



Cisco Nexus 1000V for Microsoft Hyper-V Troubleshooting Guide, Release 5.2(1)SM3(2.1)

December 21, 2018

Americas Headquarters

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

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. (1721R)

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 for Microsoft Hyper-V Troubleshooting Guide, Release 5.2(1)SM3(2.1)
© 2018 Cisco Systems, Inc. All rights reserved.



CHAPTER 1**Overview 1-1**

- Overview of the Troubleshooting Process 1-1
- Overview of Best Practices 1-1
- Troubleshooting Basics 1-1
- Overview of Symptoms 1-3
- System Messages 1-4
- Troubleshooting with Logs 1-6
- Cisco Support Communities 1-7
- Contacting Cisco Customer Support 1-7

CHAPTER 2**Tools Used in Troubleshooting 2-1**

- Commands 2-1
- Ping 2-1
- Traceroute 2-1
- Monitoring Processes and CPUs 2-2
- RADIUS 2-4
- Syslog 2-5

CHAPTER 3**Installation 3-1**

- Host Is in the Not Responding State in the Microsoft SCVMM 3-1
- VMs populated on the Microsoft SCVMM are missing 3-1
- Adding Hosts to a Logical Switch in Microsoft SCVMM Fails 3-1
- Host Changes to a Non-Responding State in Microsoft SCVMM 3-4
- Installation Failure When the Microsoft SCVMM Fails to Resolve Hostnames 3-5
- Refreshing the Connection Between the Cisco Nexus 1000V and Microsoft SCVMM Server 3-6
- Updating the Cisco Nexus 1000V Configuration Data on Hyper-V Hosts 3-6
- Verifying That the Cisco Provider Installed Correctly 3-7
- Cleaning Up Switch Extension Fails 3-8
- Removing of vem.msi Throws an Error 3-8
- Installing the Install-Nexus1000V-VSMCertificate.ps1 Fails 3-8

Refreshing Switch Extension Manager Fails	3-9
Verifying Logical Switch Compliance	3-9
Verifying the Logical Switch Extension	3-10
Verifying the Logical Switch Uplink Mode	3-10
Creating or Deleting a Switch on a Host Management Adapter	3-10
Exporting VM Templates When a Hard Disk Fails	3-11
Deleting Temporary Templates	3-11
Verifying Host Compliance in the Microsoft SCVMM	3-11
Creating a Switch on a Management NIC When a Static IP Address Fails on a Server Core	3-12
Problems with Management NICs	3-12

CHAPTER 4

Upgrade	4-1
Information About Upgrades	4-1
Problems with ISSU	4-1
Problems with the VEM Upgrade	4-4

CHAPTER 5

Licenses	5-1
Information About Licenses	5-1
Prerequisites to License Troubleshooting	5-2
Problems with Licenses	5-2
License Troubleshooting Commands	5-4

CHAPTER 6

High Availability	6-1
Information About High Availability	6-1
Problems with High Availability	6-2
Failover Clusters and the Microsoft SCVMM	6-5
Selecting Storage During VM Deployment on Failover Clusters from the Microsoft SCVMM	6-6
Live Migration Fails Due to Network Bandwidth	6-6
Cluster IP Resource Fails to Come Up	6-6
High Availability Troubleshooting Commands	6-7

CHAPTER 7

VSM and VEM Modules	7-1
Information About Modules	7-1
Troubleshooting a Module That Is Not Coming Up on the VSM	7-1

VSM and VEM Troubleshooting Commands 7-14

CHAPTER 8

Ports 8-1

- Information About Interface Characteristics 8-1
- Information About Interface Counters 8-1
- Information About Link Flapping 8-2
- Information About Port Security 8-2
- Port Diagnostic Checklist 8-2
- Problems with Ports 8-3
- Port Troubleshooting Commands 8-6

CHAPTER 9

Port Profiles 9-1

- Information About Port Profiles 9-1
- Problems with Port Profiles 9-2
- Recovering a Quarantined Offline Interface 9-6
- Verifying the Maximum Number of Ports 9-7
- Port Profile Logs 9-8
- Port Profile Troubleshooting Commands 9-9

CHAPTER 10

Port Channels and Trunking 10-1

- Port Channel Overview 10-1
- Port Channel Restriction 10-1
- Trunking Overview 10-1
- Initial Troubleshooting Checklist 10-2
- Troubleshooting Asymmetric Port Channels 10-3
- Troubleshooting LACP Port Channels 10-4
- Cannot Create a Port Channel 10-4
- Newly Added Interface Does Not Come Online in a Port Channel 10-4
- VLAN Traffic Does Not Traverse Trunk 10-5

CHAPTER 11

Layer 2 Switching 11-1

- Information About Layer 2 Ethernet Switching 11-1
- Viewing Ports from the VEM 11-2
- Viewing Ports from the VSM 11-3
- Problems with Layer 2 Switching 11-4
- Layer 2 Switching Troubleshooting Commands 11-7

Troubleshooting Microsoft NLB Unicast Mode 11-11

CHAPTER 12

VLANs 12-1

- Information About VLANs 12-1
- Initial Troubleshooting Checklist 12-2
- Cannot Create a VLAN 12-2

CHAPTER 13

Private VLANs 13-1

- Information About Private VLANs 13-1
- Troubleshooting Guidelines 13-2
- Private VLAN Troubleshooting Commands 13-2

CHAPTER 14

NetFlow 14-1

- Information About NetFlow 14-1
- NetFlow Troubleshooting Commands 14-1
- Problems with NetFlow 14-4

CHAPTER 15

Access Control Lists 15-1

- Information About ACLs 15-1
- ACL Configuration Limits 15-1
- ACL Restrictions 15-1
- ACL Troubleshooting Commands 15-2
- Displaying ACL Policies on the VEM 15-2
- Debugging Policy Verification Issues 15-3

CHAPTER 16

Quality of Service 16-1

- Information About Quality of Service 16-1
- QoS Configuration Limits 16-1
- QoS VSM Troubleshooting Commands 16-2
- QoS VEM Troubleshooting Commands 16-2
- Debugging Policing Verification Errors 16-3

CHAPTER 17

SPAN 17-1

- Information About SPAN 17-1
- Problems with SPAN 17-2
- SPAN Troubleshooting Commands 17-3

CHAPTER 18**Multicast IGMP Snooping 18-1**

- Information About Multicast IGMP Snooping 18-1
- Problems with Multicast IGMP Snooping 18-2
- Troubleshooting Guidelines 18-2
- Multicast IGMP Snooping Troubleshooting Commands 18-2

CHAPTER 19**DHCP, DAI, and IPSG 19-1**

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

CHAPTER 20**System 20-1**

- Information About the System 20-1
- VEM Troubleshooting Commands 20-1
- VEM Log Commands 20-3

CHAPTER 21**Network Segmentation Manager 21-1**

- Information About Network Segmentation Manager 21-1
- Problems with Network Segmentation Manager 21-1
- Updating VM Fails 21-1
- Network Segment Not Visible on the Microsoft SCVMM 21-2
- Network Segment Is Not Available on the Microsoft SCVMM 21-2
- Network Segmentation Manager Troubleshooting Commands 21-2

CHAPTER 22**Ethalyzer 22-1**

- Information About Ethalyzer 22-1

CHAPTER 23**Before Contacting Technical Support 23-1**

- Cisco Support Communities 23-1
- Gathering Information for Technical Support 23-1

Obtaining a File of Core Memory Information	23-2
Copying Files	23-3



Overview

This chapter introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that might occur when you configure and use the Cisco Nexus 1000V.

Overview of the Troubleshooting Process

To troubleshoot your network, follow these 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 that 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 1-3](#).
- Verify and troubleshoot any new configuration changes after implementing the change.

Troubleshooting Basics

This section introduces questions to ask when troubleshooting a problem with the 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.

Troubleshooting Guidelines

By answering the questions in the following subsections, you can determine the paths that 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, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Gather information about the 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 might use to troubleshoot your specific problem.

Each chapter in this guide includes additional tools and commands that are 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.

Enter 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 enter commands with the **internal** keyword, you must log in with the 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, enter 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 Virtual Supervisor Module Virtual Machine (VSM VM) are assigned to the right port groups and if all of the network adapters are connected from the Microsoft System Center Virtual Machine Manager (SCVMM) User Interface.

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, and trunk mode?

Enter the **show vlan brief** command to check the status of a VLAN. The status should be up.

Enter the **show port-profile** command to check a port profile configuration.

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

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

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

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 serves users who might 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 Cisco 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.

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*.

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 process 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# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
n1000v(config)# logging server 192.0.2.1 6 facility local1
```

This example shows how 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 as follows:

- 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 the syslog.

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

- d. Verify the 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 the 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

The 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 might have led up to the current problem condition that you are facing.

Viewing Logs

This example shows how to access and view logs in the Cisco Nexus 1000V:

```
n1000v# show logging ?

<CR>
>          Redirect it to a file
>>        Redirect it to a file in append mode
console    Show console logging configuration
info       Show logging configuration
internal   Logging internal information
ip         IP configuration
last       Show last few lines of logfile
level      Show facility logging configuration
logfile    Show contents of logfile
module     Show module(linecard) logging configuration
monitor    Show monitor logging configuration
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 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 Cisco Nexus 1000V software that you are running
- Version of the Microsoft SCVMM server software that you are running
- Contact phone number.
- Brief description of the problem
- Brief explanation of the steps that you have already taken to isolate and resolve the problem

If you purchased the product and support contract from Cisco, contact Cisco for support. Cisco provides Layer 1, Layer 2, and Layer 3 support.

For more information on steps to take before calling Technical Support, see the [“Gathering Information for Technical Support”](#) section on page 23-1.



Tools Used in Troubleshooting

This chapter describes the troubleshooting tools available for the Cisco Nexus 1000V.

Commands

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

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

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

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

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

Ping

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

The ping utility 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 the destination.

Traceroute

Use traceroute to do the following:

- Trace the route followed by the data traffic.

- Compute inter-switch (hop-to-hop) latency.

Traceroute identifies the path taken on a hop-by-hop basis and includes a time stamp 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.

Enter the **traceroute** command to access this feature.

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

Monitoring Processes and CPUs

The CLI has features that enable you to monitor switch processes and CPU status and utilization.

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 the following:

- 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 “-” (hyphen) usually means a daemon that is not running on any particular TTY.
- Process—Name of the process.

Process states are as follows:

- 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, which causes the system to classify it as faulty and disable it.

Example 2-1 *show processes Command*

```
n1000v# show processes ?
>      Redirect it to a file
>>    Redirect it to a file in append mode
cpu    Show processes CPU Info
log     Show information about process logs
memory Show processes Memory Info
```

```

vdc      Show processes in vdc
|        Pipe command output to filter

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

Enter the **show processes cpu** command to display CPU utilization (see [Example 2-2](#)). The command output includes the following:

- Runtime(ms)—CPU time that the process has used, expressed in milliseconds.
- Invoked—Number of times that the process has been invoked.
- uSecs—Microseconds of CPU time as an average for each process invocation.
- 1Sec—CPU utilization as a 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         2754          458     6013   0.0%   init
2          0          166      4      0.0%   kthreadd
3          0           2       0      0.0%   migration/0
4         239       51386      4      0.0%   ksoftirqd/0
5          2          72     27      0.0%   watchdog/0
6         12       1798      6      0.0%   events/0
7          0          27      7      0.0%   khelper
8         39       2278     17      0.0%   kblockd/0
9          0           2       0      0.0%   kacpid
10         0           2       0      0.0%   kacpi_notify
11         1           9    200      0.0%   kseriod
12         0           2       0      0.0%   ata/0
13         0           2       0      0.0%   ata_aux
14         0           2       0      0.0%   ksuspend_usbd
15         0           2       1      0.0%   khubd
16         9        356     27      0.0%   pdflush

```

```

17          0          5          2      0.0%  pdflush
18          0          2          1      0.0%  kswapd0
19          0          2          1      0.0%  aio/0
20          0          2          1      0.0%  nfsiod
21          0         19          4      0.0%  rpciod/0
331        24         63        385      0.0%  kjournald
...

```

Displaying CPU and Memory Information

Enter the **show system resources** command to display system-related CPU and memory statistics (see [Example 2-3](#)). The output includes the following:

- The load is defined as the 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 processes are actually running when the command is entered.
- CPU states shows the CPU usage percentage in the 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 kilobytes. 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: 0.00   5 minutes: 0.14   15 minutes: 0.16
Processes   :   295 total, 4 running
CPU states  :   0.0% user,    2.0% kernel,   98.0% idle
Memory usage: 2064844K total,  1379800K used,   685044K free
Current memory status: OK

```

RADIUS

The RADIUS protocol is 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 a Cisco Nexus 1000V device. When you try to log into a device, the 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 with a list of actual devices that the user should have access to. Once the user has been authenticated, the switch can refer to the RADIUS server to determine the extent of access that 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 can be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

This example shows how to display accounting log entries:

```
n1000v# show accounting log
Sun Dec 7 04:02:27 2002:start:/dev/pts/0_1039924947:admin
Sun Dec 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 shows only the beginning and ending (start and stop) times 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.
- Selecting the types of logging information to be captured.
- Selecting the destination of the captured logging information.

The syslog software allows you to store a chronological log of system messages locally or send 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 seven 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

The 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. You 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.

You can disable logging to the console or enable logging to a given Telnet or SSH session as follows:

- To disable console logging, enter the **no logging console** command in global configuration mode.
- To enable logging for Telnet or SSH, enter the **terminal monitor** command in EXEC mode.



Note

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

The **no logging console** command that is shown in [Example 2-4](#) disables console logging and is enabled by default.

Example 2-4 no logging console Command

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

The **terminal monitor** command that is shown in [Example 2-5](#) enables logging for Telnet or SSH and is disabled by default.

Example 2-5 terminal monitor Command

```
n1000v# terminal monitor
```

For more information about configuring syslogs, see the *Cisco Nexus 1000V for Microsoft Hyper-V System Management Configuration Guide*.



Installation

This chapter describes how to identify and resolve installation problems.

Host Is in the Not Responding State in the Microsoft SCVMM

You can refresh the host that is in the Not Responding state.

-
- | | |
|---------------|--|
| Step 1 | Launch the Microsoft SCVMM UI. |
| Step 2 | Choose the server that is in the Not Responding state. |
| Step 3 | Refresh the host. |
-

VMs populated on the Microsoft SCVMM are missing

At times, the VMs that are populated in the Microsoft SCVMM are found to be missing, but are visible under the View Dependent Resources option that is displayed when you right-click the logical switch and are accessible from the Cisco Nexus 1000V HyperV manager in the host.

In this case, perform the steps in the following document to get the VMs back in Microsoft SCVMM and refresh both, the host and the SCVMM agent.

<https://social.technet.microsoft.com/forums/en-us/f9ce9fa8-e509-40ed-b6ee-9c352f44c5c5/vdi-created-vms-vm-fails-to-show-up-in-the-system-center-2012-virtual-machine-manager>

Adding Hosts to a Logical Switch in Microsoft SCVMM Fails

When you try to add hosts to a logical switch in Microsoft SCVMM, it fails with ERROR 2912.

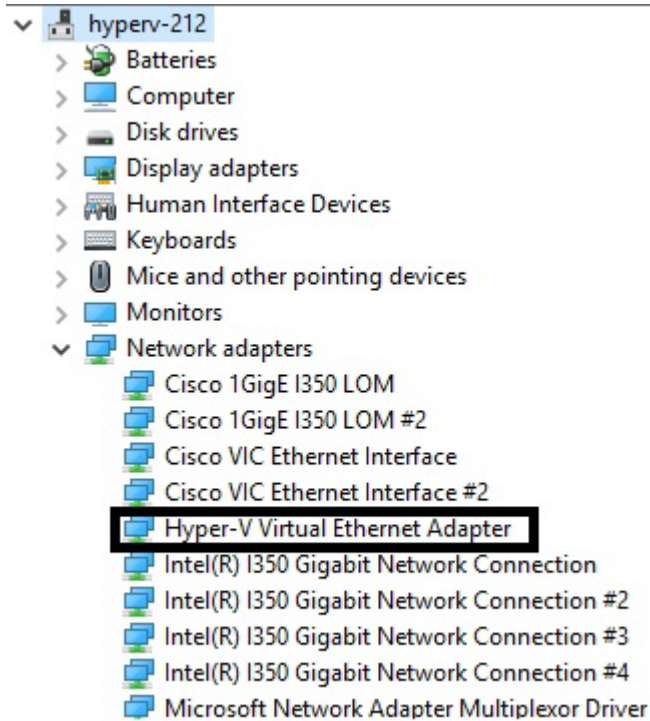
To add hosts to a logical switch in the Microsoft SCVMM, perform the following steps:

-
- | | |
|---------------|---|
| Step 1 | On the affected HyperV host, navigate to Control Panel > Hardware > Device Manager . |
|---------------|---|
-

Step 2 Under **Network Adapters**, check for the following adapters:

- Hyper-V Virtual Ethernet Adapter
- Hyper-V Switch Extension Adapter

Figure 3-1 Adapters Under Network Adapters



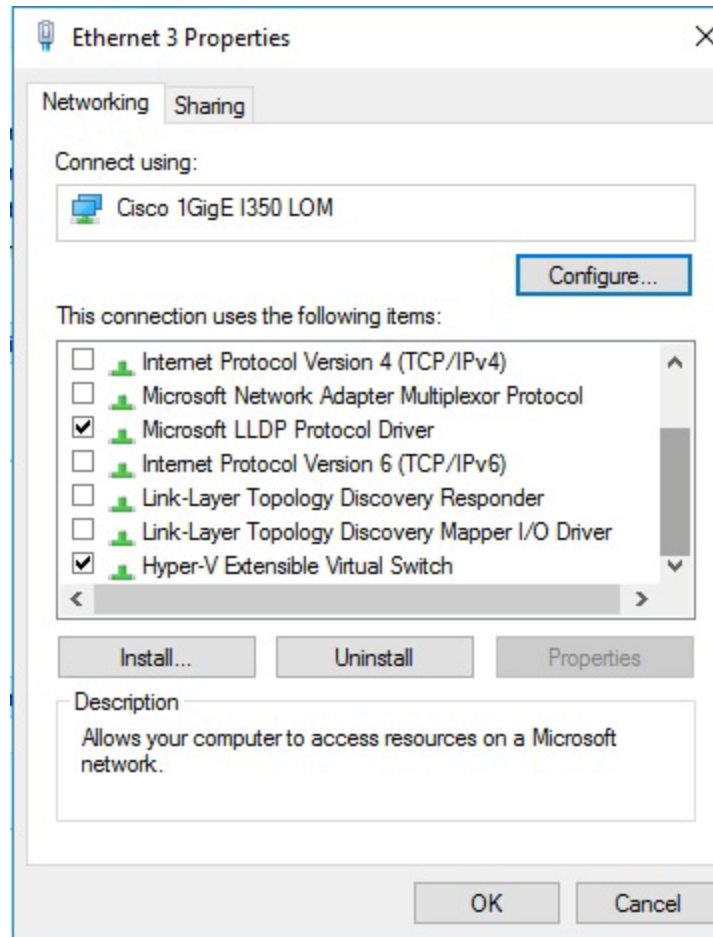
Step 3 Right-click and select **uninstall** to uninstall the above adapters if available.

Step 4 From the **Hyper-V manager**, navigate to **Server Manager > Control Panel > Adapter Settings** to verify if the adapters have been removed successfully.

Step 5 From the **Hyper-V manager**, navigate to **Server Manager > Control Panel > Adapter Settings** to create a new virtual switch with a physical external adapter.

Step 6 Right-click **Adapter Settings** to open the properties of the physical adapter over which the Virtual Switch needs to be created.

Figure 3-2 Properties Dialog Box



- Step 7** Uncheck **Hyper-V Extensible Switch** and recreate the virtual switch.
- Step 8** On the **Hyper-V host**, open **PowerShell**, and run the *Get-NetSwitchTeam* command.
- Step 9** Remove any NIC team if available by running the *Remove-NetSwitchTeam -Name <name of the team present already>* command.

Figure 3-3 PowerShell - Get and Remove NetSwitchTeam Commands

PowerShell

```

PowerShell
C) 2016 Microsoft Corporation. All rights reserved.

\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV> Get-NetSwitchTeam

Logical Switch
Ethernet 5

\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV> Remove-NetSwitchTeam -Name "Logical
\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV>
\administrator.HYPERV>

```

Step 10 Verify if the NIC team is removed successfully by running the *Get-NetSwitchTeam* command and then create the logical switch from the Microsoft SCVMM.

Host Changes to a Non-Responding State in Microsoft SCVMM

This occurs when the update patches are installed on the SCVMM server. When there is a mismatch of these update patches on the SCVMM server and host, with reference to the patch dated May 8, 2018—KB4103723 (OS Build 14393.2248) which includes the CredSSP updates for CVE-2018-0886. Hence, if the SCVMM server is updated to patch, OS Build 14393.2248 or above when the host is below this patch, then the host changes to the non-responding state.

For the Microsoft SCVMM server and host communication to function properly, both the Microsoft SCVMM and the host should be either below or above patch OS Build 14393.2248.

To resolve this issue, update the host to patch OS Build 14393.2248 or above. If you cannot update the host, perform the following steps:

-
- Step 1** Open **Run** and enter **gpedit.msc**.
- Step 2** From the Local Group Policy window that appears, navigate to **Computer Configuration > Administrative Templates > System > Credentials Delegation**.
- Step 3** Right-click **Encryption Oracle Remediation**, set it to **Enable**, set the value to **Vulnerable**, and click **OK**.
- Step 4** Open **Run** and enter **regedit**.
- Step 5** Under the **HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System\CredSSP\Parameters** registry path, set the **Allow Encryption Oracle** value to 2.

The following error is displayed when the host is refreshed:

Table 3-1 Error Message - Refreshing Host

Error	Solution
Error (2912) An internal error has occurred trying to contact the 'hyperv-117.n1kqa.com' server: : . WinRM: URL: [http://hyperv-117.n1kqa.com:5985], Verb: [INVOKE], Method: [GetVersion], Resource: [http://schemas.microsoft.com/wbem/wsman/1/wmi/root/scvmm/AgentManagement]	Verify if the WS-Management service is installed and is running on the hyperv-117.n1kqa.com server. For more information, run the winrm helpmsg hresult command. If hyperv-117.n1kqa.com' is a host/library/update server or a PXE server role then ensure that the VMM agent is installed and running. For more information, see http://support.microsoft.com/kb/2742275
The request is not supported (0x80070032)	

Installation Failure When the Microsoft SCVMM Fails to Resolve Hostnames

The Microsoft SCVMM might fail to resolve the hostnames of the managed Cisco Nexus 1000V for Microsoft Hyper-V servers. Which might result in the failure of pushing Cisco Nexus 1000V for Microsoft Hyper-V VEM MSI to the Microsoft SCVMM server hosts from the Microsoft SCVMM server.

