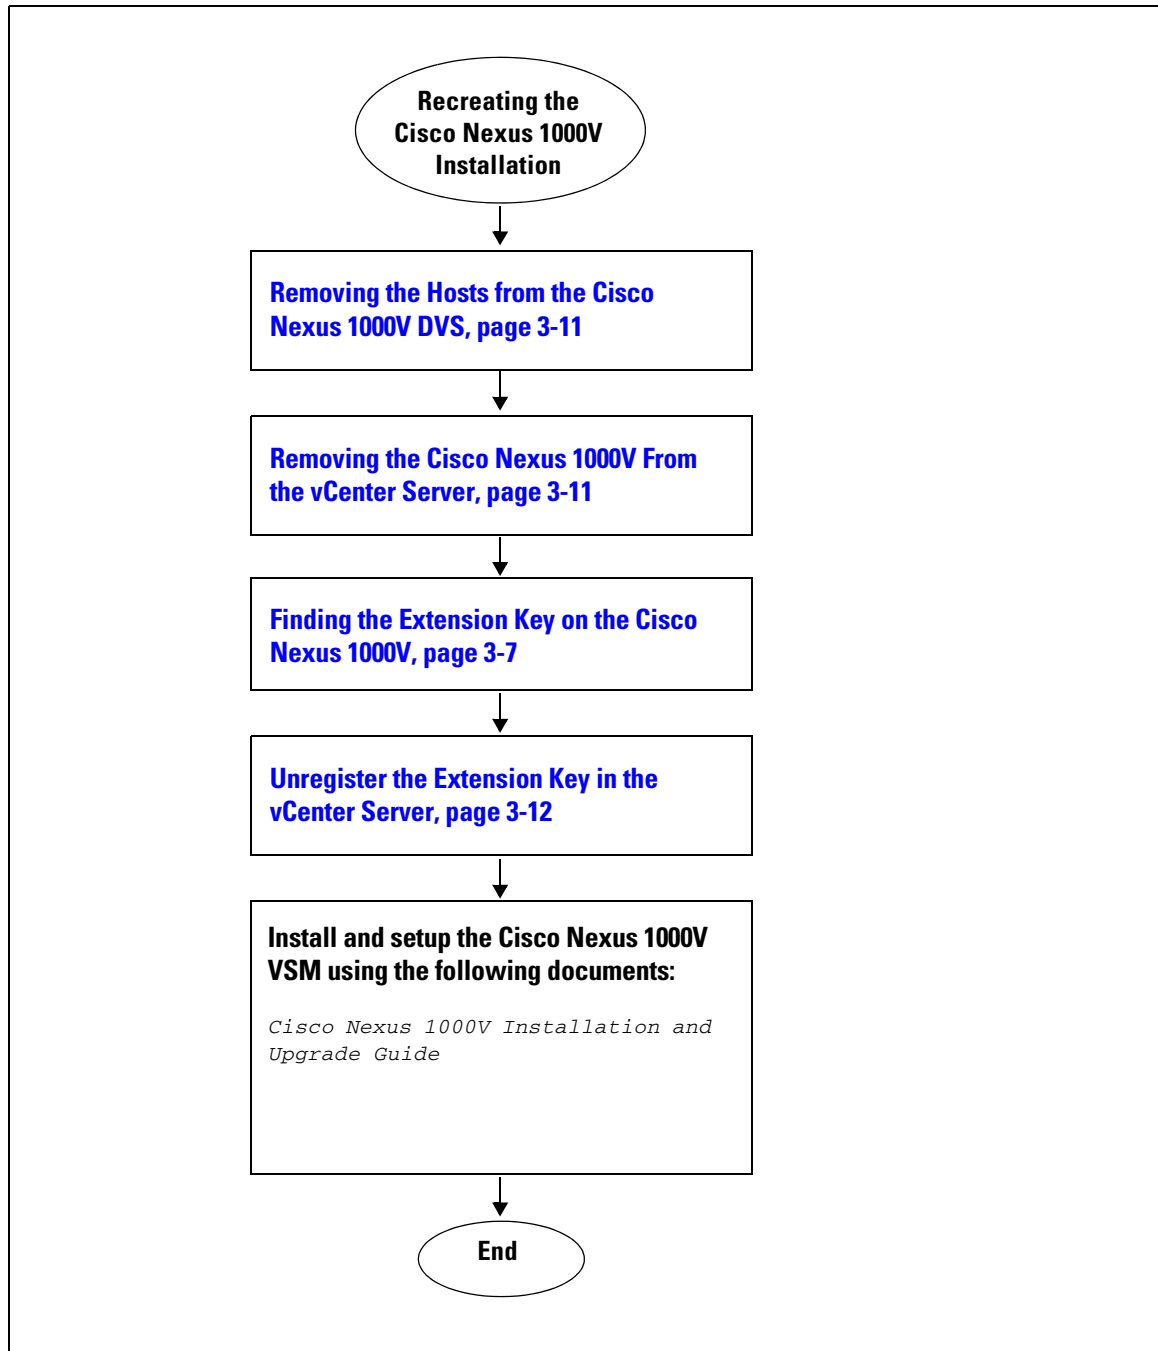
-
- | | |
|---------------|---|
| Step 1 | Find the extension key used on the Cisco Nexus 1000V using the Finding the Extension Key on the Cisco Nexus 1000V, page 3-7 . |
| Step 2 | Find the extension key used on the vCenter Server using the Finding the Extension Key Tied to a Specific DVS, page 3-8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

Resolving Extension Key Conflicts

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



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](#).
-

Unregister the Extension Key in the vCenter Server

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

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"]

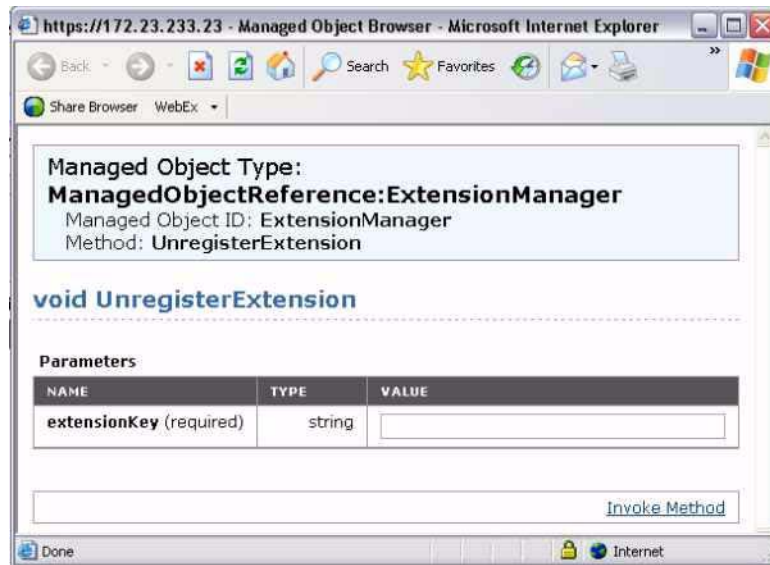
Methods

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>

A dialog box opens to unregister 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.

- Step 4** You have completed this procedure.

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

Problems with the Nexus 1000V Installation Management Center

The following are possible problems and their solutions.

Symptom	Problem	Recommended Action
Port migration fails.	VSM to VEM migration fails in Layer 2 / Layer 3 mode installation.	<ul style="list-style-type: none"> • Check if there is any VM running on the vSwitch. You need to power off all such VMs before migration. • Check if the vCenter is VUM enabled. Before migration, the host is added to the DVS by using VUM. • Verify that the native VLAN in the upstream switch configuration is correct. • Ensure that the VUM repositories are up to date and accurate.
The VEM is missing on the VSM after the migration.	<ul style="list-style-type: none"> • The installer application finishes successfully with port migration in Layer 3 mode. • The VEM is added to the vCenter but does not display when the show module command is entered on the VSM. 	<ul style="list-style-type: none"> • Verify that the Layer 3 control profile VLAN is configured as a system VLAN. • Verify that the uplink profile is allowing the Layer 3 control vmknic VLAN and that it is a system VLAN. • From the ESX host (VEM), issue a vmkping to the mgmt0/control0 IP address. It should be successful. If not, check the intermediate switches for proper routes between the subnets. • The vmknic should be pingable from the VSM. • Check the vCenter mob for opaque data propagation.
Configuration file issue.	After loading the previously saved configuration file, the installation application does not complete.	<ul style="list-style-type: none"> • Check the configuration file for appropriate contents. <p>Note You might need to change a few of the fields before reusing the previously saved files.</p> <ul style="list-style-type: none"> • Check if a VM with the same name already exists in the DC. <p>This can be identified by reviewing the Virtual Machine field in the configuration file.</p>



Licenses

This chapter describes how to identify and resolve problems related to licenses, and includes the following sections.

- [Information About Licenses, page 4-1](#)
- [Prerequisites to License Troubleshooting, page 4-2](#)
- [Problems with Licenses, page 4-3](#)
- [License Troubleshooting Commands, page 4-4](#)

Information About Licenses

The name for the Cisco Nexus 1000V license package is NEXUS1000V_LAN_SERVICES_PKG. By default, 16 licenses are installed with the VSM. These default licenses are valid for 60 days. You can purchase permanent licenses that do not expire.

Licensing is based on the number of CPU sockets on the ESX servers attached as VEMs to the VSM.

A module is either licensed or unlicensed.

- **Licensed module**—A VEM is licensed if it acquires licenses for all of its CPU sockets from the pool of available licenses installed on the VSM.
- **Unlicensed module**—A VEM is unlicensed if it does not acquire licenses for all of its CPU sockets from the pool of available licenses installed on the VSM.

If a VEM is unlicensed, the virtual Ethernet ports corresponding to the virtual machines (VMs) are kept down, and are shown as unlicensed.



Note

The server administrator has no information about VEM licenses. The VEM licensed state must be communicated to server administrators so they are aware that vEthernet interfaces on unlicensed modules cannot pass traffic.

For additional information about licensing, including how to purchase or install a license, or how to remove an installed license, see the *Cisco Nexus 1000V License Configuration Guide*.

Contents of the License File

The contents of the Cisco Nexus 1000V license file indicates the number of licenses purchased and the host-ID. To display the contents of a license file, use the **show license file** *license_name* command.

```
n1000v# show license file sample.lic
sample.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT NEXUS1000V_LAN_SERVICES_PKG cisco 1.0 permanent 16 \
  HOSTID=VDH=8449368321243879080 \
  NOTICE="<LicFileID>sample.lic</LicFileID><LicLineID>0</LicLineID> \
  <PAK>dummyPak</PAK>" SIGN=34FCB2B24AE8
```

The host-ID that appears in the license file must match that shown on the VSM. To verify this, use the **show license host-id** command. See [Example 4-6 on page 4-6](#).



Caution

Do not edit the contents of the license file. The license is invalidated if its contents are altered. If you have already done so, please contact your Cisco Customer Support Account Team.

Prerequisites to License Troubleshooting

Before you begin troubleshooting licenses, verify the information in this checklist:

- Make sure the name of the license file is less than 32 characters.
Use the **show license usage** command. See [Example 4-3 on page 4-5](#).
- 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.
Use the **show license usage** command. See [Example 4-3 on page 4-5](#).
- 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 the host-ID on the switch, using the following commands:
show license host-id See [Example 4-6 on page 4-6](#).
show license file See [Example 4-7 on page 4-6](#)

Problems with Licenses

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

Symptom	Possible Causes	Solution
<p>When you power on a virtual machine with ports on a Cisco Nexus 1000V port group, the interfaces do not come up, but display the following status:</p> <pre>VEM Unlicensed</pre>	<p>Not enough licenses were obtained to license the CPU sockets of all VEMs connected to the VSM.</p>	<ol style="list-style-type: none"> 1. Verify license usage. show license usage <i>license_name</i> See Example 4-4 on page 4-5 2. Determine the number of licenses required by viewing the sockets installed on the VEM. show module vem license-info See Example 4-2 on page 4-5 3. Contact your Cisco Customer Support Account Team to acquire additional licenses.
<p>You see the following system message:</p> <pre>PLATFORM-2-PFM_LIC_WARN_EXP Syslog 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.</pre>	<p>The default or evaluation license in use is about to expire.</p> <p>Note Permanent licenses do not expire.</p>	<ol style="list-style-type: none"> 1. Verify license usage. show license usage <i>license_name</i> See Example 4-4 on page 4-5 2. Contact your Cisco Customer Support Account Team to acquire additional licenses.
<p>You see the following system message:</p> <pre>%LICMGR-2-LOG_LIC_USAGE: Feature NEXUS1000V_LAN_SERVICES_PKG is using 17 licenses, only 16 licenses are installed.</pre>	<p>More licenses are being used than are installed.</p>	<ol style="list-style-type: none"> 1. Verify license usage. show license usage <i>license_name</i> See Example 4-4 on page 4-5 2. Contact your Cisco Customer Support Account Team to acquire additional licenses.

License Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to licenses.

Command	Purpose
<code>show module</code>	Displays display module information including license status (unlicensed or active). See Example 4-1 on page 4-4 .
<code>show module vem license info</code>	Displays the VEM license information including license status, license version, socket count. See Example 4-2 on page 4-5 .
<code>show license usage [license_name]</code>	Displays information about licenses and where they are used. If displayed for a specific license, indicates VEM and socket information. See Example 4-3 on page 4-5 . See Example 4-4 on page 4-5 .
<code>show interface veth</code>	Displays the messages logged about port profile events within the Cisco Nexus 1000V. See Example 4-5 on page 4-5 .
<code>show license host-id</code>	Displays the serial number for your Cisco Nexus 1000V license. See Example 4-6 on page 4-6 .
<code>show license file</code>	Displays the contents of a named license file. See Example 4-7 on page 4-6 .
<code>show license brief</code>	Displays a list of list of license files installed on the VSM. See Example 4-8 on page 4-6 .
<code>svs license transfer src-vem <vem no> license_pool</code>	Transfers the licenses from a VEM to the license pool. See Example 4-9 on page 4-6 .

For detailed information about show command output, see the *Cisco Nexus 1000V Command Reference*.

EXAMPLES

Example 4-1 show module

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

```

Example 4-2 show module vem license-info

```

n1000v# show module vem license-info
Licenses are Sticky
Mod  Socket Count  License Usage Count  License Version  License Status
---  -
3    2              -                    -                unlicensed
n1000v#

```

Example 4-3 show license usage

```

n1000v# show license usage
Feature                               Ins Lic  Status  Expiry Date  Comments
                               Count
---  -
NEXUS_VSG_SERVICES_PKG                No    0  Unused  -            -
NEXUS1000V_LAN_SERVICES_PKG           Yes   16  In use  Never       -
n1000v#

```

Example 4-4 show license usage <license_name>

```

n1000v# show license usage NEXUS1000V_LAN_SERVICES_PKG
-----
Feature Usage Info
-----
      Installed Licenses :    10
      Eval Licenses :      0
      Max Overdraft Licenses :    16
      Installed Licenses in Use :    4
      Overdraft Licenses in Use :    0
      Eval Licenses in Use :      0
      Licenses Available :    22
-----
Application
-----
VEM 3 - Socket 1
VEM 3 - Socket 2
VEM 4 - Socket 1
VEM 4 - Socket 2
-----
n1000v#

```

Example 4-5 show interface vethernet

```

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

```

Example 4-6 show license host-id

```

n1000v# show license host-id
License hostid: VDH=8449368321243879080
n1000v#

```

Example 4-7 show license file

```

n1000v# show license file sample.lic
sample.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT NEXUS1000V_LAN_SERVICES_PKG cisco 1.0 permanent 16 \
    HOSTID=VDH=8449368321243879080 \
    NOTICE="<LicFileID>sample.lic</LicFileID><LicLineID>0</LicLineID> \
    <PAK>dummyPak</PAK>" SIGN=34FCB2B24AE8

n1000v#

```

Example 4-8 show license brief

```

n1000v# show license brief
license_file.lic
n1000v#

```

Example 4-9 svcs license transfer src-vem <vem no> license_pool

```

n1000v# svcs license transfer src-vem 3 license_pool
n1000v#

```



Upgrade

This chapter describes how to identify and resolve problems related to upgrading the VSM software, and includes the following sections:

- [Information about Upgrades, page 5-1](#)
- [Problems with the In Service Software Upgrade, page 5-1](#)
- [Problems with the VEM Upgrade, page 5-5](#)
- [Problems with the GUI Upgrade, page 5-6](#)
- [Problems with VSM-VEM Layer 2 to 3 Conversion Tool, page 5-16](#)
- [Upgrade Troubleshooting Commands, page 5-16](#)

Information about Upgrades

The upgrade for the Cisco Nexus 1000V involves upgrading software on both the VSM and the VEM. An in service software upgrade (ISSU) is available for a stateful upgrade of the Cisco Nexus 1000V image(s) running on the VSM. A stateful upgrade is one without noticeable interruption of data plane services provided by the switch.

For detailed information, see the following document:

- *Cisco Nexus 1000V Installation and Upgrade Guide*

Problems with the In Service Software Upgrade

The following are symptoms, possible causes, and solutions for problems with software upgrade using the manual in service software upgrade (ISSU).

Table 5-1 **Problems with the ISSU**

Symptom	Possible Causes	Solution
Error Message: Pre-Upgrade check failed. Return code 0x40930062 (free space in the filesystem is below threshold).	This error indicates that there is not enough space in the /var/sysmgr partition.	1. Reboot the system.

Table 5-1 Problems with the ISSU (continued)


Symptom	Possible Causes	Solution
<p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x4093000A (SRG collection failed)</p>	A module is removed during the upgrade.	<ol style="list-style-type: none"> 1. Make sure the module removal is complete. 2. Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x40930076 (Standby sup is offline. ISSU will not proceed)</p>	The standby VSM is not present or is not synchronized with the active VSM, and the VSMs do not form a stable HA pair.	<ol style="list-style-type: none"> 1. Verify the HA synchronization state. show system redundancy status The output of the show command must indicate the following: Active VSM: Active with HA standby Standby VSM: HA standby 2. If the output of the show command indicates that the VSMs are not synchronized, then see the “Problems with High Availability” section on page 6-2. 3. When the VSMs are synchronized, restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x807B0002 (No such file or directory)</p> <p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x4093000F (Failed to copy image)</p>	<p>The software image files required for the upgrade are not present or were not copied to the bootflash: repository.</p> <p>There may not be enough room in bootflash: for the files to be copied.</p>	<ol style="list-style-type: none"> 1. Verify there is enough space in bootflash for the image files. dir 2. Do one of the following: <ul style="list-style-type: none"> – If additional space is needed, delete other files from the bootflash repository to make room for the software image files. delete <p> Caution Do not delete kickstart or system image files from bootflash. If there are no image files in bootflash, the system cannot reboot if required.</p> <ul style="list-style-type: none"> – If not, continue with the next step. <ol style="list-style-type: none"> 3. Download the required images from www.cisco.com to the bootflash: repository. 4. Verify that the correct images are in the bootflash: repository. show boot 5. When the correct software images are in the bootflash: repository, restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.

Table 5-1 Problems with the ISSU (continued)

Symptom	Possible Causes	Solution
<p>The install command fails with the following error:</p> <pre>Return code 0x4045001F (image MD5 checksum error) Pre-Upgrade check failed. Return code 0x40930011 (Image verification failed)</pre>	<p>The software image file(s) required for the upgrade do not pass the MD5 checksum verification, indicating that the correct file(s) are not present in bootflash: for the upgrade to proceed.</p> <p>A file can be truncated when copied.</p>	<ol style="list-style-type: none"> Using the README file from the upgrade zip folder at www.cisco.com, verify the MD5 checksum for each of the image files. show file bootflash: filename md5sum Replace the file(s) that do not match. Verify that the correct images are in the bootflash: repository and that checksums match. show file bootflash: filename md5sum When the correct software images are in the bootflash: repository, restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>Error message:</p> <pre>Install has failed. Return code 0x40970001 (Incompatible image)</pre>	<p>You may have used an incorrect filename when entering the install all command.</p>	<p>Restart the software upgrade using the correct filenames for the new software images.</p> <pre>install all kickstart filename1 system filename2</pre>
<p>After upgrading, the VSMs are not running the new software version.</p>	<p>The boot variables were not set properly.</p>	<ol style="list-style-type: none"> Verify that the running images and boot variables match the upgrade version. show version show boot If needed, download the required images from www.cisco.com to your local bootflash: repository. Verify that the correct images are in the bootflash: repository. show boot Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>. If the problem persists, collect details of the upgrade and open a support case. show system internal log install details
<p>Performing the configuration copy process fails and stops the upgrade.</p> <pre>Performing configuration copy. [####-----] 30%</pre>	<p>Service or system errors.</p>	<ol style="list-style-type: none"> Manually copy the configuration. copy running-config startup-config Do one of the following: <ul style="list-style-type: none"> If the progress bar gets stuck before 100% for over one minute, collect details of the upgrade and open a support case. show system internal log install details If the copy succeeds without delays, restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.

Table 5-1 Problems with the ISSU (continued)

Symptom	Possible Causes	Solution
<p>Error message:</p> <p>Another install procedure may be in progress. (0x401E0007)</p>	<p>Another upgrade session is in progress from a VSM console or SSH/Telnet.</p>	<p>Do one of the following:</p> <ul style="list-style-type: none"> Continue the first upgrade session in progress. Stop the upgrade and restart one session only using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>Install command fails with following error message:</p> <pre>-- FAIL. Return code 0x4093001E (Standby failed to come online) Install has failed. Return code 0x4093001E (Standby failed to come online)</pre>	<p>The standby VSM fails to boot with the new image.</p>	<p>Do one of the following:</p> <ul style="list-style-type: none"> Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>. Postpone the upgrade and reset the boot variables to the original filenames. <p>boot kickstart filename [sup-1] [sup-2]</p>
<p>Install command fails with following error message:</p> <pre>Install has failed. Return code 0x4093001F (Standby installer failed to take over the installation). Please identify the cause of the failure, and try "install all" again"</pre>	<p>The standby VSM takes more than 10 minutes to come up and form a stable HA pair with the active VSM.</p>	<ol style="list-style-type: none"> Reset the boot variables to the original filenames. <p>boot kickstart filename [sup-1] [sup-2]</p> <ol style="list-style-type: none"> If the standby is still running the new software version, reload it. <p>reload</p> <p>The standby synchronizes with the active, so that both are running the original software version.</p>
<p>Install command fails with following error message:</p> <pre>Module 2: Waiting for module online. -- SUCCESS -- Install has failed. Return code 0x40930000 (Current operation failed to complete within specified time)</pre>	<p>A failure at the standby VSM caused it to reload again after the Continuing with installation, please wait message and before the switchover.</p>	<ol style="list-style-type: none"> Inspect logs. <p>show logging</p> <ol style="list-style-type: none"> Look for standby reloads caused by process failures. <p>show cores</p> <p>If a process crash is observed, collect details of the upgrade and open a support case.</p> <p>show system internal log install details</p> <ol style="list-style-type: none"> Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>Pre-Upgrade check failed:</p> <p>Return code 0x40930062 (free space in the filesystem is below threshold).</p>		<ol style="list-style-type: none">

Problems with the VEM Upgrade

The following are symptoms, possible causes, and solutions for problems with VEM software upgrade.

Table 5-2 Problems with the VEM Upgrade

Symptom	Possible Causes	Solution
After starting a VEM upgrade from the VSM console, VUM skips upgrading the hosts with the new VEM.	<p>One or more of the following are enabled on the host cluster.</p> <ul style="list-style-type: none"> VMware High Availability (HA) VMware Fault Tolerance (FT) VMware Distributed Power Management (DPM) 	<ol style="list-style-type: none"> Verify the upgrade failure. show vmware vem upgrade status From vCenter server, disable HA, FT, and DPM for the cluster. Restart the VEM software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
VEM upgrade fails.	An incorrect VUM version is in use.	<ol style="list-style-type: none"> Identify the VUM version required for the upgrade using the <i>Cisco Nexus 1000V Compatibility Information</i>. Upgrade to the correct VUM version. Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
After upgrading, the host is not added to the VSM.	An incorrect VEM software version is installed on the host.	<ol style="list-style-type: none"> Identify the VEM software version required for the upgrade using the <i>Cisco Nexus 1000V Compatibility Information</i>. Proceed with the upgrade using the correct VEM software version and the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
A message on the ESX/ESXi command line shell and VMkernel logs notifies you that the loading and unloading of modules failed.	The modules were not placed in maintenance mode (all VMs vmotioned over) before starting the upgrade.	<ol style="list-style-type: none"> Place the host in maintenance mode. Proceed with the upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
	<p>The host does not have enough memory to load new modules.</p> <p>A host requires a minimum of 2 GB of physical RAM. If it also hosts a Cisco Nexus 1000V VSM VM, it needs a minimum of 4 GB of physical RAM. If it also hosts the vCenter Server VM, additional memory may be needed.</p>	<ol style="list-style-type: none"> Verify that the host has sufficient memory to load the new modules. For more information about allocating RAM and CPU, see the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>. Proceed with the upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.

Problems with the GUI Upgrade

The following are symptoms, possible causes, and solutions for problems with software upgrade using the GUI upgrade application.


Note

If you are upgrading directly from SV1(4) to SV1(4a) the GUI is not used and this section does not apply. This section is only applicable if you use the GUI for an intermediate upgrade from a SV1(3x) release to SV1(4), prior to upgrading to SV1(4a).

Table 5-3 *Problems with the GUI Upgrade*

Symptom	Possible Causes	Solution
<p>The upgrade GUI stops and times out after 10 minutes and displays the following message:</p> <p>Error: Could not contact the upgraded VSM at <i>n.n.n.n</i>. Please check the connection.</p>	<p>During the upgrade, you configured an unreachable IP address for the mgmt0 interface.</p> <p>In this case, one VSM in the redundant pair has new software installed and is unreachable; while the other VSM has the original pre-upgrade original pre-upgrade software version installed and is reachable.</p>	<ol style="list-style-type: none"> Use one of the following sets of procedures to return your VSM pair to the previous software version: <ul style="list-style-type: none"> “Recovering a Secondary VSM with Active Primary” section on page 5-7 “Recovering a Primary VSM with Active Secondary” section on page 5-12 Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.
<p>The upgrade GUI stops and times out after 10 minutes and displays the following message:</p> <p>Error: Could not contact the upgraded VSM at 10.104.244.150. Please check the connection.</p> <p>After timing out, one VSM comes up in switch(boot) mode.</p>	<p>You have selected incompatible or incorrect VSM software images for the upgrade.</p> <p>The software images you selected from the GUI selection list included a system image for one software version and a kickstart image for another software version. These images must be for the same software version.</p> <p>For an example of how software images are selected during the upgrade, see Example 5-1.</p>	<ol style="list-style-type: none"> To continue the upgrade, first recover the VSM using one of the following: <ul style="list-style-type: none"> “Recovering a Secondary VSM with Active Primary” section on page 5-7 “Recovering a Primary VSM with Active Secondary” section on page 5-12 Restart the software upgrade using the detailed instructions in the <i>Cisco Nexus 1000V Installation and Upgrade Guide</i>.

Example 5-1 Upgrade: Enter Upgrade Information

This example shows how you specify system and kickstart images during the upgrade process. In this example, the images specified are from the same release, SV1.4. If you specify a kickstart image from one release, and a system image from another, then the upgrade cannot proceed.

Recovering a Secondary VSM with Active Primary

You can use the following process to recover a secondary VSM when the primary VSM is active.

**Note**

The information in this section does not apply when upgrading from Release 4.2(1)SV1(4) to Release 4.2(1)SV2(1.1).

- Step 1** Stop the upgrade on the VSM, using the [“Stopping a VSM Upgrade” procedure on page 5-8](#)
- Step 2** Change the boot variables back to the previous version using the [“Changing Boot Variables” procedure on page 5-9](#)
- Step 3** From the vCenter Server left-hand panel, right-click the secondary VSM and then choose Delete from Disk.
The secondary VSM is deleted.
- Step 4** Create a new VSM by reinstalling the software using the vSphere client Deploy OVF Template wizard, specifying the following:
- Nexus 1000V Secondary configuration method (Configures the secondary VSM in an HA pair using a GUI setup dialog.)
 - The host or cluster of the primary VSM
 - The same domain ID and password as that of the primary VSM.

For a detailed procedure, see the *Cisco Nexus 1000V Installation and Upgrade Guide*.

The VSM comes up and forms an HA pair with the newly-created standalone VSM. The VSMs have the previous version of the software installed.

Stopping a VSM Upgrade

You can use this procedure to stop a VSM upgrade that is in progress.

BEFORE YOU BEGIN

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



Note

The information in this section does not apply when upgrading from Release 4.2(1)SV1(4) to Release 4.2(1)SV2(1.1).

DETAILED STEPS

Step 1 Display upgrade status.

show svcs upgrade status

Example:

```
n1000v# show svcs upgrade status
Upgrade State: Start
Upgrade mgmt0 ipv4 addr: 1.1.1.1
Upgrade mgmt0 ipv6 addr:
Upgrade control0 ipv4 addr:
```

Step 2 Stop the upgrade.

configure terminal

no svcs upgrade start

Example:

```
n1000v# configure terminal
n1000v#(config)# no svcs upgrade start
WARNING! VSM upgrade process is aborted
n1000v#(config)#
```

Step 3 Display upgrade status.

show svcs upgrade status

Example:

```
n1000v#(config)# show svcs upgrade status
Upgrade State: Abort
Upgrade mgmt0 ipv4 addr:
Upgrade mgmt0 ipv6 addr:
Upgrade control0 ipv4 addr:
```

Step 4 You have completed this procedure. Return to the process that pointed you here:

- [“Recovering a Secondary VSM with Active Primary” section on page 5-7](#)
- [“Recovering a Primary VSM with Active Secondary” section on page 5-12](#)

Changing Boot Variables

You can use this procedure to replace the software images used to boot the VSM.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You know the filenames of the pre-upgrade system and kickstart image files to apply.

DETAILED STEPS

Step 1 Display the current boot variables.

show boot

Example:

```
n1000v# show boot
sup-1
kickstart variable = bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin
system variable = bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin
sup-2
kickstart variable = bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4.bin
system variable = bootflash:/nexus-1000v-mzg.4.2.1.SV1.4.bin
No module boot variable set
n1000v(config)#
```

Step 2 Remove the current system and kickstart boot variables.

configure terminal

no boot system

no boot kickstart

Example:

```
n1000v# configure terminal
n1000v(config)# no boot system
n1000v(config)# no boot kickstart
n1000v#(config)#
```

Step 3 Restore the system and kickstart boot variables to the original pre-upgrade filenames.

boot system bootflash:*system-boot-variable-name*

boot system bootflash:*kickstart-boot-variable-name*

Example:

```
n1000v#(config)# boot system bootflash:nexus-1000v-mz.4.0.4.SV1.3a.bin
n1000v#(config)# boot kickstart bootflash:nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin
n1000v#(config)#
```

Step 4 Copy the running configuration to the startup configuration.

copy run start

Example:

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

Step 5 Verify the change in the system and kickstart boot variables.

show boot

Example:

```
n1000v#(config)# show boot
sup-1
kickstart variable = bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin
system variable = bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin
sup-2
kickstart variable = bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin
system variable = bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin
No module boot variable set
n1000v#(config)#
```

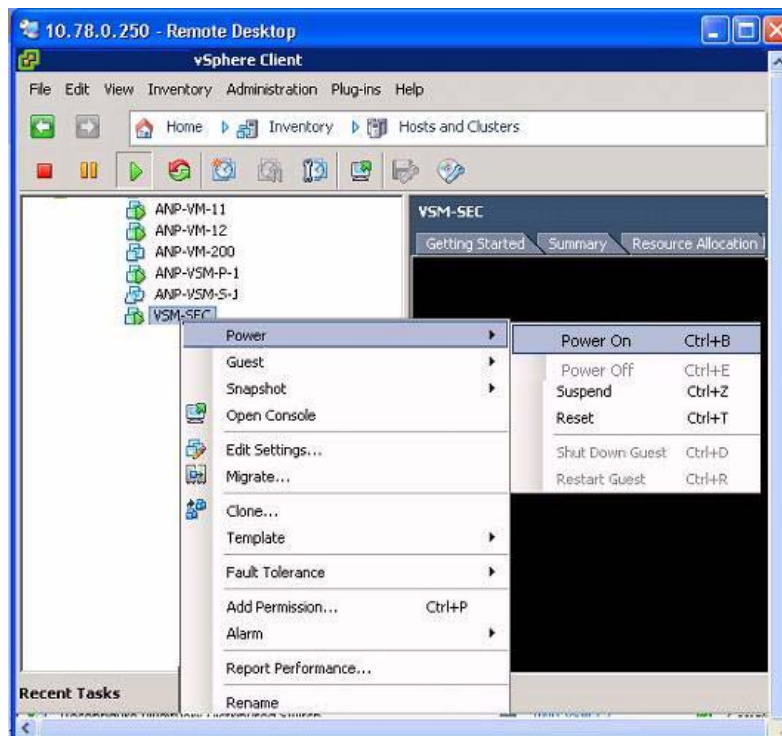
Step 6 You have completed this procedure. Return to the process that pointed you here:

- [“Recovering a Secondary VSM with Active Primary”](#) section on page 5-7
- [“Recovering a Primary VSM with Active Secondary”](#) section on page 5-12

Powering On the VSM

Use this procedure to power on the newly-created VSM.

Step 1 From the vCenter Server left-hand panel, right-click the VSM and then choose Power > Power On. The VSM starts.



Step 2 You have completed this procedure. Return to the [“Recovering a Primary VSM with Active Secondary”](#) section on page 5-12.

Changing the HA Role

You can use this procedure to change the HA role of the VSM.

BEFORE YOU BEGIN

- You are logged in to the CLI in EXEC mode.
- You know the domain ID of the existing VSM.

DETAILED STEPS

Step 1 Go to the domain of the existing VSM.

configure terminal

svs-domain

domain id *domain-id*

Example:

```
n1000v# config t
n1000v(config)# svs-domain
n1000v(config-svs-domain)# domain id 1941
Warning: Config saved but not pushed to vCenter Server due to inactive connection!
```

Step 2 Change the HA role.

system redundancy role [primary | secondary | standalone]

Example:

```
n1000v(config-svs-domain)# system redundancy role secondary
Setting will be activated on next reload.
n1000v(config-svs-domain)#
```

Example:

```
n1000v(config-svs-domain)# system redundancy role primary
Setting will be activated on next reload.
n1000v(config-svs-domain)#
```

Step 3 Copy the running configuration to the startup configuration.

copy run start

Example:

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

Step 4 You have completed this procedure. Return to the [“Recovering a Primary VSM with Active Secondary” section on page 5-12.](#)

Recovering a Primary VSM with Active Secondary

You can use the following process to recover a primary VSM when the secondary VSM is active.

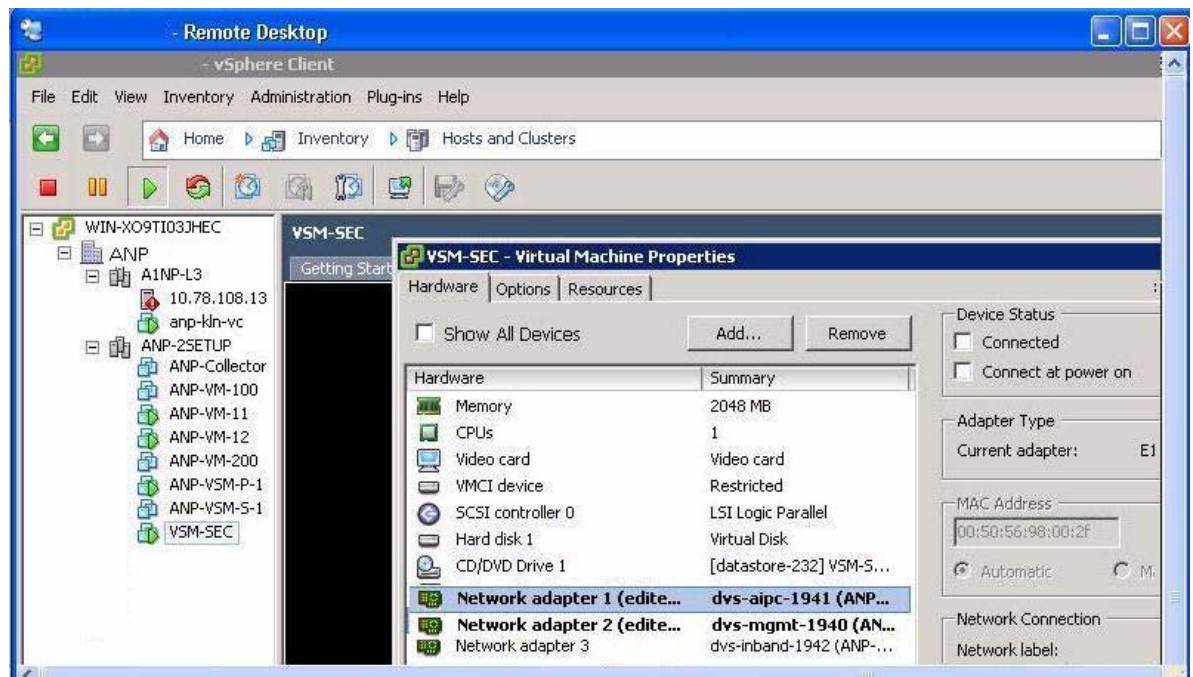
-
- Step 1** Stop the upgrade on the secondary VSM, using the [“Stopping a VSM Upgrade” procedure on page 5-8](#)
- Step 2** Change the boot variables back to the previous version using the [“Changing Boot Variables” procedure on page 5-9](#)
- Step 3** From the vCenter Server left-hand panel, right-click the primary VSM and then choose Delete from Disk.
The primary VSM is deleted.
- Step 4** Create a new VSM by reinstalling the software from the OVA and specifying the following:
- Manual (CLI) configuration method instead of GUI.
 - The host or cluster of the existing secondary VSM
- For a detailed installation procedure, see the *Cisco Nexus 1000V Installation and Upgrade Guide*.
- Step 5** Make sure the port groups between the host server and VSM are not connected when the new VSM is powered on, using the [“Disconnecting the Port Groups” procedure on page 5-12](#).
- Step 6** Power on the newly-created VSM using the [“Powering On the VSM” procedure on page 5-10](#).
The VSM comes up with the standalone HA role.
- Step 7** Change the HA role of the newly-created standalone VSM to primary and save the configuration, using the [“Changing the HA Role” procedure on page 5-11](#).
- Step 8** Power off the newly-created VSM, using the [“Powering Off the VSM” procedure on page 5-14](#).
- Step 9** Make sure the port groups between the host server and VSM are connected when the new VSM is powered on, using the [“Connecting the Port Groups” procedure on page 5-14](#).
- Step 10** Power on the newly-created VSM using the [“Powering On the VSM” procedure on page 5-10](#).
The VSM comes up, connects with the host server, and forms an HA pair with the existing primary VSM.
-

Disconnecting the Port Groups

Use this procedure to disconnect and prevent port groups to the VSM from connecting to the host server.

-
- Step 1** In vCenter Server, select the VSM and then choose Edit > Settings.
The Virtual Machine Properties dialog box opens.
- Step 2** Select the Control port group and uncheck the following Device Settings:
- Connected
 - Connect at Power On

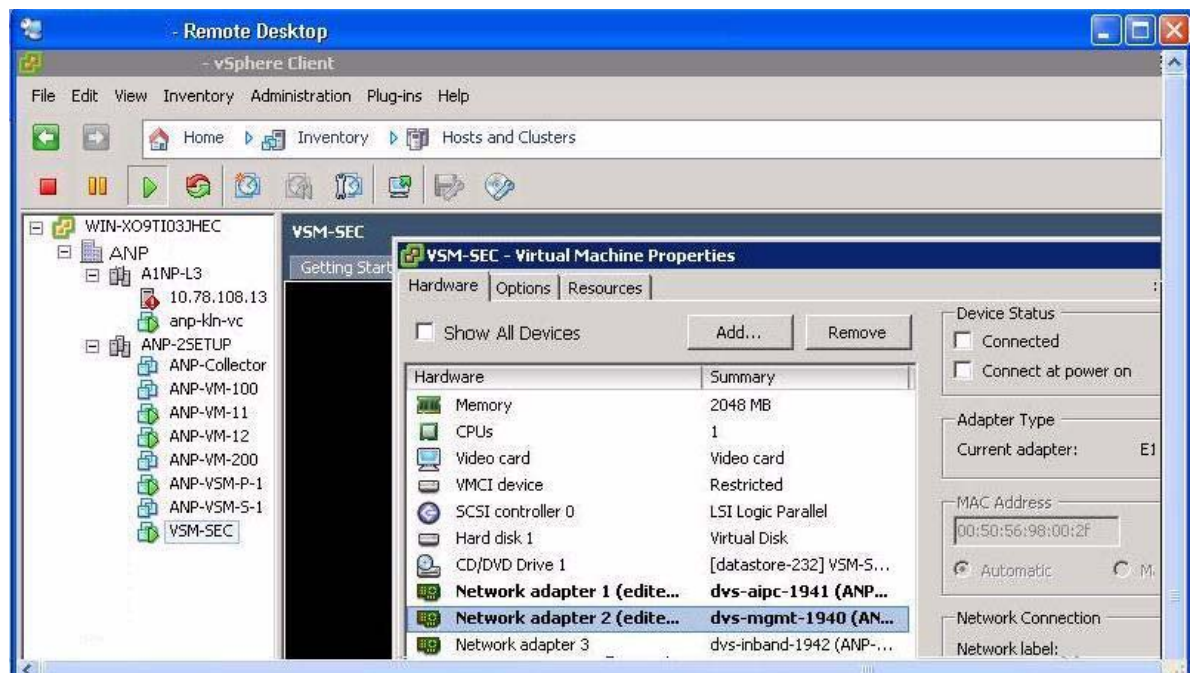
The connection from the VSM to the host server through the control port is dropped and is not restored when you power on the VSM.



Step 3 Select the Management port group and uncheck the following Device Settings:

- Connected
- Connect at Power On

The connection from the VSM to the host server through the management port is dropped and is not restored when you power on the VSM.

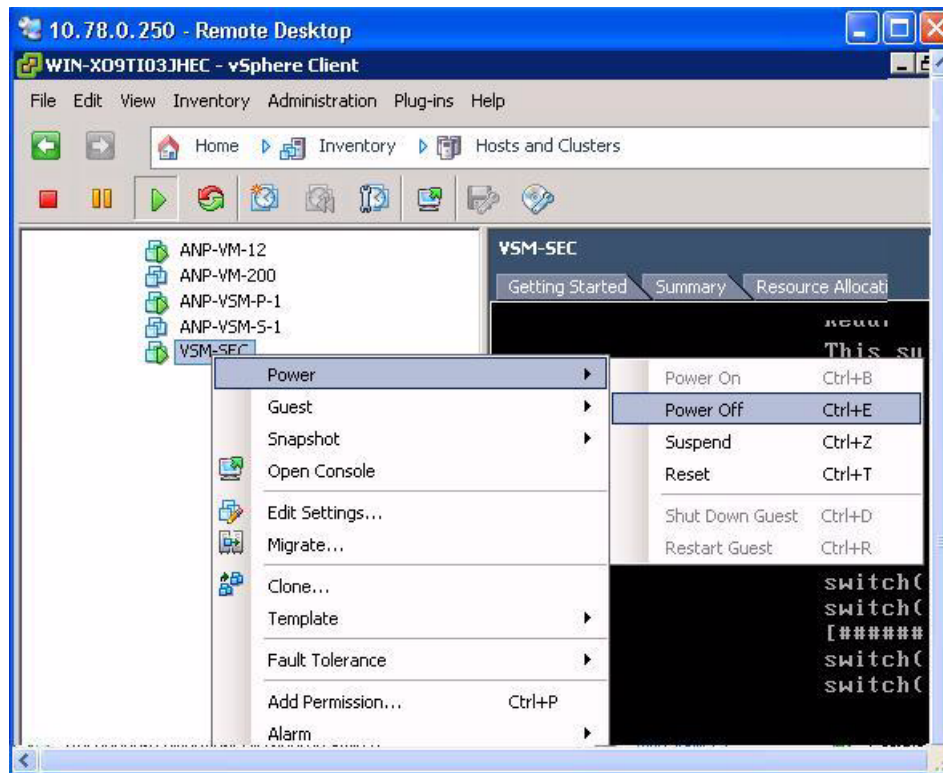


- Step 4** You have completed this procedure. Return to the [“Recovering a Primary VSM with Active Secondary”](#) section on page 5-12.

Powering Off the VSM

Use this procedure to power off the newly-created VSM.

- Step 1** From the vCenter Server left-hand panel, right-click the VSM and then choose Power > Power Off. The VSM shuts down.



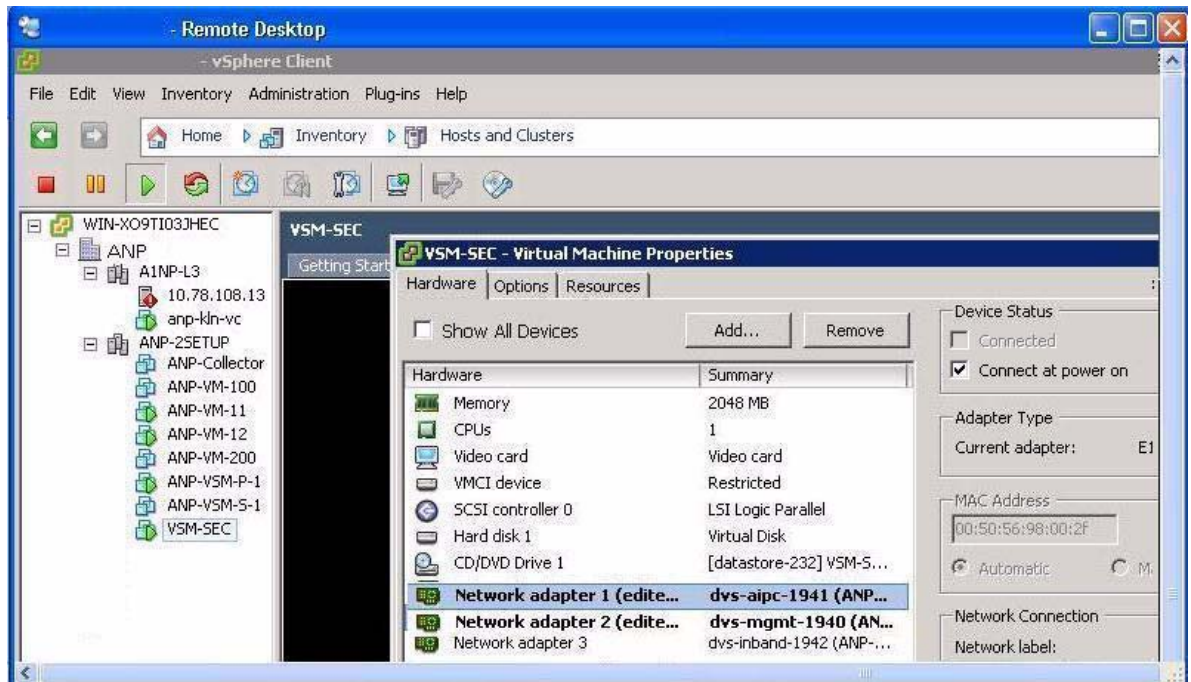
- Step 2** You have completed this procedure. Return to the [“Recovering a Primary VSM with Active Secondary”](#) section on page 5-12.

Connecting the Port Groups

Use this procedure to make sure the port groups to the host connect when you power on the VSM.

- Step 1** In vCenter Server, select the VSM and then choose Edit > Settings. The Virtual Machine Properties dialog box opens.
- Step 2** Select the Control port group and check the following Device Settings:
- Connect at Power On

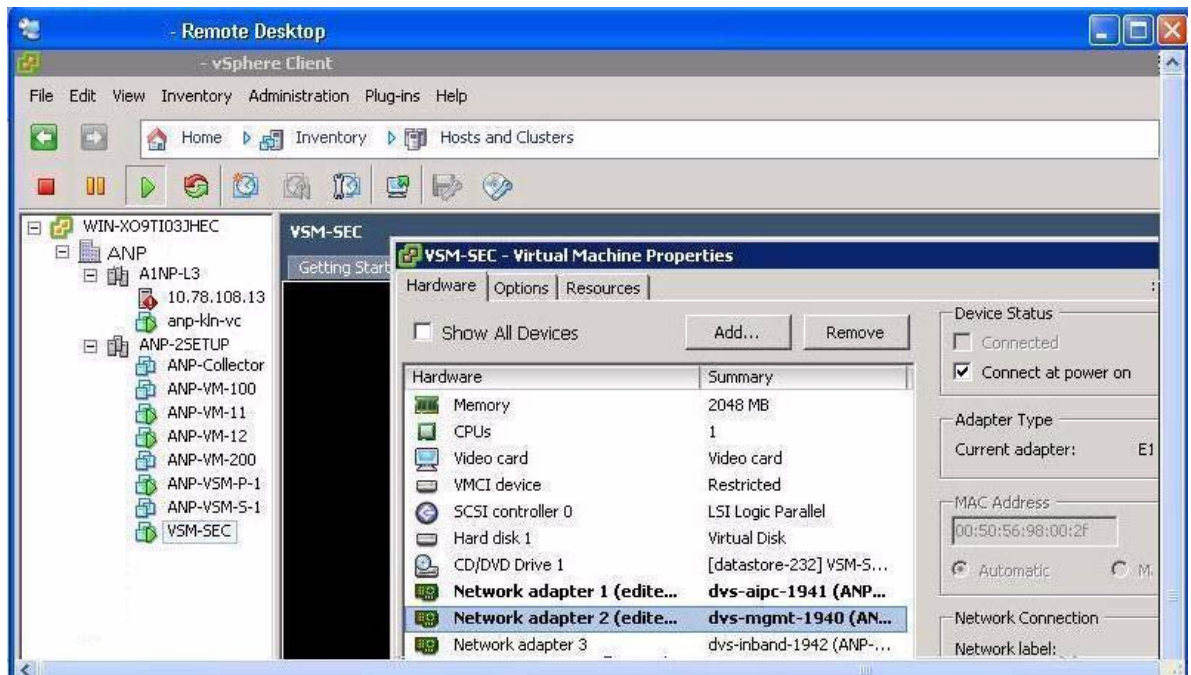
When you power on the VSM, it will connect to the host server through the control port.



Step 3 Select the Management port group and check the following Device Setting:

- Connect at Power On

When you power on the VSM, it will connect to the host server through the management port.



- Step 4** You have completed this procedure. Return to the [“Recovering a Primary VSM with Active Secondary” section on page 5-12.](#)

Problems with VSM-VEM Layer 2 to 3 Conversion Tool

The following is a symptom and solution for a problem with logging in to VSM when using the conversion tool:

Symptom	Solution
<p>When you enter your VSM and VC login credentials for the first time, the VSM-VEM Layer 2 to 3 Conversion Tool may display:</p> <pre>Timeout error. Is device down or unreachable?? ssh_expect</pre>	<ol style="list-style-type: none"> 1. Open a command line window and run an ssh command on the VSM (ssh username@vsmIPAddress). 2. When prompted, Are you sure you want to continue connecting?, enter yes. 3. Rerun the VSM-VEM Layer 2 to 3 Conversion Tool by reopening the .bat file. Ensure that the error does not reappear.

Upgrade Troubleshooting Commands

Should all examples comply w/issu upgrade? At least for this release, there is still gui upgrade to 4a from pre-sv1(4) releases, correct? Troubleshooting guide needn't only address sv1(4) upgrade to sv1(4a), correct?

You can use the commands in this section to troubleshoot problems related to upgrade.

Command	Description
show boot	Displays boot variable definitions, showing the names of software images used to boot the VSM. See Example 5-2 on page 5-17 .
show module	Displays module status for active and standby VSMs. See Example 5-4 on page 5-17 (ISSU) . See Example 5-4 on page 5-17
show running-config in boot	Displays the boot variables currently in the running configuration. See Example 5-5 on page 5-18 .
show startup-config in boot	Displays the boot variables currently in the startup configuration. See Example 5-6 on page 5-18 .

Command	Description
show svcs connections	Displays the current connections between the VSM and the VMware host server. See Example 5-7 on page 5-18 .
show svcs upgrade status	Displays the upgrade status. See Example 5-8 on page 5-18 .
show system redundancy status	Displays the current redundancy status for the VSM. See Example 5-9 on page 5-19 .
show vmware vem upgrade status	Displays the upgrade status. See Example 5-10 on page 5-19 .

Example 5-2 show boot

```
n1000v# show boot
sup-1
kickstart variable = bootflash:/nexus-1000v-kickstart-mz.4.0.4.SV1.3a.bin
system variable = bootflash:/nexus-1000v-mz.4.0.4.SV1.3a.bin
sup-2
kickstart variable = bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4.bin
system variable = bootflash:/nexus-1000v-mzg.4.2.1.SV1.4.bin
No module boot variable set
n1000v#
```

Example 5-3 show module (VSM upgraded first with ISSU, VEM upgrade pending)

```
n1000v# show module
Mod Ports Module-Type Model Status
-----
1 0 Virtual Supervisor Module Nexus1000V ha-standby
2 0 Virtual Supervisor Module Nexus1000V active *
3 248 Virtual Ethernet Module NA ok
Mod Sw Hw
-----
1 4.2(1)SV1(4a) 0.0
2 4.2(1)SV1(4a) 0.0
3 4.2(1)SV1(4) 1.9
Mod MAC-Address(es) Serial-Num
-----
1 00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8 NA
2 00-19-07-6c-5a-a8 to 00-19-07-6c-62-a8 NA
3 02-00-0c-00-03-00 to 02-00-0c-00-03-80 NA
Mod Server-IP Server-UUID Server-Name
-----
1 10.78.109.43 NA NA
2 10.78.109.43 NA NA
3 10.78.109.51 4220900d-76d3-89c5-17d7-b5a7d1a2487f 10.78.109.51
n1000v#
```

Example 5-4 show module (VEM and VSM upgraded)

```
n1000v# show module
Mod Ports Module-Type Model Status
-----
1 0 Virtual Supervisor Module Nexus1000V ha-standby
```

```

2      0      Virtual Supervisor Module      Nexus1000V      active *
3      248     Virtual Ethernet Module      NA              ok

Mod Sw          Hw
---
1      4.0(4)SV1(3)      0.0
2      4.0(4)SV1(3)      0.0
3      4.2(1)SV1(4)      1.9

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

Mod Server-IP      Server-UUID          Server-Name
---
1      10.78.109.43      NA                    NA
2      10.78.109.43      NA                    NA
3      10.78.109.51      4220900d-76d3-89c5-17d7-b5a7d1a2487f      10.78.109.51
n1000v#

```

Example 5-5 show running-config | include boot

```

n1000v# show running-config | include boot
boot kickstart bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4a.bin sup-1
boot system bootflash:/nexus-1000v-mzg.4.2.1.SV1.4a.bin sup-1
boot kickstart bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4a.bin sup-2
boot system bootflash:/nexus-1000v-mzg.4.2.1.SV1.4a.bin sup-2
n1000v#

```

Example 5-6 show startup-config | include boot

```

switch# show startup-config | include boot
boot kickstart bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4a.bin sup-1
boot system bootflash:/nexus-1000v-mzg.4.2.1.SV1.4a.bin sup-1
boot kickstart bootflash:/nexus-1000v-kickstart-mzg.4.2.1.SV1.4a.bin sup-2
boot system bootflash:/nexus-1000v-mzg.4.2.1.SV1.4a.bin sup-2
n1000v#

```

Example 5-7 show svcs connections

```

n1000v# show svcs connections

connection vc:
  hostname: 172.23.232.139
  remote port: 80
  protocol: vmware-vim https
  certificate: default
  datacenter name: Hamilton-DC
  DVS uuid: 9b dd 36 50 2e 27 27 8b-07 ed 81 89 ef 43 31 17
  config status: Enabled
  operational status: Connected
  sync status: -
  version: -
n1000v#

```

Example 5-8 show svcs upgrade status

```

n1000v# show svcs upgrade status
Upgrade State: Start

```

```
Upgrade mgmt0 ipv4 addr: 1.1.1.1
Upgrade mgmt0 ipv6 addr:
Upgrade control0 ipv4 addr:
n1000v#
```

Example 5-9 show system redundancy status

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

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

This supervisor (sup-1)
-----
      Redundancy state:  Active
      Supervisor state:  Active
      Internal state:    Active with HA standby

Other supervisor (sup-2)
-----
      Redundancy state:  Standby
      Supervisor state:  HA standby
      Internal state:    HA standby

switch#
```

Example 5-10 show vmware vem upgrade status

```
n1000v# 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: VEM400-201007101-BG
      DVS: VEM400-201007101-BG

n1000v#
```




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 6-1](#)
- [Problems with High Availability, page 6-2](#)
- [High Availability Troubleshooting Commands, page 6-5](#)

Information About High Availability

The purpose of High Availability (HA) is to limit the impact of failures—both hardware and software—within a system. The Cisco NX-OS operating system is designed for high availability at the network, system, and service levels.

The following Cisco NX-OS features minimize or prevent traffic disruption in the event of a failure:

- Redundancy— redundancy at every aspect of the software architecture.
- Isolation of processes— isolation between software components to prevent a failure within one process disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- Supervisor stateful switchover— Active/standby dual supervisor configuration. 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)

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.

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*.

Problems with High Availability

Symptom	Possible Causes	Solution
The active VSM does not see the standby VSM.	Roles are not configured properly. <ul style="list-style-type: none"> Check the role of the two VSMs using the show system redundancy status command. 	<ol style="list-style-type: none"> Confirm that the roles are the primary and secondary role, respectively. If needed, use the system redundancy role command to correct the situation. Save the configuration if roles are changed.
	Network connectivity problems. <ul style="list-style-type: none"> Check the control and management VLAN connectivity between VSM at the upstream and virtual switches. 	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.

Symptom	Possible Causes	Solution
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs. <ul style="list-style-type: none"> Check that primary and secondary VSM are using the same image version using show version command. 	If the active and the standby VSM software versions differ, 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 using the reload module <i>module-number</i> command, where <i>module-number</i> is the module number for the standby VSM.
	<ul style="list-style-type: none"> The VSM has connectivity only through the management interface. Check the output of the show system internal redundancy info command and verify if the <i>degraded_mode</i> flag is set to <i>true</i>. 	Check control VLAN connectivity between the primary and the secondary VSMs.
The standby VSM reboots periodically.	The VSM has connectivity only through the management interface. <ul style="list-style-type: none"> Check the output of the show system internal redundancy info command and verify if the <i>degraded_mode</i> flag is set to true. 	Check control VLAN connectivity between the primary and the secondary VSMs.
	The 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: 2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.	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.

Symptom	Possible Causes	Solution
Active-Active detected and resolved	<p>When control and management connectivity between the active and the standby goes down for 6 seconds, the standby VSM transitions to the active state.</p> <p>Upon restoration of control and management connectivity, both VSMs detect an active-active condition.</p>	<ol style="list-style-type: none"> Once the system detects active-active VSMs, one of the VSM is automatically reloaded based on various parameters like VEMs attached, vCenter connectivity, last configuration time, and last active time. To see any configuration changes that are performed on the rebooted VSM during the active-active condition, execute show system internal active-active remote accounting logs CLI command on the active VSM.
VSM Role Collision	<p>If another VSM is configured/provisioned with the same role (primary or secondary) in the system, the new VSM collides with the existing VSM.</p> <p>The show system redundancy info command displays the MAC addresses of the VSM(s) that collide with the working VSM.</p>	<p>If the problems exist:</p> <ol style="list-style-type: none"> Execute show system redundancy status command on the VSM console. Identify the VSM(s) that owns the MAC addresses that are displayed in the output of show system redundancy status command. Move the identified VSM(s) out of the system to stop role collision.

Symptom	Possible Causes	Solution
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.
	Different domain IDs in the two VSMS Check <i>domain</i> value using show system internal redundancy info command.	If needed, update the domain ID and save it to the startup configuration. <ul style="list-style-type: none"> Upgrading the domain ID in a dual VSM system must be done following a certain procedure. <ul style="list-style-type: none"> 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 the isolated VSM and power it on.

High Availability Troubleshooting Commands

This section lists commands that can be used troubleshoot problems related to High Availability.

To list process logs and cores, use the following commands:

- show cores**

```
n1000V# show cores
VDC No Module-num      Process-name      PID      Core-create-time
-----
1      1      private-vlan      3207     Apr 28 13:29
```

- 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
```

```
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.

CWD: /var/sysmgr/work
...

```

To check redundancy status, use the following commands:

- **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

```

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
  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: ha1
  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

```

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.

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
```

When a role collision is detected, a warning is highlighted in the CLI output. Use the following command to display the CLI output:

- **n1000v# show system redundancy status**

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



```
Supervisor state: HA standby
Internal state: HA standby
WARNING! Conflicting sup-2(s) detected in same domain
-----
```

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

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

If no collisions are detected, the highlighted output is not displayed.

Use the following command to display the accounting logs that are stored on a remote VSM.

- n1000V# **show system internal active-active remote accounting logs**

To reload a module, use the following command:

- **reload module**

```
n1000V# reload module 2
```

This command reloads the secondary VSM.



Note Issuing the **reload** command without specifying a module will reload the whole system.

To attach to the standby VSM console, use the following command.

- **attach module**

The standby VSM console is not accessible externally, but can be accessed from the active VSM through the **attach module** *module-number* command.

```
n1000V# attach module 2
```

This command attaches to the console of the secondary VSM.