Any host side operation might fail when DNS is not resolved and could resolve in the following:

- Refresh failure of the host from the Microsoft SCVMM
- Failure to create a Cisco Nexus 1000V logical switch on the host


-
- Step 1** Launch the Microsoft SCVMM UI.
 - Step 2** At the command prompt, enter the **ping** *hostname*, where the *hostname* is the name of the DNS host.
 - Step 3** Enter the **winrm id -r<hostname>** command.
 - Step 4** Repeat [Step 2](#) and [Step 3](#) from the host and replace the *hostname* with the name of the Microsoft SCVMM DNS server.
 - Step 5** If there is more than one DNS server associated with the host, make sure that the management NIC contains only the DNS server that points to the Active Directory (AD).
 - Step 6** Using your browser, navigate to **Tools > Internet Options > Connections** to relocate your alternate DNS server (if any).
-

Refreshing the Connection Between the Cisco Nexus 1000V and Microsoft SCVMM Server

You can refresh the connection between the Cisco Nexus 1000V and Microsoft SCVMM server.

-
- Step 1** Launch the Microsoft SCVMM UI.
 - Step 2** For the SCVMM 2016 SP1 server, choose **Fabric Management > Networking > Switch Extension Manager**.
 - Step 3** For the SCVMM 2016 R2 server, choose **Fabric Management > Networking > Network Service**.
 - Step 4** Choose **Cisco Nexus 1000V** and right click to refresh. See [Figure 3-4](#).

Figure 3-4 Refresh Cisco Nexus 1000V Connection with the Microsoft SCVMM Server



-
- Step 5** Verify that the job is complete by checking the **Jobs** section.
-

Updating the Cisco Nexus 1000V Configuration Data on Hyper-V Hosts

You can update the Cisco Nexus 1000V configuration data on the Hyper-V hosts.

-
- Step 1** Launch the Microsoft SCVMM UI.
-

Step 2 Choose **Fabric > Logical Switches** to display the screen. See [Figure 3-5](#).

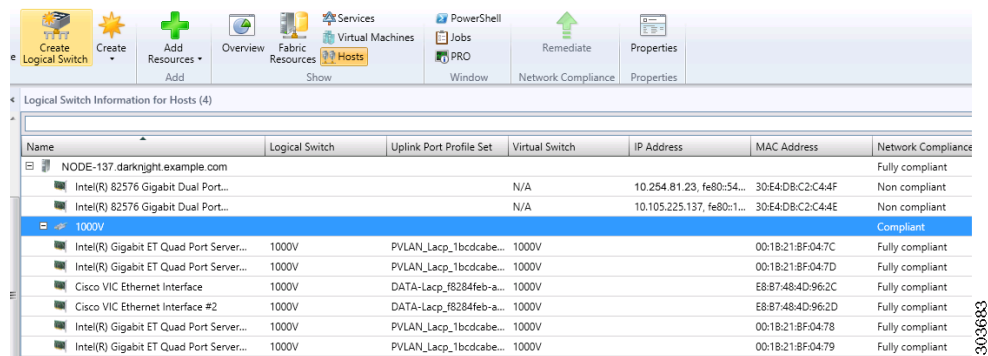
Figure 3-5 *Displaying Logical Switches*



Step 3 From the toolbar, choose **Hosts**.

Step 4 Choose the host and choose **1000V**. See [Figure 3-6](#).

Figure 3-6 *Choosing the Cisco Nexus 1000V Switch*



Step 5 From the toolbar, choose **Remediate**.

Step 6 Verify that the job was completed by checking the **Jobs** section.

Verifying That the Cisco Provider Installed Correctly

You can verify that the Cisco pCleaning up the switch extension might fail when you are deploying a VM that uses a static IP address from the static IP address pools that are published by the Cisco Nexus 1000V VSM. rovider has been installed correctly.

Step 1 Launch the Microsoft SCVMM UI.

Step 2 Navigate to **Settings**.

Step 3 Choose **Configuration Providers**.

Cleaning Up Switch Extension Fails

Cleaning up the switch extension might fail when you are deploying a VM that uses a static IP address from the static IP address pools that are published by the Cisco Nexus 1000V VSM.

**Note**

This problem is a known Microsoft issue.

Because the error is due to unrevoked IP addresses, the error shown by the Microsoft SCVMM is not specific.

Step 1 Launch the Microsoft SCVMM UI.

Step 2 Using a PowerShell window, enter the following commands, in sequence, to revoke the static IP addresses:

```
$vsem = Get-SCVirtualSwitchExtensionManager -VirtualSwitchExtensionManagerConnectionString http://<VSM-IP-address>
$poools = Get-SCStaticIPAddressPool | where { $_.VirtualSwitchExtensionManager.ID -eq $vsem.ID }
$poools | ForEach-Object { get-scipaddress -UnAssigned -StaticIPAddressPool $_ } | Revoke-SCIPAddress
```

Step 3 The configuration provider details appear on the Microsoft SCVMM.

Removing of vem.msi Throws an Error

Removing of vem.msi in WSUS throws an error if the WSUS is a part of SCVMM.

To resolve this issue, perform the following steps:

Step 1 Remove the WSUS server from SCVMM.

Step 2 Remove vem.msi from the WSUS using the powershell script.

Step 3 Readd the WSUS server to SCVMM.

Step 4 In the WSUS server, provide the cisco-get update command to successfully remove the image.

Installing the Install-Nexus1000V-VSMCertificate.ps1 Fails

If the proxy settings are enabled on the SCVMM virtual machine, you must disable these proxy settings from the virtual machine.

To disable the proxy settings, do the following:

Step 1 Open **Control panel**.

Step 2 Navigate to **Network and Internet > Internet Options**.

Step 3 In the **Internet Properties** dialog box, navigate to **Connections > LAN Settings**.

Step 4 In **Local Area Network(LAN) Settings**, Uncheck the **Proxy Server** check box.

Step 5 Click **OK**.

Now try to run the Install-Nexus1000V-VSMCertificate.ps1. It will work.

Refreshing Switch Extension Manager Fails

The following are symptoms, possible causes, and solutions for problems when refreshing the Switch Extension Manager or the Network Service.

Table 3-2 *Problems - Refresh Switch Extension*

Symptoms		
You are unable to refresh the Switch Extension Manager from the Microsoft SCVMM.	There is a problem with the connection between the Microsoft SCVMM and the VSM.	<ol style="list-style-type: none"> 1. Verify that you can navigate to the VSM <i>http://vsm_ip_address</i> from the server where the Microsoft SCVMM service is running. 2. Verify that your proxy settings and firewall settings are not impacting on the Microsoft SCVMM to VSM connectivity.
	There is an error in the VSM configuration.	On the VSM, verify the configuration by entering the show vs domain command.

Verifying Logical Switch Compliance

The Microsoft SCVMM might report a non compliant warning when you are deploying or changing port profiles on the Cisco Nexus 1000V logical switch. This problem is a result of a mismatch of the opaque data stored on Microsoft SCVMM and that of the individual hosts.



Note This issue is only a warning; it is not an error.

Step 1 Launch the Microsoft SCVMM UI.

Step 2 Navigate to **Fabric > Logical Switches > Hosts**.

Step 3 Using a Microsoft SCVMM PowerShell window, enter the following:

```
Get-SCVirtualNetwork | where-object {$_.LogicalSwitch -like "1000V"} | select VMHost,
HighlyAvailable, LogicalNetworks, VMHostNetworkAdapters | LogicalSwitchComplianceStatus
```

To remove the Logical Switch Compliance Warning, perform the following steps:

-
- Step 1** Refresh the Virtual Switch Extension Manager
 - Step 2** Choose **Fabric > Logical Switches > Hosts**.
 - Step 3** Select the appropriate logical switch and choose **Remediate the Host**.
-

Verifying the Logical Switch Extension

The Cisco Nexus 1000V logical switch extension is always a forwarding extension. You can verify the logical switch extension.

-
- Step 1** Launch the Microsoft SCVMM UI.
 - Step 2** Choose **Fabric > Logical Switches > *switch_name* > Properties > Extensions**.
 - Step 3** Verify that the extension type is **Forwarding**.
-

Verifying the Logical Switch Uplink Mode

The Cisco Nexus 1000V logical switch uplink mode should be **team**. You can verify the logical switch uplink mode.

-
- Step 1** Launch the Microsoft SCVMM UI.
 - Step 2** Choose **Fabric > Logical Switches > *switch_name* > Properties > Uplink**.
 - Step 3** Verify that the Uplink mode is **Team**.
-

Creating or Deleting a Switch on a Host Management Adapter

While you are deploying a Cisco Nexus 1000V switch or cleaning up a Cisco Nexus 1000V on a host management adapter, the operation might fail if there are network flaps or a DNS resolution. This problem might cause host connectivity loss because the failure occurs on the host management adapter.

-
- Step 1** Log in to the host using the remote console.
 - Step 2** Open an elevated PowerShell window and enter the **Remove-VMSwitch -name *switchname*** command.
 - Step 3** Remove the NetSwitch Team from the host and restore connectivity by entering the **Get-NetswitchTeam | Remove-NetSwitchTeam** command.
 - Step 4** Refresh the host from the Microsoft SCVMM.

**Note**

If [Step 2](#) fails when the WMI on the host is stuck in an inconsistent state, manually delete the switch from the registry, and perform a system reboot and proceed to [Step 3](#).

Exporting VM Templates When a Hard Disk Fails

When you are exporting a VM template and the hard disk selected fails, the problem is probably caused by the internet proxy settings.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** Verify that the internet Connection Settings field is blank.
-

Deleting Temporary Templates

You can delete temporary templates that are created by the Microsoft SCVMM.

Symptom	Possible Causes	Solution
Unable to delete Cisco Nexus 1000V objects in Microsoft SCVMM.	The Microsoft SCVMM creates temporary templates that are linked to the Cisco Nexus 1000V objects.	Delete the temporary templates by entering the following commands in a PowerShell window: <ul style="list-style-type: none"> Get-VMMServer Get-SCVMMTemplate where {\$_.Name -like "Tempoorary*"} Remove-SCVMMTemplate

Verifying Host Compliance in the Microsoft SCVMM

You can verify host compliance in the Microsoft SCVMM; all hosts should show as fully compliant.

-
- Step 1** Choose **Fabric > Logical Switches > Hosts**.
- Step 2** Choose the host from list.
- Step 3** From the toolbar, choose **Remediate**.
- Step 4** Verify that the job was completed by checking the **Jobs** section
-

Creating a Switch on a Management NIC When a Static IP Address Fails on a Server Core

Creating a switch fails when using a Cisco Nexus 1000V on a management NIC with a static IP address on a server core.



Note

This problem is a Microsoft issue with Server Core versions of Windows Server 2016.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** Log in to the host using the remote console.
- Step 3** Using a Microsoft SCVMM PowerShell window, delete the switch from the host by entering the **Remove-VMSwitch -name *switchname*** command.
- Step 4** Remove the NetSwitch Team from the host and restore connectivity by entering the **Get-NetSwitchTeam | Remove-NetSwitchTeam** command.
- Step 5** Refresh the host from the Microsoft SCVMM.
-

Problems with Management NICs

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

Symptom	Possible Causes	Solution
You are unable to push opaque data (OD) on VEMs.	The VSM IP address has changed.	<ol style="list-style-type: none"> 1. Change the IP address of the management interface (mgmt0) on the VSM. 2. Change the connection string of the Switch Manager Extension on the Microsoft SCVMM to the new VSM IP address. 3. Refresh the Switch Extension Manager/Network Service in the Microsoft SCVMM. 4. Verify the information on all screens before you choose OK. 5. Choose Fabric > Logical Switches > Hosts. 6. Choose the host from the list. 7. From the toolbar, choose Remediate. 8. Verify that the job was completed by checking the Jobs section.



Upgrade

This chapter describes how to identify and resolve problems related to upgrading the VSM and VEM software.

Information About Upgrades

The upgrade for the Cisco Nexus 1000V involves upgrading the 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 more information, see the *Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide*.

Problems with ISSU

The following are symptoms, possible causes, and solutions for problems while upgrading the software using the manual ISSU method. For information on the installation process, see the *Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide*.

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.	Reboot the system.
Error message: Pre-Upgrade check failed. Return code 0x4093000A (SRG collection failed).	A module is removed during the upgrade.	<ol style="list-style-type: none">1. Ensure that the module removal is complete.2. Restart the software upgrade.

Symptom	Possible Causes	Solution
<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. 2. The output of the show system redundancy status command must indicate the following: <ul style="list-style-type: none"> – Active VSM—Active with HA standby – Standby VSM—HA standby 3. When the VSMs are synchronized, restart the software upgrade.
<p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x807B0002 (No such file or directory).</p>	The software image files required for the upgrade are not available or were not copied to the bootflash: repository.	<ol style="list-style-type: none"> 1. Verify there is enough space in bootflash: repository for the image files. 2. If additional space is needed, delete other files from the bootflash: repository to make space for the software image files.
<p>Error message:</p> <p>Pre-Upgrade check failed. Return code 0x4093000F (Failed to copy image).</p>	There may not be enough space in bootflash: repository for the files to be copied.	<div>  <div> <p>Caution Do not delete the kickstart or system image files from the bootflash: repository. You cannot reboot the system if image files are unavailable in the bootflash: repository.</p> </div> </div> <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. 5. When the correct software images are in the bootflash: repository, restart the software upgrade
<p>The install command fails with the following error:</p> <ul style="list-style-type: none"> • Return code 0x4045001F (image MD5 checksum error). • Pre-Upgrade check failed. Return code 0x40930011 (Image verification failed). 	<ul style="list-style-type: none"> • 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. • A file can be truncated when copied. 	<ol style="list-style-type: none"> 1. Verify the MD5 checksum for each of the image files using the README file from the upgrade zip folder at www.cisco.com. 2. Replace the file(s) that do not match. 3. Verify that the correct images are in the bootflash: repository and that the checksums match. 4. When the correct software images are in the bootflash: repository, restart the software upgrade.
<p>Error message:</p> <p>Install has failed. Return code 0x40970001 (Incompatible image)</p>	While entering the install all command, you could have used an incorrect filename.	<p>Restart the software upgrade using the correct filenames for the new software images.</p> <p>Example:</p> <p>install all kickstart filename1 system filename2</p>

Symptom	Possible Causes	Solution
After upgrading, the VSMs do not run the new software version.	The boot variables were not set properly.	<ol style="list-style-type: none"> 1. Verify that the running images and boot variables match the upgrade version. 2. Download the required images from www.cisco.com to your local bootflash: repository, if required. 3. Verify that the correct images are in the bootflash: repository. 4. Restart the software upgrade. 5. If the problem persists, collect details of the upgrade and open a support case.
Performing the configuration copy process fails and stops the upgrade. Performing configuration copy. [####-----] 30%	Service or system errors.	<ol style="list-style-type: none"> 1. Manually copy the configuration. 2. 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. – If the copy succeeds without delays, restart the software upgrade.
Error message: Another install procedure may be in progress. (0x401E0007)	Another upgrade session is in progress from a VSM console or SSH/Telnet.	<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
Install command fails with following error message: -- FAIL. Return code 0x4093001E (Standby failed to come online) Install has failed. Return code 0x4093001E (Standby failed to come online).	The standby VSM fails to boot with the new image.	<p>Do one of the following:</p> <ul style="list-style-type: none"> • Restart the software upgrade. • Postpone the upgrade and reset the boot variables to the original filenames.
Install command fails with following error message: Install has failed. Return code 0x4093001F (Standby installer failed to take over the installation). Please identify the cause of the failure, and try the install all command.	The standby VSM takes more than 10 minutes to come up and form a stable HA pair with the active VSM.	<ol style="list-style-type: none"> 1. Reset the boot variables to the original filenames. 2. If the standby is still running the new software version, reload it. The standby synchronizes with the active, so that both are running the original software version.

Symptom	Possible Causes	Solution
Install command fails with following error message: Module 2: Waiting for module online. -- SUCCESS -- Install has failed. Return code 0x40930000 (Current operation failed to complete within specified time).	A failure at the standby VSM caused it to reload again after the <i>Continuing with installation, please wait</i> message and before the switchover.	Verify whether the standby VSM has come up using the show module command. After it comes up, retry the software upgrade using the install all kickstart filename1 system filename2 command.
Install command fails with following error message: Pre-upgrade check failed. Return code 0x42380014 (license validation failed).	An upgrade of the VSM, which is in advanced edition was performed without installing the evaluation or the permanent license.	<ol style="list-style-type: none"> 1. Before the upgrade process, either install a valid license (evaluation/permanent) or switch over to the Essential Edition. 2. Restart the software upgrade using the install all kickstart filename1 system filename2 command.

Problems with the VEM Upgrade

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

Symptom	Possible Causes	Solution
After upgrading the VEM, the host is not online in the VSM.	The VSM upgrade was not performed before the VEM upgrade.	Perform the VSM upgrade and check if the host comes online in the VSM.
During the VEM upgrade, the initial compliance scan fails.	The Update server is not responding.	Remove the Update server from the VMM management and re-add it to the VMM management.
The Supported VEM version range does not update after the VSM upgrade in the VMM.	The Switch Extension Manager has not been refreshed after the VSM upgrade.	Refresh the Switch Extension Manager after the VSM upgrade.
The vemcmd command on the host fails with an error initializing <code>sf_dpa_api_init()</code> .	Powershell/Command Prompt was run without the Admin privileges.	Run the Powershell/Command prompt with the Admin privileges.
The provided VEM upgrade script does not respond on the PowerShell/PowerShell ISE.	PowerShell	Press Enter when the script is not responding on the PowerShell/PowerShell ISE.



Licenses

This chapter describes how to identify and resolve problems related to licenses.

Information About Licenses

The name for the Cisco Nexus 1000V license version 3.0 package is NEXUS1000V_LAN_SERVICES_PKG. By default, 1024 licenses are installed with the Virtual Supervisor Modules (VSM). These default licenses expire after 60 days. You can purchase permanent licenses that do not expire.

Licensing is based on the number of CPU sockets on the Microsoft Hyper-V servers attached as Virtual Ethernet Modules (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 that 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 for Microsoft Hyper-V License Configuration Guide*.

Contents of the License File

The contents of the Cisco Nexus 1000V license file contents indicate 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 3.0 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 the ID that is shown on the VSM. To verify that the IDs match, use the **show license host-id** command. See [Example 5-6 on page 5-6](#).

**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, contact your Cisco Customer Support Account Team.

Prerequisites to License Troubleshooting

Before you begin troubleshooting licenses, verify the following information:

- Make sure that the name of the license file is less than 32 characters.
Check the name by entering the **show license** command. See [Example 5-3 on page 5-5](#).
- Make sure that 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.
Check the name by entering the **show license brief** command. See [Example 5-3 on page 5-5](#).
- Do not edit the contents of the license file. If you have already done so, contact your Cisco Customer Support Account Team.
- Make sure that the host ID in the license file is the same as the host ID on the switch by entering the following commands:
 - **show license host-id**. See [Example 5-6 on page 5-6](#).
 - **show license file**. See [Example 5-7 on page 5-6](#).

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> <p>VEM Unlicensed</p>	<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 by entering the show license usage <i>license_name</i> command. See Example 5-4 on page 5-5. 2. Determine the number of licenses required. View the sockets installed on the VEM and then enter the show module vem license-info command. See Example 5-2 on page 5-5. 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 2014 Nov 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 evaluation license in use is about to expire.</p> <p>Note Permanent and default licenses do not expire.</p>	<ol style="list-style-type: none"> 1. Verify license usage by entering the show license usage <i>license_name</i> command. See Example 5-4 on page 5-5. 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 by entering the show license usage <i>license_name</i> command. See Example 5-4 on page 5-5. 2. Contact your Cisco Customer Support Account Team to acquire additional licenses.
<p>You see the following system message:</p> <p>A license with a later/different expiry date already exists for this feature.</p>	<p>Multiple entries in the license file for the same feature with different expiry dates. Use show file license <i>file_name</i> command to verify.</p>	<ol style="list-style-type: none"> 1. Contact your Cisco Customer Support Account Team to get license file re-issued.

License Troubleshooting Commands

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

Command	Purpose
show module	Displays display module information including the license status (unlicensed or active). See Example 5-1 on page 5-4 .
show module vem license info	Displays the VEM license information including the license status, license version, and socket count. See Example 5-2 on page 5-5 .
show license usage [<i>license_name</i>]	Displays information about licenses and where they are used. You use this command to display information for a specific license, which indicates VEM and socket information. See Example 5-3 on page 5-5 . See Example 5-4 on page 5-5 .
show interface veth	Displays the messages logged about port profile events within the Cisco Nexus 1000V. See Example 5-5 on page 5-6 .
show license host-id	Displays the serial number for your Cisco Nexus 1000V license. See Example 5-6 on page 5-6 .
show license file	Displays the contents of a named license file. See Example 5-7 on page 5-6 .
show license brief	Displays a list of license files installed on the VSM. See Example 5-8 on page 5-6 .
show switch edition	Displays the current edition of the Cisco Nexus 1000V switch and a list of advanced features. See Example 5-9 on page 5-6 .

For detailed information about **show** command output, see the *Cisco Nexus 1000V for Microsoft Hyper-V Command Reference Guide*.