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 7-1](#)
- [Troubleshooting a Module Not Coming Up on the VSM, page 7-1](#)
- [Problems with the VSM, page 7-4](#)
- [VSM and VEM Troubleshooting Commands, page 7-18](#)

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 in Nexus 1000V 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 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 7-2](#)
- [Flow Chart for Troubleshooting Modules, page 7-3](#)
- [Verifying the VSM Is Connected to the vCenter Server, page 7-6](#)
- [Verifying the VSM Is Configured Correctly, page 7-7](#)
- [Checking the vCenter Server Configuration, page 7-10](#)
- [Checking Network Connectivity Between the VSM and the VEM, page 7-10](#)

- [Recovering Management and Control Connectivity of a Host when a VSM is Running on a VEM, page 7-12](#)
- [Checking the VEM Configuration, page 7-14](#)
- [Collecting Logs, page 7-17](#)

Guidelines for Troubleshooting Modules

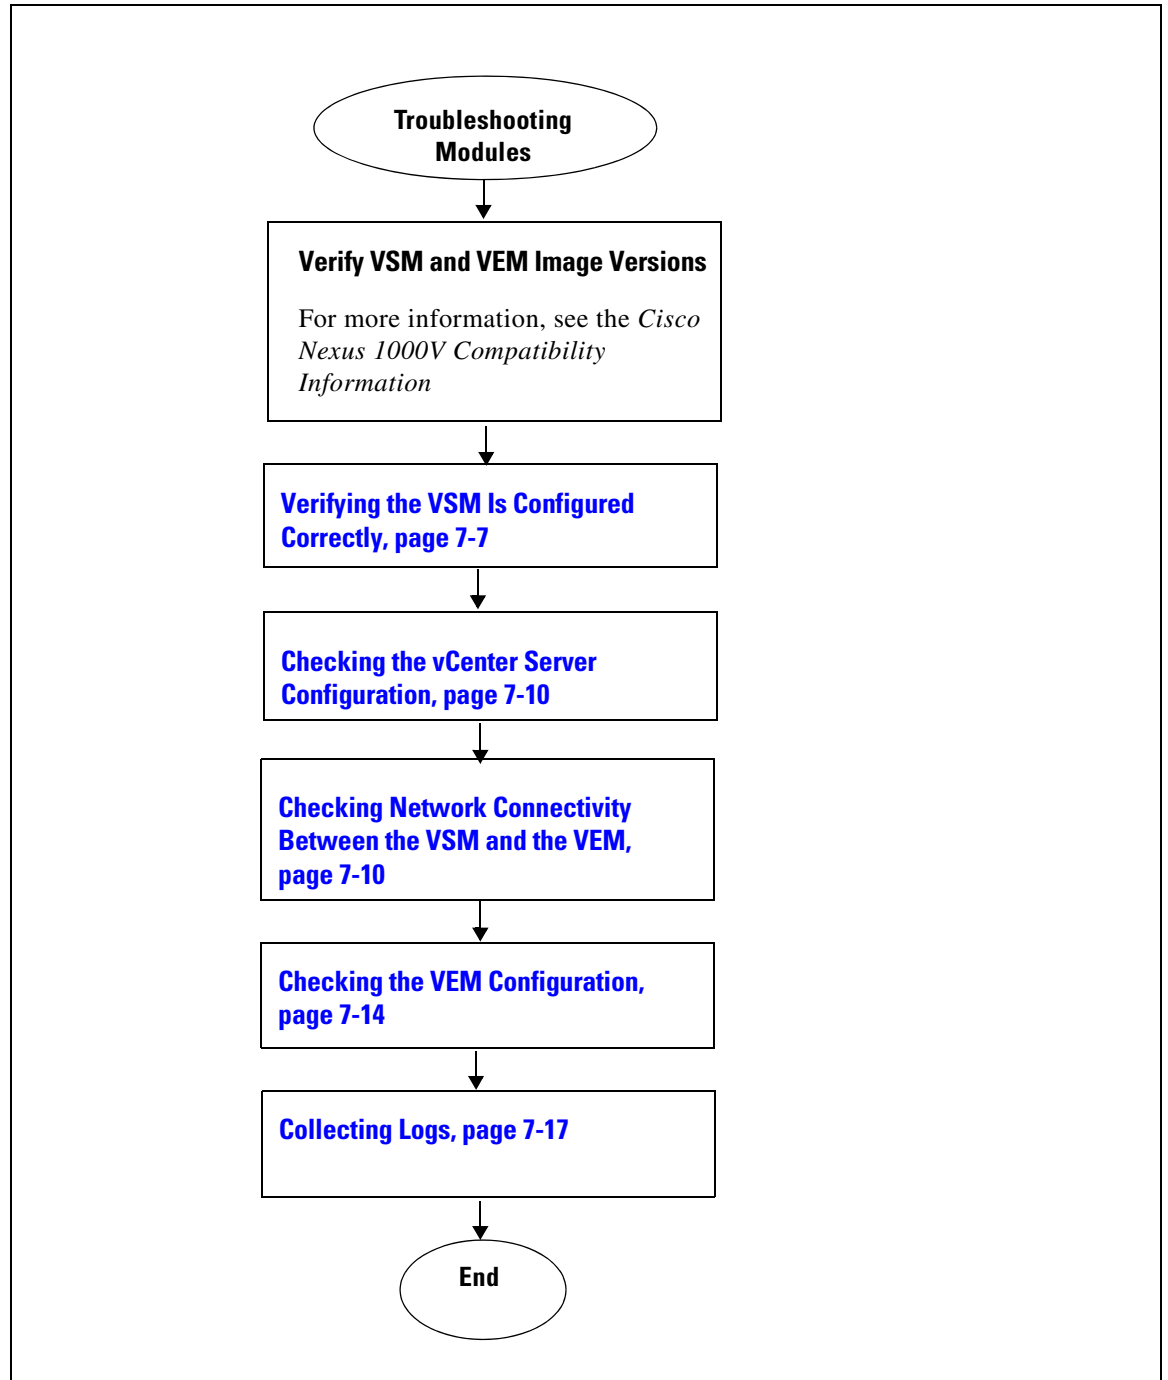
Follow these guidelines when troubleshooting a module controlled by the VSM.

- 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*.
- 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 useTo 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 7-10](#).
 - ERROR: Datacenter not found
- For a list of terms used with Cisco Nexus 1000V, see the *Cisco Nexus 1000V Getting Started Guide*.

Flow Chart for Troubleshooting Modules

Use the following flowchart to troubleshoot modules.

Flowchart: Troubleshooting Modules



Problems with the VSM

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

Table 7-1 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>	A extension or plug-in was not created for the VSM.	<ol style="list-style-type: none"> Verify that the extension or plugin was created. “Finding the Extension Key Tied to a Specific DVS” procedure on page 3-8 If the plug-in is not found, then create one using the following procedure in the <i>Cisco Nexus 1000V Getting Started Guide</i>: Creating a Cisco Nexus 1000V Plug-In on the vCenter Server
<p>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.</p>		
After boot, VSM in loader prompt.	Corrupt VSM kickstart image.	<ol style="list-style-type: none"> Boot the VSM from the CD ROM. From the CD Boot menu, choose Option 1, Install Nexus1000v and bring up new image. Follow the VSM installation procedure.
	Boot variables are not set.	<ol style="list-style-type: none"> Boot the VSM from the CD ROM. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk unformatted and bring up new image. Set the boot variables used to boot the VSM: boot system bootflash:system-boot-variable-name boot kickstart bootflash:kickstart-boot-variable-name Reload the VSM. reload
After boot, VSM in boot prompt.	Corrupt VSM system image.	<ol style="list-style-type: none"> Boot the VSM from the CD ROM. From the CD Boot menu, choose Option 1, Install Nexus1000v and bring up new image. Follow the VSM installation procedure.

Table 7-1 Problems with the VSM (continued)

Symptom	Possible Causes	Solution
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</i>: Setting Up the Software
After boot, VSM stopped at “Loader Loading.”	Corrupt boot menu file.	<ol style="list-style-type: none"> Boot the VSM from the CD ROM. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk unformatted and bring up new image. 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</i>: Setting Up the Software
After boot, secondary VSM reboots continuously.	Control VLAN or control interface down	Check control connectivity between the active and standby VSM.
	Active and standby VSMs fail to synchronize.	<p>From active VSM, check gsyncstats to identify which application caused the failure.</p> <p>show logging</p>
After a host reboot, the absence of a VLAN, or the wrong system VLAN on the VSM management port profile, the control and management connectivity of the VSM is lost.	The VSM is running on a VEM that it manages, but the VSM ports are not configured with system port profiles.	<p>Run the VEM connect script locally in the ESX host where the VEM is running. Then go to the VSM and configure the system VLANs in the port profile used for management.</p> <p>“Recovering Management and Control Connectivity of a Host when a VSM is Running on a VEM” section on page 7-12</p>

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 7-3](#).
- If not, then continue with the next step.

Step 3 Connect to the vCenter Server.

config t

svcs connection *connection_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
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 7-1](#).
- If not, go to [Step 6](#).

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

- Unregister the extension key using the “[Unregister 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*.
 - [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 7-3](#).

Verifying the VSM Is Configured Correctly

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

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

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.
- The output of the `show svcs 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 svcs domain**Example:**

```
n1000v# show svcs 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 the following procedure to verify the port profile 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.
- The output of the **show port-profile name** command should indicate the following:
 - The control and packet VLANs are assigned.
 - The port profile is enabled.
 - If you have configured a non-default system mtu setting, then it is of the correct size.

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

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

Example:

```
n1000v# show port-profile name SystemUplink
port-profile SystemUplink
  description:
  type: ethernet
  status: enabled
  capability 13control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: 114,115
  port-group: SystemUplink
  max ports: 32
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan all
    system mtu 1500
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan all
    no shutdown
  assigned interfaces:
```

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.



Note

The procedure documented is for troubleshooting VSM and VEM connectivity with layer 2 mode.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- The output of the **show running-config** command should show control and packet VLAN ID numbers among the VLANs configured,

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

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

n1000v#
.
.
.
```

Step 2 Find the AIPC MAC address of the VSM by running **show svcs neighbors** on the VSM.

```
switch(config-svs-domain)# show svcs neighbors

Active Domain ID: 27

AIPC Interface MAC: 0050-56bc-74f1 <-----
Inband Interface MAC: 0050-56bc-62bd

Src MAC          Type   Domain-id   Node-id     Last learnt (Sec. ago)
-----
0050-56bc-6a3d   VSM    27          0201        771332.97
0002-3d40-1b02   VEM    27          0302         51.60
0002-3d40-1b03   VEM    27          0402         51.60
```

Step 3 Find the DPA MAC address of the VEM by running **vemcmd show card** on the ESX host.

```
# vemcmd show card
Card UUID type 2: 24266920-d498-11e0-0000-00000000000f
Card name:
Switch name: Nexus1000v
Switch alias: DvsPortset-0
Switch uuid: ee 63 3c 50 04 b1 6d d6-58 61 ff ba 56 05 14 fd
Card domain: 27
Card slot: 3
VEM Tunnel Mode: L2 Mode
VEM Control (AIPC) MAC: 00:02:3d:10:1b:02
VEM Packet (Inband) MAC: 00:02:3d:20:1b:02
VEM Control Agent (DPA) MAC: 00:02:3d:40:1b:02 <-----
VEM SPAN MAC: 00:02:3d:30:1b:02
Primary VSM MAC : 00:50:56:bc:74:f1
Primary VSM PKT MAC : 00:50:56:bc:62:bd
Primary VSM MGMT MAC : 00:50:56:bc:0b:d5
Standby VSM CTRL MAC : 00:50:56:bc:6a:3d
Management IPv4 address: 14.17.168.1
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
Primary L3 Control IPv4 address: 0.0.0.0
Secondary VSM MAC : 00:00:00:00:00:00
Secondary L3 Control IPv4 address: 0.0.0.0
Upgrade : Default
Max physical ports: 32
Max virtual ports: 216
Card control VLAN: 168
```

```

Card packet VLAN: 168
Control type multicast: No
Card Headless Mode : No
    Processors: 16
    Processor Cores: 8
    Processor Sockets: 2
    Kernel Memory: 25102148
Port link-up delay: 5s
Global UUPB: DISABLED
Heartbeat Set: True
PC LB Algo: source-mac
Datapath portset event in progress : no
Licensed: Yes

```

Step 4 Check the upstream switches for these MAC addresses in the correct VLANs.

```

switch1 # show mac address-table | grep 1b02
* 168      0002.3d20.1b02    dynamic    20          F      F      Veth854
* 168      0002.3d40.1b02    dynamic    0           F      F      Veth854
* 1        0002.3d40.1b02    dynamic    1380       F      F      Veth854

switch2 # show mac address-table | grep 74f1
* 168      0050.56bc.74f1    dynamic    0           F      F      Eth1/1/3

```

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 Cisco Nexus 1000V 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.
 - Step 3** Confirm that the three VSM vnics are assigned to the port groups containing the control VLAN, packet VLAN, and management network.
-

Checking 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.

-
- 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 vsm 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*

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
n1000v(config)# module vem 3 execute vemcmd show l2 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 l2 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 7-14](#).
- Otherwise, you have completed this procedure.

Recovering Management and Control Connectivity of a Host when a VSM is Running on a VEM

When the VSM is running on a VEM that it manages, but the VSM ports are not configured with system port profiles, the control and management connectivity of the VSM can be lost after a host reboot or similar event. To recover from the loss, you can run the VEM connect script locally in the ESX host where the VEM is running. Then go to the VSM and configure the system VLANs in the port profile used for management.

Using the VEM Connect Script

The VEM connect script sets a given VLAN as a system VLAN on the vmknic that has the given IP address, and also sets the VLAN on all the required uplinks.

If no uplink is carrying this VLAN, you also need to specify the uplink (vmmicN) on which this VLAN needs to be applied. The uplink can be a single port or a port-channel member. If it is the latter, then the script will apply the VLANs as a system VLAN to all member uplinks of that port channel.

vem-connect -i <ip_address> -v <vlan> [-p <vmmicN>]

The -p parameter to the script is optional. If you run the script without the -p parameter, it will try to locate an uplink that carries this VLAN. If no such uplink exists, it will report this as an error. You need to specify the -p parameter and re-run the script.

You can use the following procedure to recover management and control connectivity of a host when a VSM is running on a VEM.

SUMMARY

Step 1 Enter the following command to display the VEM ports:

vemcmd show port

Example:

```
~ # vemcmd show port
LTL    VSM Port  Admin Link  State  PC-LTL  SGID  Vem Port  Type
  18    Eth9/2    UP   UP    F/B*   305    1    vmmic1
  20    Eth9/4    UP   UP    F/B*   305    3    vmmic3
  49    Veth1     UP   UP    FWD    0      3    VM-T-125.eth0
  50    Veth10    UP   UP    FWD    0      1    vmk1
 305    Po2       UP   UP    F/B*   0
```

* F/B: The port is blocked on some of the VLANs.



Note

The output *F/B The port is blocked on some of the VLANs means that the trunk is not forwarding all vlans. This may be normal depending on the port profile allowed VLAN list. Compare the output of the `vemcmd show port vlans` against the list of allowed VLANs in the trunk port profile. If the lists match, then all of the expected VLANs are forwarding and the Cisco Nexus 1000V is blocking non-allowed VLANs.

Step 2 Enter the following command to display details about the system VLANs:

vemcmd show port vlans system

Example:

```
~ # vemcmd show port vlans system
LTL    VSM Port  Mode  Native VLAN  Allowed
      VSM Port  Mode  VLAN/  State  Vlans/SegID
      SegID
  6    Internal  A      1    FWD    1
  8    Internal  A     3969  FWD    3969
  9    Internal  A     3969  FWD    3969
 10    Internal  A     210   FWD    210
 11    Internal  A     3968  FWD    3968
 12    Internal  A     211   FWD    211
 13    Internal  A      1    BLK    1
 14    Internal  A     3971  FWD    3971
 15    Internal  A     3971  FWD    3971
 16    Internal  A      1    FWD    1
 18    Eth9/2    T      1    FWD    210-211
 20    Eth9/4    T      1    FWD    210-211
 49    Veth1     A      1    FWD    1
 50    Veth10    A      1    FWD    1
```

```
305      Po2  T          1  FWD  210-211
```

Step 3 Enter the following command to recover connectivity:

```
vem-connect -i <ip_address> -v <vlan> [ -pnic <vmnicN> ]
```

Example:

```
~ # vem-connect -i 172.23.232.67 -v 232 -p vmnic3
ltl 50 and veth Veth10 vmk1
Uplink port Po2 carries vlan 232
Set System Vlan 232 port Po2 305
Uplink port Eth9/2 carries vlan 232
Set System Vlan 232 port Eth9/2 18
Uplink port Eth9/4 carries vlan 232
Set System Vlan 232 port Eth9/4 20
Set System 232 for vmk
```

Step 4 Enter the following command to confirm management connectivity:

```
vemcmd show port vlans system
```

Example:

```
~ # vemcmd show port vlans system
```

LTL	VSM Port	Mode	Native VLAN/ SegID	VLAN State	Allowed Vlans/SegID
6	Internal	A	1	FWD	1
8	Internal	A	3969	FWD	3969
9	Internal	A	3969	FWD	3969
10	Internal	A	210	FWD	210
11	Internal	A	3968	FWD	3968
12	Internal	A	211	FWD	211
13	Internal	A	1	BLK	1
14	Internal	A	3971	FWD	3971
15	Internal	A	3971	FWD	3971
16	Internal	A	1	FWD	1
18	Eth9/2	T	1	FWD	210-211,232
20	Eth9/4	T	1	FWD	210-211,232
49	Veth1	A	1	FWD	1
50	Veth10	A	232	FWD	232
305	Po2	T	1	FWD	210-211,232

Checking the VEM Configuration

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

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 value ltl port-LTL-number
```

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
```



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    1  Access  120
9     0    3969   0     2      2  VIRT   UP    UP    1  Access  121
10    0    3002   0     2      2  VIRT   UP    UP    1  Access  122
11    0    3968   0     2      2  VIRT   UP    UP    1  Access  123
12    0    3003   0     2      2  VIRT   UP    UP    1  Access  124
13    0     1     0     2      2  VIRT   UP    UP    0  Access  125
```

```

14          0 3967          0 2          2 VIRT    UP    UP    1 Access 126
16 1a030100          1 T    0 2          2 PHYS    UP    UP    1 Trunk vmnic1

```

The last line of output indicates that `vmnic1` should be in Trunk mode, with the CBL value of 1. The CBL value of the native VLAN does not have to be 1. 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. The Admin state and Port state should be UP.

Step 5 Use the following commands to verify 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

```

Example:

```

~ # vemcmd show bd 3002
BD 3002, vdc 1, vlan 3002, 2 ports
Portlist:
    10 122
    16 vmnic1
~ # vemcmd show bd 3003
BD 3003, vdc 1, vlan 3003, 2 ports
Portlist:
    12 124
    16 vmnic1

```

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 1vlan(1) cbl 1, vlan(3002) cbl 1, vlan(3003) cbl 1,

```

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 Use the following commands to 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

```

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    0 Access vmk1
~ # vemcmd set system-vlan 99 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.

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 module vem mapping

Example:

```
n1000v# show module vem mapping
Mod      Status      UUID                                     License Status
-----
60       absent      33393935-3234-5553-4538-35314e355400  unlicensed
66       powered-up  33393935-3234-5553-4538-35314e35545a  licensed
n1000v#
```

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](#).

VSM and VEM Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to VSM.

Command	Description
show svcs neighbors	Displays all svcs neighbors. See Example 7-1 on page 7-19 .
show svcs connections	Displays the Cisco Nexus 1000V connections. See Example 7-2 on page 7-20 .
show svcs domain	Displays the domain configuration. See Example 7-3 on page 7-20 .
show port-profile name <i>name</i>	Displays the configuration for a named port profile. See Example 7-4 on page 7-20 .
show running-config vlan <i>vlanID</i>	Displays the VLAN information in the running configuration. See Example 7-5 on page 7-20 .
vem-health check <i>vsm_mac_address</i>	Displays the cause of a connectivity problem and recommends next steps in troubleshooting. See Example 7-6 on page 7-21 .
show mac address-table interface	Displays the MAC address table on an upstream switch to verify the network configuration. See Example 7-7 on page 7-21 .
module vem <i>module_number</i> execute vemcmd show l2 [<i>control_vlan_id</i> <i>packet_vlan_id</i>]	Displays the VLAN configuration on the VEM to verify that the VSM MAC appears in the control and packet VLANs. See Example 7-8 on page 7-21 .
vem status	Displays the VEM status to confirm that the VEM Agent is running, and that the correct host uplinks are added to the DVS. See Example 7-9 on page 7-21 .
vemcmd show card	Displays information about cards on the VEM to verify that the domain ID, control VLANs, and packet VLANs are configured correctly on the host. See Example 7-10 on page 7-21 .
vemcmd show port [<i>port-LTL-number</i>]	Displays information about ports on the VEM 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. See Example 7-11 on page 7-22 . See Example 7-12 on page 7-22 .

Command	Description
vemcmd show bd [<i>control_vlan_id</i> <i>packet_vlan_id</i>]	Displays configured information on the VEM to verify that the VM NIC port that is supposed to carry the control VLAN and packet VLAN is present. See Example 7-14 on page 7-23 .
vemcmd show trunk	Displays configured information on the VEM to verify that the DV port groups are successfully pushed from the vCenter Server to the host and that the correct physical trunk port VM NIC is used. See Example 7-15 on page 7-23 .
vem-connect -i < <i>ip_address</i> > -v < <i>vlan</i> > [-pnic < <i>vmnicN</i> >]	Recovers management and control connectivity of a host when a VSM is running on a VEM.
show module vem mapping	Displays information about the VEM a VSM maps to, including VEM module number, status, UUID, and license status See Example 7-16 on page 7-23 .
show platform internal event-history module <i>module-number</i>	Displays platform FSM event information.
show module internal event-history module <i>module-number</i>	Displays the event log for a module.
show system internal im event-history module <i>module-number</i>	Displays the module IM event logs for the system.
show system internal vmm event-history module <i>module-number</i>	Displays the module VMM event logs for the system.
show system internal ethpm event-history module <i>module-number</i>	Displays the module Ethernet event logs for the system.
show system internal ethpm event-history int <i>type slot</i>	Displays the Ethernet interface logs for the system.

Example 7-1 show svcs neighbors

```
n1000v# show svcs neighbors

Active Domain ID: 113

AIPC Interface MAC: 0050-56b6-2bd3
Inband Interface MAC: 0050-56b6-4f2d

Src MAC           Type   Domain-id   Node-id   Last learnt (Sec. ago)
-----
0002-3d40-7102    VEM    113         0302     71441.12
0002-3d40-7103    VEM    113         0402     390.77

n1000v#
```

Example 7-2 show svcs connections

```
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
```

Example 7-3 show svcs domain

```
n1000v# show svcs 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
```

Example 7-4 show port-profile

```
n1000v# show port-profile name SystemUplink
port-profile SystemUplink
  description:
  type: ethernet
  status: enabled
  capability l3control: no
  pinning control-vlan: -
  pinning packet-vlan: -
  system vlans: 114,115
  port-group: SystemUplink
  max ports: 32
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan all
    system mtu 1500
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan all
    no shutdown
  assigned interfaces:
```

Example 7-5 show running-configuration vlan

```
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 7-6 vem-health check

```
~ # 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.

Example 7-7 show mac address-table interface

```
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

Example 7-8 module vem execute vemcmd show l2

```
n1000v# config t
n1000v(config)# module vem 3 execute vemcmd show l2 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 l2 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
```

Example 7-9 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	

Example 7-10 vemcmd show card

```
~ # 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

```

Example 7-11 vsmcmd show port

```

~ # vsmcmd 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    1 Access 120
9     0    3969   0     2     2    VIRT  UP    UP    1 Access 121
10    0    3002   0     2     2    VIRT  UP    UP    1 Access 122
11    0    3968   0     2     2    VIRT  UP    UP    1 Access 123
12    0    3003   0     2     2    VIRT  UP    UP    1 Access 124
13    0     1     0     2     2    VIRT  UP    UP    0 Access 125
14    0    3967   0     2     2    VIRT  UP    UP    1 Access 126
16   1a030100  1 T    0     2     2    PHYS  UP    UP    1 Trunk vmnic1

```

Example 7-12 vsmcmd show port

```

~ # vsmcmd 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

```

Example 7-13 vsmcmd show port

```

~ # module vem 5 e vsmcmd show port
LTL   VSM Port  Admin Link  State  PC-LTL  SGID  Vem Port
17    Eth5/1   UP    UP    FWD    305    0    vmnic0
18    Eth5/2   UP    UP    FWD    305    1    vmnic1
49    Veth11   UP    UP    FWD    0       0    vmk0
50    Veth14   UP    UP    FWD    0       1    vmk1
51    Veth15   UP    UP    FWD    0       0    vswif0
305   Po1      UP    UP    FWD    0       0

```

* F/B: Port is BLOCKED on some of the vlans.
Please run "vsmcmd show port vlans" to see the details.

```

~ # module vem 5 e vsmcmd show port vlans
Native VLAN  State  Allowed
LTL   VSM Port  Mode  VLAN  State  Vlans
17    Eth5/1   T     1     FWD    1,100,119,219,319
18    Eth5/2   T     1     FWD    1,100,119,219,319
49    Veth11   A     119   FWD    119
50    Veth14   A     119   FWD    119
51    Veth15   A     119   FWD    119
305   Po1      T     1     FWD    1,100,119,219,319

```

**Note**

The output *F/B The port is blocked on some of the VLANs means that the trunk is not forwarding all vlans. This may be normal depending on the port profile allowed VLAN list. Compare the output of the vsmcmd show port vlans against the port profile trunk allowed VLANs. If the lists match, then all of the expected VLANs are forwarding and the Cisco Nexus 1000V is blocking non-allowed VLANs.

Example 7-14 vemcmd show bd

```

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

```

Example 7-15 vemcmd show trunk

```

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

```

Example 7-16 show module vem mapping

```

n1000v# show module vem mapping
Mod      Status          UUID                                     License Status
---      -
60       absent          33393935-3234-5553-4538-35314e355400  unlicensed
66       powered-up      33393935-3234-5553-4538-35314e35545a  licensed
n1000v#

```




Ports

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

- [Information About Ports, page 8-1](#)
- [Port Diagnostic Checklist, page 8-3](#)
- [Problems with Ports, page 8-3](#)
- [Port Troubleshooting Commands, page 8-8](#)

Information About Ports

This section includes the following topics:

- [Information About Interface Characteristics, page 8-1](#)
- [Information About Interface Counters, page 8-2](#)
- [Information About Link Flapping, page 8-2](#)
- [Information About Port Security, page 8-2](#)

Information About Interface Characteristics

Before a switch can relay frames from one data link to another, you must define the characteristics of the interfaces through which the frames are received and sent. 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*.

Information About Interface Counters

Port counters are used to identify synchronization problems. Counters can show a significant disparity between received and transmitted frames. To display interface counters, use the following command:

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

See [Example 8-11 on page 8-12](#).

Values stored in 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 the present time. Create a baseline first by clearing the counters.

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

Information About Link Flapping

When a port continually goes up and down, it is said to be flapping, sometimes called link flapping. 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.

To troubleshoot link flapping, see the [“Information About Link Flapping” section on page 8-2](#).

Information About 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.

For detailed information about port security, see the *Cisco Nexus 1000V Security Configuration Guide*.

Type of port	Is Port Security Supported?
vEthernet access	Yes
vEthernet trunk	Yes
vEthernet SPAN destination	No
Standalone Ethernet interfaces	No
Port channel members	No

To troubleshoot problems with port security, see the following:

- [“VM Cannot Ping a Secured Port” section on page 8-6](#)
- [“Port Security Violations” section on page 8-7](#)

Port Diagnostic Checklist

Use the following checklist to diagnose port interface activity.

For more information about port states, see the *Cisco Nexus 1000V Interface Configuration Guide*.

Table 8-1 Port Diagnostic Checklist

Checklist	Example	✓
Verify that the module is active. show module	See Example 8-1 on page 8-10 .	
Verify that the VSM is connected to the vCenter Server. show vsys connections	See Example 8-3 on page 8-10 .	
On the vSphere Client connected to vCenter Server, verify that required port profiles are assigned to the physical NICS and the virtual NICS.		
Verify that the ports have been created. show interface brief	See Example 8-8 on page 8-11 .	
Verify the state of the interface. show interface ethernet	See Example 8-10 on page 8-12 .	

Problems with Ports

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

- [Cannot Enable an Interface, page 8-4](#)
- [Port Link Failure or Port Not Connected, page 8-4](#)
- [Link Flapping, page 8-5](#)
- [Port ErrDisabled, page 8-5](#)
- [VM Cannot Ping a Secured Port, page 8-6](#)
- [Port Security Violations, page 8-7](#)

Cannot Enable an Interface

Use these guidelines to troubleshoot an interface that cannot be enabled.

Possible Cause	Solution
Layer 2 port is not associated with an access VLAN or the VLAN is suspended.	<ol style="list-style-type: none"> 1. Verify that the interface is configured in a VLAN. show interface brief 2. If not already, associate the interface with an access VLAN. 3. Determine the VLAN status. show vlan brief 4. If not already active, configure the VLAN as active. config t vlan vlan-id state active

Port Link Failure or Port Not Connected

Use these guidelines to troubleshoot a port that remains in link failure or not connected.

Table 8-2 Troubleshooting Ports in Link Failure or Not Connected

Possible Cause	Solution
Port connection is bad.	<ol style="list-style-type: none"> 1. Verify the port state. show system internal ethpm info 2. Disable and then enable the port. shut no shut 3. Move the connection to a different port on the same module or a different module. 4. Collect the ESX side NIC configuration. vss-support
Link is stuck in initialization state or the link is in a point-to-point state.	<ol style="list-style-type: none"> 1. Check for a link failure system message. Link Failure, Not Connected show logging 2. Disable and then enable the port. shut no shut 3. Move the connection to a different port on the same module or a different module. 4. Collect the ESX side NIC configuration. vss-support

Link Flapping

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

- Who initiated the link flap.
- The actual reason for the link being down.
- For a definition of link flapping, see the “[Link Flapping](#)” section on page 8-5.

Table 8-3 Troubleshooting link flapping

Possible Cause	Solution
The bit rate exceeds the threshold and puts the port into an error disabled state.	Disable and then enable the port. shut no shut The port should return to the normal state.
A hardware failure or intermittent hardware error causes a packet drop in the switch. A software error causes a packet drop. A control frame is erroneously sent to the device.	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. <ol style="list-style-type: none"> 1. Determine the reason for the link flap as indicated by the MAC driver. 2. Use the debug facilities on the end device to troubleshoot the problem.
ESX errors, or link flapping on the upstream switch.	Use the troubleshooting guidelines in the documentation for your ESX or upstream switch.

Port ErrDisabled

Use the guidelines in this section to troubleshoot ports that are error disabled.

Table 8-4 Troubleshooting error disabled ports

Possible Cause	Solution
Defective or damaged cable.	<ol style="list-style-type: none"> 1. Verify the physical cabling. 2. Replace or repair defective cables. 3. Re-enable the port. shut no shut

Table 8-4 Troubleshooting error disabled ports (continued)

Possible Cause	Solution
You attempted to add a port to a port channel that was not configured identically; and the port is then errdisabled.	<ol style="list-style-type: none"> 1. Display the switch log file and identify the exact configuration error in the list of port state changes. show logging logfile 2. Correct the error in the configuration and add the port to the port channel. 3. Re-enable the port. shut no shut
VSM application error	<ol style="list-style-type: none"> 1. Identify the component which errored while bringing up the port. show logging logfile grep interface_number See Example 8-7 on page 8-11. 2. Identify the error transition. show system internal ethpm event-history interface interface_number 3. Open a support case and submit the output of the above commands. For more information see the “Before Contacting Technical Support” section on page 28-1.

VM Cannot Ping a Secured Port

Use these troubleshooting guidelines when you cannot ping a secured port from a VM.

Table 8-5 Troubleshooting VM Cannot Ping a Secured Port

Possible Cause	Solution
The vEthernet interface is not up.	<ol style="list-style-type: none"> 1. Verify the state of the vEthernet interface. show interface vethernet number 2. If the interface is down, enable it. shut no shut
Drop on Source Miss (DSM) is set. New MAC addresses cannot be learned by this port.	<ol style="list-style-type: none"> 1. Verify the port security configuration. module vem 3 execute vemcmd show portsec stats 2. If DSM is set, clear the DSM bit on the VSM. no port-security stop learning

Table 8-5 Troubleshooting VM Cannot Ping a Secured Port (continued)

Possible Cause	Solution
The packet VLAN is not allowed on the port.	<ol style="list-style-type: none"> Identify the packet VLAN ID. show svs domain Verify that the packet VLAN is allowed on VEM uplink ports. show port-profile na uplink-all If the packet VLAN is not allowed on the uplink port profile, add it to the allowed VLAN list.
The packet VLAN is not allowed on the upstream switch port.	<ol style="list-style-type: none"> Identify the upstream neighbors connected to the interface. show cdp neighbors Log in to the upstream switch and verify that the packet VLAN is allowed on the port. show running-config interface gigabitEthernet slot/port If the packet VLAN is not allowed on the port, add it to the allowed VLAN list.

Port Security Violations

Use these troubleshooting guidelines when a vEthernet port is disabled because of a security violation. For detailed information about port security, see the *Cisco Nexus 1000V Security Configuration Guide*.

Table 8-6 Troubleshooting Port Security Violations

Possible Cause	Solution
The configured maximum number of secured addresses on the port is exceeded.	<ol style="list-style-type: none"> Display the secure addresses. show port -security address vethernet number show port-security Identify ports with a security violation. show logging inc "PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLATION_MAX_MAC_VLAN" Correct the security violation. Enable the interface. shut no shut
A MAC that is already secured on one port is then seen on another secure port.	

Port Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to ports.

Table 8-7 Port Troubleshooting Commands

Command	Purpose
<code>show module <i>module-number</i></code>	Displays the state of a module. See Example 8-1 on page 8-10 .
<code>show svcs domain</code>	Displays the domain configuration. See Example 8-2 on page 8-10 .
<code>show svcs connections</code>	Displays the Cisco Nexus 1000V connections. See Example 8-3 on page 8-10 .
<code>show cdp neighbors</code>	Displays the neighbors connected to an interface. See Example 8-4 on page 8-10 .
<code>show port internal event-history interface</code>	Displays information about the internal state transitions of the port. See Example 8-5 on page 8-11 .
<code>show logging logfile</code>	Displays logged system messages. See Example 8-6 on page 8-11 .
<code>show logging logfile lgrep <i>interface_number</i></code>	Displays logged system messages for a specified interface. See Example 8-7 on page 8-11 .
<code>show interface brief</code>	Displays a table of interface states. See Example 8-8 on page 8-11 .
<code>show interface ethernet</code>	Displays the configuration for a named Ethernet interface, including the following: <ul style="list-style-type: none"> • Administrative state • Speed • Trunk VLAN status • Number of frames sent and received • Transmission errors, including discards, errors, CRCs, and invalid frames See Example 8-9 on page 8-11 . See Example 8-10 on page 8-12 .
<code>show interface ethernet counters</code>	Displays port counters for identifying synchronization problems. For information about counters, see “ Information About Interface Counters ” section on page 8-2. See Example 8-11 on page 8-12 .

Table 8-7 Port Troubleshooting Commands (continued)

Command	Purpose
show interface vethernet	Displays the vEthernet interface configuration. See Example 8-12 on page 8-12 .
show interface status	Displays the status of the named interface.
show interface capabilities	Displays tabular view of all configured port profiles. See Example 8-13 on page 8-13 .
show interface virtual port mapping	Displays the virtual port mapping for all vEthernet interfaces: See Example 8-14 on page 8-14 .
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	
module vem execute vemcmd show portsec status	Displays the port security status of the port. If enabled, the output shows an LTL connected to the VM network adapter. See Example 8-15 on page 8-14
show port-security	Displays information about the secured MAC addresses in the system. See Example 8-16 on page 8-15 .
show port-security interface veth	Displays secure vEthernet interfaces.
show port -security address vethernet	Displays information about secure addresses on an interface. See Example 8-17 on page 8-15 .
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	

For detailed information about show command output, see the *Cisco Nexus 1000V Command Reference*.

EXAMPLES

Example 8-1 show module

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

Example 8-2 show svcs domain