EXAMPLES

Example 5-1 show module Command

```
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    288     Virtual Ethernet Module    NA                   ok
4    288     Virtual Ethernet Module    NA                   ok

Mod  Sw                Hw
---  -

```

```

1    5.2(1)SM3(1.1)      0.0
2    5.2(1)SM3(1.1)      0.0
3    5.2(1)SM3(1.1)      Windows Server 2016 R2 - Datacenter (6.3.9600, 6.40)
4    5.2(1)SM3(1.1)      Windows Server 2016 R2 - Datacenter (6.3.9600, 6.40)

```

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
4	02-00-0c-00-04-00 to 02-00-0c-00-04-80	NA

Mod	Server-IP	Server-UUID	Server-Name
1	10.105.225.89	NA	NA
2	10.105.225.89	NA	NA
3	10.105.225.72	D1B16F89-9982-DF11-A3D0-D0D0FD09586A	HOST-200
4	10.105.225.74	B2D389D7-C082-DF11-BECC-D0D0FD0959CC	HOST-203

* this terminal session

Example 5-2 *show module vem license-info Command*

```

N1000v# show module vem license-info
Licenses are Sticky

```

Mod	Socket	Count	License Usage	Count	License Version	License Status
3	1		1		3.0	licensed
4	1		1		3.0	licensed
5	1		1		3.0	licensed
6	2		2		3.0	licensed

Example 5-3 *show license usage Command*

```

n1000v#
VSM-A-1# show license usage
Feature                               Ver  Ins  Lic   Status    Expiry Date  Comments  Count
-----
NEXUS1000V_LAN_SERVICES_PKG  3.0   No 1024   In use     09 Feb 2015   -

```

Example 5-4 *show license usage license_name Command*

```

N1000v# show license usage NEXUS1000V_LAN_SERVICES_PKG

```

```

-----
Feature Usage Info
-----
Version : 3.0
Installed Licenses : 0
Default Eval Licenses : 1024
Max Overdraft Licenses : 0
Installed Licenses in Use : 0
Overdraft Licenses in Use : 0
Default Eval Lic in Use : 5
Default Eval days left : 54
Licenses Available : 1019
Shortest Expiry : 24 Jan 2015
-----
Application
-----
VEM 3 - Socket 1
VEM 4 - Socket 1

```

```
VEM 5 - Socket 1
VEM 6 - Socket 1
VEM 6 - Socket 2
-----
```

Example 5-5 *show interface vethernet Command*

```
n1000v# show interface veth1
Vethernet1 is up
  Port description is WNV-01
  Hardware: Virtual, address: 001d.d8b7.1e5f (bia 001d.d8b7.1e5f)
  Owner is VM "WNV-01"
  Active on module 4
  DVS port 8002325c-413c-4288-b9bc-de2b1f5b848c--e9a6faa1-a7ea-4cf9-8727-582e02
88e691
  Port-Profile is dynpp_da0abdb9-b83a-44d3-b1f5-d90ce66f4d00_55e79594-417c-42b2
-a574-8f28779dfbb2
  Port mode is access
  5 minute input rate 0 bits/second, 0 packets/second
  5 minute output rate 368 bits/second, 0 packets/second
Rx
  3514 Input Packets 81 Unicast Packets
  2208 Multicast Packets 1225 Broadcast Packets
  297447 Bytes
Tx
  77067 Output Packets 94 Unicast Packets
  46827 Multicast Packets 30151 Broadcast Packets 76977 Flood Packets
  6552793 Bytes
  0 Input Packet Drops 0 Output Packet Drops
```

Example 5-6 *show license host-id Command*

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

Example 5-7 *show license file Command*

```
n1000v# show license file sample.lic
sample.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT NEXUS1000V_LAN_SERVICES_PKG cisco 1.0 permanent 512\
  HOSTID=VDH=8449368321243879080 \
  NOTICE="<LicFileID>sample.lic</LicFileID><LicLineID>0</LicLineID> \
  <PAK>dummyPak</PAK>" SIGN=34FCB2B24AE8

n1000v#
```

Example 5-8 *show license brief Command*

```
n1000v# show license brief
license_file.lic
n1000v#
```

Example 5-9 *show switch edition Command*

```
n1000v# show switch edition
Switch Edition: ADVANCED (3.0)
```

Advanced Features

Feature Name	Feature State
--------------	---------------

-------	--

dhcp-snooping	enabled
---------------	---------

Licenses Available: 1022

Licenses In Use: 2

License Expiry Date: 09 Feb 2015



High Availability

This chapter describes how to identify and resolve problems related to high availability.

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 that is disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- Supervisor stateful switchover— Active/standby dual supervisor configuration. The state and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide a seamless and stateful switchover in the event of a VSM failure.

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

- Virtual Ethernet Modules (VEMs) that run within virtualization servers. The VEMs are represented as modules within the VSM.
- A remote management component, for example, the Microsoft System Center Virtual Machine Manager (SCVMM).
- One or two VSMs that run within Virtual Machines (VMs).

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 by entering 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, enter 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 the VSM at the upstream and virtual switches. 	If network problems exist, do the following: <ol style="list-style-type: none"> From the Microsoft SCVMM UI, shut down the VSM, which should be in standby mode. From the Microsoft SCVMM UI client, bring up the standby VSM after network connectivity is restored.
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs. <ul style="list-style-type: none"> Check that the primary and secondary VSM are using the same image version by entering the show version of the command. 	If the active and 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 by entering the show system internal log sysmgr gsyncctrl command and look for fatal errors. 	

Symptom	Possible Causes	Solution
The standby VSM reboots periodically.	<p>The VSM has connectivity only through the management interface.</p> <ul style="list-style-type: none"> When a VSM is able to communicate through the management interface, but not through the control interface, the active VSM detects the situation and resets the standby VSM to prevent the two VSMs from being in HA mode and out of sync. 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 the control VLAN connectivity between the primary and secondary VSMs.
	<p>VSMs have different versions.</p> <p>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:</p> <pre>2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.</pre>	<p>Isolate the standby VSM and boot it.</p> <p>Enter the show version command to check the software version in both VSMs.</p> <p>Install the image matching the active VSM on the standby.</p>

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, do the following: <ol style="list-style-type: none"> From the Microsoft SCVMM UI client, shut down the VSM, which should be in standby mode. From the Microsoft SCVMM UI client, bring up the standby VSM after network connectivity is restored.
	Different domain IDs in the two VSMs. Check the <i>domain</i> value by entering the 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 this 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.

System-Level High Availability

The Cisco Nexus 1000V supports redundant VSM VMs—a primary and a secondary—that run as an HA pair. Dual VSMs operate in an active/standby capacity in which only one of the VSMs is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two VSMs to provide a stateful switchover if the active VSM fails.

Single or Dual Supervisors

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

- VEMs that run within virtualization servers (these VEMs are represented as modules within the VSM)
- A remote management component, such as the Microsoft SCVMM.
- One or two VSMs that run within VMs.

Single VSM Operation	Dual VSM Operation
<ul style="list-style-type: none"> Stateless—Service restarts from the startup configuration Stateful—Service resumes from previous state. 	<ul style="list-style-type: none"> One active VSM and one standby VSM. The active VSM runs all the system applications and controls the system. On the standby VSM, the applications are started and initialized in standby mode. They are also synchronized and kept up to date with the active VSM in order to maintain the runtime context of “ready to run.” On a switchover, the standby VSM takes over for the active VSM.

Network-Level High Availability

The Cisco Nexus 1000V HA at the network level includes port channels and the Link Aggregation Control Protocol (LACP). A port channel bundles physical links into a channel group to create a single logical link that provides the aggregate bandwidth of up to eight physical links. If a member port within a port channel fails, the traffic that was previously carried over the failed link switches to the remaining member ports within the port channel.

Additionally, the LACP allows you to 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 the LACP, see the *Cisco Nexus 1000V for Microsoft Hyper-V Layer 2 Switching Configuration Guide*.

Failover Clusters and the Microsoft SCVMM

Failover clustering is a hostside feature that provides high availability and scalability to multiple server workloads. In order for a Cisco Nexus 1000V switch to be considered a high availability device, the switch must meet the following criteria:

- The VM must be set to **High Availability > True** to be considered part of a failover cluster. That is, the VM can be moved automatically by the cluster in the event of a host failure.
- The high availability VM should be stored in one of the following types of Internet Protocol (IP) based storage facilities to accommodate live migration for a failover cluster:
 - Shared SMB storage
 - Clustered shared volumes (iSCSI, and so on)

When clusters are managed by the Microsoft SCVMM, certain criteria must be met for the Microsoft SCVMM to manage the VM as part of a failover cluster. That is, the logical switch that is part of the hosts of the failover clusters should be configured for high availability.

High Availability Logical Switch Criteria and Behavior

- A logical switch is considered to be highly available when it carries the same uplink networks on all the nodes of the cluster.
- If certain adapters carry the same uplink in each logical switch across all nodes and other uplinks do not then the adapters that carry the same uplink networks become high availability.

- A VM that is not configured for high availability can be connected to any switch in the failover cluster (logical or standard switch).
- A high availability VM can only be connected to uplinks that are high availability and are part of a logical switch.

Selecting Storage During VM Deployment on Failover Clusters from the Microsoft SCVMM

The failover cluster managed by the Microsoft SCVMM has more than one associated storage device. By default, the Microsoft SCVMM chooses the storage based on the deployment algorithm of the Microsoft SCVMM, which might not be what you want.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** In the **Migrate VM Wizard** screen, change the storage of the VM and the VM hard disk to the appropriate storage.
- Step 3** Pin the selection to the Microsoft SCVMM UI.
-

Live Migration Fails Due to Network Bandwidth

When a workload VM is carrying high traffic, VM live migration might not be allowed by the Microsoft SCVMM. The Microsoft SCVMM performs checks during live migration and decides the feasibility of moving the VM based on many factors, one of which is VM port traffic. From the perspective of the Microsoft SCVMM, when a VM is transmitting or receiving large amounts of traffic, it is not feasible to move the VM because it might result in a loss of bandwidth.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** In a Microsoft SCVMM PowerShell window, enter **Move-SCVirtualMachine**.
-

Cluster IP Resource Fails to Come Up

Cluster validation is an important tool used by large deployments to validate cluster configurations. When a virtual switch is deployed on the management NIC of the host with a static IP address, and the failover cluster already exists, the cluster IP resource might fail to come up. When this problem occurs, although the cluster IP address and DNS are reachable by conventional means (ping), the cluster validation tool fails.



Note

This problem is a known issue with the Microsoft SCVMM and is seen only with static IP addresses, not when the host management IP address is distributed over DHCP.

There is no known workaround for this issue. We recommended that you create clusters after you deploy the Cisco Nexus1000V on the management IP address.

High Availability Troubleshooting Commands

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

To list process logs and cores, enter these commands:

- **show cores**

```
switch# show cores
Module Instance Process-name PID Date (Year-Month-Day Time)
-----
1 1 private-vlan 3207 Apr 28 13:29
```

- **show processes log [pid pid]**

```
switch# show processes log
Process PID Normal-exit Stack Core Log-create-time
-----
private-vlan 3207 N Y N Tue Apr 28 13:29:48 2009

switch# 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: switchh-dk9.5.2.1.SM15.0.1.bin
System image version: 5.2(1)SM1(5.1)

PID: 3207
Exit code: signal 6 (core dumped) <-- Indicates that a cores for the process was
generated.

CWD: /var/sysmgr/work
...
```

To check the redundancy status, enter this command:

- **show system redundancy status**

```
switch# 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, enter this command:

- **show system internal redundancy info**

```

switch# 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_drops_wrong_ver: 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_drops_wrong_ver: 0
rx_unknown_pkts: 0

```

To check the system internal sysmgr state, enter this command:

- **show system internal sysmgr state**

```
switch# 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.
```

```
Process restart capability is currently disabled.
```

```
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.
Swover threshold settings: 20 switchovers within 1200 seconds
Switchovers within threshold interval: 0
Last switchover time: 0 seconds after system start time
Cumulative time between last 0 switchovers: 0
Start done received for 2 plugins, Total number of plugins = 2

```

Statistics:

```

Message count:          0
Total latency:          0           Max latency:          0
Total exec:             0           Max exec:             0

```

To reload a module, enter this command:

- **reload module**

```
switch# reload module 2
```

This command reloads the secondary VSM.



Note Entering the **reload** command without specifying a module reloads the whole system.

To attach to the standby VSM console, enter this 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.

```
switch# 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.

Information About Modules

The Cisco Nexus 1000V manages a logical switch that is defined by a Virtual Center. Each server with the Cisco Nexus 1000V logical switch can be managed as if it were a module in a physical Cisco switch.

The Cisco Nexus 1000V implementation has two parts:

- **Virtual Supervisor Module (VSM)**—This is the control software of the Cisco Nexus 1000V distributed virtual switch. It runs on a Virtual Machine (VM) and is based on Cisco NX-OS software.
- **Virtual Ethernet Module (VEM)**—This is the part of Cisco Nexus 1000V that actually switches data traffic. It runs on a Microsoft Hyper-V server. Several VEMs are controlled by one VSM.

Troubleshooting a Module That Is Not Coming Up on the VSM

This section describes the process that you can use when a module does not come up on the VSM.

Troubleshooting Guidelines

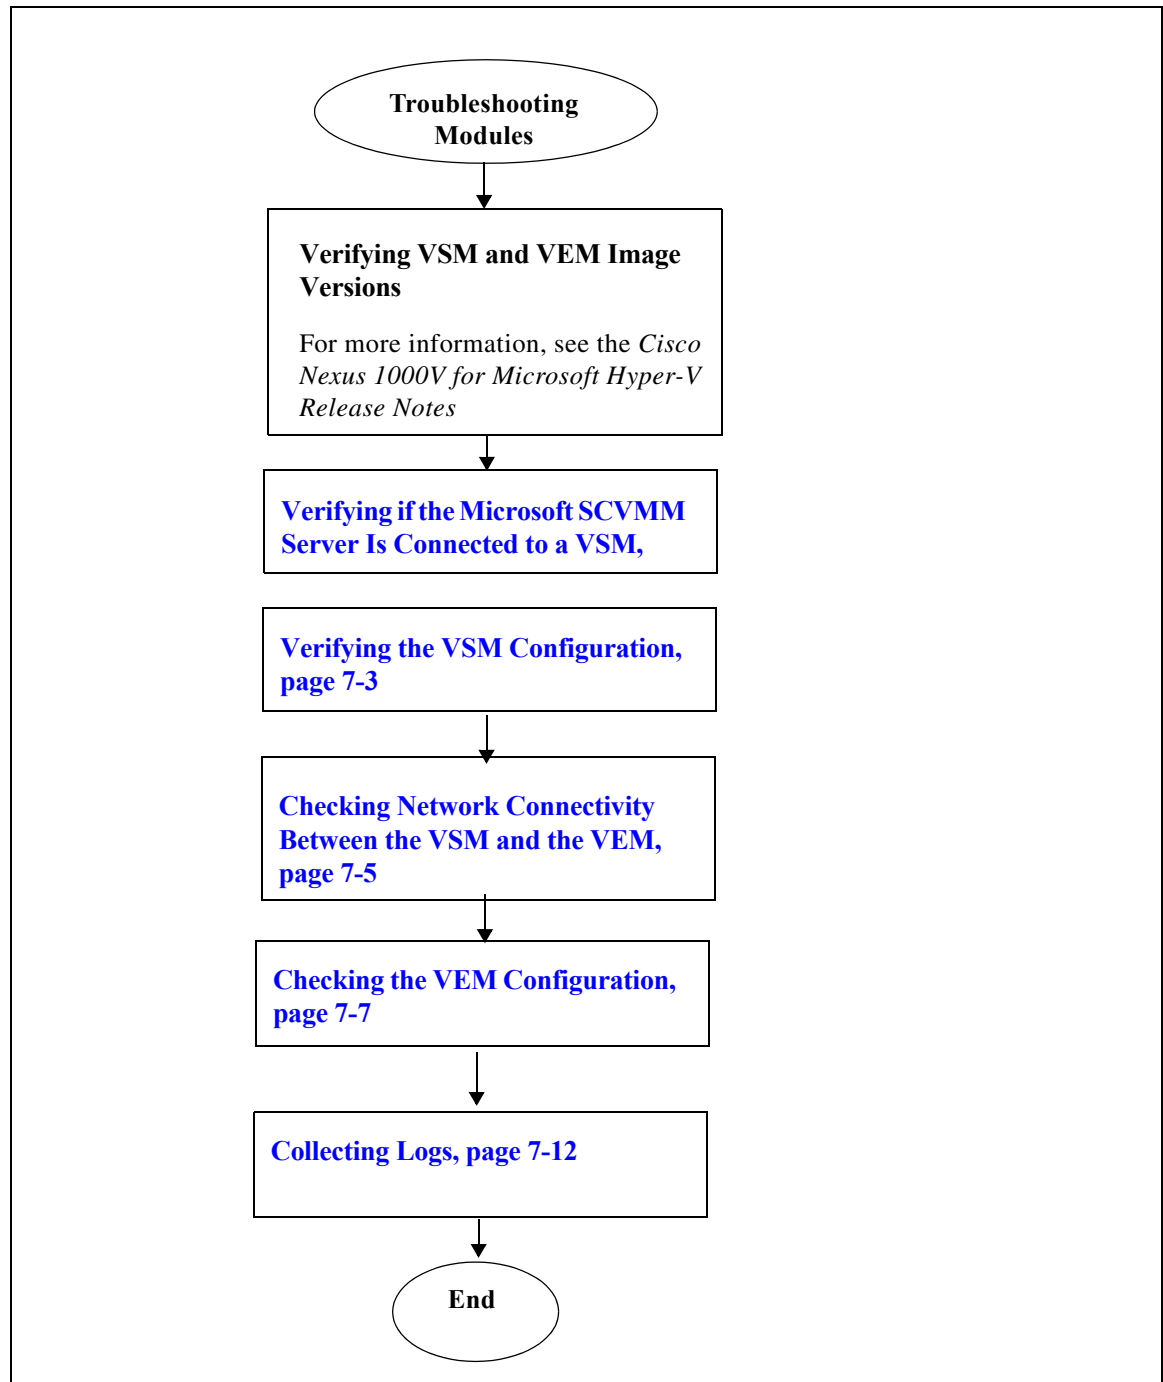
Follow these guidelines when troubleshooting a module that is controlled by the VSM:

- You must have a VSM VM and a VEM up and running.
- Make sure that you are running compatible versions of the Microsoft SCVMM server and VSM.
For more information, see the *Cisco Nexus 1000V for Microsoft Hyper-V Release Notes*.
- Make sure that `http://VSM-IP` is visible from the Microsoft SCVMM server. If it is not visible, do the following:
 1. Check the Internet Explorer proxy settings.
 2. Make sure that port 80 and port 443 are open.

Flowchart for Troubleshooting Modules

Use the following flowchart to troubleshoot modules.

Flowchart: Troubleshooting Modules



Verifying if the Microsoft SCVMM Server Is Connected to a VSM

You can verify that the Microsoft SCVMM server is connected to a VSM as follows:

- Verify that you can navigate to VSM `http://vsm_ip_address` from the server where the Microsoft SCVMM service is running.
- Verify that your proxy settings are not impacting the Microsoft SCVMM to VSM connectivity.

Verifying the VSM Configuration

You can verify the domain configuration.

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.

Step 1 On the VSM, verify the domain configuration by entering this command:

show svcs domain

```
n1000v# show svcs domain
SVS domain config:
  Domain id:      1000
  Control vlan:   1
  Packet vlan:    1
  Control mode:   L3
  Switch guid:    d8a2ad7f-a0e3-41d2-9cb9-0599946dd1e9
  L3 control interface: mgmt0
```

Problems with the VSM

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

Symptom	Possible Causes	Solution
After a VSM is rebooted, the system stops functioning in one of the following states and does not recover on its own. Attempts to debug fail.		
After boot, the VSM has a loader prompt.	The VSM kickstart image has been corrupted.	<p>Boot the VSM from the CD ROM.</p> <p>From the CD Boot menu, choose Option 1, Install Nexus1000v and bring up the new image.</p> <p>Follow the VSM installation procedure.</p>
	Boot variables are not set.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk is unformatted and bring up the new image. 3. Set the boot variables used to boot the VSM: boot system bootflash:system-boot-variable-name boot kickstart bootflash:kickstart-boot-variable-name 4. Reload the VSM by entering the reload command.
After boot, the VSM has a boot prompt.	The VSM system image has been corrupted.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 1, Install Nexus1000v and bring up the new image. 3. Follow the VSM installation procedure.
After boot, the VSM has been reconfigured.	The startup configuration has been 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 by entering the copy source filesystem: filename system:running-config command. • If you have not saved a backup copy of your configuration file, reconfigure the VSM by reading the <i>Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide</i>.

Symptom	Possible Causes	Solution
After boot, the VSM has stopped at “Loader Loading.”	The boot menu file has been corrupted.	<ol style="list-style-type: none"> 1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose Option 3, Install Nexus1000v only if the disk is unformatted and bring up new image. 3. Do one of the following: <ul style="list-style-type: none"> • If you have a saved backup copy of your configuration file, restore the configuration on the VSM by entering the copy source filesystem: filename system:running-config command. • If you have not saved a copy of your configuration file, reconfigure the VSM by reading the <i>Cisco Nexus 1000V for Microsoft Hyper-V Installation and Upgrade Guide</i>.
After boot, the secondary VSM reboots continuously.	The control VLAN or control interface is down.	Check the control connectivity between the active and standby VSM.
	Active and standby VSMs are failing to synchronize.	From the active VSM, check gsyncstats to identify which application caused the failure by entering the show logging command.
The vEthernet interface is down after the host is reloaded.	The VM is set to auto start without any start delay.	The VM should have a start delay of at least 30 seconds to allow the VEM to connect to the VSM before the VM is powered on.

Checking Network Connectivity Between the VSM and the VEM

You can verify Layer 2 network connectivity between the VSM and VEM.

Step 1 On the VSM, find its MAC address by entering this command:

show svs 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.

```
n1000v# show svs neighbors
```

```
Active Domain ID: 1000
```

```
AIPC Interface MAC: 0015-5df6-2748
```

```
Inband Interface MAC: 0015-5df6-2746
```

```
Src MAC Type Domain-id Node-id Last learnt (Sec. ago)
```

```
-----
```

```
0015-5df6-274b VSM 1000 0201 1.01
```

```
0002-3d4b-b802 VEM 1000 0302 0.76
0002-3d4b-b803 VEM 1000 0402 0.76
0002-3d4b-b804 VEM 1000 0502 0.76
```

Step 2 Do one of the following:

- If the output of the **show svs neighbors** command in [Step 1](#) does not display the Virtual Ethernet Modules (VEMs), there might be a problem with the VSM network connectivity. Proceed to the next step.
- If only some VEMs are missing, the problem might be on the VEM. See the “[Checking the VEM Configuration](#)” section on page 7-7.

Step 3 On the upstream switch, display the MAC address table to verify the network configuration by entering this command:

show mac address-table interface *int_id* vlan *vlan_id*



Note The MAC address table should be checked on the VLAN where the VSM is connected.

```
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 4 If the output from [Step 3](#) does not display the MAC address of the VSM, there might be a problem with the VSM network connectivity. Check the VSM VM adapter configuration on the Hyper-V manager or from the Microsoft SCVMM server.

Verifying the VEM Installation

Step 1 Verify the Cisco Nexus 1000V service is running by entering this command in a PowerShell window:

```
PS C:\Users\Administrator.HPV> net start | Select-String Nexus
```

Step 2 Verify that the VEM has created a team with all PNICs connected to the Cisco Nexus 1000V by entering the following command in a PowerShell window:

```
PS C:\Users\Administrator.HPV> Get-NetSwitchTeam
```

```
PS C:\Users\Administrator.HPV> Get-NetSwitchTeam
Name      : HPVec6055e5-6b3c-4cb5-9af4-d6ea373284f0
Members   : {Ethernet 4, Ethernet 2}
```

Step 3 Verify that Opaque Data (OD) is set correctly by entering this command in a PowerShell window:

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show data
```

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show data
data-version 1.0
switch-domain 1000
```

```

switch-name VSM
cp-version 5.2(1)SM1(5.1) [build 5.2(1)SM1(5.0.333)]
control-vlan 1
<truncated output>
standby-vsm ctrl mac 0015-5df6-274f
inband-vlan 1
svs-mode L3
l3control-ipaddr 12.0.45.100
upgrade state 0 mac 0015-5df6-274f l3control-ipv4 null
sequence-number 8644
end-version 1.0

```



Note The important values that you should verify are the Switch-Domain and L3control-ipaddr fields. The value for L3control-ipaddr should match the Layer 3 mode configuration of either the mgmt0 or control0.

Checking the VEM Configuration

You can verify the VEM configuration on the Hyper-V host.

Step 1 Verify the domain ID by entering this command:

vemcmd show card

```

PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show card
Card UUID type 2: 542B14B5-0CBD-E011-BD1D-30E4DBC2C6DE
Card name: WIN-35
Switch name: VSM
Switch alias:
Switch uuid: 62347037-DED0-4C10-A987-93D4EC75E555
Card domain: 1000
Card slot: 3
VEM Tunnel Mode: L3 Mode
L3 Ctrl Index: 0
L3 Control IPv4 address: 10.105.225.52
VEM Control (AIPC) MAC: 00:02:3d:1b:b8:02
VEM Packet (Inband) MAC: 00:02:3d:2b:b8:02
VEM Control Agent (DPA) MAC: 00:02:3d:4b:b8:02
VEM SPAN MAC: 00:02:3d:3b:b8:02
Primary VSM MAC : 00:15:5d:f6:27:48
Primary VSM PKT MAC : 00:15:5d:f6:27:46
Primary VSM MGMT MAC : 00:15:5d:f6:27:47
Standby VSM CTRL MAC : 00:15:5d:f6:27:4b
Management IPv4 address: 10.105.225.35
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
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: 256
Card control VLAN: 1
Card packet VLAN: 1
Card Headless Mode : No
Processors: 24
Processor Cores: 24
Processor Sockets: 2
Kernel Memory: 1073741824

```

```

Port link-up delay: 5s
Global UUFB: DISABLED
Heartbeat Set: True
PC LB Algo: source-mac
System Profile Check Enabled : Yes
Datapath portset event in progress : no
Batch Speed Duplex : yes
Licensed: Yes

```

The domain ID should match the value shown in the VSM when you enter the **show svcs domain** command.

- Step 2** Verify that the ports of the host added to the logical switch are listed and that the ports are correctly configured as access or trunk on the host by entering this command:

vemcmd show port

```

VSM# module vem 3 execute vemcmd show port-old
LTL IfIndex Vlan/ Bndl SG_ID Pinned_SGID Type Admin State CBL Mode Name
SegId
1 15f00010 3968 0 32 32 VIRT UP UP 1 Access inban
6 0 1 T 0 32 32 VIRT UP UP 1 Trunk vns
11 0 3968 0 32 32 VIRT UP UP 1 Access
16 0 1 T 0 32 32 VIRT UP UP 1 Trunk ar
20 250080c0 1 T 305 3 32 PHYS UP UP 1 Trunk Intel(R) Gigabit ET Quad Port Server Adapter
#2
pvlan promiscuous trunk port
357 --> 356
358 --> 356
21 25008100 1 T 305 4 32 PHYS UP UP 1 Trunk Intel(R) Gigabit ET Quad Port Server Adapter
#3
pvlan promiscuous trunk port
357 --> 356
358 --> 356
22 25008140 1 T 305 5 32 PHYS UP UP 1 Trunk Intel(R) Gigabit ET Quad Port Server Adapter
#4
pvlan promiscuous trunk port
357 --> 356
358 --> 356
49 1c000020 342 0 32 3 VIRT UP UP 1 Access Win2008-3-1
50 1c000000 342 0 32 4 VIRT UP UP 1 Access Win2008-2-1
305 16000001 1 T 0 32 32 CHAN UP UP 1 Trunk
pvlan promiscuous trunk port
357 --> 356
358 --> 356

```

The last line of the output indicates that vmnic1 should be in trunk mode with a color blocking logic (CBL) value of 1. The CBL value of the native VLAN does not have to be 1. It can be 0 if it is not allowed or 1 if it is VLAN 1 and not allowed. If the CBL value is 0 it is not a problem unless the native VLAN is the control VLAN. The Admin state and Port state should be UP.

- Step 3** Check if the VSM is reachable from the Hyper-v Host by entering this command:

ROUTE PRINT

arp -a

Problems with the VEM

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

Symptom	Possible Causes	Solution
A VEM that you created has failed.	The maximum transmission unit (MTU) of the NICs that are connected to the Cisco Nexus 1000V do not match.	<p>The MTU can be set using one of the following steps:</p> <ul style="list-style-type: none"> Set the MTU value in the Advanced Properties of the NIC on the host in the Control Panel. Enter the following command in a PowerShell window: PS C:\Program Files (x86)\Cisco\Nexus1000V> Get-NetAdapterAdvancedProperty-RegistryKeyword *jumbo* -Name <i>adapter name</i> Set-NetAdapterAdvancedProperty -RegistryValue <i>mtu_value</i>
	A previous deletion of the Cisco Nexus 1000V has left a stale team.	<ol style="list-style-type: none"> Find the team name by entering the following command in a PowerShell window: PS C:\> Get-NetSwitchTeam Delete the team by entering the following command in a PowerShell window: PS C:\> Remove-NetSwitchTeam -Name <i>name</i>

Symptom	Possible Causes	Solution
	A previous deletion of the Cisco Nexus 1000V has left a stale switch.	<ol style="list-style-type: none"> 1. Delete the stale switch by doing one of the following: <ul style="list-style-type: none"> – Delete the switch from the Microsoft SCVMM UI. – Delete the switch from the Hyper V manager. – Delete the switch from the registry location: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\VMSMP\Parameters\SwitchList 2. Reboot the host.
A module is not attaching to the VSM.	The switch extension is not enabled.	<ol style="list-style-type: none"> 1. Launch the Hyper-V manager. 2. Choose Virtual Switch manager. 3. Verify that the Cisco Nexus 1000V extension is enabled.
	The VEM has stale opaque data that does not have the correct VSM information.	<ol style="list-style-type: none"> 1. Launch the Microsoft SCVMM UI. 2. Choose Fabric Management > Networking > Switch Extension Manager. 3. Get the latest opaque data (OD) by refreshing the Switch Extension Manager in the Microsoft SCVMM. 4. Choose Host from the toolbar and then choose Remediate from the toolbar.

Verifying the Logical Switch Instance

You can verify the properties of a logical switch instance.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** Verify that the uplink NICs are correct.
- Step 3** Verify that the uplink networks are assigned to the logical switch.
-

Problems with Logical Switches

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

Symptom	Possible Causes	Solution
A compliant logical switch cannot be found.	The VM deployment or the port profiles might have changed.	<ol style="list-style-type: none"> View the compliance status by one of the following methods: <ul style="list-style-type: none"> Launch Microsoft SCVMM UI and choose Fabric > Logical Switches > Hosts. Enter the following command in a PowerShell window: <pre>PS C:\> Get-SCVirtualNetwork where-object {\$_.LogicalSwitch -like "1000V"} select VMHost, HighlyAvailable, LogicalNetworks, VMHostNetworkAdapters, LogicalSwitchComplianceStatus</pre> Remove the warning by choosing Fabric > Logical Switches > Hosts. Choose the logical switch and choose Remediate.
A logical switch has not been created.	A stale NetSwitch team is already present on the host.	<ol style="list-style-type: none"> Find the team name by entering this command in a PowerShell window: <pre>PS C:\> Get-NetSwitchTeam</pre> Delete the team by entering this command in a PowerShell window: <pre>PS C:\> Remove-NetSwitchTeam -Name name</pre>
	The Switch Extension Manager is not allowed on the host group.	<ol style="list-style-type: none"> Launch the Microsoft SCVMM UI. Choose Switch Extension Manager properties. Select the appropriate host group for the Switch Extension Manager.

WinRM Connectivity Problems

The Windows Remote Management (WinRM) is the Microsoft implementation of the WS-MAN protocol. WinRM is enabled by default only on Windows 2016.

Symptom	Possible Causes	Solution
WinRM error connectivity is lost.	A stale NetSwitch team is already present on the host.	<ol style="list-style-type: none"> Find the team name by entering the following command in a PowerShell window: PS C:\> Get-NetSwitchTeam Delete the team by entering the following command in a PowerShell window: PS C:\> Remove-NetSwitchTeam -Name <i>name</i>
You are unable to perform an operation on the host.	The host management connectivity is lost.	For more information, see the following URL: http://blogs.technet.com/b/jonjor/archive/2009/01/09/winrm-windows-remote-management-troubleshooting.aspx

Verifying Failover Clustering

Failover clustering provides high availability for VMs. The following restrictions apply:

- When a Failover cluster is created the VM must be made highly available for the VM to be a part of the cluster.
- A logical switch created on a host that is part of a cluster becomes highly available if and only if all the uplinks of the logical switch instances carry the same network segment pools.
- A VM on remote storage that is highly available that resides on a host that is a part of a failover cluster, can only be connected to a highly available switch.
- A VM on local storage that resides on a host that is part of a failover cluster can be connected to any switch.

Collecting Logs

After you verify the network connectivity between the VEM and the VSM, you can collect log files to help identify the problem.

Step 1 On the VEM, verify its universally unique identifier (UUID) by entering this command:

vemcmd show card info

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show card info
Card UUID type 2: 542B14B5-0CBD-E011-BD1D-30E4DBC2C6DE
Card name: WIN-35
Switch name: HPV
Switch alias:\
Switch uuid: 8A1AB26E-06E2-4F64-9258-4560C2BCB82D
Card domain: 1000
Card slot: 3
VEM Tunnel Mode: L3 Mode
L3 Ctrl Index: 0
L3 Control IPv4 address: 10.105.225.52
VEM Control (AIPC) MAC: 00:02:3d:1b:b8:02
```

```

VEM Packet (Inband) MAC: 00:02:3d:2b:b8:02
VEM Control Agent (DPA) MAC: 00:02:3d:4b:b8:02
VEM SPAN MAC: 00:02:3d:3b:b8:02
Primary VSM MAC : 00:15:5d:f6:27:48
Primary VSM PKT MAC : 00:15:5d:f6:27:46
Primary VSM MGMT MAC : 00:15:5d:f6:27:47
Standby VSM CTRL MAC : 00:15:5d:f6:27:4b
Management IPv4 address: 10.105.225.35
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
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: 256
Card control VLAN: 1
Card packet VLAN: 1
Card Headless Mode : No
Processors: 24
Processor Cores: 24
Processor Sockets: 2
Kernel Memory: 1073741824
Port link-up delay: 5s
Global UUFB: DISABLED
Heartbeat Set: True
PC LB Algo: source-mac
System Profile Check Enabled : Yes
Datapath portset event in progress : no
Batch Speed Duplex : yes
Licensed: Yes
PS C:\Program Files (x86)\cisco\Nexus1000V>

```

Step 2 On the VSM, verify the module number to which the corresponding UUID entry is mapped by entering this command:

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#

```

Step 3 Using the module number from [Step 2](#), collect the output of these 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 must have 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 SVS neighbors. See Example 7-1 on page 7-15 .
show svcs domain	Displays the domain configuration. See Example 7-2 on page 7-15 .
show port-profile name <i>name</i>	Displays the configuration for a named port profile. See Example 7-3 on page 7-16 .
show running-config vlan <i>vlanID</i>	Displays the VLAN information in the running configuration. See Example 7-4 on page 7-16 .
show mac address-table interface	Displays the MAC address table on an upstream switch to verify the network configuration. See Example 7-5 on page 7-16 .
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 address appears in the control and packet VLANs. See Example 7-6 on page 7-17 .
vemcmd show card	Displays information about the cards on the VEM to verify that the domain ID, control VLANs, and packet VLANs are configured correctly on the host. See Example 7-7 on page 7-17 .
vemcmd show vmq allocation	Displays Virtual Machine Queue (VMQ) allocation for vEthernet interfaces. See Example 7-8 on page 7-17 .
vemcmd show vmq resources	Displays information about the VMQ-enabled PNICs. See Example 7-9 on page 7-18 .
vemcmd show attach	Displays information about the platform port attach. See Example 7-10 on page 7-18 .
vemcmd show vem internal info	Displays information about the VEM queue status. See Example 7-11 on page 7-18 .
vemcmd show port [<i>port-LTL-number</i>]	Displays information about the 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 7-12 on page 7-18 .

Command	Description
vemcmd show bd [<i>control_vlan_id</i> <i>packet_vlan_id</i>]	Displays the list of ports that belong to the VLAN. Note The bd number is not the same as a VLAN number. You can display a listing of bd numbers by entering the vemcmd show bd command. See Example 7-13 on page 7-19 .
vemcmd show trunk	Displays configured information about the VEM to verify that the DV port groups are successfully pushed from the Microsoft SCVMM server to the host and that the correct physical trunk port VM NIC is used. See Example 7-14 on page 7-19 .
show module vem mapping	Displays information about the VEM that a VSM maps to, including the VEM module number, status, UUID, and license status. See Example 7-15 on page 7-19 .
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. See Example 7-16 on page 7-19 .
show system internal im event-history module <i>module-number</i>	Displays the module IM event logs for the system. See Example 7-17 on page 7-20 .
show system internal vmm event-history module <i>module-number</i>	Displays the module VMM event logs for the system. See Example 7-18 on page 7-20 .
show system internal ethpm event-history module <i>module-number</i>	Displays the module Ethernet event logs for the system. See Example 7-19 on page 7-21 .
show system internal ethpm event-history interface <i>type slot</i>	Displays the Ethernet interface logs for the system. See Example 7-20 on page 7-22 .

Example 7-1 *show svcs neighbors Command*

```

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 domain Command*

```

n1000v# show svcs domain
SVS domain config:

```

```

Domain id:      942
Control vlan: 1
Packet vlan: 1
Control mode: L3
Switch guid: dafe01ef-dfe0-4277-9533-300a2347555e
L3 control interface: mgmt0
Status: Config not pushed to Management Server.

```

Example 7-3 *show port-profile Command*

```

n1000v# show port-profile name SystemUplink
port-profile SystemUplink
type: Ethernet
description: NSM created profile. Do not delete.
status: enabled
max-ports: 512
min-ports: 1
inherit: PortChannelProfile
config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 173
evaluated config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 173
  channel-group auto
  no shutdown
assigned interfaces:
port-group:
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vn-service: no
port-profile role: none
port-binding: static

```

Example 7-4 *show running-config vlan Command*

```

n1000v# show running-config vlan 260-261
version 5.2(1)SM1(5.1)
vlan 260
  name cp_control
vlan 261
  name cp_packet

n1000v#

```

Example 7-5 *show mac address-table interface Command*

```

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 7-6 module vem execute vemcmd show l3 Command

```

n1000v# configure terminal
n1000v(config)# module vem 3 execute vemcmd show l3 1002
Bridge domain 1002 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 l3 1003
Bridge domain 1002 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 7-7 vemcmd show card info Command

```

PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show card info
Card UUID type 2: 542B14B5-0CBD-E011-BD1D-30E4DBC2C6DE
Card name: WIN-35
Switch name: HPV
Switch alias:
Switch uuid: 8A1AB26E-06E2-4F64-9258-4560C2BCB82D
Card domain: 1000
Card slot: 3
VEM Tunnel Mode: L3 Mode
L3 Ctrl Index: 0
L3 Control IPv4 address: 10.105.225.52
VEM Control (AIPC) MAC: 00:02:3d:1b:b8:02
VEM Packet (Inband) MAC: 00:02:3d:2b:b8:02
VEM Control Agent (DPA) MAC: 00:02:3d:4b:b8:02
VEM SPAN MAC: 00:02:3d:3b:b8:02
Primary VSM MAC : 00:15:5d:f6:27:48
Primary VSM PKT MAC : 00:15:5d:f6:27:46
Primary VSM MGMT MAC : 00:15:5d:f6:27:47
Standby VSM CTRL MAC : 00:15:5d:f6:27:4b
Management IPv4 address: 10.105.225.35
Management IPv6 address: 0000:0000:0000:0000:0000:0000:0000:0000
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: 256
Card control VLAN: 1
Card packet VLAN: 1
Card Headless Mode : No
Processors: 24
Processor Cores: 24
Processor Sockets: 2
Kernel Memory: 1073741824
Port link-up delay: 5s
Global UUFB: DISABLED
Heartbeat Set: True
PC LB Algo: source-mac
System Profile Check Enabled : Yes
Datapath portset event in progress : no
Batch Speed Duplex : yes
Licensed: Yes

```

Example 7-8 vemcmd show vmq allocation Command

```

~ # vemcmd show vmq allocation
LTL   VSM Port  Phy LTL  Queue id  Team queue id
49     Veth13    17         1         49

```

		18	2		49
50	Veth14		17	2	50
		18	3		50
51	Veth16		19	1	51
		20	1		51

Example 7-9 vemcmd show vmq resources Command

```
~ # vemcmd show vmq resources
LTL VSM Port Max queues Free queues
17 Eth3/1 16 10
18 Eth3/2 16 10
19 Eth3/3 8 7
```

Example 7-10 vemcmd show attach Command

```
~ # vemcmd show attach
-----
LTL: 17
-----
Port ID: 1
NIC Index: 1
Port UUID: BC9C4957-88B0-4292-879A-A4109A5A345B
NIC Instance ID: {239C8D0D-43AD-4DB7-94E1-1D90D265D21F}
MAC address: d0:d0:fd:09:31:f8
Port profile: uplink-trunk
VM/NIC name: Intel(R) 82576 Gigabit Dual Port Network Connection
VM UUID:
MTU: 1514
Link state: UP
Duplex: Full
Tx speed: 1000000000
Rx speed: 1000000000
Autoneg: Enabled
Link Params pending: No
Speed Capability 0x13
Duplex Capability 0x7
```

Example 7-11 vemcmd show vem internal info Command

```
~ # vemcmd show vem internal info
-----
VEM Internal counters
-----
# Tx pkts pending: 0
# Timer events queued: 0
# Internal pkts queued: 0
# DPA notifications queued: 0
```

Example 7-12 vemcmd show port Command

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd show port
LTL VSM Port Admin Link State PC-LTL SGID Vem Port Type
20 Eth3/4 UP UP FWD 305 3 Intel(R) Gigabit ET Quad Port Server Adapter #2
21 Eth3/5 UP UP FWD 305 4 Intel(R) Gigabit ET Quad Port Server Adapter #3
22 Eth3/6 UP UP FWD 305 5 Intel(R) Gigabit ET Quad Port Server Adapter #4
49 Veth3 UP UP FWD 0 3 Win2008-3-1
50 Veth1 UP UP FWD 0 4 Win2008-2-1
305 Po2 UP UP FWD 0
```

Example 7-13 vemcmd show bd Command

```
VSM# module vem 3 execute vemcmd show bd 8
BD 8, vdc 1, vlan 342, swbd 342, 6 ports, ""
Portlist:
20 Intel(R) Gigabit ET Quad Port Server Adapter #2
21 Intel(R) Gigabit ET Quad Port Server Adapter #3
22 Intel(R) Gigabit ET Quad Port Server Adapter #4
49 Win2008-3-1
50 Win2008-2-1
305

Multicast Group Table:
Group 0.0.0.0 Multicast LTL 4410
305
```

Example 7-14 vemcmd show trunk Command

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

Example 7-15 show module vem mapping Command

```
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#
```

Example 7-16 show module internal event-history module Command

```
n1000v# show module internal event-history module 1
>>>>FSM: <ID(257): Slot 1, node 0x0101> has 4logged transitions<<<<<

1) FSM:<ID(257): Slot 1, node 0x0101> Transition at 920000 usecs after Wed Apr
24 10:01:45 2013
   Previous state: [LCM_ST_LC_NOT_PRESENT]
   Triggered event: [LCM_EV_PFM_MODULE_SUP_INSERTED]
   Next state: [LCM_ST_SUPERVISOR_INSERTED]

2) FSM:<ID(257): Slot 1, node 0x0101> Transition at 920000 usecs after Wed Apr
24 10:01:45 2013
   Previous state: [LCM_ST_SUPERVISOR_INSERTED]
   Triggered event: [LCM_EV_START_SUP_INSERTED_SEQUENCE]
   Next state: [LCM_ST_CHECK_INSERT_SEQUENCE]

3) Event:ESQ_START length:38, at 920000 usecs after Wed Apr 24 10:01:45 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   Seq Type:SERIAL

4) FSM:<ID(257): Slot 1, node 0x0101> Transition at 0 usecs after Wed Apr 24 1
0:02:10 2013
   Previous state: [LCM_ST_LC_ONLINE]
   Triggered event: [LCM_EV_LC_ONLINE]
   Next state: [No transition found]

   Curr state: [LCM_ST_LC_ONLINE]n1000v#
n1000v#
```

Example 7-17 *show system internal im event-history module Command*

```

n1000v# show system internal im event-history module 1
>>>FSM: <Module NodeID(0x101)> has 4 logged transitions<<<<