```
n1000v# show svcs domain
SVS domain config:
  Domain id: 559
  Control vlan: 3002
  Packet vlan: 3003
  L2/L3 Aipc mode: L2
  L2/L3 Aipc interface: mgmt0
  Status: Config push to VC successful.
n1000v#
```

Example 8-3 show svcs connections

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

Example 8-4 show cdp neighbors

```
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
n1000v#
```

Example 8-5 show port internal event-history interface

```
n1000v# show port internal event-history interface e1/7
>>>>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]
```

Example 8-6 show logging logfile

```
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
n1000v#
```

Example 8-7 show logging logfile | grep interface_number

```
n1000v# show logging logfile | grep Vethernet3626
2011 Mar 25 10:56:03 n1k-bl %VIM-5-IF_ATTACHED: Interface Vethernet3626
is attached to Network Adapter 8 of gentoo-pxe-520 on port 193 of module
13 with dvport id 6899
2011 Mar 25 11:10:06 n1k-bl %ETHPORT-2-IF_SEQ_ERROR: Error ("Client data
inconsistency") while communicating with component MTS_SAP_ACLMGR for
opcode MTS_OPC_ETHPM_PORT_PRE_CFG (RID_PORT: Vethernet3626)
2011 Mar 25 11:10:06 n1k-bl %ETHPORT-2-IF_DOWN_ERROR_DISABLED: Interface
Vethernet3626 is down (Error disabled. Reason:Client data inconsistency)
```

Example 8-8 show interface brief

```
n1000v# show int brief
-----
Port VRF Status IP Address Speed MTU
-----
mgmt0 -- up 172.23.232.141 1000 1500
-----
Ethernet VLAN Type Mode Status Reason Speed Port
Interface Ch #
-----
Eth3/2 1 eth trunk up none 1000(D) --
Eth3/3 1 eth access up none 1000(D) --
n1000v#
```

Example 8-9 show interface ethernet

```
n1000v# show interface e1/14
e1/7 is down (errDisabled)
```

Example 8-10 show interface ethernet**Example:**

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

Example 8-11 show interface ethernet counters

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

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

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

Example 8-12 show interface vEthernet

```
n1000v# show interface veth1
Vethernet1 is up
  Port description is gentool, Network Adapter 1
  Hardware is Virtual, address is 0050.56bd.42f6
  Owner is VM "gentool", adapter is Network Adapter 1
  Active on module 33
  VMware DVS port 100
  Port-Profile is vlan48
  Port mode is access
  Rx
    491242 Input Packets 491180 Unicast Packets
    7 Multicast Packets 55 Broadcast Packets
    29488527 Bytes
  Tx
    504958 Output Packets 491181 Unicast Packets
    1 Multicast Packets 13776 Broadcast Packets 941 Flood Packets
    714925076 Bytes
    11 Input Packet Drops 0 Output Packet Drops
```

```
n1000v#
```

Example 8-13 show interface capabilities

```
n1000v# show interface capabilities
mgmt0
  Model:                --
  Type:                 --
  Speed:                10,100,1000,auto
  Duplex:               half/full/auto
  Trunk encap. type:    802.1Q
  Channel:              no
  Broadcast suppression: none
  Flowcontrol:          rx-(none),tx-(none)
  Rate mode:            none
  QOS scheduling:       rx-(none),tx-(none)
  CoS rewrite:          yes
  ToS rewrite:          yes
  SPAN:                 yes
  UDLD:                 yes
  Link Debounce:        no
  Link Debounce Time:   no
  MDIX:                 no
  Port Group Members:   none

port-channel1
  Model:                unavailable
  Type:                 unknown
  Speed:                10,100,1000,10000,auto
  Duplex:               half/full/auto
  Trunk encap. type:    802.1Q
  Channel:              yes
  Broadcast suppression: percentage(0-100)
  Flowcontrol:          rx-(off/on/desired),tx-(off/on/desired)
  Rate mode:            none
  QOS scheduling:       rx-(none),tx-(none)
  CoS rewrite:          yes
  ToS rewrite:          yes
  SPAN:                 yes
  UDLD:                 no
  Link Debounce:        no
  Link Debounce Time:   no
  MDIX:                 no
  Port Group Members:   none

port-channel2
  Model:                unavailable
  Type:                 unknown
  Speed:                10,100,1000,10000,auto
  Duplex:               half/full/auto
  Trunk encap. type:    802.1Q
  Channel:              yes
  Broadcast suppression: percentage(0-100)
  Flowcontrol:          rx-(off/on/desired),tx-(off/on/desired)
  Rate mode:            none
  QOS scheduling:       rx-(none),tx-(none)
  CoS rewrite:          yes
  ToS rewrite:          yes
  SPAN:                 yes
  UDLD:                 no
  Link Debounce:        no
  Link Debounce Time:   no
  MDIX:                 no
```

```

Port Group Members:    none

port-channel12
Model:                 unavailable
Type:                 unknown
Speed:                10,100,1000,10000,auto
Duplex:               half/full/auto
Trunk encap. type:    802.1Q
Channel:              yes
Broadcast suppression: percentage(0-100)
Flowcontrol:          rx-(off/on/desired),tx-(off/on/desired)
Rate mode:            none
QOS scheduling:       rx-(none),tx-(none)
CoS rewrite:          yes
ToS rewrite:          yes
SPAN:                 yes
UDLD:                 no
Link Debounce:        no
Link Debounce Time:   no
MDIX:                 no
Port Group Members:   none

control0
Model:                --
Type:                 --
Speed:                10,100,1000,auto
Duplex:               half/full/auto
Trunk encap. type:    802.1Q
Channel:              no
Broadcast suppression: none
Flowcontrol:          rx-(none),tx-(none)
Rate mode:            none
QOS scheduling:       rx-(none),tx-(none)
CoS rewrite:          yes
ToS rewrite:          yes
SPAN:                 yes
UDLD:                 yes
Link Debounce:        no
Link Debounce Time:   no
MDIX:                 no
Port Group Members:   none

n1000v#

```

Example 8-14 show interface virtual port-mapping

```
n1000v# show interface virtual port-mapping
```

```

-----
Port      Hypervisor Port      Binding Type      Status      Reason
-----
Veth1     DVPort5747           static            up          none
Veth2     DVPort3361           static            up          none
n1000v#

```

Example 8-15 module vem execute vemcmd show portsec status

```

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

```


n1000V#

Example 8-16 show port security

```
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
=====
```

Example 8-17 show port security address interface vethernet

```
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
=====
```




Port Profiles

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

- [Information About Port Profiles, page 9-1](#)
- [Problems with Port Profiles, page 9-2](#)
- [Port Profile Logs, page 9-5](#)
- [Port Profile Troubleshooting Commands, page 9-5](#)

Information About 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 are 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.

vEthernet port profiles can be assigned by the server administrator to physical ports (a VMNIC or a PNIC). Port profiles not configured as vEthernet 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 to physical or virtual ports, see your VMware documentation.

To verify that the profiles are assigned as expected to physical or virtual ports, use the following show commands:

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

To verify port profile inheritance, use the following command:

- **show running-config interface** *interface-id*

**Note**

Inherited port profiles cannot be changed or removed from an interface from the Cisco Nexus 1000V CLI. This can only be done from vCenter Server.

**Note**

Inherited port profiles are automatically configured by the Cisco Nexus 1000V 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*.

Problems with Port Profiles

The following are symptoms, possible causes, and solutions for problems with port profiles.

Table 9-1 Problems with Port Profiles

Symptom	Possible Causes	Solution
You do not see the port group on vCenter Server or the following message is displayed: Warning: Operation succeeded locally but update failed on vCenter server. Please check if you are connected to vCenter Server.	The connection to vCenter server is down.	<ol style="list-style-type: none"> 1. Verify that the connection to the vCenter Server is Enabled and Connected. show svcs connections 2. Reconnect to vCenter server. For detailed instructions, see the <i>Connecting to vCenter Server</i> procedure in the <i>Cisco Nexus 1000V System Management Configuration Guide</i>.
	The domain configuration was not successfully pushed to vCenter server.	<ol style="list-style-type: none"> 1. Verify that the domain configuration was successfully pushed to vCenter Server. show svcs domain 2. Fix any problems with the domain configuration. For information about configuring the domain, see the <i>Cisco Nexus 1000V System Management Configuration Guide</i>.
	The port profile is configured incorrectly.	<ol style="list-style-type: none"> 1. Verify that the vmware port-group is configured for the port profile and that the port profile is enabled. show port profile name <i>name</i> 2. Fix the port profile using the procedures in the <i>Cisco Nexus 1000V Port Profile Configuration Guide</i>.

Table 9-1 Problems with Port Profiles (continued)

Symptom	Possible Causes	Solution
<p>A port configuration is not applied to an interface.</p>	<p>Management connectivity between the vCenter server and the VSM has prevented the port profile assignment from being sent or received.</p>	<ol style="list-style-type: none"> 1. Display the port profile usage by interface. show port-profile virtual usage 2. Verify that the interface level configuration did not overwrite the port profile configuration. show run show port-profile expand-interface 3. If the show command output is incorrect, then on vCenter server, reassign the port group to the interface.
<p>An Ethernet interface or vEthernet interface is administratively down.</p> <p>A system message similar to the following is logged:</p> <pre>%VMS-3-DVPG_NICS_MOVED: '1' nics have been moved from port-group 'Access483' to 'Unused_Or_Quarantine_Veth'.</pre>	<p>The interface is inheriting a quarantined port profile.</p> <p>A configuration was not saved prior to rebooting the VSM, the configuration was lost, and the interfaces were moved to one of the following port profiles:</p> <ul style="list-style-type: none"> • Unused_Or_Quarantine_Uplink for ethernet types • Unused_Or_Quarantine_Veth for Vethernet types 	<ol style="list-style-type: none"> 1. Verify the port profile-to-interface mapping. show port-profile virtual usage 2. Reassign the VMNIC or PNIC to a non-quarantined port group to enable the interface to be up and forwarding traffic. This requires changing the port group on vCenter Server.
<p>After applying a port profile, an online interface is quarantined.</p> <p>A system message similar to the following is logged:</p> <pre>%PORT-PROFILE-2-INTERFACE_QUARANTINED: Interface Ethernet3/3 has been quarantined due to Cache Overrun</pre>	<p>The assigned port profile is incorrectly configured. The incorrect command fails when the port profile is applied to an interface.</p> <p>Although a specific command fails, the port profile-to-interface mapping is created.</p>	<ol style="list-style-type: none"> 1. Identify the command that failed. show accounting log grep FAILURE 2. Verify the interface is quarantined. show port-profile sync-status 3. Verify the port profile-to-interface mapping. show port-profile virtual usage 4. Fix the error in the port profile using the procedures in the <i>Cisco Nexus 1000V Port Profile Configuration Guide</i>. 5. Bring the interface out of quarantine. no shutdown The interface comes back online. 6. Return shutdown control to the port-profile. default shutdown

Table 9-1 Problems with Port Profiles (continued)

Symptom	Possible Causes	Solution
<p>After modifying a port profile, an assigned offline interface is quarantined.</p> <p>A system message similar to the following is logged:</p> <pre>%PORT-PROFILE-2-INTERFACE_QUARANTINED: Interface Ethernet4/3 has been quarantined due to Cache Overrun</pre>	<p>The interface has been removed from the DVS.</p>	<p>To bring the interface back online, use the “Recovering a Quarantined Offline Interface” procedure on page 9-4.</p>
<p>A module and all associated interfaces are offline.</p> <p>A system message similar to the following is logged:</p> <pre>2011 Mar 2 22:28:50 n1000v %VEM_MGR-2-VEM_MGR_REMOVE_NO_HB: Removing VEM 3 (heartbeats lost) 2011 Mar 2 22:29:00 n1000v %VEM_MGR-2-MOD_OFFLINE: Module 3 is offline</pre>	<p>The interface carrying system VLANs for the module has gone down for one of the following reasons:</p> <ul style="list-style-type: none"> • System interfaces were removed from the DVS on the vCenter Server. • The module was powered down. • There is general loss of connectivity to the module. 	<p>Follow VEM troubleshooting guide to bring module back online</p> <p>To bring the interface back online, use the “Recovering a Quarantined Offline Interface” procedure on page 9-4.</p>

Recovering a Quarantined Offline Interface

You can use this procedure to recover and bring online an interface that is offline and has been quarantined.

BEFORE YOU BEGIN

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

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

DETAILED STEPS

-
- Step 1** Verify the interface has is quarantined. The interface appears in the show command output.
- show port-profile sync-status**
- Step 2** On the vCenter server, add or associate the PNIC to a port profile (either the original port profile or a different port profile).
- The interface comes back online.
- Step 3** Verify that the interface has come back online.
- show interface brief**
- Step 4** Verify the port profile-to-interface mapping.
- show port-profile virtual usage**

Step 5 Verify the interface has come out of quarantine automatically. The interface should no longer appear in the show command output.

show port-profile sync-status

Step 6 Return shutdown control to the port-profile.

default shutdown

Port Profile Logs

To enable and collect detailed logs for port profiles, use the following commands:

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

After enabling the debug log, the results of any subsequent port profile configuration are captured in the log file.

Port Profile Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to port profiles.

Command	Purpose
show port-profile	Displays the port profile configuration. See Example 9-1 on page 9-6 .
show port-profile name <i>name</i>	Displays the configuration for a named port profile. See Example 9-2 on page 9-7 .
show port-profile brief	Displays tabular view of all configured port profiles. See Example 9-3 on page 9-8 .
show port-profile expand-interface	Displays all configured port profiles expanded to include the interfaces assigned to them. See Example 9-4 on page 9-8 .
show port-profile expand-interface name <i>name</i>	Displays a named port profile expanded to include the interfaces assigned to it. See Example 9-5 on page 9-10 .

Command	Purpose
show port-profile-role [name <i>port-profile-role-name</i>]	Displays the port profile role configuration, including role names, descriptions, assigned users, and assigned groups. See Example 9-7 on page 9-11 .
show running-config port-profile [<i>profile-name</i>]	Displays the port profile configuration. See Example 9-6 on page 9-11 .
show port-profile-role	Displays the port profile role configuration. See Example 9-7 on page 9-11 .
show port-profile-role users	Displays available users and groups. See Example 9-8 on page 9-11 .
show port-profile sync-status [interface <i>if-name</i>]	Displays interfaces that are out of sync with the port profile. See Example 9-9 on page 9-11 .
show port-profile virtual usage [name <i>profile-name</i>]	Displays the port profile usage by interface. See Example 9-10 on page 9-12 .
show msp internal info	Displays port profile mappings on vCenter server and configured roles. See Example 9-11 on page 9-12 .
show system internal port-profile profile-fsm	Displays port profile activity on the Cisco Nexus 1000V, including transitions such as inherits and configurations. If the following displays, then all inherits are processed: Curr state: [PPM_PROFILE_ST_SIDLE] See Example 9-12 on page 9-16 .
show system internal port-profile event-history msgs	Displays the messages logged about port profile events within the Cisco Nexus 1000V. See Example 9-13 on page 9-16 .

For detailed information about show command output, see the *Cisco Nexus 1000V Command Reference*.

EXAMPLES

Example 9-1 show port-profile

```
n1000v# show port-profile
port-profile vEthProfile1
  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 mac-pinning
evaluated config attributes:
  channel-group auto mode on mac-pinning
assigned interfaces:
port-profile vEthProfile2
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 vEthProfile3
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#

```

Example 9-2 show port-profile name

```

n1000v# show port-profile name vEthProfile3
port-profile vEthProfile3
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#

```

Example 9-3 show port-profile brief

```
n1000v# show port-profile brief
-----
Port Profile           Profile Type   Profile State  Conf Items  Eval Items  Assigned Intfs  Child Profs
-----
AccessProf            Vethernet 0       0       0       0       0       0       0
PP1027                Vethernet 1       0       0       0       0       0
PP1028                Vethernet 1       0       0       0       0       0
Unused_Or_Quarantine_Uplink  Ethernet 1       1       0       0       0       0
Unused_Or_Quarantine_Veth  Vethernet 1       1       0       0       0       0
accessprof            Vethernet 0       3       3       0       0       0
portpl                Vethernet 0       0       0       0       0       0
-----
Profile   Total
Type      Intfs
-----
Vethernet 8
Ethernet  10
n1000v#
```

Example 9-4 show port-profile expand-interface

```
n1000v# show port-profile expand-interface
port-profile AccessProf
  id: 1
  capability: 0x0
  state: 0x0
  type: 0x0
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 32
  used ports: 0
  vmware config information
    pg name: AccessProf
    dvs: (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
port-profile PP1027
  id: 4
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 1
  max ports: 64
  min ports: 32
  used ports: 0
  vmware config information
    pg name: PP1027
    dvs: (ignore)
    reserved ports: 16
  port-profile role:
  alias information:
    pg id: PP1027
    dvs uuid:
```

```
    type: 1
    pg id: dvportgroup-3180
    dvs uuid: 12 98 0e 50 6b 78 6f c5-74 af b2 3a 16 6e 45 10
    type: 2
port-profile PP1028
  id: 3
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: dynamic
  bind_opts: 1
  max ports: 64
  min ports: 32
  used ports: 0
  vmware config information
    pg name: PP1028
    dvs: (ignore)
    reserved ports: 16
  port-profile role:
  alias information:
    pg id: PP1028
    dvs uuid:
    type: 1
    pg id: dvportgroup-3181
    dvs uuid: 12 98 0e 50 6b 78 6f c5-74 af b2 3a 16 6e 45 10
    type: 2
port-profile Unused_Or_Quarantine_Uplink
  id: 6
  capability: 0x1
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: dynamic
  bind_opts: 1
  max ports: 32
  min ports: 32
  used ports: 0
  vmware config information
    pg name: Unused_Or_Quarantine_Uplink
    dvs: (ignore)
    reserved ports: 0
  port-profile role:
  alias information:
    pg id: Unused_Or_Quarantine_Uplink
    dvs uuid:
    type: 1
    pg id: dvportgroup-3182
    dvs uuid: 12 98 0e 50 6b 78 6f c5-74 af b2 3a 16 6e 45 10
    type: 2
port-profile Unused_Or_Quarantine_Veth
  id: 7
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: dynamic
  bind_opts: 1
  max ports: 32
  min ports: 32
  used ports: 0
```

```

vmware config information
  pg name: Unused_Or_Quarantine_Veth
  dvs: (ignore)
  reserved ports: 32
port-profile role:
alias information:
  pg id: Unused_Or_Quarantine_Veth
  dvs uuid:
  type: 1
  pg id: dvportgroup-3183
  dvs uuid: 12 98 0e 50 6b 78 6f c5-74 af b2 3a 16 6e 45 10
  type: 2
port-profile accessprof
  id: 5
  capability: 0x0
  state: 0x0
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: dynamic
  bind_opts: 1
  max ports: 32
  min ports: 32
  used ports: 0
  vmware config information
    pg name: accessprof
    dvs: (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
    pg id: accessprof
    dvs uuid:
    type: 1
port-profile portp1
  id: 2
  capability: 0x0
  state: 0x0
  type: 0x0
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 32
  used ports: 0
  vmware config information
    pg name: portp1
    dvs: (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
pending binds:
n1000v#

```

Example 9-5 *show port-profile expand-interface name UplinkProfile1*

```

n1000v# show port-profile expand-interface name UplinkProfile1
port-profile EthProfile1
Ethernet2/2
  switchport mode trunk
  switchport trunk allowed vlan 110-119
  no shutdown
n1000v#

```

Example 9-6 show running-config port-profile

```
n1000v# show running-config port-profile
port-profile type ethernet UplinkProfile1
  description "Profile for critical system ports"
  vmware port-group
  switchport mode access
  switchport access vlan 113
  switchport trunk native vlan 113
  channel-group auto mode on
  no shutdown
port-profile type vethernet vEthProfile2
  vmware port-group
  vmware max-ports 5
  switchport mode trunk
  switchport trunk native vlan 112
  channel-group auto mode on sub-group cdp
  no shutdown
n1000v#
```

Example 9-7 show port-profile-role

```
n1000v# show port-profile-role name adminUser

Name: adminUser
Description: adminOnly
Users:
  hdbaar (user)
Assigned port-profiles:
  allaccess2
n1000v#
```

Example 9-8 show port-profile-role users

```
switch# show port-profile-role users
Groups:
  Administrators
  TestGroupB
Users:
  hdbaar
  fgreen
  suchen
  mariofr
switch#
```

Example 9-9 show port-profile sync-status

```
n1000v# show port-profile sync-status interface ethernet 3/2
Ethernet3/2
  port-profile: uplink
  interface status: quarantine
  sync status: out of sync
  cached commands:
  errors:
    command cache overrun
  recovery steps:
    bring interface online
n1000v#
```

Example 9-10 show port-profile virtual usage

```
n1000v# show port-profile virtual usage
-----
Port Profile          Port          Adapter      Owner
-----
n1kv-uplink0         Po1
                    Eth3/2        vmn1c1       localhost.
                    Eth3/3        vmn1c2       localhost.
vlan1767              Veth7         Net Adapter 1 all-tool-7
                    Veth8         Net Adapter 1 all-tool-8
aicp1765              Veth4         Net Adapter 1 bl-h-s
inband1766           Veth6         Net Adapter 3 bl-h-s
mgmt1764              Veth5         Net Adapter 2 bl-h-s
vpc-mac-uplink       Po7
                    Eth5/2        vmn1c1       localhost.
                    Eth5/3        vmn1c2       localhost.
ch-vpc-mac-uplink    Po2
                    Po3
                    Eth4/2        vmn1c1       VDANIKLNCOS
                    Eth4/3        vmn1c2       VDANIKLNCOS
ch-aicp1765          Veth1         Net Adapter 1 bl-h-p
ch-mgmt1764          Veth2         Net Adapter 2 bl-h-p
ch-inband1766        Veth3         Net Adapter 3 bl-h-p
n1000v#
```

Example 9-11 show msp internal info

```
n1000v# show msp internal info
port-profile Access484
  id: 5
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  max ports: 256
  vmware config information
    pg name: Access484
    dvs: (ignore)
  port-profile role:
  alias information:
    pg id: Access484
    dvs uuid:
    type: 1
    pg id: dvportgroup-3285
    dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
    type: 2
    pg id: dvportgroup-3292
    dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
    type: 2
port-profile Unused_Or_Quarantine_Uplink
  id: 1
  capability: 0x1
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  max ports: 32
  vmware config information
    pg name: Unused_Or_Quarantine_Uplink
```

```

    dvs: (ignore)
port-profile role:
alias information:
  pg id: Unused_Or_Quarantine_Uplink
  dvs uuid:
    type: 1
  pg id: dvportgroup-2444
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile Unused_Or_Quarantine_Veth
id: 2
capability: 0x0
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 32
vmware config information
  pg name: Unused_Or_Quarantine_Veth
  dvs: (ignore)
port-profile role:
alias information:
  pg id: Unused_Or_Quarantine_Veth
  dvs uuid:
    type: 1
  pg id: dvportgroup-2445
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile eth-break-deinherit
id: 10
capability: 0x1
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 32
vmware config information
  pg name: eth-break-deinherit
  dvs: (ignore)
port-profile role:
alias information:
  pg id: eth-break-deinherit
  dvs uuid:
    type: 1
  pg id: dvportgroup-3286
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3293
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile eth-break-inherit
id: 9
capability: 0x1
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 32
vmware config information
  pg name: eth-break-inherit
  dvs: (ignore)

```

```

port-profile role:
alias information:
  pg id: eth-break-inherit
  dvs uuid:
  type: 1
  pg id: dvportgroup-3287
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3294
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile uplink
id: 3
capability: 0x3
state: 0x1
type: 0x1
system vlan mode: trunk
system vlans: 480-481
port-binding: static
max ports: 32
vmware config information
  pg name: uplink
  dvs: (ignore)
port-profile role:
alias information:
  pg id: uplink
  dvs uuid:
  type: 1
  pg id: dvportgroup-3283
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile uplink-quar
id: 12
capability: 0x1
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 32
vmware config information
  pg name: uplink-quar
  dvs: (ignore)
port-profile role:
alias information:
  pg id: uplink-quar
  dvs uuid:
  type: 1
  pg id: dvportgroup-3288
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3295
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile veth-break-deinherit
id: 8
capability: 0x0
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 256
vmware config information

```



```

    pg name: veth-break-deinherit
    dvs: (ignore)
port-profile role:
alias information:
  pg id: veth-break-deinherit
  dvs uuid:
  type: 1
  pg id: dvportgroup-3289
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3296
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile veth-break-inherit
id: 7
capability: 0x0
state: 0x1
type: 0x1
system vlan mode: -
system vlans:
port-binding: static
max ports: 256
vmware config information
  pg name: veth-break-inherit
  dvs: (ignore)
port-profile role:
alias information:
  pg id: veth-break-inherit
  dvs uuid:
  type: 1
  pg id: dvportgroup-3290
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3297
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
port-profile vpc-uplink
id: 6
capability: 0x3
state: 0x1
type: 0x1
system vlan mode: trunk
system vlans: 480-481
port-binding: static
max ports: 32
vmware config information
  pg name: vpc-uplink
  dvs: (ignore)
port-profile role:
alias information:
  pg id: vpc-uplink
  dvs uuid:
  type: 1
  pg id: dvportgroup-3291
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
  pg id: dvportgroup-3298
  dvs uuid: 44 dc 3b 50 53 11 b7 ac-ef ed ef 46 ee df c2 d5
  type: 2
pending binds:
port-profile-role adfd
id: 0
desc:
num users: 1

```

```
group GROUP

n1000v#
```

Example 9-12 show system internal port-profile profile-fsm

```
n1000v# show system internal port-profile profile-fsm
>>>FSM: <PROFILE_FSM:1> has 4 logged transitions<<<<<

1) FSM:<PROFILE_FSM:1> Transition at 856903 usecs after Tue Mar  8 19:11:47 2011
   Previous state: [PPM_PROFILE_ST_SIDLE]
   Triggered event: [PPM_PROFILE_EV_EIF_STATUS_CHANGE]
   Next state: [PPM_PROFILE_ST_SIDLE]

2) FSM:<PROFILE_FSM:1> Transition at 858442 usecs after Tue Mar  8 19:11:47 2011
   Previous state: [PPM_PROFILE_ST_SIDLE]
   Triggered event: [PPM_PROFILE_EV_ELEARN]
   Next state: [PPM_PROFILE_ST_SIF_CREATE]

3) FSM:<PROFILE_FSM:1> Transition at 842710 usecs after Tue Mar  8 19:12:04 2011
   Previous state: [PPM_PROFILE_ST_SIF_CREATE]
   Triggered event: [PPM_PROFILE_EV_EACKNOWLEDGE]
   Next state: [FSM_ST_NO_CHANGE]

4) FSM:<PROFILE_FSM:1> Transition at 873872 usecs after Tue Mar  8 19:12:04 2011
   Previous state: [PPM_PROFILE_ST_SIF_CREATE]
   Triggered event: [PPM_PROFILE_EV_ESUCCESS]
   Next state: [PPM_PROFILE_ST_SIDLE]

   Curr state: [PPM_PROFILE_ST_SIDLE]
n1000v#
```

Example 9-13 show system internal port-profile event-history msgs

```
n1000v# show system internal port-profile event-history msgs
1) Event:E_MTS_RX, length:60, at 538337 usecs after Tue Mar  8 19:13:02 2011
   [NOT] Opc:MTS_OPC_IM_IF_CREATED(62467), Id:0X0000B814, Ret:SUCCESS
   Src:0x00000101/175, Dst:0x00000101/0, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:120
   Payload:
   0x0000:  00 00 00 02 00 00 00 02 00 00 00 0c 00 00 00 29

2) Event:E_MTS_RX, length:60, at 515030 usecs after Tue Mar  8 19:13:02 2011
   [NOT] Opc:MTS_OPC_LC_ONLINE(1084), Id:0X0000B7E8, Ret:SUCCESS
   Src:0x00000101/744, Dst:0x00000101/0, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:234
   Payload:
   0x0000:  02 00 00 03 00 00 00 00 00 00 03 02 03 02 00 00

3) Event:E_MTS_RX, length:60, at 624319 usecs after Tue Mar  8 19:12:05 2011
   [NOT] Opc:MTS_OPC_PPM_INTERFACE_UPDATE(152601), Id:0X00003908, Ret:SUCCESS
   Src:0x00000101/489, Dst:0x00000101/0, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:107
   Payload:
   0x0000:  00 00 00 02 00 00 00 02 00 00 00 0c 00 00 00 26

4) Event:E_MTS_RX, length:60, at 624180 usecs after Tue Mar  8 19:12:05 2011
   [NOT] Opc:MTS_OPC_PPM_INTERFACE_UPDATE(152601), Id:0X00003905, Ret:SUCCESS
   Src:0x00000101/489, Dst:0x00000101/0, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:107
   Payload:
   0x0000:  00 00 00 02 00 00 00 02 00 00 00 0c 00 00 00 26
```