1) FSM:<Module NodeID(0x101)> Transition at 970000 usecs after Wed Apr 24 10:02
:09 2013
    Previous state: [IM_MOD_ST_MODULE_NOT_EXISTENT]
    Triggered event: [IM_MOD_EV_MOD_INSERTED]
    Next state: [IM_MOD_ST_WAIT_CONFIG_FLUSH]

2) FSM:<Module NodeID(0x101)> Transition at 970000 usecs after Wed Apr 24 10:02
:09 2013
    Previous state: [IM_MOD_ST_WAIT_CONFIG_FLUSH]
    Triggered event: [IM_MOD_EV_CONFIG_FLUSH_BYPASSED]
    Next state: [IM_MOD_ST_WAIT_PLATFORM_INIT]

3) Event:ESQ_START length:38, at 970000 usecs after Wed Apr 24 10:02:09 2013
    Instance:257, Seq Id:0x1, Ret:SUCCESS
    Seq Type:SERIAL
>>>FSM: <Module NodeID(0x101)> has 13 logged transitions<<<<

4) FSM:<Module NodeID(0x101)> Transition at 980000 usecs after Wed Apr 24 10:0
2:09 2013
    Previous state: [IM_MOD_ST_WAIT_INTERFACE_BIND]
    Triggered event: [IM_MOD_EV_INTERFACE_BIND_BYPASSED]
    Next state: [IM_MOD_ST_MODULE_INIT_DONE]

    Curr state: [IM_MOD_ST_MODULE_INIT_DONE]

```

Example 7-18 *show system internal vmm event-history module Command*

```

n1000v# show system internal vmm event-history module 1
>>>FSM: <ID(257): Module 1> has 8 logged transitions<<<<

1) FSM:<ID(257): Module 1> Transition at 950000 usecs after Wed Apr 24 10:02:15
2013
    Previous state: [VMM_ST_IDLE]
    Triggered event: [VMM_EV_IF_BIND]
    Next state: [VMM_ST_CHECK_INSERT_SEQUENCE]

2) Event:ESQ_START length:38, at 950000 usecs after Wed Apr 24 10:02:15 2013
    Instance:257, Seq Id:0x1, Ret:SUCCESS
    Seq Type:SERIAL

3) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
    Instance:257, Seq Id:0x1, Ret:SUCCESS
    [E_MTS_TX] Dst:MTS_SAP_ETH_PORT_CHANNEL_MGR(378), Opc:MTS_OPC_IM_IF_VDC_BIN
D(62488)
    RRtoken:0x000019F0

4) Event:ESQ_RSP length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
    Instance:257, Seq Id:0x1, Ret:SUCCESS
    [E_MTS_RX] Src:MTS_SAP_ETH_PORT_CHANNEL_MGR(378), Opc:MTS_OPC_IM_IF_VDC_BIN
D(62488)
    RRtoken:0x000019F0

5) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013

```

```

Instance:257, Seq Id:0x1, Ret:SUCCESS
[E_MTS_TX] Dst:MTS_SAP_TEST_ETHPM(175), Opc:MTS_OPC_IM_IF_VDC_BIND(62488)
RRtoken:0x000019F5

6) Event:ESQ_RSP length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
[E_MTS_RX] Src:MTS_SAP_TEST_ETHPM(175), Opc:MTS_OPC_IM_IF_VDC_BIND(62488)
RRtoken:0x000019F5

7) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
Type: 0

8) FSM:<ID(257): Module 1> Transition at 990000 usecs after Wed Apr 24 10:02:15
2013
Previous state: [VMM_ST_CHECK_INSERT_SEQUENCE]
Triggered event: [VMM_EV_INSERT_SEQ_DONE]
Next state: [VMM_ST_IDLE]

Curr state: [VMM_ST_IDLE]
n1000v#

```

Example 7-19 *show system internal ethpm event-history module Command*

```

n1000v# show system internal ethpm event-history module 1
>>>FSM: <Module NodeID(0x101)> has 8 logged transitions<<<<

1) FSM:<Module NodeID(0x101)> Transition at 990000 usecs after Wed Apr 24 10:02:15 2013
Previous state: [ETHPM_MODULE_ST_MODULE_NOT_EXISTENT]
Triggered event: [ETHPM_MODULE_EV_IF_BIND_CMD]
Next state: [FSM_ST_NO_CHANGE]

2) FSM:<Module NodeID(0x101)> Transition at 990000 usecs after Wed Apr 24 10:02:15 2013
Previous state: [ETHPM_MODULE_ST_MODULE_NOT_EXISTENT]
Triggered event: [ETHPM_MODULE_EV_SUP_INSERT]
Next state: [ETHPM_MODULE_ST_AWAIT_SUP_INSERT]

3) Event:ESQ_START length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
Seq Type:SERIAL

4) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
Seq:SUP_INTERNAL_INIT

5) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
[E_MTS_TX] Dst:MTS_SAP_REGISTRY(0), Opc:MTS_OPC_PSSHELPER_PUB_WRITE(28673)

6) Event:ESQ_RSP length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
[E_MTS_RX] Src:MTS_SAP_REGISTRY(0), Opc:MTS_OPC_PSSHELPER_PUB_WRITE(28673)

7) Event:ESQ_REQ length:38, at 990000 usecs after Wed Apr 24 10:02:15 2013
Instance:257, Seq Id:0x1, Ret:SUCCESS
Seq:Update_Sup_Module_PSS

8) FSM:<Module NodeID(0x101)> Transition at 990000 usecs after Wed Apr 24 10:02:15 2013

```

```

Previous state: [ETHPM_MODULE_ST_AWAIT_SUP_INSERT]
Triggered event: [ETHPM_MODULE_EV_SUP_INSERT_DONE]
Next state: [ETHPM_MODULE_ST_MODULE_PRESENT]

```

```

Curr state: [ETHPM_MODULE_ST_MODULE_PRESENT]
n1000v#

```

Example 7-20 *show system internal ethpm event-history module Command*

```

n1000v# show system internal ethpm event-history module 1

>>>FSM: <Module NodeID(0x101)> has 8 logged transitions<<<<

1) FSM:<Module NodeID(0x101)> Transition at 120000 usecs after Thu Apr 25 11:35
:21 2013
   Previous state: [ETHPM_MODULE_ST_MODULE_NOT_EXISTENT]
   Triggered event: [ETHPM_MODULE_EV_IF_BIND_CMD]
   Next state: [FSM_ST_NO_CHANGE]

2) FSM:<Module NodeID(0x101)> Transition at 120000 usecs after Thu Apr 25 11:35
:21 2013
   Previous state: [ETHPM_MODULE_ST_MODULE_NOT_EXISTENT]
   Triggered event: [ETHPM_MODULE_EV_SUP_INSERT]
   Next state: [ETHPM_MODULE_ST_AWAIT_SUP_INSERT]

3) Event:ESQ_START length:38, at 120000 usecs after Thu Apr 25 11:35:21 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   Seq Type:SERIAL

4) Event:ESQ_REQ length:38, at 120000 usecs after Thu Apr 25 11:35:21 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   Seq:SUP_INTERNAL_INIT

5) Event:ESQ_REQ length:38, at 120000 usecs after Thu Apr 25 11:35:21 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   [E_MTS_TX] Dst:MTS_SAP_REGISTRY(0), Opc:MTS_OPC_PSSHELPER_PUB_WRITE(28673)

6) Event:ESQ_RSP length:38, at 120000 usecs after Thu Apr 25 11:35:21 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   [E_MTS_RX] Src:MTS_SAP_REGISTRY(0), Opc:MTS_OPC_PSSHELPER_PUB_WRITE(28673)

7) Event:ESQ_REQ length:38, at 120000 usecs after Thu Apr 25 11:35:21 2013
   Instance:257, Seq Id:0x1, Ret:SUCCESS
   Seq:Update_Sup_Module_PSS

8) FSM:<Module NodeID(0x101)> Transition at 120000 usecs after Thu Apr 25 11:35
:21 2013
   Previous state: [ETHPM_MODULE_ST_AWAIT_SUP_INSERT]
   Triggered event: [ETHPM_MODULE_EV_SUP_INSERT_DONE]
   Next state: [ETHPM_MODULE_ST_MODULE_PRESENT]

Curr state: [ETHPM_MODULE_ST_MODULE_PRESENT]

```



Ports

This chapter describes how to identify and resolve problems with ports.

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 might not be valid when the interface is down (such as the operational speed).

For a complete description of port modes, administrative states, and operational states, see the *Cisco Nexus 1000V for Microsoft Hyper-V 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, enter this command:

show interface ethernet *mod/port* counters

See [Example 8-10 on page 8-10](#).

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 by entering this command:

clear counters interface ethernet *mod/port*

Information About Link Flapping

A port that continually goes up and down is called flapping or a link-flapping port. 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

Port security enables you to secure a port by limiting and identifying the MAC addresses that can access the port. Secure MAC addressees can be manually configured or dynamically learned.

For detailed information about port security, see the *Cisco Nexus 1000V for Microsoft Hyper-V 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 topics:

- [VMs Cannot Ping a Secured Port, page 8-5](#)
- [Port Security Violations, page 8-6](#)

Port Diagnostic Checklist

Use the following checklist to diagnose the port interface activity.

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

Checklist	Example	✓
Verify that the module is active by entering the show module command.	See Example 8-1 on page 8-8 .	

Checklist (continued)	Example	✓
On the VMM client that is connected to the Microsoft SCVMM server, verify that required port profiles are assigned to the physical NICs and the virtual NICs.		
Verify that the ports have been created and the state of the interface by entering the show interface brief command.	See Example 8-7 on page 8-10 .	

Problems with Ports

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

- [Cannot Enable an Interface, page 8-3](#)
- [Port Link Failure or Port Not Connected, page 8-4](#)
- [Link Flapping, page 8-4](#)
- [Port is ErrDisabled, page 8-5](#)
- [VMs Cannot Ping a Secured Port, page 8-5](#)
- [Port Security Violations, page 8-6](#)

Cannot Enable an Interface

Possible Cause	Solution
A 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 by entering the show interface brief command. 2. If not already associated, associate the interface with an access VLAN. 3. Determine the VLAN status by entering the show vlan brief command. 4. If the VLAN is not already active, configure the VLAN as active by entering these commands: <ul style="list-style-type: none"> – config terminal – vlan <i>vlan-id</i> – state active

Port Link Failure or Port Not Connected

Possible Cause	Solution
The port connection is bad.	<ol style="list-style-type: none"> 1. Verify the port state by entering the show system internal ethpm info command. 2. Disable and then enable the port by entering these commands: <ul style="list-style-type: none"> – shut – no shut 3. Collect the Hyper-V side NIC configuration. In a PowerShell window, enter this command: PS C:\Program Files (x86)\Cisco\Nexus1000V\Support> .\vem-support.ps1
The link is stuck in the initialization state or the link is in a point-to-point state.	<ol style="list-style-type: none"> 1. Check for the link failure system message “Link Failure, Not Connected” by entering the show logging command. 2. Disable and then enable the port by entering these commands: <ul style="list-style-type: none"> – shut – no shut 3. Collect the Hyper-V side NIC configuration. In a PowerShell window, enter this command: PS C:\Program Files (x86)\Cisco\Nexus1000V\Support> .\vem-support.ps1

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.

Possible Cause	Solution
The bit rate exceeds the threshold and puts the port into an error-disabled state.	Disable and then enable the port by entering these commands: <ul style="list-style-type: none"> • 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.	An external device might 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.
A software error causes a packet drop.	
A control frame is erroneously sent to the device.	
A Windows error or link flapping on the upstream switch has occurred.	Use the troubleshooting guidelines in the documentation for your Windows or upstream switch.

Port is ErrDisabled

Possible Cause	Solution
The cable is defective or damaged.	<ol style="list-style-type: none"> 1. Verify the physical cabling. 2. Replace or repair defective cables. 3. Reenable the port by entering these commands: <ul style="list-style-type: none"> • shut • no shut
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 by entering the show logging logfile command. 2. Correct the error in the configuration and add the port to the port channel. 3. Reenable the port by entering these commands: <ul style="list-style-type: none"> • shut • no shut
A VSM application error has occurred.	<ol style="list-style-type: none"> 1. Identify the component that had the error while bringing up the port by entering this command: show logging log file grep interface_number See Example 8-6 on page 8-9. 2. Identify the error transition by entering this command: 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 23-1.

VMs 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 by entering the show interface vethernet number command.
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 by entering the module vem number execute vemcmd show portsec stats command. 2. If the DSM is set, clear the DSM bit on the VSM by entering the no port-security stop learning command.

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 for Microsoft Hyper-V Security Configuration Guide*.

Possible Cause	Solution
The configured maximum number of secured addresses on the port is exceeded.	<ol style="list-style-type: none"> 1. Display the secure addresses by entering these commands: <ul style="list-style-type: none"> – show port-security address vethernet <i>number</i> – show port-security 2. Identify ports with a security violation as follows: <pre>show logging inc "PORT-SECURITY-2-ETH_PORT_SEC_SECURITY_VIOLAT ION_MAX_MAC_VLAN"</pre> 3. Correct the security violation. 4. Enable the interface by entering these commands: <ul style="list-style-type: none"> – shut – no shut
A MAC address 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.

Command	Purpose
show module <i>module-number</i>	Displays the state of a module. See Example 8-1 on page 8-8 .
show svs domain	Displays the domain configuration. See Example 8-2 on page 8-8 .
show cdp neighbors	Displays the neighbors connected to an interface. See Example 8-3 on page 8-9 .
show system internal ethpm event-history interface <i>interface</i>	Displays information about the internal state transitions of the port. See Example 8-4 on page 8-9 .
show logging logfile	Displays logged system messages. See Example 8-5 on page 8-9 .
show logging logfile grep <i>interface_number</i>	Displays logged system messages for a specified interface. See Example 8-6 on page 8-9 .
show interface brief	Displays a table of interface states. See Example 8-7 on page 8-10 .

Command	Purpose
show interface ethernet <i>mod/port</i>	Displays the status of a named interface. See Example 8-8 on page 8-10 .
show running-config interface ethernet <i>mod/port expand-port-profile</i>	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 8-10 .
show interface ethernet counters	Displays port counters for identifying synchronization problems. For information about counters, see the “ Information About Interface Counters ” section on page 8-1. See Example 8-10 on page 8-10 .
show interface vethernet <i>number</i>	Displays the vEthernet interface configuration. See Example 8-11 on page 8-11 .
show interface <i>mod/port status</i>	Displays the status of the named interface.
show interface capabilities	Displays a tabular view of all configured port profiles. See Example 8-12 on page 8-11 .
show interface virtual attach binding	Displays the virtual port mapping for all vEthernet interfaces. See Example 8-13 on page 8-12 .
show system internal ethpm errors	Displays the ethpm error logs. See Example 8-14 on page 8-12 .
show system internal ethpm event-history errors	Displays the ethpm event logs. See Example 8-15 on page 8-13 .
show system internal ethpm info	Displays the internal data structure information. See Example 8-16 on page 8-13 .
show system internal ethpm mem-stats	Displays the ethpm memory allocation statistics. See Example 8-17 on page 8-13 .
show system internal ethpm msgs	Displays the ethpm message logs. See Example 8-18 on page 8-13 .
show system internal vim errors	Displays VIM error logs. See Example 8-19 on page 8-14 .
show system internal vim event-history	Displays various VIM event logs. See Example 8-20 on page 8-14 .
show system internal vim info	Displays internal data structure information. See Example 8-21 on page 8-15 .

Command	Purpose
show system internal vim mem-stats	Displays memory allocation statistics of ethpm. See Example 8-22 on page 8-15 .
show system internal vim msgs	Displays various message logs of ethpm. See Example 8-23 on page 8-15 .
module vem execute vemcmd show portsec status	Displays the port security status of the port. If enabled, the output shows a LTL connected to the VM network adapter. See Example 8-24 on page 8-16 .
show port-security	Displays information about the secured MAC addresses in the system. See Example 8-25 on page 8-16 .
show port-security address vethernet	Displays information about the secured addresses on an interface. See Example 8-26 on page 8-16 .
show system internal port-security msgs	Displays various message logs of eth_port_sec. See Example 8-27 on page 8-17 .
show system internal port-security errors	Displays error logs of eth_port_sec. See Example 8-28 on page 8-17 .
show system internal pktmgr interface brief	Displays a summary of the pktmgr interface status and configuration. See Example 8-29 on page 8-17 .
show system internal pktmgr client detail	Displays detailed filter information. See Example 8-30 on page 8-18 .

For detailed information about the **show** command output, see the *Cisco Nexus 1000V for Microsoft Hyper-V Command Reference Guide*.

Example 8-1 show module Command

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

Mod  Sw  Hw
---  ---  ---
3    5.2(1)SM1(5.1)  Windows Server 8 - Datacenter (6.2.9200, 6.30)

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 Command

```
n1000v# show svcs domain
SVS domain config:
  Domain id: 942
```

```
Control vlan: 1
Packet vlan: 1
Control mode: L3
L3 control interface: mgmt0
Status: Config push to Management Server successful.
n1000v#
```

Example 8-3 *show cdp neighbors Command*

```
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
Nexus1000V(1)
                        mgmt0                145      B T B T S Nexus1000V  mgmt0
n1000V#
```

Example 8-4 *show system internal ethpm event-history interface Command*

```
n1000v# show system internal ethpm 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 8-5 *show logging logfile Command*

```
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 8-6 *show logging logfile | grep interface_number Command*

```
n1000v# show logging logfile | grep Vethernet3626
2011 Mar 25 10:56:03 nlk-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 nlk-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 nlk-bl %ETHPORT-2-IF_DOWN_ERROR_DISABLED: Interface
Vethernet3626 is down (Error disabled. Reason:Client data inconsistency)
```

Example 8-7 *show interface brief Command*

```
n1000v# show interface brief
-----
Port VRF Status IP Address Speed MTU
-----
mgmt0 -- up 172.23.232.141 1000 1500
-----

Ethernet      VLAN   Type Mode   Status Reason          Speed   Port
Interface                                           Ch #
-----
Eth3/1        1      eth  pvlan  up      none            10G     2
Eth3/2        1      eth  pvlan  up      none            10G     2
-----

Port-channel  VLAN   Type Mode   Status Reason          Speed   Protocol
Interface                                          
-----
Po1           1      eth  trunk  up      none            a-1000(D) none
Po2           1      eth  pvlan  up      none            a-10G(D)  none
-----

Vethernet     VLAN   Type Mode   Status Reason          Speed
-----
Veth1         262    virt access up      none            auto
Veth2         262    virt access up      none            auto
-----

Port      VRF          Status IP Address          Speed   MTU
-----
control0  --          up      --                  --      1500

n1000v#
```

Example 8-8 *show interface ethernet Command*

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

Example 8-9 *show running-config interface ethernet mod/port expand-port-profile Command*

```
n1000v# show running-config interface ethernet 3/2 expand-port-profile

!Command: show running-config interface Ethernet3/2 expand-port-profile
!Time: Thu Feb 14 17:33:21 2013

version 5.2(1)SM1(5.1)

interface Ethernet3/2
  switchport mode private-vlan trunk promiscuous
  switchport private-vlan trunk allowed vlan 214,224,234,244,254,260,284
  switchport private-vlan trunk allowed vlan add 294,298
  switchport private-vlan mapping trunk 264 10,20,30,40,50
  channel-group auto mode on mac-pinning
  no shutdown

n1000v#
```

Example 8-10 *show interface ethernet counters Command*

```
n1000v# show interface eth3/3 counters
```


Port	InOctets	InUcastPkts
Eth3/3	167944438	154350
Port	InMcastPkts	InBcastPkts
Eth3/3	68452	298184
Port	OutOctets	OutUcastPkts
Eth3/3	1789120	8738
Port	OutMcastPkts	OutBcastPkts
Eth3/3	1461	3172

Example 8-11 *show interface vethernet Command*

```

n1000v# show interface veth1
Vethernet1 is up
  Port description is LINUX-RHEL-01
  Hardware: Virtual, address: 001d.d8b7.1f81 (bia 001d.d8b7.1f81)
  Owner is VM "LINUX-RHEL-01"
  Active on module 3
  DVS port 4903633b-2994-4cc1-859d-f18030927ac4--ff1bb4d4-9edb-4784-b3cc-239908355ea5
  Port-Profile is
dynpp_03ac7d00-933d-4fc6-89c6-83bdaadf4248_96d80bc3-aa5b-43b8-a784-ca023f59759a
  Port mode is access
  5 minute input rate 1975576 bits/second, 3723 packets/second
  5 minute output rate 89381728 bits/second, 7378 packets/second
Rx
  23088599 Input Packets 23088568 Unicast Packets
  4 Multicast Packets 27 Broadcast Packets
  1530981234 Bytes
Tx
  45745626 Output Packets 45744734 Unicast Packets
  382 Multicast Packets 511 Broadcast Packets 893 Flood Packets
  69252149146 Bytes
  0 Input Packet Drops 0 Output Packet Drops
n1000v#

```

Example 8-12 *show interface capabilities Command*

```

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:                yes
TDR capable:         no
FabricPath capable:  no
Port mode:           Unknown

port-channel1
Model:               --
Type (Non SFP):      --
Speed:               10,100,auto
Duplex:              half/full/auto
Trunk encap. type:   802.1Q
Channel:             no
Broadcast suppression: no
Flowcontrol:         rx- (none) ,tx- (none)
Rate mode:           none
QOS scheduling:       rx- (none) ,tx- (none)
CoS rewrite:         no
ToS rewrite:         no
SPAN:               no
UDLD:               no
Link Debounce:       no
Link Debounce Time:  no
MDIX:               no
TDR capable:         no
FabricPath capable:  no
Port mode:           Unknown

```

```
n1000v#
```

Example 8-13 show interface virtual attach binding Command

```
n1000v# show interface virtual attach binding
```

```

-----
Port          Bind-Type Hypervisor-Port
-----
Veth1         static      21e2dfd4-660e-4aa1-9813-bb02db4a5b6a--7c3f9397-c2c0-42bb-
a19d-1f3737e06b8f
Veth2         static      21e2dfd4-660e-4aa1-9813-bb02db4a5b6a--f4444d8b-5e95-41de-
a750-91fe500b221a
Veth3         static      21e2dfd4-660e-4aa1-9813-bb02db4a5b6a--284227ec-e395-4203-
aa10-67d40271c184
Veth4         static      21e2dfd4-660e-4aa1-9813-bb02db4a5b6a--0ee0a71e-5b55-4fdf-
a03f-7860ed6d1011

```

```
n1000v#
```

Example 8-14 show system internal ethpm errors Command

```
n1000V# show system internal ethpm errors
```

```

1) Event:E_DEBUG, length:59, at 620000 usecs after Mon May 27 16:56:27 2013
   [102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0

```

```

2) Event:E_DEBUG, length:59, at 620000 usecs after Mon May 27 16:56:27 2013

```

```
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0
```

```
3) Event:E_DEBUG, length:59, at 160000 usecs after Mon May 27 16:56:26 2013
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0
```

Example 8-15 *show system internal ethpm event-history errors Command*

```
n1000V# show system internal ethpm event-history errors
1) Event:E_DEBUG, length:59, at 900000 usecs after Mon May 27 16:56:25 2013
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0

2) Event:E_DEBUG, length:59, at 900000 usecs after Mon May 27 16:56:25 2013
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0

3) Event:E_DEBUG, length:59, at 830000 usecs after Mon May 27 16:56:25 2013
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0

4) Event:E_DEBUG, length:59, at 830000 usecs after Mon May 27 16:56:25 2013
[102] ethpm_shared_port_down_notif(616): seqno = 1 const= 0
```

Example 8-16 *show system internal ethpm mem-stats Command*

```
n1000V# show system internal ethpm mem-stats
ETHPM Log Buffer info:
[Mon May 27 16:57:58 2013] PORT_FSM_ACTION_INIT      fsm->prev_state:22, eve
nt_id: 65, if_index:0x250080c0 (Ethernet3/4), oper_port_state:0x1, layer:0x2
[Mon May 27 16:57:58 2013] PORT_FSM_ACTION_INIT      fsm->prev_state:22, eve
nt_id: 65, if_index:0x25008140 (Ethernet3/6), oper_port_state:0x1, layer:0x2
[Mon May 27 16:57:58 2013] PORT_FSM_ACTION_INIT      fsm->prev_state:22, eve
nt_id: 65, if_index:0x25008180 (Ethernet3/7), oper_port_state:0x1, layer:0x2
[Mon May 27 16:57:58 2013] PORT_FSM_ACTION_INIT      fsm->prev_state:22, eve
nt_id: 65, if_index:0x250081c0 (Ethernet3/8), oper_port_state:0x1, layer:0x2
```

Example 8-17 *show system internal ethpm mem-stats Command*

```
n1000V# show system internal ethpm mem-stats
Private Mem stats for UUID : Malloc track Library(103) Max types: 5
-----
-
Curr alloc: 1587 Curr alloc bytes: 108176(105k)

Private Mem stats for UUID : Non mtrack users(0) Max types: 150
-----
-
Curr alloc: 1150 Curr alloc bytes: 94275(92k)
```

Example 8-18 *show system internal ethpm msgs Command*

```
n1000V# show system internal ethpm msgs
1) Event:E_MTS_RX, length:60, at 770000 usecs after Mon May 27 16:57:58 2013
```

```

[NOT] Opc:MTS_OPC_IM_L2_BUNDLED_PHY_PORT_STATE_CHANGE(62485), Id:0X0001646D
, Ret:SUCCESS
Src:0x00000101/175, Dst:0x00000101/0, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:1269
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 770000 usecs after Mon May 27 16:57:58 2013
[NOT] Opc:MTS_OPC_IM_L2_BUNDLED_PHY_PORT_STATE_CHANGE(62485), Id:0X00016468
, Ret:SUCCESS
Src:0x00000101/175, Dst:0x00000101/0, Flags:None
HA_SEQNO:0X00000000, RRtoken:0x00000000, Sync:UNKNOWN, Payloadsize:1269
Payload:
0x0000: 00 00 00 02 00 00 00 02 00 00 00 0c 00 00 00 29

```

Example 8-19 *show system internal vim errors Command*

```

n1000V# show system internal vim errors
1) Event:E_DEBUG, length:82, at 940000 usecs after Mon May 27 16:40:10 2013
[102] vim_mod_handle_att_ack(447): Received attach ack status 0x418f000c on
Eth3/3

2) Event:E_DEBUG, length:139, at 990000 usecs after Mon May 27 16:10:36 2013
[102] vim_vem_handle_veth_attach_pending_info(731): Dropping attach req for
Veth2<->lveth4/1: removal of active attach on Veth2 in progress

3) Event:E_DEBUG, length:128, at 410000 usecs after Mon May 27 16:10:36 2013
[102] vim_vem_bq_profile_bind_resp(235): Profile bind error on Veth2: statu
s=no port-profile found matching request (0x41b00014)

4) Event:E_DEBUG, length:117, at 230000 usecs after Mon May 27 16:10:36 2013
[102] vim_vem_handle_vem_lic_state_chg_notif(1285): Received License State
Change Notification for non-existing VEM 4

```

Example 8-20 *show system internal vim event-history all Command*

```

n1000V# show system internal vim event-history all
>>>>FSM: <Vethernet1> has 3 logged transitions<<<<

1) FSM:<Vethernet1> Transition at 410000 usecs after Mon May 27 15:46:33 2013
Previous state: [VIM_VETH_FSM_ST_NOT_EXISTENT]
Triggered event: [VIM_VETH_FSM_EV_CREATE]
Next state: [VIM_VETH_FSM_ST_WAIT_INIT]

2) Event:ESQ_START length:38, at 410000 usecs after Mon May 27 15:46:33 2013
Instance:-1073741824, Seq Id:0x1, Ret:SUCCESS
Seq Type:SERIAL

3) Event:ESQ_REQ length:38, at 410000 usecs after Mon May 27 15:46:33 2013
Instance:-1073741824, Seq Id:0x1, Ret:SUCCESS
[E_MTS_TX] Dst:MTS_SAP_IFMGR(179), Opc:MTS_OPC_IM_IF_IOD_ASSIGN_RELEASE(624
66)

```

Example 8-21 *show system internal vim info Command*

```
n1000V# show system internal vim info
is_vmfex_enabled: false
auto_setup: true
auto_delete: true
issu_in_progress: false
auto_config_purge: false
module 3:
  ports: ETH 32, LVETH 256
  node_addr: 0x00000302
  fsm_state: VIM_MOD_FSM_ST_INSERTED
  srv_license_state: licensed
  num_atts_in_progress: 0
  flags: 0x00000040
  lveth3/1:
    if_index: 0x1b020000
    attached: Veth1
    flags: 0x00000048
    attach_cookie: 0x00000002
  Eth3/1:
    if_index: 0x25008000
    pp_alias: DATA-Macpin (11)
    ds_id: none
    ds_port_uuid: (5)
    mac: 00:25:b5:aa:ab:4f
    flags: 0x00000094
    attach_cookie: 0x00000011
    speed_cap: 0x00000010 0x00000004
    pp_guid: 33701631-C4FD-47E0-9C30-5AF1CAC4ACE8
```

Example 8-22 *show system internal vim mem-stats Command*

```
n1000V# show system internal vim mem-stats
Private Mem stats for UUID : Malloc track Library(103) Max types: 5
-----
-
Curr alloc: 1727 Curr alloc bytes: 114526(111k)

Private Mem stats for UUID : Non mtrack users(0) Max types: 164
-----
-
Curr alloc: 474 Curr alloc bytes: 51576(50k)

Private Mem stats for UUID : libsdwrap(115) Max types: 22
-----
-
Curr alloc: 28 Curr alloc bytes: 715264(698k)

Private Mem stats for UUID : Associative_db library(175) Max types: 14
-----
-
Curr alloc: 210 Curr alloc bytes: 5912(5k)
```

Example 8-23 *show system internal vim msgs Command*

```
n1000V# show system internal vim msgs
1) Event:E_MTS_RX, length:60, at 190000 usecs after Wed May 29 14:24:20 2013
```

```

[REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X00106411, Ret:SUCCESS
Src:0x00000101/13924, Dst:0x00000101/403, Flags:None
HA_SEQNO:0X00000000, RRToken:0x00106411, Sync:UNKNOWN, Payloadsize:216
Payload:
0x0000:  01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 32 32

2) Event:E_MTS_RX, length:60, at 770000 usecs after Wed May 29 14:22:13 2013
[REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X00105EC5, Ret:SUCCESS
Src:0x00000101/13909, Dst:0x00000101/403, Flags:None
HA_SEQNO:0X00000000, RRToken:0x00105EC5, Sync:UNKNOWN, Payloadsize:208
Payload:
0x0000:  01 00 2f 74 6d 70 2f 64 62 67 64 75 6d 70 32 32

3) Event:E_MTS_RX, length:60, at 30000 usecs after Wed May 29 14:20:19 2013
[REQ] Opc:MTS_OPC_VSH_CMD_TLV(7679), Id:0X00104DF2, Ret:SUCCESS
Src:0x00000101/13798, Dst:0x00000101/403, Flags:None
HA_SEQNO:0X00000000, RRToken:0x00104DF2, Sync:UNKNOWN, Payloadsize:244
Payload:
0x0000:  04 03 02 01 f4 00 00 00 00 00 00 00 00 00 00 00

```

Example 8-24 module vem execute vemcmd show portsec status Command

```

n1000V# module vem 3 execute vemcmd show portsec status

LTL   if_index  Max      Aging   Aging   DSM   Sticky   VM
      Secure    Secure   Time    Type    Bit    Enabled  Name
      Addresses
    49  1b020000      1       0  Absolute  Clr     Yes     î
    50  1b020010      1       0  Absolute  Clr     Yes     î

n1000V#

```

Example 8-25 show port security Command

```

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 8-26 show port security address interface vethernet Command

```

n1000v# show port-security address interface vethernet 1
Secure Mac Address Table
-----
Vlan    Mac Address          Type          Ports          Configured Age
      (mins)
-----
262     001D.D8B7.1F81      STICKY        Vethernet1      0
-----

```

Example 8-27 *show system internal port-security msgs Command*

```
n1000v# show system internal port-security msgs
1) Event:E_MTS_RX, length:60, at 120000 usecs after Wed May 29 14:39:36 2013
   [REQ] Opc:MTS_OPC_SDWRAP_DEBUG_DUMP(1530), Id:0X00109749, Ret:SUCCESS
   Src:0x00000101/14132, Dst:0x00000101/191, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00109749, 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_RX, length:60, at 430000 usecs after Wed May 29 14:38:53 2013
   [REQ] Opc:MTS_OPC_VSH_CMD_TLV(7679), Id:0X0010936C, Ret:SUCCESS
   Src:0x00000101/13798, Dst:0x00000101/191, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x0010936C, Sync:UNKNOWN, Payloadsize:288
   Payload:
   0x0000: 04 03 02 01 20 01 00 00 00 00 00 00 00 00 00 00

3) Event:E_MTS_RX, length:60, at 120000 usecs after Wed May 29 14:38:03 2013
   [REQ] Opc:MTS_OPC_VSH_CMD_TLV(7679), Id:0X00108FA0, Ret:SUCCESS
   Src:0x00000101/13798, Dst:0x00000101/191, Flags:None
   HA_SEQNO:0X00000000, RRtoken:0x00108FA0, Sync:UNKNOWN, Payloadsize:232
   Payload:
   0x0000: 04 03 02 01 e8 00 00 00 00 00 00 00 00 00 00 00
```

Example 8-28 *show system internal port-security errors Command*

```
n1000v# show system internal port-security errors
1) Event:E_DEBUG, length:62, at 550000 usecs after Wed May 29 14:37:06 2013
   [102] psec_get_interface_info(8666): Port security not enabled

2) Event:E_DEBUG, length:63, at 550000 usecs after Wed May 29 14:37:06 2013
   [102] eth_port_sec_interface_port_status(409):admin_status is 1

3) Event:E_DEBUG, length:62, at 230000 usecs after Wed May 29 14:37:03 2013
   [102] psec_get_interface_info(8666): Port security not enabled

4) Event:E_DEBUG, length:63, at 230000 usecs after Wed May 29 14:37:03 2013
   [102] eth_port_sec_interface_port_status(409):admin_status is 1
```

Example 8-29 *show system internal pktmgr interface brief Command*

```
n1000v# show system internal pktmgr interface brief
Interface      Type      Interface Status
mgmt0          protocol-up/link-up/admin-up
control0       protocol-up/link-up/admin-up
sup-eth1       protocol-up/link-up/admin-up
sup-eth2       protocol-up/link-up/admin-up
sup-eth3       protocol-up/link-up/admin-up
port-channel1  protocol-up/link-up/admin-up
port-channel2  protocol-up/link-up/admin-up
```

Example 8-30 *show system internal pktmgr client detail Command*

```
n1000v# show system internal pktmgr client detail
Client uuid: 268, 3 filters, pid 2422
  Filter 1: EthType 0x0806,
    Rx: 62537, Drop: 0
  Filter 2: EthType 0xffff0, Exc 8,
    Rx: 0, Drop: 0
  Filter 3: EthType 0x8841, Snap 34881,
    Rx: 0, Drop: 0
Options: TO 0, Flags 0x18040, AppId 0, Epid 0
Ctrl SAP: 278, Data SAP 337 (1)
Total Rx: 125074, Drop: 0, Tx: 2906, Drop: 0
Recirc Rx: 0, Drop: 0
Rx pps Inst/Max: 0/60
Tx pps Inst/Max: 0/1
COS=0 Rx: 0, Tx: 0    COS=1 Rx: 0, Tx: 0
COS=2 Rx: 0, Tx: 0    COS=3 Rx: 0, Tx: 0
COS=4 Rx: 0, Tx: 0    COS=5 Rx: 0, Tx: 0
COS=6 Rx: 0, Tx: 2906  COS=7 Rx: 62537, Tx: 0
```




Port Profiles

This chapter describes how to identify and resolve problems with port profiles.

Information About Port Profiles

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

In the Microsoft Hyper-V Server, a port profile is represented as a port group. The vEthernet or Ethernet interfaces are assigned in the Microsoft SCVMM server to a port profile to do the following:

- Define the port configuration by policy.
- Apply a single policy across a large number of ports.
- Support both vEthernet and Ethernet ports.

Port profiles can be assigned by the server administrator to physical ports (a VMNIC or a PNIC). Port profiles that are configured as vEthernet can be assigned only to a vNIC port while port profiles that are configured as Ethernet can be assigned only to physical adapters.



Note

While a manual interface configuration overrides that of the port profile, it is not recommended. A manual interface configuration is used only, 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 *Cisco Nexus 1000V for Microsoft Hyper-V Port Profile Configuration Guide*.

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

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

To verify port profile inheritance, enter this 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 process can be done from the Microsoft SCVMM server.

**Note**

Inherited port profiles are automatically configured by the Cisco Nexus 1000V when the ports are attached on the hosts. This process occurs when the Microsoft Hyper-V port GUID that is assigned by the system administrator is matched with the port profile that created it.

For detailed information about port profiles, see the *Cisco Nexus 1000V for Microsoft Hyper-V Port Profile Configuration Guide*.

Problems with Port Profiles

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

Symptom	Possible Causes	Solution
You do not see the port profile/uplink network/network segment on the Microsoft SCVMM server.	The connection to the Microsoft SCVMM server is down.	<ol style="list-style-type: none"> 1. Connect to <code>http://VSM-IP</code> to verify that the VSM is reachable from the Microsoft SCVMM server. 2. Launch the Microsoft SCVMM user interface (UI). 3. Choose Fabric Management > Networking > Switch Extension Manager. 4. Refresh the Cisco Nexus 1000V.
	The Microsoft SCVMM server has not pulled the new configuration from the VSM.	<ol style="list-style-type: none"> 1. Launch the Microsoft SCVMM user interface (UI). 2. Choose Fabric Management > Networking > Switch Extension Manager. 3. Refresh the Cisco Nexus 1000V to force Microsoft SCVMM to pull new configuration from the VSM. 4. Fix any problems with the domain configuration. <p>For information about configuring the domain, see the <i>Cisco Nexus 1000V for Microsoft Hyper-V System Management Configuration Guide</i>.</p>
	The port profile is configured incorrectly.	<ol style="list-style-type: none"> 1. To verify that publish port profile is configured for the port profile/Network segment/uplink network, enter this command: show network segment name <i>name</i> 2. Fix the port profile using the procedures in the <i>Cisco Nexus 1000V for Microsoft Hyper-V Port Profile Configuration Guide</i>

Symptom	Possible Causes	Solution
<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>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 by entering the show port-profile virtual usage command. 2. Reassign the VMNIC or PNIC to a nonquarantined port group to enable the interface to be up and forwarding traffic. You must change the port group on the Microsoft SCVMM 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 by entering the show accounting log grep FAILURE command. 2. Verify the port profile-to-interface mapping by entering the show port-profile virtual usage command. 3. Fix the error in the port profile using the procedures in the <i>Cisco Nexus 1000V for Microsoft Hyper-V Port Profile Configuration Guide</i>. 4. Bring the interface out of quarantine by entering the no shutdown command. The interface comes back online. 5. Return shutdown control to the port profile by entering the default shutdown command.
<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, see the “Recovering a Quarantined Offline Interface” section on page 9-6.</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 that causes 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 Microsoft SCVMM Server. The module was powered down. There is general loss of connectivity to the module. 	<p>See the VEM troubleshooting guide to bring the module back online.</p> <p>To bring the interface back online, see the “Recovering a Quarantined Offline Interface” section on page 9-6.</p>

Symptom	Possible Causes	Solution
The interface is in the NoPortProfile state.	The port profile or uplink networks have been deleted from the VSM but are still on the VMM. If the port profiles are used to attach Ethernet and vEthernet interfaces, the interface will go into the NoPortProfile state.	<ol style="list-style-type: none"> Check if the GUID of the Ethernet port profile and GUID of the Microsoft SCVMM server uplink port profile set match as follows: <ol style="list-style-type: none"> Display the GUID of the Ethernet port profile by entering the show running-config port-profile uplinkname command on the VSM. Display the GUID of the Microsoft SCVMM server uplink port profile set by entering the PS C:\Program files(x86)\Cisco\Nexus1000V\Support > Get-SCExtensionUplinkPortProfile -name uplinkname command in a PowerShell window. Check if the GUID of the vEthernet port profile matches the GUID of the Microsoft SCVMM server port profile under port classification as follows: <ol style="list-style-type: none"> Display the GUID of the vEthernet port profile by entering the show running-config port-profile vEthernetName command on the VSM. Display the GUID of the Microsoft SCVMM server port profile by entering the PS C:\Program files(x86)\Cisco\Nexus1000V\Support > Get-SCVirtualNetworkAdapterExtensionPortProfile -name vEthernetName command in a PowerShell window. Detach the vEthernet interface from the Microsoft SCVMM server Delete the vEthernet interface manually from the VSM. Refresh the Switch Extension Manager in the Microsoft SCVMM server. Create a new uplink port profile set or a virtual port classification. Attach the original combination of Ethernet and vEthernet interfaces.

Symptom	Possible Causes	Solution
	In Microsoft SCVMM 2016 you can attach a vEthernet port to a Cisco Nexus 1000V, by using two mandatory parameters VMNetwork and PortClassification. If the PortClassification is not applied, the vEthernet port will go into NoPortProfile state.	<ol style="list-style-type: none"> 1. In the Microsoft SCVMM UI, use the VM properties to detach the vEthernet/vNIC from the Cisco Nexus 1000V logical switch. 2. Attach the vEthernet/vNIC to the Cisco Nexus 1000V logical switch with both of the parameters.
The interface is in NoPortProfile and an update failure has occurred at the Microsoft SCVMM UI.	The vEthernet port profile is a system and the network segment is a nonsystem. This scenario is not supported.	<ol style="list-style-type: none"> 1. In the Microsoft SCVMM UI, repair the VM and click ignore. 2. Detach the vEthernet interface from the Cisco Nexus 1000V host in the Microsoft SCVMM UI by moving it to a disconnected state. 3. Attach the vEthernet interface back to the Cisco Nexus 1000V by using one of the following supported combinations: <ul style="list-style-type: none"> – A vEthernet interface that is attached with a system vEthernet port profile and a system network segment. – A vEthernet interface that is attached with a nonsystem vEthernet port profile and a nonsystem network segment – A vEthernet interface that is attached with a nonsystem vEthernet port profile and a system network segment.
You are unable to add a PVLAN network segment to an uplink port that contains system VLANs.	If uplink networks are marked as system, the VSM does not allow PVLAN network segment pools.	<ol style="list-style-type: none"> 1. Remove the system VLAN configuration from the uplink network. 2. Attach the PVLAN network segment pool to the uplink network. 3. Convert the uplink network back to the system.

Symptom	Possible Causes	Solution
You are unable to remove the system keyword from the network segment.	Sometimes, removing a system keyword is not allowed.	<ol style="list-style-type: none"> 1. Check if the segment is used by any dynamic profile by entering the show dynamic-port-profile network segment name command. 2. If it is used, delete the dynamic profile. 3. If Step 2 fails, do the following: <ol style="list-style-type: none"> a. Check if any vEthernet interfaces are using the profile by entering show port-profile name name command. b. If vEthernet interfaces are present, detach the vEthernet interfaces from the Cisco Nexus 1000V by putting the interfaces in the NotConnected state. c. Detach the vEthernet interfaces to move the vEthernet interfaces to a Nonparticipating state so that they can be deleted without any impact. d. After all the vEthernet interfaces are associated with the dynamic port profile are deleted, repeat Step 2. 4. Remove the system keyword from the profile.
The network segment is not available for attaching to a vEthernet interface.	A new network segment is associated with an existing network segment pool. However, when you try to attach a vEthernet interface to this new segment, it might not appear the same on the Microsoft SCVMM UI.	<ol style="list-style-type: none"> 1. Check if the uplink profiles allow the network segment pool that contains the new network segment on the VSM and the Microsoft SCVMM UI. 2. If the segment is missing on the VSM, allow the network segment pool under the uplink network. Refresh the Switch Extension Manager to retrieve the new information. 3. If the segment is missing on the Microsoft SCVMM UI, perform a refresh of the Switch Extension Manager to retrieve the new information.
VM network not available on Cisco Nexus 1000V.	The VM network is created on the Microsoft SCVMM but when using this VM network to assign to a VM it fails.	<ol style="list-style-type: none"> 1. Identify the uplink network of the host where the VM is residing. 2. Verify that the uplink network carries the VM network. 3. Use the appropriate uplink network.