```
5) Event:E_MTS_RX, length:60, at 624041 usecs after Tue Mar 8 19:12:05 2011
   [NOT] Opc:MTS_OPC_PPM_INTERFACE_UPDATE(152601), Id:0X00003903, Ret:SUCCESS
   Src:0x00000101/489, Dst:0x00000101/0, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:107
   Payload:
   0x0000: 00 00 00 02 00 00 00 02 00 00 00 0c 00 00 00 26
   ...
```




Port Channels and Trunking

Use this chapter to troubleshoot port channels and trunking.

This chapter includes the following topics:

- [Overview, page 10-1](#)
- [Initial Troubleshooting Checklist, page 10-2](#)
- [Troubleshooting Asymmetric Port Channels, page 10-3](#)
- [Cannot Create Port Channel, page 10-4](#)
- [Newly Added Interface Does Not Come Online In a Port Channel, page 10-4](#)
- [VLAN Traffic Does Not Traverse Trunk, page 10-5](#)

Overview

This section includes the following topics:

- [Port Channel Overview, page 10-1](#)
- [Trunking Overview, page 10-2](#)

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.

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) feature in Nexus 1000V, 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***

- **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 10-1 shows output of the **show port-channel summary** command.

Example 10-1 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.

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 10-4 , 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.

Example:

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

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.



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 11-1](#)
- [Port Model, page 11-1](#)
- [Layer 2 Switching Problems, page 11-4](#)
- [Verifying Layer 2 Switching, page 11-7](#)
- [Troubleshooting Microsoft NLB Unicast Mode, page 11-12](#)

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

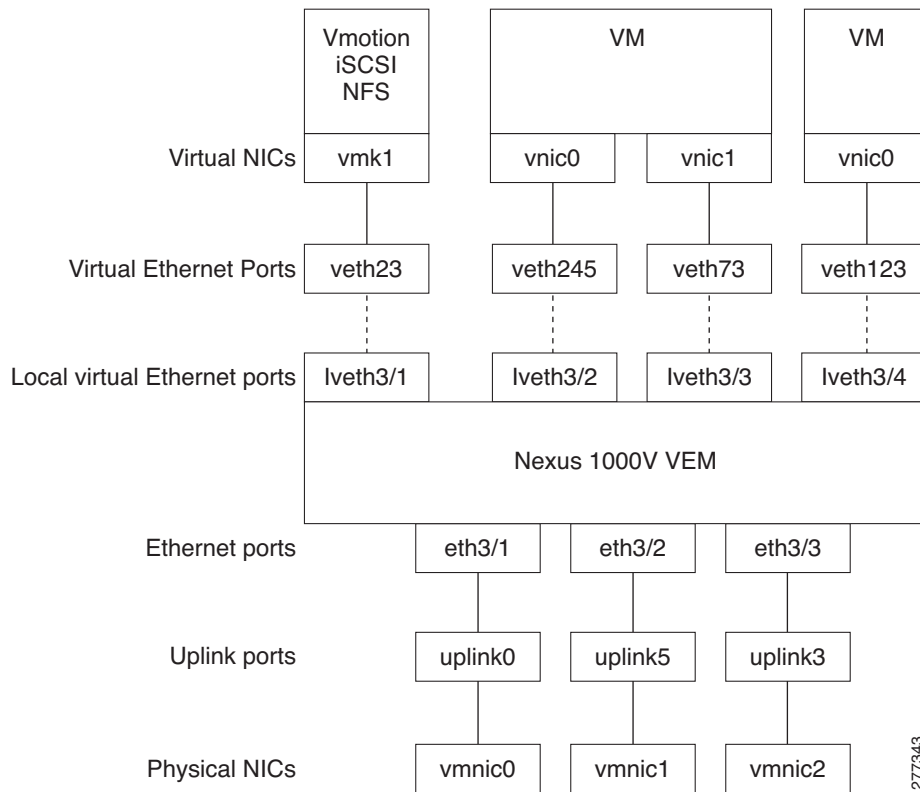
This section describes the following port perspectives:

- [Viewing Ports from the VEM, page 11-2](#)
- [Viewing Ports from the VSM, page 11-3.](#)

Viewing Ports from the VEM

The Nexus1000V differentiates between virtual and physical ports on each of the VEMs. [Figure 11-1](#) shows how ports on the Nexus1000V switch are bound to physical and virtual VMware ports within a VEM.

Figure 11-1 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.

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 11-2 shows the VSM view ports.

Figure 11-2 VSM View of Ports



277344

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*.

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 11-4](#)
- [Verifying a Connection Between VEMs, page 11-5](#)
- [Isolating Traffic Interruptions, page 11-6](#)
- [Verifying Layer 2 Switching, page 11-7](#)

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.

- 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 1 for forwarding (active).
- 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
```

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
```


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 11-1 on page 11-8
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 11-2 on page 11-9
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 mac address-table static inc veth	Displays the static MAC address of vEthernet interfaces in case a VEM physical port learns a dynamic MAC and the packet source is in another VEM on the same VSM. See Example 11-3 on page 11-9
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 11-4 on page 11-9 .
show vlan summary	Displays a summary of VLAN information.

Command	Purpose
show interface brief	Displays a table of interface states. See Example 11-5 on page 11-10 .
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 11-6 on page 11-10 .
module vem <i>module-number</i> execute vemcmd show bd command	For the specified VEM, displays its VLANs and their port lists. See Example 11-7 on page 11-11 .
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 forwarding (active) on a port, then its CBL state should be 1. • If a VLAN is blocked, then its CBL state is 0. See Example 11-8 on page 11-11 .
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 11-9 on page 11-11 .
show interface <i>interface_id</i> mac	Displays the MAC addresses and the burn-in MAC address for an interface.

Example 11-1 show mac address-table

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



Tip Module indicates the VEM on which 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
VLAN      MAC Address      Type      Age      Port              Module
-----+-----+-----+-----+-----+-----
1         0002.3d11.5502   static    0        N1KV Internal Port  3
1         0002.3d21.5500   static    0        N1KV Internal Port  3
1         0002.3d21.5502   static    0        N1KV Internal Port  3
1         0002.3d31.5502   static    0        N1KV Internal Port  3
1         0002.3d41.5502   static    0        N1KV Internal Port  3
1         0002.3d61.5500   static    0        N1KV Internal Port  3
1         0002.3d61.5502   static    0        N1KV Internal Port  3
1         0002.3d81.5502   static    0        N1KV Internal Port  3
3         12ab.47dd.ff89   static    0        Eth3/3             3
342      0002.3d41.5502   static    0        N1KV Internal Port  3
342      0050.568d.5a3f   dynamic   0        Eth3/3             3
343      0002.3d21.5502   static    0        N1KV Internal Port  3
```

```

343      0050.568d.2aa0    dynamic 9          Eth3/3          3
Total MAC Addresses: 13
n1000v#

```

Example 11-2 show mac address-table address



Tip This command shows all interfaces on which a MAC is learned dynamically. In this example, the same MAC appears on Eth3/3 and Eth4/3.

```

n1000v# show mac address-table address 0050.568d.5a3f
VLAN      MAC Address      Type   Age      Port          Module
-----+-----+-----+-----+-----+-----
342      0050.568d.5a3f  dynamic 0          Eth3/3          3
342      0050.568d.5a3f  dynamic 0          Eth4/3          4
Total MAC Addresses: 1
n1000v#

```

Example 11-3 show mac address-table static | inc veth

```

n1000v# show mac address-table static | inc veth
460      0050.5678.ed16    static 0          Veth2          3
460      0050.567b.1864    static 0          Veth1          4
n1000v#

```

Example 11-4 show vlan



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
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 11-5 show interface brief

```

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 11-6 module vem module-number execute vemcmd show port

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    1 Access 120
  9     0     3969   0     2     2     VIRT  UP    UP    1 Access 121
 10     0     115    0     2     0     VIRT  UP    UP    1 Access 122
 11     0     3968   0     2     2     VIRT  UP    UP    1 Access 123
 12     0     116    0     2     0     VIRT  UP    UP    1 Access 124
 13     0     1      0     2     2     VIRT  UP    UP    0 Access 125
 14     0     3967   0     2     2     VIRT  UP    UP    1 Access 126
 16  1a030100  1 T    0     0     2     PHYS  UP    UP    1 Trunk
vmnic1
 17  1a030200  1 T    0     2     2     PHYS  UP    UP    1 Trunk
vmnic2

```

Example 11-7 *module vem module-number execute vemcmd show bd*

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 11-8 *module vem module-number execute vemcmd show trunk*

Tip If a VLAN is active on a port, then its CBL state should be 1.
If a VLAN is blocked, then its CBL state is 0.

```

~ # module vem 5 execute vemcmd show trunk
Trunk port 16 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(110) cbl 1, vlan(111) cbl 1, vlan(112) cbl 1, vlan(113) cbl 1,
vlan(114) cbl 1,vlan(115) cbl 1, vlan(116) cbl 1, vlan(117) cbl 1, vlan(118) cbl 1,
vlan(119) cbl 1,
Trunk port 17 native_vlan 1 CBL 0
vlan(1) cbl 1, vlan(117) cbl 1,
~ #

```

Example 11-9 *module vem module-number execute vemcmd show l2*

```

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

```

Troubleshooting Microsoft NLB Unicast Mode

Microsoft Network Load Balancing (MS-NLB) is a clustering technology offered by Microsoft as part of the Windows server operating systems. Clustering enables a group of independent servers to be managed as a single system for higher availability, easier manageability, and greater scalability.

For more information about Microsoft Network Load Balancing. See this URL:

<http://technet.microsoft.com/en-us/library/bb742455.aspx>



Note

Access to third-party websites identified in this document is provided solely as a courtesy to customers and others. Cisco Systems, Inc. and its affiliates are not in any way responsible or liable for the functioning of any third-party website, or the download, performance, quality, functioning or support of any software program or other item accessed through the website, or any damages, repairs, corrections or costs arising out of any use of the website or any software program or other item accessed through the website. Cisco's End User License Agreement does not apply to the terms and conditions of use of a third-party website or any software program or other item accessed through the website.

Limitations and Restrictions

A syslog is generated if one of the following configurations exists when you try to disable automatic static MAC learning for MS-NLB because they do not support this feature:

- PVLAN port
- Ports configured with unknown unicast flood blocking (UUFb)
- Ports configured with switchport port-security mac-address sticky

Disabling Automatic Static MAC Learning on vEthernet

You must disable automatic static MAC learning before you can successfully configure NLB on vEthernet (vEth).

In interface configuration mode use the following commands:

```
switch(config)# int veth 1
switch(config-if)# no mac auto-static-learn
```

In port profile configuration mode use the following commands:

```
switch(config)# port-profile type vethernet ms-nlb
switch(config-port-prof)# no mac auto-static-learn
```

Checking Status on a VSM

If the NLB unicast mode configuration does not function, check the status of the Virtual Supervisor Module (VSM).

Confirm that **no mac auto-static-learn** is listed in the vEth and/or port profile configurations.

In interface configuration mode use the following command to generate the VSM status:

switch(config-if)# show running-config int veth1

```
interface Vethernet1
  inherit port-profile vm59
  description Fedora117, Network Adapter 2
  no mac auto-static-learn
  vmware dvport 32 dvswitch uuid "ea 5c 3b 50 cd 00 9f 55-41 a3 2d 61 84 9e 0e c4"
```

In port profile configuration mode use the following command to generate the VSM status:

switch(config-if)# show running-config port-profile ms-nlb

```
port-profile type vethernet ms-nlb
  vmware port-group
  switchport mode access
  switchport access vlan 59
  no mac auto-static-learn
  no shutdown
  state enabled
```

Checking Status on a VEM

If the NLB unicast mode configuration does not function, check the status of the Virtual Ethernet Module (VEM). Check the following:

- Confirm that MS-NLB veths are disabled.
- Confirm that MS-NLB shared-MAC (starting with 02:BF) is not listed in the layer 2 (L2) MAC table

Use the following command to generate the VEM status:

~ # vemcmd show port auto-smac-learning

```
LTL   VSM Port  Auto Static MAC Learning
49    Veth4    DISABLED
50    Veth5    DISABLED
51    Veth6    DISABLED
```

Use the following command to generate the L2 MAC table for VLAN59:

~ # vemcmd show l2 59

```
Bridge domain 15 brtmax 4096, brtcnt 6, timeout 300
VLAN 59, swbd 59, ""
Flags: P - PVLAN S - Secure D - Drop
      Type          MAC Address  LTL  timeout  Flags  PVLAN
Dynamic 00:15:5d:b4:d7:02 305    4
Dynamic 00:15:5d:b4:d7:04 305    25
Dynamic 00:50:56:b3:00:96  51     4
Dynamic 00:50:56:b3:00:94 305    5
Dynamic 00:0b:45:b6:e4:00 305    5
Dynamic 00:00:5e:00:01:0a  51     0
```

Configuring MS NLB for Multiple VM NICs in the Same Subnet

When MS NLB VMs have more than one port on the same subnet, a request is flooded, which causes both ports to receive it. The server cannot manage this situation.

To workaroud this situation, enable Unknown Unicast Flood Blocking (UUFB).

Enabling UUFB

To enable UUFB, enter these configuration commands, one on each line. At the end, enter Cntl + Z.

```
n1000v# configure terminal
n1000v (config)# uufb enable
n1000v (config)#
```

This configuration conceals the requests from the non-NLB ports and allows the system to function as it is expected.

Disabling UUFB for VMs that use Dynamic MAC Addresses

Issues might occur for VMs that use dynamic MAC addresses, other than those assigned by VMware. For ports hosting these types of VMs, disable UUFB. To disable UUFB, enter the following commands:

```
n1000v(config)# int veth3
n1000v(config-if)# switchport uufb disable
n1000v(config-if)#
```




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 12-1](#)
- [Initial Troubleshooting Checklist, page 12-2](#)
- [Cannot Create a VLAN, page 12-3](#)

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.

- 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**

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.



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 13-1](#)
- [Troubleshooting Guidelines, page 13-2](#)
- [Private VLAN Troubleshooting Commands, page 13-2](#)

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.

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*.

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
```

```

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											

```
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**



NetFlow

This chapter describes how to identify and resolve problems that relate to Netflow.

This chapter includes the following sections:

- [Information About NetFlow, page 14-1](#)
- [NetFlow Troubleshooting Commands, page 14-2](#)
- [Common NetFlow Problems, page 14-3](#)

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 Nexus 1000V flow record.

For detailed information about configuring NetFlow, see the *Cisco Nexus 1000V System Management Configuration Guide*.

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**

- **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.

- Issue the **show flow exporter** command to view statistics for the exporter and identify any drops on the VSM.



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 15-1](#)
- [ACL Configuration Limits, page 15-1](#)
- [ACL Restrictions, page 15-2](#)
- [Troubleshooting ACLs, page 15-2](#)
- [Displaying ACL Policies on the VEM, page 15-2](#)
- [Debugging Policy Verification Issues, page 15-3](#)
- [Troubleshooting ACL Logging, page 15-3](#)

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*.

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.

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
```

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.

Troubleshooting ACL Logging

This section includes the following topics:

- [Using the CLI to Troubleshoot ACL Logging on a VEM, page 15-4](#)
- [ACL Logging Troubleshooting Scenarios, page 15-5](#)

Using the CLI to Troubleshoot ACL Logging on a VEM

The commands in this section will help you troubleshoot ACL logging by examining ACL flows.

Viewing Current Flows

You can troubleshoot ACL logging by viewing the current flows on a VEM. Enter the following command:

```
vemcmd show aclflows stats
```

EXAMPLE

The following shows an example of the output when you enter this command:

```
[root@esx /]# vemcmd show aclflows stats
Current Flow stats:
  Permit Flows:          1647
  Deny Flows:           0
  Current New Flows:    419      --- current new flows yet to be reported.
```

Viewing Active Flows

You can display all the active flows on a VEM by entering the following command:

```
vemcmd show aclflows [permit | deny]
```

If you do not specify **permit** or **deny**, the command displays both.

EXAMPLE

The following shows an example of the output when you enter this command:

```
[root@esx /]# vemcmd show aclflows [permit | deny]
If      SrcIP      DstIP      SrcPort  DstPort  Proto  Direction  Action  Stats
Veth4   192.168.1.20  192.168.1.10  5345    8080     6      Ingress    permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5769     6      Egress     permit  1
Veth4   192.168.1.20  192.168.1.10  6256    8080     6      Ingress    permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5801     6      Egress     permit  1
Veth4   192.168.1.20  192.168.1.10  5217    8080     6      Ingress    permit  1
Veth4   192.168.1.10  192.168.1.20  8080    57211    6      Egress     permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5865     6      Egress     permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5833     6      Egress     permit  1
Veth4   192.168.1.20  192.168.1.10  5601    8080     6      Ingress    permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5705     6      Egress     permit  1
Veth4   192.168.1.10  192.168.1.20  8080    5737     6      Egress     permit  1
Veth4   192.168.1.20  192.168.1.10  5473    8080     6      Ingress    permit  1
Veth4   192.168.1.20  192.168.1.10  57211   8080     6      Ingress    permit  1
```

Flushing All ACL Flows

You can use this command to detect any new flows affecting the VEM. Clear all the existing flows, then you can detect new flows that match any expected traffic. Syslog messages are not sent when you do this. Enter the following command:

```
vemcmd flush aclflows
```


Showing Flow Debug Statistics

You can show ACL debug statistics.

To display internal ACL flow statistics, enter the following command:

```
vemcmd show aclflows dbgstats
```

To clear all internal ACL flow debug statistics, enter the following command:

```
vemcmd clear aclflows dbgstats
```

ACL Logging Troubleshooting Scenarios

This section describes situations that you might encounter when you are using ACL logging.

Troubleshooting a Syslog Server Configuration

If syslog messages are not being sent from the VEM, you can check the syslog server configuration and check if ACL logging is configured by entering the commands shown in the following procedure.

BEFORE YOU BEGIN

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

- You are logged in to the VSM and VEM CLI.

SUMMARY STEPS

- show logging ip access-list status
- vemcmd show acllog config
- vemcmd show aclflows dbgstats

PROCEDURE

	Command	Description
Step 1	<pre>show logging ip access-list status</pre> <p>Example</p> <pre>n1000v # show logging ip access-list status n1000v #</pre>	Verifies that the remote syslog server is configured properly.
Step 2	<pre>vemcmd show acllog config</pre> <p>Example:</p> <pre>n1000v# vemcmd show acllog config n1000v #</pre>	Verifies ACL logging on the VEM.
Step 3	<pre>vemcmd show aclflows dbgstats</pre> <p>Example:</p> <pre>n1000v# vemcmd show aclflows dbgstats n1000v #</pre>	Checks to see if any errors occurred.

Troubleshooting an ACL Rule That Does Not Have a Log Keyword

If the ACL rule does not have a log keyword, any flow matching the ACL is not reported although the ACL statistics continue to advance. You can verify a log keyword by entering the commands shown in the following procedure.

BEFORE YOU BEGIN

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

- You are logged in to the VSM and VEM CLI.

SUMMARY STEPS

1. **show running-config aclmgr**
2. **show logging ip access-list status**
3. **vemcmd show acllog config**

PROCEDURE

	Command	Description
Step 1	show running-config aclmgr Example n1000v # show running-config aclmgr n1000v #	Verify that the log keyword is enabled
Step 2	show logging ip access-list status Example: n1000v # show logging ip access-list status n1000v #	Verify that ACL logging is configured properly
Step 3	vemcmd show acllog config Example: n1000v # vemcmd show acllog config n1000v #	Verifies ACL logging on the VEM.

Troubleshooting a Maximum Flow Limit Value That is Too Low

If the number of flows does not reach 5000 for either permit or deny flows, you can increase the maximum flows by entering the commands shown in the following procedure.

BEFORE YOU BEGIN

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

- You are logged in to the VSM and VEM CLI.

SUMMARY STEPS

1. **show logging ip access-list status**

2. `vemcmd show aclog config`
3. `logging ip access-list cache max-deny- flows <num>`

PROCEDURE

	Command	Description
Step 1	<pre>show logging ip access-list status</pre> <p>Example:</p> <pre>n1000v # show logging ip access-list status n1000v #</pre>	Verifies that ACL logging is configured properly.
Step 2	<pre>vemcmd show aclog config</pre> <p>Example:</p> <pre>n1000v # vemcmd show aclog config n1000v #</pre>	Verifies ACL logging on the VEM.
Step 3	<pre>logging ip access-list cache max-deny- flows <num></pre> <p>Example:</p> <pre>n1000v # logging ip access-list cache max-deny- flows <num> n1000v #</pre>	Increases maximum flows to the desired value.

Troubleshooting a Mismatched Configuration between a VSM and a VEM

If syslog messages are not being sent and the flow information counters are invalid, the configuration between a VSM and a VEM might be mismatched. Enter the commands shown in this procedure.

Modify any mismatched configurations using the appropriate configuration command. If the problem persists, enable aclog debugging on both VSM and the VEM and retry the commands.

BEFORE YOU BEGIN

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

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

SUMMARY STEPS

1. `show logging ip access-list status`
2. `vemcmd show aclog config`

PROCEDURE

	Command	Description
Step 1	show logging ip access-list status Example: n1000v # show logging ip access-list status n1000v #	Verifies that ACL logging is configured properly.
Step 2	vemcmd show acllog config Example: n1000v # vemcmd show acllog config n1000v #	Verifies ACL logging on the VEM.



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 16-1](#)
- [QoS Configuration Limits, page 16-1](#)
- [QoS Troubleshooting Commands, page 16-2](#)
- [Troubleshooting the VEM, page 16-2](#)
- [Debugging Policing Verification Errors, page 16-3](#)

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*.

QoS Configuration Limits

[Table 16-1](#) and [Table 16-2](#) list the configuration limits for QoS.

Table 16-1 QoS Configuration Limits

Item	DVS Limit	Per Server Limit
Class map	1000	64 (with policies)
Policy map	128	16
Service policy	—	128

Table 16-2 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
-----
```

```

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.

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 ltl 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
```



SPAN

This chapter describes how to identify and resolve problems that relate to SPAN and includes the following topics:

- [Information About SPAN, page 17-1](#)
- [Problems with SPAN, page 17-2](#)
- [SPAN Troubleshooting Commands, page 17-3](#)

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 local SPAN or ERSPAN, see the *Cisco Nexus 1000V System Management Configuration Guide*.

SPAN Session Guidelines

The following are SPAN session guidelines:

- 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. Examples include the following:
 - 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.

- The following are required for a running SPAN session:
 - The limit of 64 SPAN sessions is not exceeded.
 - At least one operational source is configured.
 - At least one operational destination is configured.
 - The configured source and destination are on the same host.
 - The session is enabled with the **no shut** command.
- A session is stopped if any of the following 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.

Problems with SPAN

Symptom	Possible Causes	Solution
You observe issues with VM traffic after configuring a session with Ethernet destinations.	—	Ensure that the Ethernet 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.
A session state is up and the packets are not received at the destination ports.	—	Verify that the correct VLANs are allowed on the trunk destination ports.
The session displays an error.	—	<ol style="list-style-type: none"> 1. Make sure that NX-OS VEM connectivity is working correctly. 2. Force reprogramming of the session on the VEM. <p>shut no shut</p>
The ERSPAN session is up, but does not see packets at the destination.	The ERSPAN ID is not configured.	Make sure that ERSPAN ID is configured at the destination.
	An ERSPAN enabled VMKernel NIC is not configured on the host or VEM.	Make sure you create a VMKernel NIC on the host using a port profile configured for ERSPAN.
	The ERSPAN enabled VMKernel NIC is not configured with a proper IP, gateway, or both.	Ping the ERSPAN IP destination from the host VMKernel NIC. vmkping dest-id

SPAN Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to SPAN.

Command	Purpose
show monitor	Displays the status of SPAN sessions. See Example 17-1 on page 17-3 .
show monitor session	Displays the current state of a SPAN session, the reason it is down, and the session configuration. See Example 17-2 on page 17-3 .
module vem <i>module-number</i> execute vemcmd show span	Displays the VEM source IP and SPAN configuration. See Example 17-3 on page 17-4 .
show monitor internal errors	
show monitor internal event-history msgs	
show monitor internal info global-info	
show monitor internal mem-stats	

Example 17-1 show monitor

```
n1000v# show monitor
Session State Reason Description
-----
17 down Session admin shut folio
```

Example 17-2 show monitor session

```
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 :
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
```

Example 17-3 module vem execute vemcmd show span

```
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
```



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 18-1](#)
- [Multicast IGMP Snooping, page 18-1](#)
- [Problems with Multicast IGMP Snooping, page 18-2](#)

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.

- 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 if the Nexus 1000V switch is configured correctly and is ready to forward multicast traffic. In the displayed output of the command, 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 means that the Nexus 1000V 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 18-1 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,
 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 18-2 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 18-3 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 18-4](#), 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.

Example 18-4 *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 vmic3
  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.



DHCP, DAI, and IPSG

This chapter describes how to identify and resolve problems related to the following security features:

- Dynamic Host Configuration Protocol (DHCP) Snooping
- Dynamic ARP Inspection (DAI)
- IP Source Guard (IPSG)

This chapter includes the following sections:

- [Information About DHCP Snooping, page 19-1](#)
- [Information About Dynamic ARP Inspection, page 19-2](#)
- [Information About IP Source Guard, page 19-2](#)
- [Guidelines and Limitations for Troubleshooting, page 19-2](#)
- [Problems with DHCP Snooping, page 19-3](#)
- [Troubleshooting Dropped ARP Responses, page 19-4](#)
- [Problems with IP Source Guard, page 19-5](#)
- [Collecting and Evaluating Logs, page 19-5](#)
- [DHCP, DAI, and IPSG Troubleshooting Commands, page 19-6](#)

Information About DHCP Snooping

DHCP snooping acts like a firewall between untrusted hosts and trusted DHCP servers by doing the following:

- Validates DHCP messages received from untrusted sources and filters out invalid response messages from DHCP servers.
- Builds and maintains the DHCP snooping binding database, which contains information about untrusted hosts with leased IP addresses.
- Uses the DHCP snooping binding database to validate subsequent requests from untrusted hosts.

Dynamic ARP inspection (DAI) and IP Source Guard also use information stored in the DHCP snooping binding database.

For detailed information about configuring DHCP snooping, see the *Cisco Nexus 1000V Security Configuration Guide*.

Information About Dynamic ARP Inspection

DAI is used to validate ARP requests and responses as follows:

- Intercepts all ARP requests and responses on untrusted ports.
- Verifies that a packet has a valid IP-to-MAC address binding before updating the ARP cache or forwarding the packet.
- Drops invalid ARP packets.

DAI can determine the validity of an ARP packet based on valid IP-to-MAC address bindings stored in a Dynamic Host Configuration Protocol (DHCP) snooping binding database. This database is built by DHCP snooping when it is enabled on the VLANs and on the device. It may also contain static entries that you have created.

For detailed information about configuring DAI, see the *Cisco Nexus 1000V Security Configuration Guide*.

Information About IP Source Guard

IP Source Guard is a per-interface traffic filter that permits IP traffic only when the IP address and MAC address of each packet matches the IP and MAC address bindings of dynamic or static IP source entries in the Dynamic Host Configuration Protocol (DHCP) snooping binding table.

For detailed information about configuring IP Source Guard, see the *Cisco Nexus 1000V Security Configuration Guide*.

Guidelines and Limitations for Troubleshooting

The following guidelines and limitations apply when troubleshooting DHCP snooping, Dynamic ARP Inspection, or IP Source Guard:

- A maximum of 2000 DHCP entries can be snooped and learned system-wide in the DVS. This is a combined total for both entries learned dynamically and entries configured statically.
- Rate limits on interfaces must be set to high values for trusted interfaces such as VSD SVM ports or vEthernet ports connecting to DHCP servers.

For detailed guidelines and limitations used in configuring these features, see the *Cisco Nexus 1000V Security Configuration Guide*.

Problems with DHCP Snooping

The following are symptoms, possible causes, and solutions for problems with DHCP snooping.

Symptom	Possible Causes	Solution
With snooping configured, DHCP client is not able to obtain an IP address from the server.	IP address was not added to binding database. Faulty connection between DHCP server and client.	<ol style="list-style-type: none"> 1. Verify the connection between the DHCP server(s) and the host connected to the client. vmkping 2. If the connection between DHCP server and the host is broken, do the following: <ul style="list-style-type: none"> – Check the configuration in the upstream switch, for example, verifying that the VLAN is allowed, etc. – Make sure the server itself is up and running.
	The interface of the DHCP server(s) connected to the DVS as a VM is not trusted.	<ol style="list-style-type: none"> 1. On the VSM, verify that the interface is trusted. show ip dhcp snooping 2. On the VSM, verify the vEthernet interface attached to the server is trusted. module vem mod# execute vemcmd show dhcps interfaces
	DHCP requests from the VM are not reaching the server for acknowledgement.	On the DHCP server, log in and use a packet capture utility to verify requests and acknowledgements in packets.
	DHCP requests and acknowledgements are not reaching the Cisco Nexus 1000V.	<ul style="list-style-type: none"> • From the client vEthernet interface, SPAN the packets to verify they are reaching the client. • On the host connected to the client, enable VEM packet capture to verify incoming requests and acknowledgements in packets.
	The Cisco Nexus 1000V is dropping packets.	On the VSM, verify DHCP statistics. show ip dhcp snooping statistics module vem mod# execute vemcmd show dhcps stats

Troubleshooting Dropped ARP Responses

The following are possible causes, and solutions for dropped ARP responses.

Possible Causes	Solution
ARP inspection is not configured on the VSM	<p>On the VSM, verify that ARP inspection is configured as expected.</p> <p>show ip arp inspection</p> <p>For detailed information about configuring DAI, see the <i>Cisco Nexus 1000V Security Configuration Guide</i></p>
DHCP snooping is not enabled globally on the VSM, or is not enabled on the VLAN.	<p>On the VSM, verify the DHCP snooping configuration.</p> <p>show ip dhcp snooping</p> <p>For detailed information about enabling DHCP, and configuring DAI, see the <i>Cisco Nexus 1000V Security Configuration Guide</i>.</p>
DHCP snooping is not enabled on the VEM, or is not enabled on the VLAN.	<ol style="list-style-type: none"> 1. From the VSM, verify the VEM DHCP snooping configuration. <p>module vem mod# execute vemcmd show dhcps vlan</p> 2. Do one of the following: <ul style="list-style-type: none"> – Correct any errors in the VSM DHCP configuration. For detailed information, see the <i>Cisco Nexus 1000V Security Configuration Guide</i>. – If the configuration appears correct on the VSM but fails on the VEM, capture and analyze the error logs from both VSM and the VEM to identify the reason for the failure.
If snooping is disabled, the binding entry is not statically configured in the binding table.	<ol style="list-style-type: none"> 1. On the VSM, display the binding table. <p>show ip dhcp snooping binding</p> 2. Correct any errors in the static binding table. <p>For detailed information about clearing entries from the table, enabling DHCP, and configuring DAI, see the <i>Cisco Nexus 1000V Security Configuration Guide</i>.</p>
The binding corresponding to the VM sending the ARP response is not present in the binding table.	<ol style="list-style-type: none"> 1. On the VSM, display the binding table. <p>show ip dhcp snooping binding</p> 2. Correct any errors in the static binding table. <p>For detailed information about clearing entries from the table, enabling DHCP, and configuring DAI, see the <i>Cisco Nexus 1000V Security Configuration Guide</i>.</p> 3. If all configurations are correct, make sure to turn on DHCP snooping before DAI or IPSG. This is to make sure the Cisco Nexus 1000V has enough time to add the binding in the snooping database. <p>For more information, see the <i>Cisco Nexus 1000V Security Configuration Guide</i>.</p>

Problems with IP Source Guard

The following are symptoms, possible causes, and solutions for problems with IP Source Guard.

Symptom	Possible Causes	Solution
Traffic disruptions	ARP inspection is not configured on the VSM.	<p>On the VSM, verify that IP Source Guard is configured as expected.</p> <pre>show port-profile name profile_name</pre> <pre>show running interface if_ID</pre> <pre>show ip verify source</pre> <p>For detailed information about configuring IP Source Guard, see the <i>Cisco Nexus 1000V Security Configuration Guide</i></p>
	The IP address corresponding to the vEthernet interface is not in the snooping binding table.	<ol style="list-style-type: none"> 1. On the VSM, display the binding table. <pre>show ip dhcp snooping binding</pre> 2. Configure the missing static entry or renew the lease on the VM. 3. On the VSM, display the binding table again to verify the entry is added correctly. <pre>show ip dhcp snooping binding</pre>

Collecting and Evaluating Logs

You can use the commands in this section from the VSM to collect and view logs related to DHCP, DAI, and IP Source Guard.

- [VSM Logging, page 19-5](#)
- [Host Logging, page 19-6](#)

VSM Logging

You can use the commands in this section from the VSM to collect and view logs related to DHCP, DAI, and IP Source Guard.

VSM Command	Description
<code>debug dhcp all</code>	Enable debug all for dhcp configuration flags
<code>debug dhcp errors</code>	Enable debugging of errors
<code>debug dhcp mts-errors</code>	Enable debugging of mts errors
<code>debug dhcp mts-events</code>	Enable debugging of mts events
<code>debug dhcp pkt-events</code>	Enable debugging of pkt events
<code>debug dhcp pss-errors</code>	Enable debugging of pss errors
<code>debug dhcp pss-events</code>	Enable debugging of pss events

Host Logging

You can use the commands in this section from the ESX host to collect and view logs related to DHCP, DAI, and IP Source Guard.

ESX Host Command	Description
<code>echo "logfile enable" > /tmp/dpafifo</code>	Enables DPA debug logging. Logs are output to /var/log/vemdpa.log file.
<code>echo "debug sfdhcpsagent all" > /tmp/dpafifo</code>	Enables DPA DHCP agent debug logging. Logs are output to /var/log/vemdpa.log file.
<code>vemlog debug sfdhcps all</code>	Enables datapath debug logging, and captures logs for the data packets sent between the client and the server.
<code>vemlog debug sfdhcps_config all</code>	Enables datapath debug logging, and captures logs for configuration coming from the VSM.
<code>vemlog debug sfdhcps_binding_table all</code>	Enables datapath debug logging, and captures logs corresponding to binding database changes.

DHCP, DAI, and IPSG Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to DHCP snooping, DAI, and IP Source Guard.

Command	Description
<code>show running-config dhcp</code>	Displays the DHCP snooping, DAI, and IP Source Guard configuration See Example 19-1 on page 19-7 .
<code>show ip dhcp snooping</code>	Displays general information about DHCP snooping. See Example 19-2 on page 19-7 .
<code>show ip dhcp snooping binding</code>	Display the contents of the DHCP snooping binding table. See Example 19-3 on page 19-7 .
<code>show feature</code>	Displays the features available, such as DHCP, and whether they are enabled. See Example 19-4 on page 19-8 .
<code>show ip arp inspection</code>	Displays the status of DAI. See Example 19-5 on page 19-8 .
<code>show ip arp inspection interface vethernet <i>interface-number</i></code>	Displays the trust state and ARP packet rate for a specific interface. See Example 19-6 on page 19-8 .

Command	Description
show ip arp inspection vlan <i>vlan-ID</i>	Displays the DAI configuration for a specific VLAN. See Example 19-7 on page 19-8 .
show ip verify source	Displays interfaces where IP source guard is enabled and the IP-MAC address bindings. See Example 19-8 on page 19-9 .
show system internal dhcp { <i>event-history</i> <i>mem-stats</i> <i>msgs</i> }	Debugs any issues in the filter-mode configuration. See Example 19-9 on page 19-9 , Example 19-10 on page 19-9 , and Example 19-11 on page 19-10 .
Debug dhcp all	Enables debug all for dhcp configuration flags on the VSM. See Example 19-12 on page 19-10 .

Example 19-1 show running-config dhcp

```
n1000v# show running-config dhcp

!Command: show running-config dhcp
!Time: Wed Feb 16 14:20:36 2011

version 4.2(1)SV1(4)
feature dhcp

no ip dhcp relay

n1000v#
```

Example 19-2 show ip dhcp snooping

```
n1000v# show ip dhcp snooping
DHCP snooping service is enabled
Switch DHCP snooping is enabled
DHCP snooping is configured on the following VLANs:
1,13
DHCP snooping is operational on the following VLANs:
1
Insertion of Option 82 is disabled
Verification of MAC address is enabled
DHCP snooping trust is configured on the following interfaces:
Interface          Trusted
-----          -
vEthernet 3        Yes

n1000v#
```

Example 19-3 show ip dhcp snooping binding

```
n1000v# show ip dhcp snooping binding
MacAddress          IPAddress          LeaseSec  Type          VLAN  Interface
-----          -
0f:00:60:b3:23:33  10.3.2.2          infinite  static        13   vEthernet 6
0f:00:60:b3:23:35  10.2.2.2          infinite  static        100  vEthernet 10

n1000v#
```

Example 19-4 show feature

```
n1000v# show feature
Feature Name           Instance  State
-----
dhcp-snooping         1        enabled
http-server           1        enabled
ippool                 1        enabled
lACP                   1        enabled
lisp                   1        enabled
lisp-helper            1        enabled
netflow                1        disabled
port-profile-roles    1        enabled
private-vlan           1        disabled
sshServer              1        enabled
tacacs                 1        enabled
telnetServer           1        enabled
n1000v#
```

Example 19-5 show ip arp inspection

```
n1000v# show ip arp inspection

Source Mac Validation      : Disabled
Destination Mac Validation : Disabled
IP Address Validation      : Disabled

Vlan : 1
-----
Configuration              : Disabled
Operation State             : Inactive

Vlan : 5
-----
Configuration              : Disabled
Operation State             : Inactive

Vlan : 100
-----
Configuration              : Disabled
Operation State             : Inactive

Vlan : 101
-----
Configuration              : Disabled
Operation State             : Inactive
n1000v#
```

Example 19-6 show ip arp inspection interface

```
n1000v# show ip arp inspection interface vethernet 6

Interface      Trust State
-----
vEthernet 6    Trusted
n1000v#
```

Example 19-7 show ip arp inspection vlan

```
n1000v# show ip arp inspection vlan 13
```



```
Source Mac Validation      : Disabled
Destination Mac Validation : Enabled
IP Address Validation     : Enabled
```

```
n1000v#
```

Example 19-8 show ip verify source

```
n1000v# show ip arp inspection vlan 13
```

```
IP source guard is enabled on the following interfaces:
```

```
-----
Vethernet1

Interface      Filter-mode      IP-address      Mac-address      Vlan
-----
Vethernet11    active           25.0.0.128     00:50:56:88:00:20  25
```

Example 19-9 show system internal dhcp event-history msgs

```
n1000v# show system internal dhcp event-history msgs
```

```
1) Event:E_MTS_RX, length:60, at 809122 usecs after Mon Oct 8 20:59:08 2012
   [RSP] Opc:MTS_OPC_PDL32(148511), Id:0X00F132AB, Ret:SUCCESS
   Src:0x00000302/747, Dst:0x00000201/360, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00009498, Sync:UNKNOWN, Payloadsize:132
   Payload:
   0x0000: 00 00 00 03 00 00 00 01 00 00 00 64 00 00 00 07

2) Event:E_MTS_RX, length:60, at 809100 usecs after Mon Oct 8 20:59:08 2012
   [RSP] Opc:MTS_OPC_PDL32(148511), Id:0X00E01555, Ret:SUCCESS
   Src:0x00000502/747, Dst:0x00000201/360, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00009497, Sync:UNKNOWN, Payloadsize:132
   Payload:
   0x0000: 00 00 00 03 00 00 00 01 00 00 00 64 00 00 00 07

3) Event:E_MTS_RX, length:60, at 809079 usecs after Mon Oct 8 20:59:08 2012
   [RSP] Opc:MTS_OPC_PDL32(148511), Id:0X006BE1FC, Ret:SUCCESS
   Src:0x00000602/747, Dst:0x00000201/360, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00009496, Sync:UNKNOWN, Payloadsize:132
   Payload:
   0x0000: 00 00 00 03 00 00 00 01 00 00 00 64 00 00 00 07

4) Event:E_MTS_RX, length:60, at 809028 usecs after Mon Oct 8 20:59:08 2012
   [RSP] Opc:MTS_OPC_PDL32(148511), Id:0X00F132AA, Ret:SUCCESS
   Src:0x00000302/747, Dst:0x00000201/360, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00009474, Sync:UNKNOWN, Payloadsize:132
   Payload:
   0x0000: 00 00 00 03 00 00 00 01 00 00 00 64 00 00 00 07
contd.
```

Example 19-10 show system internal dhcp mem-stats detail

```
VSM-N1k# show system internal dhcp mem-stats detail
```

```
Private Mem stats for UUID : Malloc track Library(103) Max types: 5
```

```
-----
TYPE NAME                                ALLOCS                                BYTES
                                CURR  MAX                                CURR  MAX
2 MT_MEM_mtrack_hdl                     33    34                                19236 19384
3 MT_MEM_mtrack_info                     588   880                                9408  14080
4 MT_MEM_mtrack_lib_name                 882   1174                               42246 56230
```

```
-----
Total bytes: 70890 (69k)
-----
```

```
Private Mem stats for UUID : Non mtrack users(0) Max types: 149
-----
```

TYPE	NAME	ALLOCS		BYTES	
		CURR	MAX	CURR	MAX
11	[r-xp]/isan/plugin/0/isan/lib/libavl.so	3421	3421	68360	68360
26	[r-xp]/isan/plugin/0/isan/lib/libddbcom	116	141	302445	308307
47	[r-xp]/isan/plugin/0/isan/lib/libindxob	6	6	456	456
50	[r-xp]/isan/plugin/0/isan/lib/libip.so.	1	1	212	212
64	[r-xp]/isan/plugin/0/isan/lib/libmpmts.	0	9	0	785
66	[r-xp]/isan/plugin/0/isan/lib/libmts.so	10	11	972	984
68	[r-xp]/isan/plugin/0/isan/lib/libnetsta	1	2	704	1350
81	[r-xp]/isan/plugin/0/isan/lib/libpss.so	158	262	101579	204281
85	[r-xp]/isan/plugin/0/isan/lib/libfdb.so	44	44	3914	3914
89	[r-xp]/isan/plugin/0/isan/lib/libmm.so	3	3	216	216
111	[r-xp]/isan/plugin/0/isan/lib/libutils.	4	7	69	349
112	[r-xp]/isan/plugin/0/isan/lib/libvdc_mg	0	1	0	20
118	[r-xp]/isan/plugin/2/isan/bin/dhcp_snoo	0	2	0	64
121	[r-xp]/isan/plugin/2/isan/lib/libpdlser	4	29	208	1016
128	[r-xp]/lib/ld-2.3.3.so	33	33	5363	5371
131	[r-xp]/lib/tls/libc-2.3.3.so	51	51	1347	1637
134	[r-xp]/lib/tls/libpthread-2.3.3.so	1	1	33	33
138	[r-xp]/usr/lib/libglib-2.0.so.0.600.1	15	16	10372	10392
145	[r-xp]/isan/plugin/1/isan/lib/libvem_mg	0	1	0	1940

```
-----
Total bytes: 496250 (484k)
-----
```

```
contd.
```

Example 19-11 show system internal dhcp msgs

```
n1000v# show system internal dhcp msgs
1) Event:E_DEBUG, length:75, at 409832 usecs after Mon Oct 8 20:57:48 2012
   [16843009] Session close, handle -767541913, sess-id 0xff0101ba02812d08, state 3

2) Event:E_DEBUG, length:62, at 399944 usecs after Mon Oct 8 20:57:48 2012
   [16843009] PPF session open session-id 0xff0101ba02812d08, msg_id 0

3) Event:E_DEBUG, length:30, at 399866 usecs after Mon Oct 8 20:57:48 2012
   [16843009] PPF goto setting state 1

4) Event:E_DEBUG, length:23, at 682346 usecs after Mon Oct 8 20:57:11 2012
   [16843009] Processed log-mts

contd
```

Example 19-12 debug dhcp all

```
n1000v# debug dhcp all
#
```



Virtual Service Domain

This chapter describes how to identify and resolve problems related to Virtual Service Domain (VSD).

This chapter includes the following sections:

- [Information about Virtual Service Domain, page 20-1](#)
- [Problems with Virtual Service Domain, page 20-1](#)
- [Collecting and Evaluating Logs, page 20-2](#)
- [Virtual Service Domain Troubleshooting Commands, page 20-3](#)

Information about Virtual Service Domain

A Virtual Service Domain (VSD) is a logical group of interfaces that is serviced by a common Service VM (SVM). With VSD the Cisco Nexus 1000V can support third party appliances such as vShield.

VSD lets you classify and separate traffic for network services such as firewalls and traffic monitoring.

Multiple VSDs can co-exist on a host; with each VSD serviced by an SVM.

For more information, to configure VSD, an example configuration, and for configuration limits, see the *Cisco Nexus 1000V System Management Configuration Guide*.

Problems with Virtual Service Domain

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

Symptom	Possible Causes	Solution
The SVM does not come online.	There is more than one SVM per VSD per host. There can be only one SVM per VSD per host. If a second SVM tries to come up, the SVM ports are error disabled.	<ol style="list-style-type: none"> 1. Check for multiple SVMs per VSD per host. show virtual-service-domain interface If output indicates Invalid SVM interface, then there are multiple SVMs per VSD per host. 2. Remove or relocate one of the SVMs.

Symptom	Possible Causes	Solution
A loop occurs.	SVM ports are not correctly attached to the inside and outside port profiles.	<ol style="list-style-type: none"> 1. Turn off the SVM looping capability or the SVM itself. 2. Display the interfaces attached to the port profiles. show port-profile usage 3. Correct configuration errors. <p>For information about configuring VSD, see the <i>Cisco Nexus 1000V System Management Configuration Guide</i>.</p>

Collecting and Evaluating Logs

You can use the commands in this section from the VSM to collect and view logs related to VSD captured as follows:

- VSM logs: /var/log/external/startupdebug
- VEM DPA logs: /var/log/vemdpa.log

Command	Description
module vem <i>module_number</i> execute vemdpalog writelogs	Enables the DPA logs and writes them to vemdpa.log.
module vem <i>module_number</i> execute vemdpalog debug sfvsimagent all	
module vem <i>module_number</i> execute vemdpalog start	Starts and stops DPA logging for viewing.
module vem <i>module_number</i> execute vemdpalog stop	
module vem <i>module_number</i> execute vemdpalog show all	Displays DPA logs.
module vem <i>module_number</i> execute vemlog debug sfvsim all	Enables DP logs.
module vem <i>module_number</i> execute vemlog start	Starts and stops DP logging for viewing.
module vem <i>module_number</i> execute vemlog stop	
module vem <i>module_number</i> execute vemlog show all	Displays DPA logs.

Example 20-1 VSM Logs

```

2011 Feb 17 10:14:01 vsm vsim: <{vsim}> [DBG]=====ZONES=====
2011 Feb 17 10:14:01 vsm vsim: <{vsim}>[DBG]Zone_id: 1, name: vsd1, is_in_use? 1,
default_action: (DROP), member_cnt: 5
2011 Feb 17 10:14:01 vsm vsim: <{vsim}> [DBG]=====INTFS=====

```

```

2011 Feb 17 10:14:01 vsm vsim: <{vsim}>[DBG]Ifindex 0x1c000000, zoneid 1, status ATTACHED,
type SVM_MEMBER (2)
2011 Feb 17 10:14:01 vsm vsim: <{vsim}>[DBG]Ifindex 0x1c000010, zoneid 1, status ATTACHED,
type SVM_MEMBER (2)
2011 Feb 17 10:14:01 vsm vsim: <{vsim}>[DBG]Ifindex 0x1c000020, zoneid 1, status ATTACHED,
type SVM_MEMBER (2)
2011 Feb 17 10:14:01 vsm vsim: <{vsim}>[DBG]Ifindex 0x1c000030, zoneid 1, status ATTACHED,
type SVM_MEMBER (2)

```

Example 20-2 VEM DPA Logs

```

Feb 17 16:11:02.645378: sfvsimagent: PDL Lite :Opening new session
Feb 17 16:11:02.723186: sfvsimagent: PDL Lite :Add policy callback
Feb 17 16:11:02.727281: sfvsimagent: PDL Lite :Add policy node callback
Feb 17 16:11:02.727293: sfvsimagent: sf_vsim_add_vzone: Entered
Feb 17 16:11:02.727303: sfvsimagent: sf_vsim_dpa_vzone_init: Entered
Feb 17 16:11:02.727324: sfvsimagent: MTS Opcode: 142337

```

Example 20-3 VEM Logs

```

Feb 17 15:58:42.924322      4411   1     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 18 dst l1l 10
Feb 17 15:58:42.924337      4412   1     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 9  dst l1l 8
Feb 17 15:58:43.038065      4413   1     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 18 dst l1l 10
Feb 17 15:58:43.038087      4414   1     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 9  dst l1l 8
Feb 17 15:58:43.038128      4415   2     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 8  dst l1l 4282
Feb 17 15:58:43.038152      4416   1     1  16   Debug sfvsimsrc: Reached vsim stage src
l1l 10 dst l1l 18
Feb 17 15:58:43.038156      4417   2     0  0     Suspending log

```

Virtual Service Domain Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to VSD.

Command	Description
show system internal ethpm event-history interface	Displays the request/response pre-configuration event. Useful when the port is error disabled. See Example 20-4 on page 20-4 .
show system internal vsim event-history msgs	Displays a log of the MTS events processed by VSIM. See Example 20-5 on page 20-4 .
module vem <i>mod-number</i> execute vemcmd show port	Displays the port state on the VEM. Useful for debugging traffic flow on interfaces. See Example 20-6 on page 20-5 .
show virtual-service-domain name <i>vsd-name</i>	Displays a specific VSD configuration. See Example 20-7 on page 20-5 .

Command	Description
show virtual-service-domain brief	Displays a summary of all VSD configurations. See Example 20-8 on page 20-5 .
show virtual-service-domain interface	Displays the interface configuration for all VSDs. See Example 20-9 on page 20-6 .
module vem <i>module_number</i> execute vemcmd show vsd	Displays the VEM VSD configuration by sending the command to the VEM from the remote Cisco Nexus 1000V. See Example 20-10 on page 20-6 .
module vem <i>module_number</i> execute vemcmd show vsd ports	Displays the VEM VSD ports configuration by sending the command to the VEM from the remote Cisco Nexus 1000V. See Example 20-11 on page 20-6 .
show port-profile name <i>profile-name</i>	Displays the port profile configuration. See

Example 20-4 show system internal ethpm event-history interface vethernet 1

```
n1000v# show system internal ethpm event-history interface vethernet 1
...
18) Event:ESQ_REQ length:34, at 725272 usecs after Thu Feb 17 15:42:13 2011
Instance:469762048, Seq Id:0x1, Ret:success
[E_MTS_TX] Dst:MTS_SAP_VSIM(716), Opc:MTS_OPC_ETHPM_PORT_PRE_CFG(61441)

19) Event:ESQ_RSP length:34, at 739984 usecs after Thu Feb 17 15:42:13 2011
Instance:469762048, Seq Id:0x1, Ret:success
[E_MTS_RX] Src:MTS_SAP_VSIM(716), Opc:MTS_OPC_ETHPM_PORT_PRE_CFG(61441)
...
n1000v#
```

Example 20-5 show system internal vsim event-history msgs

```
n1000v# show system internal vsim event-history msgs
1) Event:E_MTS_RX, length:60, at 215249 usecs after Thu Feb 17 10:16:53 2011
[REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X000C14C4, Ret:SUCCESS
Src:0x00000101/2282, Dst:0x00000101/716, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x000C14C4, Sync:UNKNOWN, Payloadsize:216
Payload:
0x0000: 01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 32 34

2) Event:E_MTS_TX, length:60, at 833885 usecs after Thu Feb 17 10:14:01 2011
[NOT] Opc:MTS_OPC_FSMUTILS_SYNC_PSS_TO_STDBY(1523), Id:0X000C05B3, Ret:SUCCESS
Src:0x00000101/716, Dst:0x00000101/0, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:380
Payload:
0x0000: 00 00 00 00 00 00 00 01 00 00 00 01 00 00 00 01

3) Event:E_FU_UNLOCK, length:36, at 820289 usecs after Thu Feb 17 10:14:01 2011
Status: 0x0
Gwrap: 0x80fa09c Cat: 0x0
Opc:MTS_OPC_VSH_CMD_TLV_SYNC(7682)
Msg id: 0X000C05A5
Lock type: 0
```

```

RID Size: 8
Val :
0x0000: 01 00 00 00 00 00 00 01
4) Event:E_FU_UNLOCK, length:36, at 818291 usecs after Thu Feb 17 10:14:01 2011
Status: 0x0
Gwrap: 0x80fa09c Cat: 0x0
Opc:MTS_OPC_VSH_CMD_TLV_SYNC(7682)
Msg id: 0X000C05A5
Lock type: 0
RID Size: 8
Val :
0x0000: 00 00 00 1c 00 00 00 02
5) Event:E_FU_UNLOCK, length:36, at 816421 usecs after Thu Feb 17 10:14:01 2011
Status: 0x0
Gwrap: 0x80fa09c Cat: 0x0
Opc:MTS_OPC_VSH_CMD_TLV_SYNC(7682)
Msg id: 0X000C05A5
Lock type: 0
RID Size: 8
Val :
0x0000: 10 00 00 1c 00 00 00 02
n1000v#

```

Example 20-6 module vem # execute vemcmd show port

```

n1000v# module vem 3 execute vemcmd show port
LTL   VSM Port  Admin Link  State  PC-LTL  SGID  Vem Port
18    Eth3/2   UP    UP    F/B*    0      vmnic1
49    Veth1     UP    UP    FWD     0      New Virtual Machine.eth0
50    Veth2     UP    UP    FWD     0      New Virtual Machine.eth1
51    Veth3     UP    UP    FWD     0      New Virtual Machine.eth2
52    Veth4     UP    UP    FWD     0      New Virtual Machine.eth3

```

* F/B: Port is BLOCKED on some of the vlans.
Please run "vemcmd show port vlans" to see the details.

```
n1000v#
```

Example 20-7 show virtual-service-domain name vsd_name

```
n1000v# show virtual-service-domain name vsd1
Default Action: drop
```

Interface	Type
Vethernet1	Member
Vethernet2	Member
Vethernet3	Member
Vethernet6	Member
Vethernet7	Inside
Vethernet8	Outside

```
n1000v#
```

Example 20-8 show virtual-service-domain brief

```
n1000v# show virtual-service-domain brief
```

Name	vsd-id	default action	in-ports	out-ports	mem-ports	Modules with VSD Enabled
zone	1	forward	1	1	2	4

```
n1000v#
```

Example 20-9 show virtual-service-domain interface

```
n1000v# sho virtual-service-domain interface
-----
Name           Interface           Type           Status
-----
vsd1           Vethernet1          Member         Active
vsd1           Vethernet2          Member         Active
vsd1           Vethernet3          Member         Active
vsd1           Vethernet6          Member         Active
vsd1           Vethernet7          Inside         Active
vsd1           Vethernet8          Outside        Active
vsd2           Vethernet9          Inside         Active
vsd2           Vethernet10         Outside        Active
```

Example 20-10 module module_number execute vemcmd show vsd

```
n1000v# module vem 4 execute vemcmd show vsd
ID Def_Act ILTL OLT L NMLTL State Member LTLs
1 FRWD 51 50 1 ENA 49
n1000v#
```

Example 20-11 module module_number execute vemcmd show vsd ports

```
n1000v# module vem 4 execute vemcmd show vsd ports
LTL IfIndex VSD_ID VSD_PORT_TYPE
49 1c000010 1 REGULAR
50 1c000040 1 OUTSIDE
51 1c000030 1 INSIDE
n1000v#
```

Example 20-12 show port-profile name UpLinkProfile

```
n1000v# show port-profile name UpLinkProfile3
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#
```




System

This chapter describes how to identify and resolve problems related to the Nexus 1000V system.

This chapter includes the following sections:

- [Information About the System, page 21-1](#)
- [General Restrictions for vCenter Server, page 21-2](#)
- [Recovering a DVS, page 21-2](#)
- [Problems Related to VSM and vCenter Server Connectivity, page 21-5](#)
- [Connection Failure After ESX Reboot, page 21-6](#)
- [VSM Creation, page 21-9](#)
- [Port Profiles, page 21-9](#)
- [Problems with Hosts, page 21-10](#)
- [Problems with VM Traffic, page 21-10](#)
- [VEM Troubleshooting Commands, page 21-11](#)
- [VEM Log Commands, page 21-12](#)
- [Error Messages, page 21-12](#)

Information About the System

Cisco Nexus 1000V provides Layer 2 switching functions in a virtualized server environment. Nexus 1000V replaces virtual switches within ESX servers and allows users to configure and monitor the virtual switch using the Cisco NX-OS command line interface. Nexus 1000V also gives you visibility into the networking components of the ESX servers and access to the virtual switches within the network.

The Nexus 1000V manages a data center defined by the vCenter Server. Each server in the Datacenter is represented as a linecard in Nexus 1000V and can be managed as if it were a line card in a physical Cisco switch. The Nexus 1000V implementation has two components:

- **Virtual supervisor module (VSM)** – This is the control software of the Nexus 1000V distributed virtual switch. It runs on a virtual machine (VM) and is based on 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.

See the *Cisco Nexus 1000V Getting Started Guide* for a detailed overview of how the Nexus 1000V works with VMware ESX software.

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 21-3](#)
- [Recovering a DVS Without a Saved Copy of the VSM, page 21-4](#)

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 21-4](#).
- 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 [“Unregister 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*.
- 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**

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 21-3](#).
- 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 [“Unregister 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*.

- 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*.

- 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
Connections are not supported between Release 4.0(4)SV1(3a) VSMs and VMware vCenter Server 5.0	Upgrade to a compatible version of the Cisco Nexus 1000V software.
The vCenter Server connection seems to succeed, but does not.	Make sure that the domain ID is configured correctly.
The svcs connection command fails.	Make sure you have configured all parameters for the svcs connection command.
	Make sure you can ping the vCenter Server IP address.
	Make sure that the proxy.xml file is correct for both the IP address and length.
	Restart the vCenter Server

Symptom	Solution
The connection fails after an ESX reboot	“Connection Failure After ESX Reboot” procedure on page 21-6
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.

[Example 21-1](#) 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 21-1 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.
```

Connection Failure After ESX Reboot

To prevent a loss of connectivity between the VSM and VEM, and preserve a non-default MTU setting for a physical NIC across reboots of the ESX, you must configure a system MTU in the system port profile.

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 result in a mismatch with the VMware kernel NIC MTU and failure of the VSM and VEM. For example, you may manually configure an MTU of other than 1500 in networks with jumbo frames. During a power cycle, the ESX reboots and the MTU of the physical NIC reverts to the default of 1500 but the VMware kernel NIC does not.

To prevent a loss of connectivity in resulting from an MTU mismatch, see the [“Setting the System MTU” procedure on page 21-7](#).

To recover connectivity if you have not configured system mtu in the system uplink port profile, see

Setting the System MTU

Use this procedure to set a system MTU in your existing system uplink port profiles.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- The system port profiles are already configured and you know the uplink profile names.
For more information, see the *Cisco Nexus 1000V Port Profile Configuration Guide*.
- The MTU size you set for the **system mtu** on the port profile must be less than the size of the **system jumbomtu** configured on the interface.
For more information about configuring MTU on the interface, see the *Cisco Nexus 1000V Interface Configuration Guide*.
- When you configure a system MTU on a system port profile, it takes precedence over an MTU you may have configured on the interface.
- To verify the ESX MTU settings for corresponding PNICs, use the **esxcfg-nics -l** command.

SUMMARY STEPS

1. **config t**
2. **port-profile** *profilename*
3. **system mtu** *mtu value*
4. **show port-profile** [**brief** | **expand-interface** | **usage**] [*name profilename*]
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Description
Step 1	config t Example: n1000v# config t n1000v(config)#	Enters global configuration mode.
Step 2	port-profile <i>name</i> Example: n1000v(config)# port-profile AccessProf n1000v(config-port-prof)#	Enters port profile configuration mode for the named system uplink port profile.
Step 3	system mtu <i>mtu-size</i> Example: n1000v(config-port-prof)# system mtu 4000 n1000v(config-port-prof)#	Designates the MTU size. <ul style="list-style-type: none"> • Must be an even number between 1500 and 9000. • Must be less than the size of the system jumbomtu on the interface.

	Command	Description
Step 4	<pre>show port-profile [brief expand-interface usage] [name profile-name]</pre> <p>Example: n1000v(config-port-prof)# show port-profile name AccessProf</p>	(Optional) Displays the configuration for verification.
Step 5	<pre>copy running-config startup-config</pre> <p>Example: n1000v(config-port-prof)# copy running-config startup-config</p>	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Recovering Lost Connectivity Due to MTU Mismatch

Use this procedure to recover lost connectivity due to an MTU mismatch between the physical NIC and the VMware kernel NIC after an ESX reboot.

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- To verify the ESX MTU settings for corresponding PNICs, use the **esxcfg-nics -l** command.



Note Use **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.

SUMMARY STEPS

- config t**
- module vem module_number execute vemcmd show port port-LTL-number**
- module vem module_number execute vemcmd set mtu size ltl port-LTL-number**

DETAILED STEPS

	Command	Description
Step 1	<pre>config t</pre> <p>Example: n1000v# config t n1000v(config)#</p>	Enters global configuration mode.
Step 2	<pre>module vem module_number execute vemcmd show port port-LTL-number</pre>	Displays the port configuration including the LTL number needed for Step 3 .

Command	Description
<p>Example:</p> <pre>n1000v(config)# module vem 3 execute 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 vmn1c1 n1000v(config)#</pre>	
<p>Step 3</p> <pre>module vem <i>module_number</i> execute vemcmd set mtu <i>size</i> ltl <i>port-LTL-number</i></pre> <p>Example:</p> <pre>n1000v(config)# module vem 3 execute vemcmd set mtu 9000 ltl 17 n1000v(config)#</pre>	Designates the MTU size for the port, using the LTL number obtained in Step 2 .

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.