Recovering a Quarantined Offline Interface

You can recover and bring online an interface that is offline and has been quarantined.

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.

DETAILED STEPS

-
- | | |
|---------------|--|
| Step 1 | Verify that the interface is quarantined. The interface appears in the show command output. |
| Step 2 | On the Microsoft SCVMM 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 by entering the show interface brief command. |
| Step 4 | Verify the port profile-to-interface mapping by entering the show port-profile virtual usage command. |
| Step 5 | Verify the interface has come out of quarantine automatically. The interface should no longer appear in the show command output. |
| Step 6 | Return shutdown control to the port profile by entering the default shutdown command. |
-

Verifying the Maximum Number of Ports

The validation for the maximum number of ports is carried out by the Microsoft SCVMM only during the following two scenarios:

- When a new VM is deployed from the Microsoft SCVMM.
- When an existing VM is moved or migrated live.

Maximum ports are classified in two categories

- Per port profile with a default value of 32 and a maximum of 1024.
- Per host or VEM with a value of 216 on the Hyper-V.

Verifying Maximum Number of Ports on a Logical Switch

-
- | | |
|---------------|--|
| Step 1 | Launch the Microsoft SCVMM UI to view the maximum ports on a logical switch. |
| Step 2 | Choose Fabric > Switch Extension Manager > Properties > Extensions . |
-

Verifying the Maximum Number of Ports on a Virtual Port Profile

-
- | | |
|---------------|---|
| Step 1 | On the VSM, verify the maximum number of ports on the virtual port profile by entering the show port-profile name <i>virtual_port_profile_name</i> command.

n1000v(config)# show port-profile name veth-policy
port-profile veth-policy |
|---------------|---|

```

type: Vethernet
description:
status: enabled
max-ports: 32
min-ports: 1
inherit:
config attributes:
no shutdown
evaluated config attributes:
no shutdown
assigned interfaces:
port-group: new
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vn-service: no
port-profile role: none
port-binding: static

```

Verifying the Maximum Number of Ports of a vEthernet Interface on the Microsoft SCVMM

-
- Step 1** Launch a PowerShell window to view the maximum ports of the vEthernet port profile on the Microsoft SCVMM.
- Step 2** In the PowerShell window, enter this command.:
- ```
PS C:\> Get-SCVirtualNetworkAdapterExtensionPortProfile -Name vEthernet_port_profile
```
- 

## Port Profile Logs

To enable and collect detailed logs for port profiles, enter these commands:

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

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                                                                                                                                                                                                                                                               |
|---------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Get-SCPortClassification</b>                                     | Displays the port profile and port classification information.                                                                                                                                                                                                        |
| <b>Get-SCVirtualNetworkAdapterExtensionPort Profile</b>             | See <a href="#">Example 9-1 on page 9-9</a> .                                                                                                                                                                                                                         |
| <b>show port-profile</b>                                            | Displays the port profile configuration.<br>See <a href="#">Example 9-2 on page 9-10</a> .                                                                                                                                                                            |
| <b>show port-profile name</b> <i>name</i>                           | Displays the configuration for a named port profile.<br>See <a href="#">Example 9-3 on page 9-10</a> .                                                                                                                                                                |
| <b>show port-profile brief</b>                                      | Displays a tabular view of all configured port profiles.<br>See <a href="#">Example 9-4 on page 9-11</a> .                                                                                                                                                            |
| <b>show port-profile expand-interface name</b> <i>name</i>          | Displays a named port profile expanded to include the interfaces assigned to it.<br>See <a href="#">Example 9-5 on page 9-11</a> .                                                                                                                                    |
| <b>show running-config port-profile</b> [ <i>profile-name</i> ]     | Displays the port profile configuration.<br>See <a href="#">Example 9-6 on page 9-11</a> .                                                                                                                                                                            |
| <b>show port-profile virtual usage</b> [ <i>name profile-name</i> ] | Displays the port profile usage by interface.<br>See <a href="#">Example 9-7 on page 9-12</a> .                                                                                                                                                                       |
| <b>show msp internal info</b>                                       | Displays port profile mappings on the Microsoft SCVMM server and configured roles.<br>See <a href="#">Example 9-8 on page 9-13</a> .                                                                                                                                  |
| <b>show system internal port-profile profile-fsm</b>                | Displays port profile activity on the Cisco Nexus 1000V, including transitions such as inherits and configurations. If the following appears, all inherits are processed:<br><br>Curr state: [PPM_PROFILE_ST_SIDLE]<br>See <a href="#">Example 9-9 on page 9-14</a> . |
| <b>show system internal port-profile event-history msgs</b>         | Displays the messages logged about port profile events within the Cisco Nexus 1000V.<br>See <a href="#">Example 9-10 on page 9-14</a> .                                                                                                                               |

For detailed information about **show** command output, see the *Cisco Nexus 1000V for Microsoft Hyper-V Command Reference Guide*.

## EXAMPLES

### Example 9-1 Get-SCPortClassification Command

```
PS C:\Users\Administrator.HyperV> Get-SCPortClassification

Name : NexusNoRestrict-2
```

```

Description :
ServerConnection : Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection
ID : 9f8819c1-8b53-42bd-a6fd-0173804e3194
IsViewOnly : False
ObjectType : PortClassification
MarkedForDeletion : False
IsFullyCached : True

```

```
PS C:\Users\Administrator\HYPERV> Get-SCVirtualNetworkAdapterExtensionPortProfile
```

```

Name : NoRest-unicast-norest
ExternalId : 308ad66b-7c42-4067-90af-13f7a6e59afe
NetworkEntityType : ExternallyManaged
VirtualSwitchExtension : n1kv-test
Tags : {}
AllowedVNicType : Both
MaxNumberOfPorts : 32
MaxNumberOfPortsPerHost : 216
ProfileData : 0
ServerConnection :
Microsoft.SystemCenter.VirtualMachineManager.Remoting.ServerConnection
ID : 8934a01c-0cb7-4ee2-ae9d-21ff5b26568f
IsViewOnly : False
ObjectType : VirtualSwitchExtensionVirtualPortProfile
MarkedForDeletion : False
IsFullyCached : True

```

#### **Example 9-2** *show port-profile name new Command*

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

```

port-profile new
type: Vethernet
description:
status: enabled
max-ports: 32
min-ports: 1
inherit:
config attributes:
no shutdown
evaluated config attributes:
no shutdown
assigned interfaces:
port-group: new
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vn-service: no
port-profile role: none
port-binding: static

```

#### **Example 9-3** *show port-profile name Command*

```

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 9-4 *show port-profile brief Command*

```

n1000v# show port-profile brief

Port Profile Profile Conf Eval Assigned Child
Profile Type State Items Items Intfs Profs

LACP Ethernet 1 2 2 0 2
LACP_PIN Ethernet 1 4 5 4 0
MAC Ethernet 1 2 2 0 1
MAC_PIN Ethernet 1 4 5 7 0
MAC_PIN_343 Ethernet 1 2 4 1 0
NSM_template_segmentation Vethernet 1 1 1 0 0
NSM_template_vlan Vethernet 1 1 1 0 0
basic Vethernet 1 1 1 0 0
default Vethernet 1 1 1 0 0
dynpp_a7ab47ce-07c3-4fc8-ae74-321a10818199_76604d2a-f62e-40a4-85d1-0ccad8d1c9c0
Vethernet 1 2 3 0 0
dynpp_a7ab47ce-07c3-4fc8-ae74-321a10818199_aa914386-bf85-48e6-98ca-541a764e7580
Vethernet 1 2 3 2 0
dynpp_a7ab47ce-07c3-4fc8-ae74-321a10818199_b4490e62-57c2-4c3d-81f9-99ca0b6a6a82
Vethernet 1 2 3 8 0
new Vethernet 1 1 1 0 3
system Vethernet 1 1 1 0 0
uplink_network_default_policy Ethernet 1 1 1 0 0

Profile Assigned Total Sys Parent Child UsedBy
Type Intfs Prfls Prfls Prfls Prfls Prfls

Vethernet 10 9 1 8 3 2
Ethernet 12 6 0 4 3 3
n1000v#

```

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

```

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 Command*

```

n1000v# show running-config port-profile
!Command: show running-config port-profile

```

```

!Time: Sun Mar 17 13:17:03 2013

version 5.2(1)SM1(5.1)
port-profile default max-ports 32
port-profile type vethernet NSM_template_vlan
no shutdown
guid 100b8834-85a7-4a9f-a942-83b8218b4fc1
description NSM default port-profile for VLAN networks. Do not delete.
state enabled
port-profile type vethernet NSM_template_segmentation
no shutdown
guid aee2046c-eb9d-4018-bae7-e1000f5b2d54
description NSM default port-profile for VXLAN networks. Do not delete.
state enabled
port-profile type ethernet MAC
channel-group auto mode on mac-pinning
no shutdown
guid 51217cb4-280d-4cbe-a73d-18299cc347c2
max-ports 512
state enabled
port-profile type ethernet LACP
channel-group auto mode active
no shutdown
guid 28a414ca-7c10-4c0d-a73e-a1af409bdb5f
max-ports 512
state enabled
port-profile type vethernet basic
no shutdown
guid bbf3ec9f-9ca3-445a-9376-630180c35250
publish port-profile basic-non-system
state enabled
port-profile type vethernet system
no shutdown
guid 2e21ff4a-e966-4432-95ae-6600e0cbe50f
publish port-profile basic-system
system port-profile
state enabled
port-profile type ethernet uplink_network_default_policy
no shutdown
guid 4cc1067c-7104-4aa1-8556-ce18ada165e8
max-ports 512
description NSM created profile. Do not delete.
state enabled
port-profile type vethernet default
no shutdown
guid 622e109d-6465-4abd-882f-d026938b830d
state enabled
port-profile type vethernet new
no shutdown
guid a7ab47ce-07c3-4fc8-ae74-321a10818199
publish port-profile
state enabled
n1000v#

```

**Example 9-7** *show port-profile virtual usage Command*

```

n1000v# show port-profile virtual usage

Port Profile Port Adapter Owner

MAC_PIN Po2
Po6
Eth3/4 vmnic3 WIN-35

```

```

Eth3/5 vmnic4 WIN-35
Eth3/6 vmnic5 WIN-35
Eth4/1 vmnic0 WIN-37
Eth4/3 vmnic2 WIN-37
LACP_PIN Po1
Po3
Eth5/1 vmnic0 WIN-39
Eth5/2 vmnic1 WIN-39
dynpp_a7ab47ce-07c3-4fc8-a
e74-321a10818199_b4490e62-
57c2-4c3d-81f9-99ca0b6a6a8
2 Veth1 Net Adapter Win2008-2-1
Veth2 Net Adapter Win2008-1-1
Veth3 Net Adapter Win2008-3-1
Veth4 Net Adapter Win2008-4-1
Veth5 Net Adapter Win2008-2-2
Veth6 Net Adapter Win2008-1-2
Veth7 Net Adapter Win2008-3-2
Veth8 Net Adapter Win2008-4-2
MAC_PIN_343 Po4
dynpp_a7ab47ce-07c3-4fc8-a
e74-321a10818199_aa914386-
bf85-48e6-98ca-541a764e758
0 Veth9 Net Adapter WIN-Legacy
Veth10 Net Adapter WIN-SPAN-3
n1000v#

```

#### **Example 9-8** *show msp internal info Command*

```

n1000v# show msp internal info
port-profile NSM_template_segmentation
 id: 2
 capability: 0x0
 state: 0x1
 type: 0x0
 system vlan mode: -
 system vlans:
 port-binding: static
 bind_opts: 0
 max ports: 32
 min ports: 1
 active used ports count: 0
 intf inherit count: 0
 Hyper-V config information
 pg name: NSM_template_segmentation
 dvs: (ignore)
 reserved ports: 32
 port-profile role:
 alias information:
 pg id: 8eebad90-fe9a-4460-b44e-9f71b8ebc88d
 dvs uuid:
 type: 11
port-profile NSM_template_vlan
 id: 1
 capability: 0x0
 state: 0x1
 type: 0x0
 system vlan mode: -
 system vlans:
 port-binding: static
 bind_opts: 0
 max ports: 32
 min ports: 1

```

```

active used ports count: 0
intf inherit count: 0
Hyper-V config information
 pg name: NSM_template_vlan
 dvs: (ignore)
 reserved ports: 32
port-profile role:
alias information:
 pg id: 83e41305-c443-4d30-a142-f1260183d974
 dvs uuid:
 type: 11
pending binds:
PPM restore_complete:TRUE
 opq_data_info.ppm_sdb_restored:1
NSMGR restore_complete:TRUE
 opq_data_info.nsm_sdb_restored:1

```

#### **Example 9-9** *show system internal port-profile profile-fsm Command*

```

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 9-10** *show system internal port-profile event-history msgs Command*

```

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

---

This chapter describes how to identify and resolve problems that relate to port channels and trunking.

## 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 failures.

## Port Channel Restriction

The following are port channel restrictions:

- Port channels do not support access control lists (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                                                                                                                                                                                                                                                                                              | ✓ |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|
| To determine port channel requirements, enter the <b>show port-channel compatibility-parameters</b> command.                                                                                                                                                                                           |   |
| Ensure that all interfaces in the port channel have the same destination device for Link Aggregation Control Protocol (LACP) channels. By using the Asymmetric Port Channel (APC) feature in the Cisco Nexus 1000V, ports in an ON mode channel can be connected to two different destination devices. |   |
|  <b>Note</b> APC is supported only in ON mode channels. It is not supported for LACP channels.                                                                                                                        |   |
| To verify that either side of a port channel is connected to the same number of interfaces, enter these commands: <ul style="list-style-type: none"> <li>• <b>show port-channel summary</b></li> <li>• <b>show ether channel summary</b></li> </ul>                                                    |   |
| To verify that each interface is connected to the same type of interface on the other side, enter the <b>show interface brief</b> command.                                                                                                                                                             |   |
| To verify that all required VLANs on a trunk port are in the allowed VLAN list, enter the <b>show interface switchport</b> command.                                                                                                                                                                    |   |
| To verify that all the members trying to form a port channel are on the same module, enter the <b>show port-channel summary</b> command.                                                                                                                                                               |   |
| To verify that the port channel configuration is present in the profile used by the physical ports, enter the <b>show port-channel name name</b> command and check for the <b>channel-group auto</b> setting.                                                                                          |   |
| 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.                                                                                                                                                                                                      |   |

The following commands can help you to 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/port-number**
- **show system internal ethpm event-history interface port-channel channel-number**
- **show system internal ethpm event-history interface ethernet slot-number/port-number**
- **show vlan internal trunk interface ethernet slot-number/port-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 port vlans**
- **module vem module-number execute vemcmd show pc**
- **module vem module-number execute vemcmd show trunk**

Example 10-1 shows the 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)
```

## Verifying a Port Channel Configuration

You can debug port channels that are configured through a port profile.

### BEFORE YOUR BEGIN

- Log in to the CLI in global configuration mode.

### DETAILED STEPS

- |        |                                                                                                                                   |
|--------|-----------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Verify that you have configured a port channel in the profile by entering the <b>show port-profile name profile-name</b> command. |
| Step 2 | Verify the port configuration by entering the <b>show port-channel summary</b> command.                                           |
| Step 3 | Configure debugging of port-channel trace by entering the <b>debug port-channel trace</b> command.                                |

## 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.
- Ports in APC only come up when they are assigned subgroup IDs.
- For MAC-pinning and MAC-pinning relative APCs, subgroup IDs are automatically assigned based on the vmnic numbers.
- For the Cisco Discovery Protocol (CDP) subgroup APCs, physical ports within an APC get assigned subgroup IDs based on CDP information received from upstream switches. Make sure that CDP is enabled on Virtual Supervisor Modules (VSM) and upstream switches.
- Verify CDP adjacency and subgroup mapping for upstream switches by entering the **show cdp neighbors** and **show port-channel cdp-map** commands on the VSM.

- For manual subgroup APCs, ensure subgroup IDs are manually configured on the physical ports in the interface configuration submenu.
- After the ports came up, check that ports are put in the correct subgroups by entering the **module vem module-number execute vemcmd show pc** command on the VSM.
- Configure debugging of port-channel trace by entering the **debug port-channel trace** command.

## Troubleshooting LACP Port Channels

When troubleshooting LACP port channels, follow these guidelines:

- All physical ports in the port channel should be connected to a single upstream switch.
- The LACP feature should be enabled on both the VSM and the upstream switch.
- The LACP channel group should be configured on all upstream ports and the channel-group ID that is assigned to the upstream ports of a single port channel should be identical.
- At least one end (the Cisco Nexus 1000V or upstream switch) of the port channel should have active LACP mode configured.
- After the ports come up, check that ports are in a LACP port channel by entering the **show lacp port-channel** and **module vem module-number execute vemcmd show pc** commands on the VSM.

## Cannot Create a Port Channel

| Symptom                       | Possible Cause                                                                                      | Solution                                                                                                                                                                                                          |
|-------------------------------|-----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cannot create a port channel. | A maximum number of port channels has been reached for the system or Virtual Ethernet Module (VEM). | Verify the number of port channels already configured by entering the <b>show port-channel summary</b> command. You can have a maximum of 256 port channels on the Cisco Nexus 1000V and 8 port channels per VEM. |

## Newly Added Interface Does Not Come Online in a Port Channel

| Symptom                                                         | Possible Cause                                                                        | Solution                                                                                                        |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| A newly added interface does not come online in a port channel. | The port channel has not been configured.                                             | Make sure that you have the port channel configuration in the port profile (port group) used by that interface. |
|                                                                 | The interface parameters are not compatible with the parameters of the existing port. | Configure compatible parameters on all physical ports in a port channel.                                        |

## VLAN Traffic Does Not Traverse Trunk

| Symptom                                       | Possible Cause                          | Solution                                                                                                                                                            |
|-----------------------------------------------|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The VLAN traffic does not traverse the trunk. | A VLAN is not in the allowed VLAN list. | Add the VLAN to the allowed VLAN list by entering the <b>switchport trunk allowed vlan add <i>vlan-id</i></b> command in the profile that is used by the interface. |





# Layer 2 Switching

---

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

## Information About Layer 2 Ethernet Switching

The Cisco Nexus 1000V 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.

# Viewing Ports from the VEM

The Cisco Nexus 1000V differentiates between virtual and physical ports on each of the VEMs. [Figure 11-1](#) shows how ports on the Cisco Nexus 1000V switch are bound to physical and virtual Microsoft Hyper-V ports within a VEM.

**Figure 11-1 VEM View of Ports**



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

- **Virtual NICs**—There are two types of Virtual NICs. The virtual NIC (`vnic`) is part of the VM and represents the physical port of the host that is plugged into the switch. Internal NICs are used by the hypervisor for internal purposes. Each type maps to a `vEth` port within the Cisco Nexus 1000V.
- **Virtual Ethernet Ports (VEth)**—A `vEth` port is a port on the Cisco Nexus 1000V distributed virtual switch. The 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 that is running the VM.
- **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.

`vEth` ports are assigned to port groups.

These local ports do not move and you can address them by the module-port number method.

Each physical NIC is represented by an interface called a `vmnic`. The `vmnic` number is allocated during Microsoft Hyper-V 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. The port acts like an local vEth 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 that is added to the Cisco Nexus 1000V switch appears as a physical Ethernet port, just as it would on a hardware-based switch.

The uplink port concept is handled entirely by the hypervisor and is used to associate port configuration with vmnics. There is no fixed relationship between the uplink number and vmnic number, and the uplink and the vmnic numbers 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 of the ports.

Figure 11-2 VSM View of Ports



## 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 hypervisor.
  - Internal NICs on the hypervisor.

- eths (physical Ethernet interfaces)—Correspond to the physical NICs on the hypervisor.
- Po (port channel interfaces)—The physical NICs of a hypervisor 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 for Microsoft Hyper-V Layer 2 Switching Configuration Guide*.

## Problems with Layer 2 Switching

This section describes how to troubleshoot Layer 2 problems and lists troubleshooting commands.

### Verifying a Connection Between VEM Ports

- 
- Step 1** View the state of the VLANs associated with the port by entering the **show vlan** command on the VSM. If the VLAN associated with a port is not active, the port might 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 the **show interface brief** command.
- Step 3** 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 by entering the **module vem module-number execute vemcmd show port** command.

The key things to look for in the output are as follows:

- State of the port.
  - CBL.
  - Mode.
  - Attached device name.
  - The LTL of the port that 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 that the state of the port is up. If not, verify the configuration of the port on the VSM.
- Step 4** View the VLANs and their port lists on a particular VEM by entering 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 that you see the port name or LTL in the port list of that VLAN.

---

### Verifying a Connection Between VEMs

- 
- Step 1** Check if the VLAN associated with the port is created on the VSM by entering the **show vlan** command.
- Step 2** Check if the ports are up in the VSM by entering the **show interface brief** command.
- Step 3** Check if the CBL state of the two ports is set to the value of 1 for forwarding (active) by entering the **module vem 3 execute vemcmd show port** command on the VEM.

- Step 4** Check if the two vEth ports are listed in the flood list of the VLAN to which they are trying to communicate by entering the **module vem 3 execute vemcmd show bd** command on the VEM.
- 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 the port on the upstream switch to which the physical NIC (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 Intrfce | 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 that 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 to 69 and 231 to 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

- Step 1** In the 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 VLANs are 3002 and 3003)
  - If the physical NIC in your configuration carries the VLAN for VM, that VLAN is also present in the allowed VLAN list.