Problems with Hosts

Symptom	Solution
You receive an error message, DVS Operation failed for one or more members.”	Issue the vem status -v command to verify if the VEM is running on the host.
	Issue the vem unload command to unload the VEM.
	In the vSphere Client, remove the stale DVS: <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.	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. 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.
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
- **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 21-2 **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 svb_switch_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 21-3 **vemcmd show info Command**

```
[root@esx-cos1 ~]# vemcmd 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 21-4 **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
- **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 21-1 lists error messages that you might see on the VSM.

Table 21-1 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.

Table 21-1 Error Messages on the VSM

Error	Description
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.



Network Segmentation Manager

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

- [Information About Network Segmentation Manager, page 22-1](#)
- [Problems with Network Segmentation Manager, page 22-2](#)
- [Network Segmentation Manager Troubleshooting Commands, page 22-7](#)

Information About Network Segmentation Manager

See the *Cisco Nexus 1000V Network Segmentation Manager Configuration Guide* for more information.

Problems with Network Segmentation Manager

This section includes symptoms, possible causes and solutions for the following problems with Network Segmentation Manager (NSM). The system message for the majority of the problems is logged in the vShield Manager or the vCloud Director.

Table 22-1 *Problems with Network Segmentation Manager*

Symptom	Possible Causes	Verification and Solution
Registration failure of vShield Manager with Network Segmentation Manager A system message is logged in the vShield Manager.	vShield Manager is unable to reach Network Segmentation Manager.	Verify that the connection between Cisco Nexus 1000V and VMware vShield Manager is enabled. Check that vShield Manager is able to ping the Cisco Nexus 1000V. If not, reestablish the L2 or L3 connectivity between vShield Manager and the Cisco Nexus 1000V. See the <i>Cisco Nexus 1000V Network Segmentation Manager Configuration Guide</i> for more information.
	vShield Manager is unable to authenticate with Network Segmentation Manager.	Verify if the username and password are accurate by checking the VSM system logs. The following system log will be displayed if the username and password are inaccurate. 2012 Jan 20 00:49:59 switch %USER-3-SYSTEM_MSG: VALIDATE: user: admin, Authentication failure - validate If not, replace the username and password on the in the networking configuration on the vShield Manager.
	The NSM feature is not enabled on the Cisco Nexus 1000V.	Verify if the NSM feature is enabled on the Cisco Nexus 1000V. show feature If not, enable the NSM feature. feature network-segmentation-manager
	HTTPS is not enabled on the Cisco Nexus 1000V.	Check if the browser can connect to <code>https://<vsm-ip>/?</code> If not, enable the HTTPS server on the VSM. feature http-server

Table 22-1 Problems with Network Segmentation Manager (continued)

Symptom	Possible Causes	Verification and Solution
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to create network segment</pre>	vCloud Director is unable to create the VLAN associated with the network.	<ol style="list-style-type: none"> 1. Verify that the resources are available to create a VLAN by checking the existing number of VLANs. show vlan summary If the number of VLANs existing exceeds the number of supported VLANs (i.e. 2048), then evaluate if there are any VLANs that can be removed from the system. 2. Verify that the VLAN pool in the vCloud Director does not contain more than 2048 available VLANs.
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Template could not be inherited on port-profile</pre>	Unable to inherit the port profile associated with the network segment policy onto the port profile created for the network.	<ol style="list-style-type: none"> 1. Verify if the port profile exists. show running-config port-profile name To identify the name of the port profile, you will need to determine the network segment policy the network was attempting to use. You will need the information about the tenant/organization UUID and the type of network pool the network was being created from (VXLAN or VLAN) to find the corresponding network segment policy that has these values configured. If no network segment policy is configured with these values, then use the default network segment policy to identify the name of the port profile. 2. Check system logs for a port profile inheritance failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to set max-ports</pre>	Unable to set the max ports on the port profile.	Check system logs for a max port failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Network already exists</pre>	A network with the same name already exists in the vCloud Director.	<ol style="list-style-type: none"> 1. Delete the existing network that has the same name. no port-profile network name 2. Delete the bridge domain with the same name if it exists. no bridge-domain name

Table 22-1 Problems with Network Segmentation Manager (continued)

Symptom	Possible Causes	Verification and Solution
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to create port-profile	Cisco Nexus 1000V is unable to create the port profile required for the network.	Check system logs for a port profile failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Template does not exist	Unable to find the port profile associated with the network segment policy associated with the network.	<ol style="list-style-type: none"> Verify if the port profile exists. show running-config port-profile name To identify the name of the port profile, you will need to determine the network segment policy the network was attempting to use. You will need the information about the tenant/organization UUID and the type of network pool the network was being created from (VXLAN or VLAN) to find the corresponding network segment policy that has these values configured. If no network segment policy is configured with these values, then use the default network segment policy to identify the name of the port profile. Check system logs for a port profile failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Alias ID not found	Unable to retrieve the port group ID associated with the port profile created for the network	Verify that the Virtual Supervisor Module (VSM) has an active SVS connection. show svvs connection When you enter the command, the output must display operational status: connected
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to set port-binding	Unable to set the port binding on the port profile associated with the network	Check system logs for a port binding failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to set vlan	Unable to set the access VLAN on the port profile associated with the network	Check system logs for a set VLAN failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.

Table 22-1 Problems with Network Segmentation Manager (continued)

Symptom	Possible Causes	Verification and Solution
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to set vmware port-group	Unable to set VMware port group property on the port profile.	Check system logs for a port group property failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to set state enabled	Unable to set the property state on the port profile to enabled.	Check system logs for a state enabled property failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to collect svcs configuration	Unable to execute the command show svcs connection	Verify that the Virtual Supervisor Module (VSM) has an active SVS connection. show svcs connection When you enter the command, the output must display operational status: connected
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Operational status is missing	Unable to locate the operational status in the SVS connection.	<ol style="list-style-type: none"> Verify that the Virtual Supervisor Module (VSM) has an active SVS connection. show svcs connection When you enter the command, the output must display operational status: connected Check system logs for a operational status failure message. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: SVS connection is disconnected	SVS connection is disconnected.	Verify that the Virtual Supervisor Module (VSM) has an active SVS connection. show svcs connection When you enter the command, the output must display operational status: connected
The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director: Failed to create bridge domain	Unable to create the bridge domain associated with the network.	Verify that the feature Segmentation is enabled. show feature If not, enable the segmentation feature. feature segmentation

Table 22-1 Problems with Network Segmentation Manager (continued)

Symptom	Possible Causes	Verification and Solution
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to set segment ID</pre>	Unable to set the segment ID associated with the network.	<p>Verify that the segment ID is not already in use by another bridge domain.</p> <p>show bridge-domain</p> <p>Check the error message on the system log to retrieve the segment ID.</p>
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to set group IP</pre>	Unable to set the group IP associated with the network.	<p>Verify that the group IP is a valid multicast IP address by checking the system logs for invalid IP address error message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.</p>
<p>The network creation triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to set port-profile description</pre>	Unable to set the description for the port profile associated with the network.	<p>Check system logs for a port profile description failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.</p>
<p>The network deletion triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to delete interface using the port-profile</pre>	Unable to delete the interfaces inheriting the port profile.	<ol style="list-style-type: none"> 1. Manually delete the interfaces. 2. In the vCenter Server ensure that the VMs associated with the vApp are powered down. 3. In the VSM execute the command <p>no interface vethernet vethernet number</p>
<p>The network deletion triggered from vCloud Directors fails. A system message similar to the following is logged in the vCloud Director:</p> <pre>Failed to delete the port-profile</pre>	Unable to delete the port profile associated with the network.	<ol style="list-style-type: none"> 1. Manually delete the port profile. 2. Check system logs for a port profile deletion failure message reported by network segmentation manager. See the <i>Cisco NX-OS System Messages Reference</i> for more information.
<p>An vEthernet interface is administratively down. The interface will be in NoPortProfile state.</p>	The vEthernet interface is in quarantine state.	<ol style="list-style-type: none"> 1. Verify the interface is quarantined. <p>show port-profile sync-status</p> 2. Bring the interface out of quarantine. <p>no shutdown</p> <p>The interface comes back online.</p> 3. Verify if the interface is online. <p>show interface vethernet</p>

Network Segmentation Manager Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to the Network Segmentation Manager.

Table 22-2 Network Segmentation Manager Troubleshooting Commands

Command	Purpose
show network-segment manager switch	Displays the Cisco Nexus 1000V configured with NSM.
show running-config port-profile	Displays the port profile configuration.
show running-config network-segment policy	Displays the NSM policy configuration.
show network-segment policy usage	Displays the network segmentation policy usage by networks.
show network-segment network	Displays the networks associated with a network segmentation policy.
show network-segment network id <i>id</i>	Displays the network ids associated with a network segmentation policy.
show network-segment network name <i>name</i>	Displays the name of the networks associated with a network segmentation policy.
show logging logfile grep NSMGR	Displays the system logs from the network segmentation manager.

For detailed information about show command output, see the *Cisco Nexus 1000V Command Reference*.



VXLANs

This chapter describes how to identify and resolve problems that might occur when implementing Virtual Extensible Local Area Networks (VXLANs).

This chapter includes the following sections:

- [Information About VXLANs, page 23-1](#)
- [VXLAN Troubleshooting Commands, page 23-4](#)
- [VEM Packet Path Debugging, page 23-9](#)
- [VEM Multicast Debugging, page 23-10](#)
- [VXLAN Datapath Debugging, page 23-11](#)

Information About VXLANs

- [Overview, page 23-1](#)
- [VXLAN Tunnel EndPoint, page 23-2](#)
- [VXLAN Gateway, page 23-2](#)
- [VXLAN Trunks, page 23-2](#)
- [Multi-MAC Capability, page 23-3](#)
- [Fragmentation, page 23-3](#)
- [Scalability, page 23-3](#)
- [Supported Features, page 23-3](#)

Overview

A Virtual Extensible LAN creates LAN segments by using an overlay approach with MAC-in-UDP encapsulation and a 24-bit segment identifier in the form of a VXLAN ID. The encapsulation carries the original Layer 2 (L2) frame from the Virtual Machine (VM) that is encapsulated from within the Virtual Ethernet Module (VEM). Each VEM is assigned a IP address that are used as the source IP address when encapsulating MAC frames to be sent on the network. You can have multiple vmknics per VEM that are used as sources for this encapsulated traffic. The encapsulation carries the VXLAN identifier used to scope the MAC address of the payload frame. The VXLAN ID to which a VM belongs is indicated within the port profile configuration of the vNIC and is applied when the VM connects to the network. A VXLAN supports three different modes for broadcast, multicast, and MAC distribution mode transport:

- **Multicast Mode**— A VXLAN uses an IP multicast network to send broadcast, multicast, and unknown unicast flood frames. When a VM joins a VXLAN segment, the server joins a multicast group. Broadcast traffic from the VM is encapsulated and is sent using the multicast outer destination IP address to all the servers in the same multicast group. Subsequent unicast packets are encapsulated and unicast directly to the destination server without multicast IP address.
- **Unicast-only Mode**— A VXLAN uses each VEM's single unicast IP address as the destination IP address to send broadcast, multicast, and unknown unicast flood frames. Broadcast traffic from the VM is replicated to each VEM by encapsulating it with a VXLAN header and the designated IP address as the outer destination IP address.
 - **MAC Distribution Mode**(supported only in unicast mode)—In this mode, the unknown unicast flooding is reduced because the Virtual Supervisor Module (VSM) learns all the MAC addresses from the VEMs in all VXLANs and distributes those MAC addresses with VXLAN Tunnel Endpoint (VTEP) IP mappings to other VEMs.

The VXLAN creates LAN segments by using an overlay approach with MAC in IP encapsulation. The encapsulation carries the original Layer 2 (L2) frame from the Virtual Machine (VM) which is encapsulated from within the Virtual Ethernet Module (VEM). Each VEM is assigned an IP address which is used as the source IP address when encapsulating MAC frames to be sent on the network. You can have multiple vmknics per VEM that are used as sources for this encapsulated traffic. The encapsulation carries the VXLAN identifier which is used to scope the MAC address of the payload frame.

VXLAN Tunnel EndPoint

Each VEM requires at least one IP/MAC pair to terminate VXLAN packets. This IP/MAC address pair is known as the VXLAN Tunnel End Point (VTEP) IP/MAC addresses. The VEM supports IPv4 addressing for this purpose. The IP/MAC address that the VTEP uses is configured when you enter the capability vxlan command. You can have a maximum of four VTEPs in a single VEM.

One VTEP per VXLAN segment is designated to receive all broadcast, multicast, and unknown unicast flood traffic for the VEM.

When encapsulated traffic is destined to a VEM that is connected to different subnet, the VEM does not use the VMware host routing table. Instead, the VTEPs initiate the Address Resolution Protocol (ARP) for remote VEM IP addresses. If the VTEPs in the different VEMs are in different subnets, you must configure the upstream router to respond by using the Proxy ARP.

VXLAN Gateway

VXLAN termination (encapsulation and decapsulation) is supported on virtual switches. As a result, the only endpoints that can connect into VXLANs are VMs that are connected to a virtual switch. Physical servers cannot be in VXLANs and routers or services that have traditional VLAN interfaces cannot be used by VXLAN networks. The only way that VXLANs can currently interconnect with traditional VLANs is through VM-based software routers.

VXLAN Trunks

A VXLAN trunk allows you to trunk multiple VXLANs on a single virtual Ethernet interface. In order to achieve this configuration, you must encapsulate a VXLAN-VLAN mapping on the virtual Ethernet interface.

VXLAN-VLAN mappings are configured through the VSM and must always be a 1:1 mapping for each Layer 2 domain. VXLAN-VLAN mappings are applied on a virtual Ethernet interface using a port-profile. A single port profile can support multiple VLAN-VXLAN mappings.

Multi-MAC Capability

You can use multi-MAC addresses to mark a virtual Ethernet interface as capable of sourcing packets from multiple MAC addresses. For example, you can use this feature if you have a virtual Ethernet port and you have enabled VXLAN trunking on it and the VM that is connected to the port bridges packets that are sourced from multiple MAC addresses.

By using this feature, you can easily identify such multi-MAC capable ports and handle live migration scenarios correctly for those ports.

Fragmentation

The VXLAN encapsulation overhead is 50 bytes. In order to prevent performance degradation due to fragmentation, the entire interconnection infrastructure between all VEMs exchange VXLAN packets must be configured to carry 50 bytes more than what the VM VNICs are configured to send. For example, if the default VNIC configuration is 1500 bytes, the VEM uplink port profile, upstream physical switch port, and interswitch links, and any routers if present, must be configured to carry a maximum transmission unit (MTU) of at least 1550 bytes. If that is not possible, we recommend that the MTU within the guest VMs you configure to be smaller by 50 bytes.

If you do not configure a smaller MTU, the VEM attempts to notify the VM if it performs Path MTU (PMTU) Discovery. If the VM does not send packets with a smaller MTU, the VM fragments the IP packets. Fragmentation occurs only at the IP layer. If the VM sends a frame that is too large, the frame will be dropped after VXLAN encapsulation and if the frame does not contain an IP packet.

Scalability

Maximum Number of VXLANs

The Cisco Nexus 1000V supports a total of 4096 VLANs or VXLANs (or a maximum of 2048 VLANs or 2048 VXLANs in any combination that totals 4096).

Supported Features

This section contains the following topics:

- [Jumbo Frames, page 23-3](#)
- [Disabling the VXLAN Feature Globally, page 23-4](#)

Jumbo Frames

Jumbo frames are supported by the Cisco Nexus 1000V if there is space on the frame to accommodate the VXLAN encapsulation overhead of at least 50 bytes, and the physical switch/router infrastructure has the capability to transport these jumbo-sized IP packets.

Disabling the VXLAN Feature Globally

As a safety precaution, do not use the no feature segmentation command if there are any ports associated with a VXLAN port profile. You must remove all associations before you can disable this feature. You can use the no feature segmentation command to remove all the VXLAN bridge domain configurations on the Cisco Nexus 1000V.

VXLAN Troubleshooting Commands

Use the following commands to display VXLAN attributes.

This section contains the following topics:

- [VSM Commands, page 23-4](#)
- [VXLAN Gateway Commands, page 23-5](#)

VSM Commands

To display ports belonging to a specific segment:

```
switch(config)# show system internal seg_bd info segment 10000
Bridge-domain: A
Port Count: 11
Veth1
Veth2
Veth3
```

To display the vEthernet bridge domain configuration:

```
switch(config)# show system internal seg_bd info port vethernet 1
Bridge-domain: A
segment_id = 10000
Group IP: 225.1.1.1
```

To display the vEthernet bridge configuration with ifindex as an argument:

```
switch(config)# show system internal seg_bd info port ifindex 0x1c000050
Bridge-domain: A
segment_id = 10000
Group IP: 225.1.1.1
```

To display the total number of bridge domain ports:

```
switch(config)# show system internal seg_bd info port_count
Number of ports: 11
```

To display the bridge domain internal configuration:

```
switch(config)# show system internal seg_bd info bd vxlan-home

Bridge-domain vxlan-home (2 ports in all)
Segment ID: 5555 (Manual/Active)
Group IP: 235.5.5.5
State: UP                               Mac learning: Enabled
is_bd_created: Yes
current state: SEG_BD_FSM_ST_READY
pending_delete: 0
port_count: 2
action: 4
```

```

hwbd: 28
pa_count: 0
Veth2, Veth5
switch(config)#

```

To display VXLAN vEthernet information:

```

switch# show system internal seg_bd info port
if_index = <0x1c000010>
Bridge-domain vxlan-pepsi
rid = 216172786878513168
swbd = 4098

if_index = <0x1c000040>
Bridge-domain vxlan-pepsi
rid = 216172786878513216
swbd = 4098

switch#

```

Additional **show** commands:

```

show system internal seg_bd info {pss | sdb | global | all}

show system internal seg_bd {event-history | errors | mem-stats | msgs}

show system internal seg_bd info (sdb | bd)

```

VXLAN Gateway Commands

To display VXLAN Gateway information attached to VSM:

```

switch# show module vem
Mod  Ports  Module-Type                Model                Status
---  ---
3    7      Virtual Service Module     VXLAN Gateway       ok

```

To display VXLAN Gateway information that is not attached to the VSM:

```

VXLANGW# attach vem
VXLANGW(vem-attach)# ?
  vemcmd      Execute vem command
  vemdpa      Execute vemdpa command
  vemdpalog   Execute vemdpalog command
  vemlog      Execute vemlog command
  vempkt      Execute vempkt command
  vemset      Execute vemset command
switch(vem-attach)#

```

To display VXLAN Gateway statistics:

```

switch(vem-attach)# vemcmd show vxlan-stats
  LTL  Ucast  Mcast/Repl  Ucast  Mcast  Total
      Encaps  Encaps  Decaps  Decaps  Drops
  17   8717   173       8334   0      242
switch(vem-attach)#

switch(vem-attach)# vemcmd show vxlan-stats ltl 17
VXLAN Port Stats for LTL 17
Unicast Encapsulations: 8756

```

```

Multicast Encapsulations/HeadEnd Replications: 173
Unicast Decapsulations: 8372
Multicast Decapsulations: 0
IP Pre-fragmentations: 0
TSO Processed Packets: 0
ICMP Pkt Too Big msgs from upstream: 0
ICMP Pkt Too Big msgs sent to VM: 0
Packets generated by Head End Replication: 172

```

To display the VXLAN Gateway packet path:

```
switch(vem-attach)# vemlog show all
```

To display the bridge-domain configuration on VSM:

```
switch# show bridge-domain
```

Note - This command is common for both gateway and VEM.

```

Global Configuration:
Mode: Unicast-only
MAC Distribution: Disable
Note - If you have enabled MAC distribution, the above command will display Enable.
Bridge-domain segment-cisco (3 ports in all)
Segment ID: 9001 (Manual/Active)
Mode: Unicast-only (default)
MAC Distribution: Disable (default)
Group IP: NULL
State: UP                               Mac learning: Enabled
Veth2, Veth3, Veth5

```

To display the vlan-vxlan mappings programmed on the VSM:

```
switch# show bridge-domain mapping
```

To display the interfaces on the VSM:

```
switch# show module vtep
```

To display the the bridge domain-to-VTEP mappings that are maintained by the VSM and are pushed to all VEMs:

```
switch# show bridge-domain vtep
```

To displays the MACs learnt on VSM through VEM distribution:

```
switch# sho bridge-domain mac
```

```

Bridge-domain: segment-cisco
MAC TABLE Version: 1
Note: You can compare with VEM output using the echo show vxlan version-table command.
MAC Address      Module      Port      VTEP-IP Address  VM-IP Address
-----
0050.5683.014e   5           Veth5     10.106.199.117  -
0050.5683.0160   4           Veth2     10.106.199.116  -
0050.5683.0161   4           Veth3     10.106.199.116  -

```

To verify the port configuration on VSM:

```
switch# show int switchport | begin Vethernet2
```

```

Name: Vethernet2
Switchport: Enabled
Switchport Monitor: Not enabled
Operational Mode: access
Access Mode VLAN: 0 (none)

```

```
Access BD name: segment-cisco
```

To verify the VTEP distribution on VSM:

```
switch# show bridge-domain segment-cisco vteps
```

```
D: Designated VTEP      I:Forwarding Publish Incapable VTEP
```

```
Bridge-domain: segment-cisco
```

```
VTEP Table Version: 2
```

Note: You can compare the VTEP table version with the echo show vxlan version-table on VEM.

```
Ifindex      Module  VTEP-IP Address
```

```
-----
```

```
Veth4        4      10.106.199.116(D)
```

```
Veth1        5      10.106.199.117(D)
```

```
switch#
```

Additional **show** commands:

```
show platform fwm errors
```

```
show platform fwm info (vtep | trace | error history)
```

```
show platform fwm info error history
```

```
show platform fwm event-history msgs
```

```
show platform fwm info vlan (all|swbd)
```

VEM Commands

To verify VXLAN vEthernet programming:

```
~ # vemcmd show port segments
```

LTL	VSM Port	Mode	Native SegID	Seg State
50	Veth5	A	5555	FWD
51	Veth9	A	8888	FWD

```
~ #
```

To verify VXLAN vmknic programming:

```
~ # vemcmd show vxlan interfaces
```

LTL	IP	Seconds since Last IGMP Query Received
49	10.3.3.3	50 *
52	10.3.3.6	50

```
~ #
```

Use "vemcmd show port vlans" to verify that the vmknics are in the correct transport VLAN.

To verify bridge domain creation on the VEM:

```
~ # vemcmd show bd bd-name vxlan-home
```

```
BD 31, vdc 1, segment id 5555, segment group IP 235.5.5.5, swbd 4098, 1 ports, "vxlan-home"
```

```
Portlist:
```

```
50 RedHat_VM1.eth0
```

```
~ #
```

To verify remote IP learning:

```
~ # vemcmd show 12 bd-name vxlan-home
Bridge domain 31 brtmax 4096, brtcnt 2, timeout 300
Segment ID 5555, swbd 4098, "vxlan-home"
Flags: P - PVLAN S - Secure D - Drop
      Type          MAC Address  LTL  timeout  Flags  PVLAN  Remote IP
      Dynamic      00:50:56:ad:71:4e  305   2        2      10.3.3.100
      Static       00:50:56:85:01:5b   50   0        0      0.0.0.0
```

```
~ #
```

To display statistics:

```
~ # vemcmd show vxlan-stats
      LTL  Ucast  Mcast  Ucast  Mcast  Total
      Encaps  Encaps  Decaps  Decaps  Drops
      49    5   14265    4    15    0
      50    6   14261    4    15   213
      51    1    15     0     0    10
      52    0    11     0     0    15
```

```
~ #
```

To display detailed per-port statistics for a VXLAN vEthernet/vmknic:

```
~ # vemcmd show vxlan-stats ltl 51
```

To display detailed per-port-per-bridge domain statistics for a VXLAN vmknic for all bridge domains:

```
~ # vemcmd show vxlan-stats ltl <vxlan_vmknic_ltl> bd-all
```

To display detailed per-port-per-bridge domain statistics for a VXLAN vmknic for a specified bridge domain:

```
~ # vemcmd show vxlan-stats ltl vxlan_vmknic_ltl bd-name bd-name
```

To verify the bridge-domain configuration on VEM:

```
switch# vemcmd show bd bd-name segment-cisco
Note - Use the module command to check the details of VEM and gateway on the VSM.

BD 26, vdc 1, segment id 9001, segment group IP 0.0.0.0, swbd 4102, 2 ports,
"segment-cisco"
Segment Mode: Unicast
Note: If MAC distribution is enabled, the above command will displays Segment mode as
Unicast MAC distribution
VTEP DSN: 1 , MAC DSN: 1
Note: You can check the VTEP and MAC download sequence numbers using the vemcmd show
vxlan-vteps and vemcmd show 12 bd bd-name commands.
Portlist:
      53 RedHat_VM1_112.eth4
      54 RedHat_VM1_112.eth5
~ #
```

To display the MAC address table that shows the MACs pushed by the VSM:

```
switch# vemcmd show 12 bd-name segment-cisco
Bridge domain 26 brtmax 4096, brtcnt 3, timeout 300
Segment ID 9001, swbd 4102, "segment-cisco"
Flags: P - PVLAN S - Secure D - Drop
      Type          MAC Address  LTL  timeout  Flags  PVLAN  Remote IP  DSN
      SwInsta      00:50:56:83:01:4e  561   0        0      10.106.199.117  1
```

```

Static 00:50:56:83:01:61 54 0 0.0.0.0 1
Static 00:50:56:83:01:60 53 0 0.0.0.0 1

```

switch#

To verify the port configuration on VEM:

```

switch# vemcmd show port
LTL   VSM Port  Admin Link  State  PC-LTL  SGID  Vem Port  Type
17    Eth4/1    UP    UP    F/B*    561    0    vmnic0
49                    DOWN  UP    BLK     0      0    RedHat_VM1_112 ethernet7
50    Veth8     DOWN  UP    BLK     0      0    RedHat_VM1_112.eth8
51    Veth4     UP    UP    FWD     0      0    vmk1    VXLAN
52                    DOWN  UP    BLK     0      0    RedHat_VM1_112.eth6
53    Veth2     UP    UP    FWD     0      0    RedHat_VM1_112.eth4
54    Veth3     UP    UP    FWD     0      0    RedHat_VM1_112.eth5
561   Po2       UP    UP    F/B*    0

```

To verify the VTEP distribution on VEM:

```

switch# vemcmd show vxlan-vteps
Bridge-Domain: segment-cisco Segment ID: 9001
Designated Remote VTEP IPs (*=forwarding publish incapable):
10.106.199.117(DSN: 1),
Note: You can compare the download sequence number against the VTEP download sequence
number using the vemcmd show bd bd-name.

```

To verify if the MAC address table displays the remote IP learning in the segment-cisco bridge domain:

```

switch# vemcmd show 12 bd-name segment-cisco
Note - Use the module command to check the details of VEM and gateway on the VSM.

Bridge domain 26 brtmax 4096, brtcnt 3, timeout 300
Segment ID 9001, swbd 4102, "segment-cisco"
Flags: P - PVLAN S - Secure D - Drop
      Type      MAC Address  LTL  timeout  Flags  PVLAN  Remote IP  DSN
      Dynamic   00:50:56:83:01:4e 561   1        10.106.199.117 0
      Static    00:50:56:83:01:61 54    0        0.0.0.0 0
      Static    00:50:56:83:01:60 53    0        0.0.0.0 0

```

To display the vlan-vxlan mappings programmed on a VEM:

```

switch# vemcmd show vlan-vxlan mapping
Note - Use the module command to check the details of VEM and gateway on the VSM.

```

To display the multi-MAC capable interfaces on a VEM:

```

Note - Use the module command to check the details of VEM and gateway on the VSM.
switch# vemcmd show multi-mac-capable interfaces

```

VEM Packet Path Debugging

Use the following commands to debug VXLAN traffic from a VM on VEM1 to a VM on VEM2.

- VEM1: Verify that packets are coming into the switch from the segment vEthernet.

```
vempkt capture ingress lt1 vxlan_veth
```

- VEM1: Verify VXLAN encapsulation.

```
vemlog debug sflisp all
vemlog debug sfvnsegment all
```

- VEM1: Verify remote IP is learned:

```
vemcmd show l2 bd-name segbdname
```

If the remote IP is not learned, packets are sent multicast encapsulated. For example, an initial ARP request from VM is sent in this manner.

- VEM1: Verify encapsulated packets go out uplink.

Use the **vemcmd show vxlan-encap ltl ltl** command or the **vemcmd show l2lisp-encap mac mac** to find out which uplink is being used.

```
vempkt capture egress ltl uplink
```

- VEM1: Look at statistics for any failures.

```
vemcmd show vxlan-stats all
vemcmd show vxlan-stats ltl veth/vxlanvmknic
```

- VEM2: Verify encapsulated packets are arriving on the uplink.

```
vempkt capture ingress ltl uplink
```

- VEM2: Verify VXLAN decapsulation.

```
vemlog debug sflisp all
vemlog debug sfvnsegment all
```

- VEM2: Verify decapsulated packets go out on VXLAN vEthernet.

```
vempkt capture egress ltl vxlan_veth
```

- VEM2: Look at statistics for any failures:

```
vemcmd show vxlan-stats all
vemcmd show vxlan-stats ltl veth/vxlanvmknic
```

Use the following commands to debug the VXLAN packet path:

```
switch# module vem 4 execute vemlog debug vssnet all
switch# module vem 4 execute vemlog debug sfsched all
switch# module vem 4 execute vemlog debug sfport all
switch# module vem 4 execute vemlog debug sflisp all
switch# module vem 4 execute vemlog debug sfvnsegment all
```

Use the following commands to debug the VXLAN packet path from the VSM:

```
switch# module vem 4 execute vemdpalog debug if_bridge_rt all
switch# module vem 4 execute vemdpalog debug sfbid all
switch# module vem 4 execute vemdpalog debug sf_dp_threads all
switch# module vem 4 execute vemdpalog debug sf12agent all
switch# module vem 4 execute vemlog debug sfporttable all
```

You can view the output for all the above logs by using the **module vem 4 execute vemlog show all** command.

VEM Multicast Debugging

Use the following command to debug VEM multicast.

- IGMP state on the VEM:

```
vemcmd show igmp vxlan_transport_vlan detail
```


**Note**

This command does not show any output for the segment multicast groups. To save multicast table space, segment groups are not tracked by IGMP snooping on the VEM.

- IGMP queries:

Use the **vemcmd show vxlan interfaces** command to verify that IGMP queries are being received.

- IGMP joins from vmknic:

Use the **vempkt capture ingress ltl *first_vxlan_vmknic_ltl*** command to see if the VMware stack is sending joins.

Use the **vempkt capture egress ltl *uplink_ltl*** command to see if the joins are being sent out to the upstream switch.

VXLAN Datapath Debugging

Use the commands listed in this section to troubleshoot VXLAN problems.

This section contains the following topics:

- [Vemlog Debugging, page 23-11](#)
- [Vempkt, page 23-12](#)
- [Statistics, page 23-12](#)
- [Show Commands, page 23-13](#)

Vemlog Debugging

To debug the bridge domain setup or configuration, use the following command:

```
vemlog debug sfbid all
```

To debug port configuration/CBL/vEthernet LTL pinning, use the following command:

```
vemlog debug sfporttable all
```

(for encap/decap setup and decisions)

```
vemlog debug sfvnsegment all
```

To debug for actual packet editing, VXLAN interface handling, and multicast handling, use the following command:

```
vemlog debug sflisp all
```

To debug multicast joins or leaves on the DPA socket, use the following command:

```
echo "debug dpa_allplatform all" > /tmp/dpafifo
```

To debug the bridge domain configuration, use the following command:

```
echo "debug sf12agent all" > /tmp/dpafifo
```

To debug port configuration, use the following command:

```
echo "debug sfportagent all" > /tmp/dpafifo
```

To debug hitless reconnect (HR) for capability l2-lisp, use the following command:

```
echo "debug sfpportl2lisp_cache all" > /tmp/dpafifo
```

To debug CBL programming.

```
echo "debug sfpixmagent all" > /tmp/dpafifo
```

To debug VXLAN agent interacting with the VSM, use the following command:

```
echo "debug sfvxlanagent all" > /tmp/dpafifo
```

To check the VTEP and MAC version, use the following command:

```
~ # echo "show vxlan version-table" > /tmp/dpafifo
```

Content written to /var/log/vemdpa.log

Slot	SWBD	VTEP Version	MAC Version
0	4096	0	0
6	4102	2	1
7	4103	0	1

Note: You can compare the MAC version output on the VSM using the show bridge-domain mac command and VTEP version output on the VSM using the show bridge-domain vtep command.

To check the MACs to be distributed on the VSM, use the following command:

```
~ # echo "show vxlan mac-table" > /tmp/dpafifo
```

Content written to /var/log/vemdpa.log

Flags: R - Report to VSM I - VSM Informed

Del* - Stale entry in VSM

No VTEP - NO VTEP. Entry to be removed from VSM

Count	BD	MAC Address	Iindex	Id	VTEP	Flags	VM-IP
0	4102	00:50:56:83:01:61	0x1c000020	5422209	10.106.199.116	I	0
0	4102	00:50:56:83:01:60	0x1c000010	5422209	10.106.199.116	I	0

Vempkt

Vempkt has been enhanced to display VLAN/SegmentID. Use vempkt to trace the packet path through VEM.

- Encap: Capture ingress on Seg-VEth LTL – Egress on uplink
- Decap: Capture ingress on uplink – Egress on Seg-VEth LTL

Statistics

To display a summary of per-port statistics, use the following command:

```
vemcmd show vxlan-stats
```

To display detailed per-port statistics for VXLAN vmknic, use the following command:

```
vemcmd show vxlan-stats lt1 vxlan_vmknic_lt1
```

To display detailed per-port statistics for vEthernet in a VXLAN, use the following command:

```
vemcmd show vxlan-stats lt1 vxlan_veth_ltl
```

To display detailed per-port-per-bridge domain statistics for a VXLAN vmknic for all bridge domains, use the following command:

```
vemcmd show vxlan-stats lt1 vxlan_vmknic_ltl bd-all
```

To display detailed per-port-per-bridge domain statistics for a VXLAN vmknic for the specified bridge domain, use the following command:

```
vemcmd show vxlan-stats lt1 vxlan_vmknic_ltl bd-name bd-name
```

To display which VXLAN vmknic used for encap and subsequent pinning to uplink PC for static MAC learned on port, use the following command:

```
vemcmd show vxlan-encap lt1 vxlan_veth_ltl
```

To display which VXLAN vmknic used for encapsulation and subsequent pinning to uplink PC, use the following command:

```
vemcmd show vxlan-encap mac vxlan_vm_mac
```

Show Commands

Table 23-1 lists available **vemcmd show** commands.

Table 23-1 vemcmd Show Commands

Command	Result
vemcmd show vxlan interfaces	Displays the VXLAN encapsulated interfaces.
vemcmd show port vlans	Checks the port programming and CBL state for the bridge domain.
vemcmd show bd	Displays the bridge domain segmentId/group/list of ports.
vemcmd show bd bd-name <i>bd-name-string</i>	Displays one segment bridge domain.
vemcmd show l2 all	Displays the remote IP being learned.
vemcmd show l2 bd-name <i>bd-name-string</i>	Displays the Layer 2 table for one segment bridge domain.
vemcmd show arp all	Displays the IP-MAC mapping for the outer encapsulated header.



VSI Discovery and Configuration Protocol

This chapter describes how to identify and resolve problems that might occur when implementing the VSI Discovery and Configuration Protocol (VDP).

This chapter includes the following sections:

- [Information About VDP, page 24-1](#)
- [Problems with VDP, page 24-2](#)
- [VDP Troubleshooting Commands, page 24-2](#)

Information About VDP

VDP on the Cisco Nexus 1000V is an implementation of the IEEE standard 802.1Qbg/D2.2 (Edge Virtual Bridging). VDP can detect and signal the presence of end-hosts and exchange capability with an adjacent VDP-capable bridge. VDP serves as a reliable first-hop protocol and communicates the presence of end-host Virtual Machines (VMs) to adjacent leaf nodes on the Cisco Dynamic Fabric Automation (DFA) architecture. In addition to detecting the MAC and IP addresses of the end-host VMs when a host comes up, or during VM mobility events, the VDP triggers auto-configuration of leaf nodes on the DFA architecture to make them ready for more VM traffic.

VDP enables network-based overlays which are a more scalable alternative when compared to the host-based overlays for segmentation and enable access to more than four thousand VLANs in a multi-tenant network. With the VDP configured on the Cisco Nexus 1000V, segmentation support for bridge-domains is extended to native encapsulated bridge-domains. The original VXLAN-based bridge domains can also coexist with these bridge domains.

For more information about the Cisco DFA architecture, see the *Cisco DFA Solutions Guide*.

Problems with VDP

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

Table 24-1 Problems with VDP

Symptom	Possible Causes	Solution
VDP packets not received by a leaf	The connected port on the VEM does not have the trunk dynamic port profile.	<ol style="list-style-type: none"> Verify that the connected port on the VEM has the trunk dynamic port profile: show interface ethernet slot/port If the output of the show interface ethernet command does not contain “Allowing Dynamic VLANs”, configure the port profile for trunk dynamic mode: <ol style="list-style-type: none"> switch# configure terminal switch(config)# port-profile name switch(config-port-prof)# switchport mode trunk switch(config-port-prof)# switchport trunk dynamic
VM is associated but it is not pinging	The encapsulation mode is not native.	<p>Verify that encapsulation mode is native and a valid VLAN value is returned by the leaf:</p> <p>module vem module_number execute vemcmd show bd</p> <p>module vem module_number execute vemcmd show segment segment_id</p>

VDP Troubleshooting Commands

This section contains the following topics:

- [VSM Commands, page 24-2](#)
- [VEM Commands, page 24-4](#)

VSM Commands

You can use the commands in this section to troubleshoot problems related to VDP.

Command	Purpose
show evb vsi interface vethernet interface-number	Displays if VDP association sequence is complete for a vEthernet interface. Identify the vEthernet port of the VM and use this command. A VSI state of 3 means it is associated See Example 24-1 on page 24-3 .
show evb	Displays configured information in the EVB process. See Example 24-2 on page 24-3 .

Command	Purpose
show run evb	Displays the running configuration for the EVB segmentation. See Example 24-3 on page 24-3 .
show ecp	Displays the configured information for ECP. See Example 24-4 on page 24-3 .

EXAMPLES

Example 24-1 show evb vsi interface

```
switch(config)# show evb vsi interface vethernet 40
LTL : 135 [module: 2]
Segment : 30000
MAC : 0050.5693.63A1
IP : 30.0.1.2
VSI State : 3
State Machine State : 7
Rwd Expiry Count : 4621
Last CMD Time : 125
Last RSP Time : 125
```

Example 24-2 show evb

```
switch(config)# show evb
Edge Virtual Bridging
Role : VDP Station
VDP Mac Address : 0180.0000.0000
VDP Resource Wait Delay : 22(66 secs)
VDP Reinit Keep Alive : 21(20 secs)
```

Example 24-3 show run evb

```
switch(config)# show run evb
evb resource-wait-delay 24
evb reinit-keep-alive 25
ecp retransmission-timer-exponent 15
ecp max-retries 6
```

Example 24-4 show ecp

```
switch(config)# show ecp
ECP Max ReTries : 3
ECP Retransmission Timer Exp : 14(163840 micro seconds)
```

VEM Commands

You can use the VEM commands in this section to troubleshoot problems related to VDP.

Command	Purpose
vemcmd show segment <i>segment-id</i>	Displays a list of VM interfaces that are a part of a segment and indicates if a segment is configured as VDP (native encapsulation mode). See Example 24-5 on page 24-4 .
vemcmd show bd <i>hwbd</i>	Displays a list of VM interfaces that are a part of an internal bridge domain and indicates if the bridge domain is configured as VDP (native encapsulation mode). See Example 24-6 on page 24-5 .
vemcmd show bd <i>bd-name</i> <i>bd-name</i>	Displays a list of VM interfaces that are a part of a configured bridge domain and indicates if the bridge domain is configured as VDP (native encapsulation mode). See Example 24-7 on page 24-5 .

EXAMPLES

Example 24-5 vemcmd show segment

```

~ # vemcmd show segment 8000
BD 21, vdc 1, segment id 8000, segment group IP 224.9.19.10, encap NATIVE, vff_mode
Anycast,swbd 4098, VLAN 0, 28 ports, "BD-Mcast"
Segment Mode: Multicast
Portlist:
    52 VM-L-13-25-10.eth7
    62 VM-L-13-25-2.eth7
    72 VM-L-13-25-1.eth7
    82 VM-L-13-25-3.eth7
    92 VM-L-13-25-7.eth7
   102 VM-L-13-25-5.eth7
   112 VM-L-13-25-4.eth7
   122 VM-L-13-25-6.eth7
   132 VM-L-13-25-8.eth7
   144 VM-L-14-25-1.eth7

   145 VM-L-14-25-2.eth7
   162 VM-L-14-25-10.eth7
   172 VM-L-14-25-3.eth7
   182 VM-L-13-25-9.eth7
   192 VM-L-14-25-4.eth7
   202 VM-L-14-25-8.eth7
   212 VM-L-14-25-7.eth7
   222 VM-L-14-25-6.eth7
   232 VM-L-14-25-5.eth7
   242 VM-L-14-25-9.eth7

   252 VM-L-15-25-10.eth7
   262 VM-L-15-25-3.eth7
   272 VM-L-15-25-2.eth7
   282 VM-L-15-25-1.eth7

```



```
294 VM-L-15-25-7.eth7
295 VM-L-15-25-4.eth7
312 VM-L-15-25-5.eth7
322 VM-L-15-25-6.eth7
```

Example 24-6 `vemcmd show bd`

```
~ # vemcmd show bd 21
BD 21, vdc 1, segment id 8000, segment group IP 224.9.19.10, encap NATIVE, vff_mode
Anycast,swbd 4098, VLAN 0, 28 ports, "BD-Mcast"
Segment Mode: Multicast
Portlist:
    52 VM-L-13-25-10.eth7
    62 VM-L-13-25-2.eth7
    72 VM-L-13-25-1.eth7
    82 VM-L-13-25-3.eth7
    92 VM-L-13-25-7.eth7
   102 VM-L-13-25-5.eth7
   112 VM-L-13-25-4.eth7
   122 VM-L-13-25-6.eth7
   132 VM-L-13-25-8.eth7
   144 VM-L-14-25-1.eth7

   145 VM-L-14-25-2.eth7
   162 VM-L-14-25-10.eth7
   172 VM-L-14-25-3.eth7
   182 VM-L-13-25-9.eth7
   192 VM-L-14-25-4.eth7
   202 VM-L-14-25-8.eth7
   212 VM-L-14-25-7.eth7
   222 VM-L-14-25-6.eth7
   232 VM-L-14-25-5.eth7
   242 VM-L-14-25-9.eth7

   252 VM-L-15-25-10.eth7
   262 VM-L-15-25-3.eth7
   272 VM-L-15-25-2.eth7
   282 VM-L-15-25-1.eth7
   294 VM-L-15-25-7.eth7
   295 VM-L-15-25-4.eth7
   312 VM-L-15-25-5.eth7
   322 VM-L-15-25-6.eth7
```

Example 24-7 `vemcmd show bd bd-name`

```
~ # vemcmd show bd bd-name BD-Mcast
BD 21, vdc 1, segment id 8000, segment group IP 224.9.19.10, encap NATIVE, vff_mode
Anycast,swbd 4098, VLAN 0, 28 ports, "BD-Mcast"
Segment Mode: Multicast
Portlist:
    52 VM-L-13-25-10.eth7
    62 VM-L-13-25-2.eth7
    72 VM-L-13-25-1.eth7
    82 VM-L-13-25-3.eth7
    92 VM-L-13-25-7.eth7
   102 VM-L-13-25-5.eth7
   112 VM-L-13-25-4.eth7
   122 VM-L-13-25-6.eth7
   132 VM-L-13-25-8.eth7
   144 VM-L-14-25-1.eth7
```

```
145 VM-L-14-25-2.eth7
162 VM-L-14-25-10.eth7
172 VM-L-14-25-3.eth7
182 VM-L-13-25-9.eth7
192 VM-L-14-25-4.eth7
202 VM-L-14-25-8.eth7
212 VM-L-14-25-7.eth7
222 VM-L-14-25-6.eth7
232 VM-L-14-25-5.eth7
242 VM-L-14-25-9.eth7

252 VM-L-15-25-10.eth7
262 VM-L-15-25-3.eth7
272 VM-L-15-25-2.eth7
282 VM-L-15-25-1.eth7
294 VM-L-15-25-7.eth7
295 VM-L-15-25-4.eth7
312 VM-L-15-25-5.eth7
322 VM-L-15-25-6.eth7
```



Cisco TrustSec

This chapter describes how to identify and resolve problems that might occur when configuring Cisco TrustSec.

This chapter includes the following sections:

- [Information About Cisco TrustSec, page 25-1](#)
- [Guidelines and Limitations for Troubleshooting Cisco TrustSec, page 25-1](#)
- [Cisco TrustSec Troubleshooting Commands, page 25-2](#)
- [Problems with Cisco TrustSec, page 25-4](#)

Information About Cisco TrustSec

The Cisco TrustSec security architecture builds secure networks by establishing clouds of trusted network devices. Each device in the cloud is authenticated by its neighbors. Communication on the links between devices in the cloud is secured with a combination of encryption, message integrity checks, and data-path replay protection mechanisms.

Cisco TrustSec also uses the device and user identification information acquired during authentication for classifying, or coloring, the packets as they enter the network. This packet classification is maintained by tagging packets on ingress to the Cisco TrustSec network so that they can be properly identified for the purpose of applying security and other policy criteria along the data path. The tag, also called the security group tag (SGT), allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic.

See the *Cisco Nexus 1000V Security Configuration Guide* for more information on the Cisco TrustSec feature on Cisco Nexus 1000V.

Guidelines and Limitations for Troubleshooting Cisco TrustSec

The following guidelines and limitations apply when troubleshooting Cisco TrustSec SXP:

- In this release, SGT Exchange Protocol (SXP) is supported for Cisco Nexus 1000V.
- Cisco Nexus 1000V VSM will always be configured as the SXP speaker in all peer connections. Listener functionality is not supported in this release.
- A maximum of 2048 IP-SGT mappings can be learned system-wide in the DVS. This is a combined total for both entries learned via DHCP snooping as well as device tracking of individual virtual machines by ARP as well as IP traffic inspection.

- The IP-SGT mappings can be communicated to up to 64 SXP peer devices.
- In order to assign a SGT to a virtual machine, SGT interactions need to be manually configured in the port profile or vEthernet interface. This is not supported on a management interface or a ethernet interface.

Cisco TrustSec Troubleshooting Commands

This section contains the following topics:

- [Debugging Commands, page 25-2](#)
- [Host Logging Commands, page 25-3](#)
- [Show Commands, page 25-4](#)

Debugging Commands

Table 25-1 lists the available debugging commands.

Table 25-1 Cisco TrustSec Debugging Commands

Command	Purpose
debug cts authentication	Collect and view logs related to Cisco TrustSec authentication.
debug cts authorization	Collect and view logs related to Cisco TrustSec authorization.
debug cts errors	Collect and view logs related to Cisco TrustSec errors and warning messages.
debug cts messages	Collect and view logs related to Cisco TrustSec messages.
debug cts packets	Collect and view logs related to Cisco TrustSec packets.
debug cts relay	Collect and view logs related to Cisco TrustSec relay functionality.
debug cts sxp	Collect and view logs related to Cisco TrustSec SXP.
debug cts sap	Collect and view logs related to Cisco TrustSec security association protocol (SAP).
debug cts trace	Collect and view logs related to Cisco TrustSec trace functionality.
show cts internal debug-info	Displays Cisco TrustSec debug information.

Host Logging Commands

Table 25-2 lists the commands from the ESX host to collect and view logs related to Cisco TrustSec.

Table 25-2 ESX Host Commands

ESX Host Command	Description
<code>echo "logfile enable" > /tmp/dpafifo</code>	Enables DPA debug logging. Logs are output to <code>/var/log/vemdpa.log</code> file.
<code>echo "debug sfctsagent all" > /tmp/dpafifo</code>	Enables TrustSec SXP agent debug logging. Logs are output to <code>/var/log/vemdpa.log</code> file.
<code>vemlog debug sfcts_config all</code>	Enables datapath debug logging, and captures logs for the data packets sent between the client and the server.
<code>vemlog debug sfdhcps_config all</code>	Enables datapath debug logging, and captures logs for DHCP snooping configuration coming from the VSM. To view the logs DHCP snooping should be enable in Cisco Nexus 1000V.
<code>vemlog debug sfdhcps_binding_table all</code>	Enables datapath debug logging, and captures logs corresponding to binding database changes. To view the logs DHCP snooping should be enabled on Cisco Nexus 1000V.
<code>vemlog debug sfipdb all</code>	Enables datapath debug logging, and captures logs corresponding to IP database that maintains the IP addresses for all the virtual machines that are being tracked using Cisco TrustSec device tracking. To view the logs Cisco TrustSec device tracking should be enabled on Cisco Nexus 1000V.
<code>vemcmd show learnt ip</code>	Displays Cisco TrustSec configuration on Cisco Nexus 1000V.
<code>vemcmd show cts global</code>	Displays if Cisco TrustSec is enabled on Cisco Nexus 1000V.
<code>vemcmd show cts ipsqt</code>	Displays Cisco TrustSec configuration on Cisco Nexus 1000V.

Example

The following examples displays Cisco TrustSec specific information on Cisco Nexus 1000V.

```
switch# vemcmd show learnt ip
IP Address LTL VLAN BD
/SegID
10.78.1.76 49 353 7
switch#
```

```
switch# vemcmd show cts global
CTS Global Configuration:
CTS is: Enabled
CTS Device Tracking is: Enabled
switch#
```

```
switch# vexec show cts ipsgt
IP Address LTL VLAN BD SGT Learnt
10.78.1.76 49 353 7 6766 Device Tracking
switch#
```

Show Commands

Table 25-3 lists available Cisco TrustSec show commands. See the *Cisco Nexus 1000V Command Reference* for more information on the show commands for Cisco TrustSec.

Table 25-3 Cisco TrustSec Show Commands

Command	Purpose
show cts	Displays Cisco TrustSec configuration.
show cts sxp	Displays the SXP configuration for Cisco TrustSec.
show feature	Displays the features available, such as CTS, and whether they are enabled.
show running-configuration cts	Displays the running configuration information for Cisco TrustSec.
show cts device tracking	Displays the Cisco TrustSec device tracking configuration.
show cts ipsgt entries	Display the SXP SGT entries for Cisco TrustSec.
show cts role-based sgt-map	Displays the mapping of the IP address to SGT for Cisco TrustSec.
show cts sxp connection	Displays SXP connections for Cisco TrustSec.
show cts interface delete-hold timer	Displays the interface delete hold timer period for Cisco TrustSec.
show cts internal event-history [error mem-stats msgs sxp]	Displays event logs for Cisco TrustSec.

Problems with Cisco TrustSec

This section includes symptoms, possible causes and solutions for the following problems with Cisco TrustSec.

Table 25-1 Problems with Cisco TrustSec

Symptom	Possible Causes	Verification and Solution
The Cisco Nexus 1000V is unable to form a SXP session with Cisco TrustSec.	There is no connection between Cisco Nexus 1000V and its peer.	Verify if the Cisco Nexus 1000V is connected to its peer. ping
	The Cisco TrustSec SXP is not enabled on the Cisco Nexus 1000V.	Verify if the Cisco TrustSec SXP is enabled on the Cisco Nexus 1000V. show cts sxp If not, enable the Cisco TrustSec SXP. cts sxp enable
	The password configured on the Cisco Nexus 1000V does not match the password configured on its peer.	Verify if the passwords configured on the Cisco Nexus 1000V matches its peer. show cts sxp
	The default source IPv4 address is not configured on the Cisco Nexus 1000V.	Verify if the default source IPv4 address is not configured on the Cisco Nexus 1000V. show cts sxp
	The SXP peer is not configured as the listener.	Verify that the SXP peer is configured as the listener. show cts sxp connection
Cisco TrustSec SXP is unable to learn any IP-SGT mappings on the Cisco Nexus 1000V.	The Cisco TrustSec device tracking is not enabled on the Cisco Nexus 1000V.	Verify if the Cisco TrustSec device tracking is enabled on the Cisco Nexus 1000V. show cts device tracking If not, enable the Cisco TrustSec device tracking. cts sxp device tracking
	The DHCP Snooping feature is not enabled globally and on a VLAN on the Cisco Nexus 1000V.	Verify if the DHCP Snooping feature is enabled globally on the Cisco Nexus 1000V. show feature If not, enable the DHCP Snooping feature globally. feature dhcp Verify if the DHCP Snooping feature is enabled on a VLAN on the Cisco Nexus 1000V. show ip dhcp snooping If not, enable the DHCP Snooping feature on a VLAN. ip dhcp snooping vlan <i>vlan-list</i>



vCenter Plug-in

Use this chapter to troubleshoot the vCenter Plug-in functionality.

This chapter includes the following topics:

- [vCenter Plug-in Overview, page 26-1](#)
- [Requirements for VMware vSphere Web Client, page 26-1](#)
- [Generating a Log Bundle, page 26-2](#)

vCenter Plug-in Overview

The Cisco Nexus 1000V virtual switch is a software based Layer 2 switch for the virtualized server environments that are running VMware ESX. The Cisco Nexus 1000V provides a consistent networking experience across the physical and the virtual environments. It consists of two components: the Virtual Ethernet Module (VEM), a software switch that is embedded in the hypervisor and a Virtual Supervisor Module (VSM) that manages the networking policies and the quality of service for the virtual machines.

With earlier releases of Cisco Nexus 1000V, the system administrators had no visibility into the networking aspects of the Cisco Nexus 1000V virtual switch. Starting with Cisco NX-OS Release 4.2(1)SV2(1.1), the Cisco Nexus 1000V Plug-in for the VMware vCenter Server (vCenter Plug-in) is supported on the Cisco Nexus 1000V virtual switch. It provides the server administrators a holistic view of the virtual network and a visibility into the networking aspects of the Cisco Nexus 1000V virtual switch.

Starting with Cisco NX-OS Release 4.2(1)SV2(1.1), the vCenter Plug-in is supported on the VMware vSphere Web Clients only. The VMware vSphere Web Client enables you to connect to a VMware vCenter Server system to manage a Cisco Nexus 1000V through a browser. The vCenter Plug-in is installed as a new tab called Cisco Nexus 1000v as part of the user interface in the vSphere Web client.

With the vCenter Plug-in, the server administrators can export the necessary networking details from the vCenter server, investigate the root cause of and prevent the networking issues, and deploy the virtual machines with the suitable policies. The server administrators can monitor and manage the resources effectively with the network details provided in the vCenter Plug-in.

Requirements for VMware vSphere Web Client

Refer to the following pre-requisites before configuring the vCenter Plug-in functionality on Cisco Nexus 1000V:

- VMware vCenter Server 5.0 and/or higher.

- VMware vCenter Web Client 5.1. The vCenter Plug-in does not work with the vSphere 5.0 Web Client.
- The following browsers are supported for version 5.1 of the vSphere Web Client:
 - Microsoft Internet Explorer 7, 8, and 9.
 - Mozilla Firefox 3.6 and later.
 - Google Chrome 14 and later.
- The vSphere Web Client requires the Adobe Flash Player version 11.1.0 or later to be installed.
- Make sure that Cisco Nexus 1000V Release 4.2(1)SV2(1.1) is installed and configured to a vCenter.

Generating a Log Bundle

You can collect the diagnostic information for VMware vCenter Server by collecting vSphere log files into a single location. Complete the following steps to generate the vSphere Web Client log bundles from a vCenter Server:

1. Log in to the Windows server where the VMware vCenter Server is installed.
2. Click **Start > All Programs > VMware > Generate vSphere Web Client Log Bundle**.

You can use this step to generate the vSphere Web Client log bundles even when you are not able to connect to the vCenter Server using the vSphere Client. The log bundle is generated as a .zip file. See VMware documentation *Collect vSphere Log Files* for more information on collecting the log files.

**Note**

Currently the login to the vCenter Plug-in is available via the administrator account only.



Ethalyzer

This chapter describes how to use Ethalyzer as a Cisco NX-OS protocol analyzer tool.

This chapter includes the following section:

- [Using Ethalyzer, page 27-1](#)

Using Ethalyzer

Ethalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark (formerly Ethereal) open source code. Ethalyzer is a command-line version of Wireshark that captures and decodes packets. You can use Ethalyzer to troubleshoot your network and analyze the control-plane traffic.

To configure Ethalyzer, use one or more of the following commands:

Table 27-1 Ethalyzer Commands Used for Configuring

Command	Purpose
switch# ethalyzer local sniff-interface interface	Captures packets sent or received by the supervisor and provides detailed protocol information. Note For all commands in this table, interface is control, ha-primary, ha-secondary, inband (packet interface) or mgmt (management interface).
switch# ethalyzer local sniff-interface interface detailed-dissection	Displays detailed protocol information
switch# ethalyzer local sniff-interface interface limit-captured-frames	Limits the number of frames to capture.
switch# ethalyzer local sniff-interface interface limit-frame-size	Limits the length of the frame to capture.
switch# ethalyzer local sniff-interface interface capture-filter	Filters the types of packets to capture.
switch# ethalyzer local sniff-interface interface display-filter	Filters the types of captured packets to display.
switch# ethalyzer local sniff-interface interface dump-pkt	Dump the packet in HEX/ASCII with possibly one line summary

Table 27-1 Ethalyzer Commands Used for Configuring

Command	Purpose
switch# ethalyzer local sniff-interface interface write	Saves the captured data to a file.
switch# ethalyzer local read file	Opens a captured data file and analyzes it.

Ethalyzer does not capture data traffic that Cisco NX-OS forwards in the hardware. Ethalyzer uses the same capture filter syntax as tcpdump. For more information, see the following URL:

http://www.tcpdump.org/tcpdump_man.html

For information about the syntax of the display filter, see the following URL:

<http://wiki.wireshark.org/DisplayFilters>

This example shows captured data (limited to four packets) on the management interface:

```
switch# ethalyzer local sniff-interface mgmt limit-captured-frames 4
Capturing on eth1
2012-10-01 19:15:23.794943 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=64
2012-10-01 19:15:23.796142 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=144
2012-10-01 19:15:23.796608 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=144
2012-10-01 19:15:23.797060 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=144
4 packets captured
switch#
```

For more information about Wireshark, see the following URL: <http://www.wireshark.org/docs/>



Before Contacting Technical Support

This chapter describes the steps to take before calling for technical support and includes the following sections:

- [Cisco Support Communities](#), page 28-1
- [Gathering Information for Technical Support](#), page 28-1
- [Obtaining a File of Core Memory Information](#), page 28-2
- [Copying Files](#), page 28-3



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>

Cisco Support Communities

For additional information, visit one of the following support communities:

- [Cisco Support Community for Server Networking](#)
- [Cisco Communities: Nexus 1000V](#)

Gathering Information for Technical Support

At some point, you may need to contact your customer support representative or Cisco TAC for some additional assistance. This section outlines the steps that you should perform prior to contacting your next level of support, as this will reduce the amount of time spent resolving the issue.



Note

Do not reload the module or the switch at least until you have completed [Step 1](#) below. Some logs and counters are kept in volatile storage and will not survive a reload.

To prepare for contacting your customer support representative, follow these steps:

-
- Step 1** Collect switch information and configuration. This should be done before and after the issue has been resolved.
- Configure your Telnet or SSH application to log the screen output to a text file. Use the **terminal length 0** CLI command and then use the **show tech-support details** CLI command.
- Step 2** Capture the exact error codes you see in CLI message logs using one of the following commands.
- **show logging log** CLI (displays the error messages)
 - **show logging last *number*** (displays the last lines of the log)
- Step 3** Answer the following questions before calling for technical support:
- 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 4** 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.

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**.



A

Access Control Lists. See ACLs.

ACLs

- commands for troubleshooting [15-2](#)
- configuration limits [15-1](#)
- debugging policy verification [15-3](#)
- description [15-1](#)
- displaying ACL policies on the VEM [15-2](#)
- restrictions [15-2](#)

automatic static MAC learning [11-12](#)

C

checking VEM status [11-13](#)

checking VSM status [11-12](#)

Cisco Support Communities [1-7, 28-1](#)

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 [7-10](#)
- VSM and vCenter Server problem symptoms and solutions [21-5](#)

core dumps [28-2](#)

CPU status

- monitoring [2-2](#)

customer support

- collecting information [28-1](#)
- contacting Cisco or VMware [1-7](#)

D

DHCP

- logs, collecting and evaluating [19-5](#)

DHCP snooping

- displaying DHCP bindings [19-6](#)
- problem symptoms and solutions [19-3](#)

disabling

- VXLAN feature globally [23-4](#)

disabling automatic static MAC learning [11-12](#)

documentation

- additional publications [ii-xvi, ii-xviii](#)
- conventions [ii-xv](#)

domain parameters [3-4](#)

DVS

- find extension key [3-8](#)
- recovering [21-2](#)

dynamic ARP inspection

- DHCP snooping binding database [19-2](#)

E

error messages

- vSphere Client [21-12](#)

extension key

- finding for specific DVS [3-8](#)
- unregister in vCenter Server [3-12](#)

F

fragmentation

- VXLANs [23-3](#)

G

GUI upgrade

problem symptoms and solutions [5-6](#)

H

HA

commands to troubleshoot [6-5](#)

description [6-1](#)

network level support [6-2](#)

problem symptoms and solutions [6-2](#)

system level support [6-2](#)

High Availability. See HA

hosts

problem symptoms and solutions [21-9, 21-10](#)

I

IGMP snooping

commands for troubleshooting [18-2](#)

description [18-1](#)

problem symptoms and solutions [18-4](#)

troubleshooting guidelines [18-2](#)

installation

recreating the installation flowchart [3-3, 3-10](#)

Internet Group Management Protocol snooping. See IGMP snooping

IP ARP Inspection

problem symptoms and solutions [19-4](#)

IP Source Guard

problem symptoms and solutions [19-5](#)

ISSU upgrade

problem symptoms and solutions [5-1](#)

J

jumbo frames and MTU [21-6](#)

K

key, extension [3-8](#)

L

Layer 2 switching

inter-VEM ping [11-5](#)

intra-VEM ping [11-4](#)

overview [11-1](#)

problem symptoms and solutions [11-7](#)

traffic interruptions [11-6](#)

license

Cisco Nexus N1000V license package [4-1](#)

contents of Cisco Nexus N1000V license file [4-2](#)

troubleshooting checklist [4-2](#)

VMware Enterprise Plus [3-1](#)

licensed module [4-1](#)

Link Aggregation Control Protocol (LACP) [6-2](#)

logging levels [2-5](#)

logs [1-6](#)

M

MAC address tables

verifying [11-7](#)

MAC learning [11-12](#)

maximum number

VXLANs [23-3](#)

module

licensed [4-1](#)

not coming up on the VSM [7-1](#)

unlicensed [4-1](#)

virtual Ethernet module (VEM) [7-1](#)

virtual supervisor module (VSM) [7-1](#)

MS-NLB [11-12](#)

MTU settings, with jumbo frames [21-6](#)

multicast

description [18-1](#)

N

NetFlow

- commands for troubleshooting [14-2](#)
- configuration problems [14-3](#)
- description [14-1](#)

network adapter [3-5](#)Network Load Balancing [11-12](#)

Nexus 1000V

- system overview [21-1](#)

Nexus 1000V switch

- copying files to or from [28-3](#)
- plug-in [3-7](#)

NLB [11-12](#)

O
online support communities [1-7, 28-1](#)

overview

- VXLANs [23-1](#)

P

port channels

- asymmetric [10-3](#)
- cannot create port channel [10-4](#)
- commands to troubleshoot [10-2](#)
- description [10-1](#)
- forcing port channel characteristics onto an interface [10-4](#)
- initial checklist [10-2, 26-2](#)
- interface does not come online [10-4](#)
- troubleshooting checklist [10-2](#)
- verifying a port channel configuration [10-5](#)

port groups

- assigning to VSM VM [3-4](#)
- virtual interfaces [3-4](#)

port profiles

- commands to troubleshoot [4-4, 9-5](#)

creating corresponding port groups on vCenter Server [21-9](#)

debug logs [9-5](#)

description [9-1](#)

information about [9-1](#)

problem symptoms and solutions [9-2](#)

ports

error disabled [8-5](#)

flapping [8-5](#)

interface description [8-1](#)

overview [8-1](#)

port counters [8-2, 8-8](#)

port enabled and port security is ErrDisabled [8-7](#)

port interface cannot be enabled [8-4](#)

port security address count exceed violation [8-7](#)

port security problems [8-2](#)

port state is link failure or not connected [8-4](#)

port types [11-4](#)

troubleshooting checklist [8-3](#)

verifying [1-3](#)

viewing port state [8-8](#)

private VLANs

commands to troubleshoot [13-2](#)

description [13-1](#)

troubleshooting guidelines [13-2](#)

Q

QoS

commands to troubleshoot [16-2](#)

configuration limits [16-1](#)

debugging policy verification errors [16-3](#)

description [16-1](#)

troubleshooting QoS policies on the VEM [16-2](#)

Quality of Service. See QoS

R

RADIUS

accounting logs [2-5](#)
 related documents [ii-xvi, ii-xviii](#)

S

scalability
 VXLANs [23-3](#)
 show port-profile sync-status command [9-6](#)
 show port-profile virtual usage command [9-6](#)
 software
 core dumps [28-2](#)
 SPAN
 commands to troubleshoot [17-3](#)
 configuration guidelines [17-1](#)
 description [17-1](#)
 problem symptoms and solutions [17-2](#)
 session requirements [17-2](#)
 static MAC learning [11-12](#)
 status, VEM [11-13](#)
 status, VSM [11-12](#)
 support communities, links to [1-7, 28-1](#)
 supported features
 jumbo frames [23-3](#)
 Switched Port Analyzer. See SPAN
 synchronization problems [8-2, 8-8](#)
 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 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 [10-2, 26-2](#)
 overview [10-2](#)

U

unlicensed module [4-1](#)
 unregister an extension key [3-12](#)
 upgrade
 GUI problems symptoms and solutions [5-6](#)
 ISSU problems symptoms and solutions [5-1](#)
 VEM problems symptoms and solutions [5-5](#)

V

vCenter Server
 refreshing connection [3-3](#)
 removing the VSM [3-11](#)
 restrictions [21-2](#)
 unregister the VSM [3-12](#)
 verifying connection to VSM [7-6](#)
 verifying correct configuration [7-10](#)
 VEM
 commands for vemlog [21-12](#)
 commands to troubleshoot [21-11](#)
 domain parameters [3-4](#)
 physical ports [11-2](#)
 status [11-13](#)
 verifying correct configuration [7-14](#)
 view of ports [11-2](#)
 virtual ports [11-2](#)

- VEM upgrade
 - problem symptoms and solutions [5-5](#)
- verifying
 - MAC address tables [11-7](#)
- Virtual Ethernet Module. See VEM.
- virtual Ethernet port (veth) [11-2](#)
- virtual NIC [11-2](#)
- Virtual Supervisor Module. See VSM.
- VLAN
 - cannot create [12-3](#)
 - traffic does not traverse trunk [10-5](#)
 - troubleshooting checklist [12-2](#)
- VLANs
 - description [12-1](#)
- VM
 - improving performance [3-4](#)
 - traffic problems [21-10](#)
- vmnic
 - number allocation [11-3](#)
- VSD
 - commands to troubleshoot [5-16, 19-6, 20-3](#)
 - logs, collecting and evaluating [20-2](#)
 - problem symptoms and solutions [20-1](#)
- VSM
 - commands to troubleshoot [7-18](#)
 - creating [21-9](#)
 - domain parameters [3-4](#)
 - identifying extension key [3-12](#)
 - removing hosts from [3-11](#)
 - status [11-12](#)
 - verifying correct configuration [7-7](#)
 - view of ports [11-3](#)
- vSphere Client
 - error messages [21-12](#)
- VXLANs
 - fragmentation [23-3](#)
 - maximum number [23-3](#)
 - overview [23-1](#)
 - scalability [23-3](#)