```
n1000v# show port-profile name alluplink
port-profile alluplink
 type: Ethernet
 description:
 status: enabled
 max-ports: 512
 min-ports: 1
 inherit:
 config attributes:
 switchport mode trunk
```

```

switchport trunk allowed vlan 1,80,3002,610,620,630-650
evaluated config attributes:
switchport mode trunk
switchport trunk allowed vlan 1,80,3002,3003,610,620,630-650
no shutdown
assigned interfaces:
Ethernet2/2
port-group:
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vn-service: no
port-profile role: none
port-binding: static

```

**Step 2** Verify that the Ethernet interface is up by entering the **ifconfig -a** command inside the VM.

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, look for the association between the IP and MAC address by entering these commands:

- **debug arp**
- **show arp**

This example shows how to debug the Address Resolution Protocol (ARP):

```

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#

```

This example shows how to display ARP:

```

n1000v_CAT6K# show 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  |

# Layer 2 Switching Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to the Layer 2 MAC address configuration.

| Command                                                                                                                             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>show mac address-table</b>                                                                                                       | Displays the MAC address table to verify all MAC addresses on all VEMs controlled by the VSM.<br>See <a href="#">Example 11-1 on page 11-8</a> .                                                                                                                                                                                                                                                                             |
| <b>show mac address-table module</b><br><i>module-number</i>                                                                        | Displays all the MAC addresses on the specified VEM.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>show mac address-table static</b><br><i>HHHH.WWWW.HHHH</i>                                                                       | Displays the MAC address table static entries.<br>See <a href="#">Example 11-2 on page 11-8</a> .                                                                                                                                                                                                                                                                                                                            |
| <b>show mac address-table address</b><br><i>HHHH.WWWW.HHHH</i>                                                                      | Displays the interface on which the MAC address specified is learned or configured. <ul style="list-style-type: none"> <li>For dynamic MAC addresses, if the same MAC address appears on multiple interfaces, then each of them is displayed separately.</li> <li>For static MAC addresses, if the same MAC address appears on multiple interfaces, then only the entry on the configured interface is displayed.</li> </ul> |
| <b>show mac address-table static   inc veth</b>                                                                                     | Displays the static MAC address of vEthernet interfaces in case a VEM physical port learns a dynamic MAC address and the packet source is in another VEM on the same VSM.<br>See <a href="#">Example 11-3 on page 11-8</a> .                                                                                                                                                                                                 |
| <b>show running-config vlan</b> <i>vlan-id</i>                                                                                      | Displays VLAN information in the running configuration.                                                                                                                                                                                                                                                                                                                                                                      |
| <b>show vlan</b> [ <b>all-ports</b>   <b>brief</b>   <b>id</b> <i>vlan-id</i>   <b>name</b> <i>name</i>   <b>dot1q tag native</b> ] | Displays VLAN information as specified. See <a href="#">Example 11-4 on page 11-9</a> .                                                                                                                                                                                                                                                                                                                                      |
| <b>show vlan summary</b>                                                                                                            | Displays a summary of VLAN information.                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>show interface brief</b>                                                                                                         | Displays a table of interface states.<br>See <a href="#">Example 11-5 on page 11-9</a> .                                                                                                                                                                                                                                                                                                                                     |
| <b>module vem</b> <i>module-number</i> <b>execute vemcmd show port</b>                                                              | On the VEM, displays the port state on a particular VEM.<br>This command can only be used from the VEM.<br>See <a href="#">Example 11-6 on page 11-10</a> .                                                                                                                                                                                                                                                                  |
| <b>module vem</b> <i>module-number</i> <b>execute vemcmd show bd</b>                                                                | For the specified VEM, displays its VLANs and their port lists.<br>See <a href="#">Example 11-7 on page 11-10</a> .                                                                                                                                                                                                                                                                                                          |
| <b>module vem</b> <i>module-number</i> <b>execute vemcmd show trunk</b>                                                             | For the specified VEM, displays the VLAN state on a trunk port. <ul style="list-style-type: none"> <li>If a VLAN is forwarding (active) on a port, its CBL state should be 1.</li> <li>If a VLAN is blocked, its CBL state is 0.</li> </ul> See <a href="#">Example 11-8 on page 11-11</a> .                                                                                                                                 |

| Command                                                                                       | Purpose                                                                                                                                 |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <b>module vem</b> <i>module-number</i> <b>execute vemcmd</b><br><b>show l2</b> <i>vlan-id</i> | For the specified VEM, displays the VLAN forwarding table for a specified VLAN.<br><br>See <a href="#">Example 11-9 on page 11-11</a> . |
| <b>show interface</b> <i>interface_id</i> <b>mac-address</b>                                  | Displays the MAC addresses and the burn-in MAC address for an interface.                                                                |

**Example 11-1** *show mac address-table command*

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



**Tip** Module indicates the VEM on which this MAC address is seen.

The N1KV Internal Port refers to an internal port that is 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 Mod
-----+-----+-----+-----+-----+-----+-----
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 command*

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

```
n1000v# show mac address-table address 0050.568d.5a3f
VLAN MAC Address Type Age Port Mod
-----+-----+-----+-----+-----+-----+-----
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 Command*

```
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 Command



**Tip** This command shows the state of each VLAN that is 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 Type Vlan-mode
```

```
Remote SPAN VLANs
```

```
Primary Secondary Type Ports
```

#### Example 11-5 show interface brief Command

```
n1000v# show interface brief
```

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

| Ethernet Interface | VLAN | Type | Mode  | Status | Reason | Speed    | Port Ch # |
|--------------------|------|------|-------|--------|--------|----------|-----------|
| 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) | --        |

| Port-channel Interface | VLAN | Type | Mode  | Status | Reason | Speed      | Protocol |
|------------------------|------|------|-------|--------|--------|------------|----------|
| Po1                    | 1    | eth  | trunk | up     | none   | a-1000 (D) | none     |
| Po2                    | 1    | eth  | pvlan | up     | none   | a-10G (D)  | none     |

| Vethernet | VLAN | Type | Mode   | Status | Reason | Speed |
|-----------|------|------|--------|--------|--------|-------|
| Veth1     | 262  | virt | access | up     | none   | auto  |

| Port     | VRF | Status | IP Address | Speed | MTU  |
|----------|-----|--------|------------|-------|------|
| control0 | --  | up     | --         | --    | 1500 |

**Example 11-6** *module vem module-number execute vemcmd show port Command*



**Tip** Look for the state of the 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    | 1   | Access | l20    |
| 9   | 0        | 3969 | 0    | 2     | 2           | VIRT | UP    | UP    | 1   | Access | l21    |
| 10  | 0        | 115  | 0    | 2     | 0           | VIRT | UP    | UP    | 1   | Access | l22    |
| 11  | 0        | 3968 | 0    | 2     | 2           | VIRT | UP    | UP    | 1   | Access | l23    |
| 12  | 0        | 116  | 0    | 2     | 0           | VIRT | UP    | UP    | 1   | Access | l24    |
| 13  | 0        | 1    | 0    | 2     | 2           | VIRT | UP    | UP    | 0   | Access | l25    |
| 14  | 0        | 3967 | 0    | 2     | 2           | VIRT | UP    | UP    | 1   | Access | l26    |
| 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 Command*



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

```
n1000v# 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 Command*



**Tip**

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

```

n1000v# 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,
n1000v#

```

**Example 11-9** *module vem module-number execute vemcmd show L2 Command*

```

n1000v# configure terminal
n1000v(config)# module vem 3 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
n1000v#

```

## 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 MS-NLB, see the following 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:

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

## Disabling Automatic Static MAC Learning on vEthernet Interfaces

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

In interface configuration mode, enter these commands:

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

In port profile configuration mode, enter these commands:

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

## Checking the 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.

This example shows how to generate the VSM status in the interface configuration mode:

```
switch(config-if)# show running-config int veth1
interface Vethernet1
 inherit port-profile vm59
 description Fedora117, Network Adapter 2
 switchport port-security mac-address 001D.D8B7.1F81
 dvport uuid "ea 5c 3b 50 cd 00 9f 55-41 a3 2d 61 84 9e 0e c4"
```

This example shows how to generate the VSM status in the port profile configuration:

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

```

port-profile type vethernet ms-nlb
 ip port access-group abhi-acl in
 ip port access-group abhi-acl out
 no shutdown
 guid a85154f4-b07a-4cc4-86ad-fac0246557fe
 publish port-profile
 max-ports 300
 state enabled

```

## Checking the 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 the MS-NLB vEthernet interfaces are disabled.
- Confirm that the MS-NLB shared-MAC address (starting with 02:BF) is not listed in the Layer 2 (L2) MAC table.

This example shows how to generate the VSM status:

```

~ # vemcmd show port auto-smac-learning
LTL VSM Port Auto Static MAC Learning
49 Veth4 DISABLED
50 Veth5 DISABLED
51 Veth6 DISABLED

```

This example shows how to generate the Layer 2 MAC address table for VLAN 59:

```

~ # 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 UUFB to Block Unwanted MS-NLB Traffic

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.

A workaround for this situation is to enable unknown unicast flood blocking (UUFB).

### Enabling UUFB

This example shows how to enable UUFB. After you enter the commands in the example, press Ctrl-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 expected.

## Disabling UUFB for VMs That Use Dynamic MAC Addresses

Issues might occur for VMs that use dynamic MAC addresses. For ports that host these types of VMs, disable UUFB.

This example shows how to disable UUFB:

```
n1000v(config)# interface 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.

## 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 to 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 the sticky Address Resolution Protocol (ARP) when you configure private VLANs. ARP entries that are learned on Layer 3 private VLAN interfaces are sticky ARP entries. For security reasons, private VLAN port sticky ARP entries do not age out.

- IGMP runs only 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.)
- You can configure a SPAN to span both primary and secondary VLANs or, alternatively, to span either one if you are 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:

|                                                                  |   |
|------------------------------------------------------------------|---|
| <b>Checklist</b>                                                 | ✓ |
| 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                                                                         |
|---------------------------|----------------------------------------------------|----------------------------------------------------------------------------------|
| You cannot create a VLAN. | The Cisco Nexus 1000V is 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.

### Information About Private VLANs

Private VLANs (PVLANS) are used to segregate Layer 2 Internet service provider (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 isolated at Layer 2. In turn, the use of larger subnets reduces address management overhead. Three separate port designations are used. Each has its own unique set of rules that regulate 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; 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 do not need to be aware of the special VLAN type and can carry frames tagged with these VLANs just as they do with 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 with the packet, it is possible to maintain consistent behavior throughout the network. Therefore, the mechanism that 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 that 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 for Microsoft Hyper-V Layer 2 Switching Configuration Guide*.

## Troubleshooting Guidelines

Follow these guidelines when troubleshooting private VLAN issues:

- Verify that a private VLAN is configured correctly by entering the **show vlan private-vlan** command.
- Verify that the interface is up by entering the **show interface slot-port** command.
- Verify that the VEM is configured correctly by entering the **module vem module-number execute vemcmd show port** command.

## Private VLAN Troubleshooting Commands

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

To verify that a private VLAN is configured correctly, enter this 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, enter this command:

- **show interface**

```
n1000V# show interface eth3/4
Ethernet3/4 is up
 Hardware: Ethernet, address: 0050.565a.ca50 (bia 0050.565a.ca50)
 Port-Profile is DATA-Macpin
 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
 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
```

To verify if a virtual Ethernet interface in private VLAN host mode is up, enter this command:

- **show interface**

```
n1000V# show interface v3
Vethernet3 is up
 Port description is fedora9
 Hardware is Virtual, address is 0050.56bb.6330
 Owner is VM "fedora9", adapter is Network Adapter 1
 Active on module 3
 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, enter this command:

- **module vem module-number execute vemcmd show port**

```
n1000V# module vem 3 execute vemcmd show port
```

| LTl  | VSM    | Admin | Link | State | PC-LTL | Vlan | SG_ID | Vem Port |
|------|--------|-------|------|-------|--------|------|-------|----------|
| Type |        |       |      |       |        |      |       |          |
| 8    | Eth3/1 | UP    | UP   | UP    | 305    | 3969 | 2     |          |
| 9    | Eth3/2 | UP    | UP   | UP    | 305    | 3969 | 2     |          |
| 10   | Eth3/3 | UP    | UP   | UP    | 306    | 150  | 2     |          |
| 11   | Eth3/4 | UP    | UP   | UP    | 306    | 3968 | 2     |          |
| 12   | Eth3/5 | UP    | UP   | UP    | 306    | 151  | 2     |          |
| 13   | Eth3/6 | UP    | UP   | UP    | 0      | 1    | 2     |          |
| 14   | Veth33 | UP    | UP   | UP    | 0      | 3967 | 2     |          |
| 16   | Veth34 | UP    | UP   | UP    | 0      | 1 T  | 2     |          |

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.

## Information About NetFlow

NetFlow allows you to 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 subinterface) that matches a set of criteria. You create a flow using a flow record to define the criteria for your flow and all criteria must match for the packet to count in the given flow. Flows are stored in the NetFlow cache. Flow information tells you the following:

- The source address tells you who is originating the traffic.
- The destination address tells who is receiving the traffic.
- Ports characterize the application that use the traffic.
- The class of service (CoS) 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 the packets in the flow and the types of counters gathered per flow. You can define new flow records or use the predefined the Cisco Nexus 1000V flow records.

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

## NetFlow Troubleshooting Commands

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

To redirect the output of the following **debug** commands to a file stored in bootflash, enter this command:

- **debug logfile** *filename*
  - **debug nfm all**

To print monitor configuration, enter this command:

- **module vem *module-number* execute vemcmd show netflow monitor**

```
n1000V# module vem 3 execute vemcmd show netflow monitor
Flow Monitor m1:
 Table ID : 1
 Monitor ID: 65537
 Use Count: 1
 Inactive Timeout: 15
 Active Timeout: 1800
 Cache Type: normal
 Cache State: allocated
```

To print interface configuration, enter this command:

- **module vem *module-number* execute vemcmd show netflow interface**

```
n1000V# module vem 3 execute vemcmd show netflow interface
Interface: LTL49
 Monitor: m1
 Direction: Input
```

To print tracked configuration features, enter this command:

- **module vem *module-number* execute vemcmd show netflow stats**

```
n1000V# module vem 3 execute vemcmd show netflow stats
Netflow DPA-DP Session statistics:
 Session Opens: 1
 Session Verify: 1
 Session Commit: 1
 Session Abort: 0
 Session Add Monitor: 5
 Session Del Monitor: 0
 Get Cache stats: 0
 Get CPU stats: 0
 Show Cache: 0
 Ager Polls: 13775
 Module Cleanup: 6

Netflow DPA-DP Session Failure statistics:
 Opens Failures: 0
 Verify Failures: 0
 Commit Failures: 0
 Abort Failures: 0
 Add Monitor Failures: 0
 Del Monitor Failures: 0
 Get Cache stats Failures: 0
 Get CPU stats Failures: 0
 Show Cache Failures: 0
 Ager Polls Failures: 0

Netflow Packet Path Failure statistics:
 No Free Flows: 0
 Lost Flows: 0
 Ingress Pak Store Missing: 0
 Ingress Feature Store Missing: 0
 Ingress Permanent Full: 0
 Ingress Memory Failure: 0
 Ingress Multicast Packets: 0
 Ingress Non-IP Packets: 0
 Ingress Lock Failure: 0
 Ingress Policy not found: 0
 Post Ingress Pak Store Missing: 0
 Post Ingress Feature Store Missing: 0
 Post Ingress Permanent Full: 0
```

```

Post Ingress Multicast Packets: 0
Post Ingress Non-IP Packets: 0
Post Ingress Lock Failure: 0
Post Ingress Policy not found: 0
Egress Permanent Full: 0
Egress Memory Failure: 0
Egress Multicast Packets: 0
Egress Non-IP Packets: 0
Egress Lock Failure: 0
Egress Policy not found: 0

Netflow Packet Store Failure statistics:
Client Ref In Use: 0
Client Ref Null: 0
Pak Ref Null: 0
Alloc Client Ref Null: 0
Clear Client Ref Null: 0
Alloc Fail: 0
Central Info Mismatch: 0

Netflow Cache failure statistics:
No Free Entry: 0
Being Deleted: 0
Emergency Age Failure: 0
Normal Ager Failure: 0
No Ager Offset: 0

```

To dump the pakstore usage for a policy on an interface, enter the following command. The output goes to a vemlog internal buffer. Make sure that the output shows the correct monitor name and interface.

- **vemdebug netflow dump pakstore**

```

PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemdebug netflow dump pakstore
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

```

To enable NetFlow debugging for policy installation on the VEM, enter the following commands. Debug messages are printed for every PDL session open, verify, and commit requests coming from the DPA.

- **vemlog debug sfnetflow\_cache all**
- **vemlog debug sfnetflow\_config all**
- **vemlog debug sfnetflow\_flowmon all**
- **vemlog debug sfnetflow\_ager all**
- **vemlog debug sfnetflow\_flowapi all**

To enable packet path debugging for NetFlow policies on the VEM, enter the following command. Debug messages are printed for every packet that hits a NetFlow policy. Use this command with caution. High traffic could result in a lot of debug messages.

- **vemlog debug sfnetflow all**

Enter these commands to collect information about NetFlow manager (NFM) process run-time configuration errors:

- **show flow internal event-history errors**
- **show flow internal event-history msgs**
- **show flow internal pdl detailed**

- **show flow internal mem-stats** (to debug memory usage and leaks)

## Problems with NetFlow

Common NetFlow configuration problems on the Virtual Supervisor Module (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 that 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.
- Configure monitors in multiple levels of a port-profile inheritance tree.

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. Enter the **no** form of the original command to clear the configuration and then reenter the command.

## Debugging a Policy Verification Error

- 
- |               |                                                                                                                                                     |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step 1</b> | Configure all debug flags of NetFlow monitor (NFM) by entering the <b>debug nfm all</b> command.                                                    |
| <b>Step 2</b> | Save the Secure Shell Telnet (SSH) session buffer to a file.                                                                                        |
| <b>Step 3</b> | Enable a flexible NFM for traffic that the router is receiving or forwarding by entering the <b>ip flow monitor monitor name direction</b> command. |

The command executes once again and the debug traces are 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 Problems

When debugging a NetFlow statistics export problem, follow these guidelines:

- Ensure that the destination IP address is reachable from the VSM and Virtual Ethernet Modules (VEMs).
- Ensure that the UDP port configured on the exporter matches that used by the NetFlow Collector.
- View statistics for the exporter and identify any drops by entering the **show flow exporter** command.





## Access Control Lists

---

This chapter describes how to identify and resolve problems that relate to access control lists (ACLs).

### Information About 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 for Microsoft Hyper-V 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 Virtual Ethernet Module (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 on ACL rules.
- Non-initial fragments are not subject to an 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.
- ACLs are not supported in port channels.

## ACL Troubleshooting Commands

The commands listed in this section can be used on the Virtual Supervisor Module (VSM) to see the policies that are configured and applied on the interfaces.

Display the configured ACLs by entering this command:

- **show access-list summary**

Display the run-time information of the ACLMGR and ACLCOMP during configuration errors and to collect ACLMGR process run-time information configuration errors by entering these commands on the VSM:

- **show system internal aclmgr event-history errors**
- **show system internal aclmgr event-history msgs**
- **show system internal aclmgr ppf control**
- **show system internal aclmgr mem-stats (to debug memory usage and leaks)**
- **show system internal aclmgr status**
- **show system internal aclmgr dictionary**

Collect ACLCOMP process run-time information configuration errors by entering these commands:

- **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)**
- **show system internal aclcomp ppf control**

## Displaying ACL Policies on the VEM

You can use the commands in this section to display configured ACL policies on the VEM.

To list the ACLs installed on that server, enter this command:

- **module vem *module-number* execute vemcmd show acl**

```
n1000v # module vem 3 execute vemcmd show acl
Acl-id Ref-cnt Type Numrules Stats Stat-id
 1 33 IPv4 127 enabled 9
 2 81 IPv4 127 enabled 10
 3 33 IPv4 127 enabled 11
```

The Acl-id is the local ACL ID for this VEM. Ref-cnt refers to the number of instances of this ACL in this VEM.

To list the interfaces on which ACLs have been installed, enter this command:

- **module vem *module-number* execute vemcmd show acl pinst**

```
n1000v# module vem 3 execute vemcmd show acl pinst
LTL Acl-id Dir
16 1 ingress
```

## Debugging Policy Verification Issues

- Step 1** Redirect the output to a file in bootflash by entering the **debug logfile *filename*** command on the VSM.
- Step 2** Configure all debug flags of aclmgr by entering the **debug aclmgr all** command.
- Step 3** Configure all debug flags of aclcomp by entering the **debug aclcomp all** command.
- Step 4** From the VSM enter the following steps:



**Note** The output goes to the console.

- Enable ACL logging on the DPA by entering these commands:
    - **module vem *module-number* execute vemcmd dpa debug sfaclagent all**
    - **module vem *module-number* execute vemcmd dpa debug sfpdlagent all**
  - Enable logging on the VEM by enter the **module vem *module-number* execute vemlog debug sfacl all** command.
  - Enable DPA logging for viewing by entering the **module vem *module-number* execute vemlog start** command.
  - Enable DPA logging for viewing by entering the **module vem *module-number* execute vemlog start** command.
- Step 5** Configure the policy that was causing the verification error.
- Step 6** Display DPA logs by entering the **module vem *module-number* execute vemlog show all** command.
- Step 7** Save the Telnet or SSH session buffer to a file.
- Step 8** Copy the logfile created in bootflash.





# Quality of Service

This chapter describes how to identify and resolve problems related to Quality of Service (QoS).

## Information About Quality of Service

QoS allows you to 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, class of service (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, see the *Cisco Nexus 1000V for Microsoft Hyper-V 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 VSM Troubleshooting Commands

You can use the commands in this section on the Virtual Supervisor Module (VSM) to troubleshoot the policies that are configured and applied on the interfaces.

Display configured policies and class maps by entering these commands:

- **show policy-map** [*policy-map-name*]
- **show class-map** [*class-map-name*]

Display installed policies by entering this command:

- **show policy-map interface brief**

Collect QOSMGR process run-time information configuration errors by entering these commands on the VSM:

- **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**

Collect ACLCOMP process run-time information configuration errors by entering these commands on the VSM:

- **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)

## QoS VEM Troubleshooting Commands

You can use the commands in this section to display configured QoS policies on the VEM.

To list all class maps and policies in use on the server, enter this command:

- **module vem module-number execute vemcmd show qos node #**
- ```
n1000v# 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

```

To list all the installed policy maps in use on the server, enter this command:

- **module vem module-number execute vemcmd show qos policy**

```

n1000v# module vem 3 execute vemcmd show qos policy
policyid classid policerid set_type value
-----
0          1          -1          dscp          5
          2          0          dscp          0

```

To list all service policies installed on the server, enter this command:

- **module vem module-number execute vemcmd show qos pinst**

```

n1000v# 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

-
- Step 1** Enter the **debug aclmgr all** command if the policy references an ACL.
- Step 2** Set all debug flags for IP QoS manager by entering the **debug ipqos all** command.
- Step 3** Configure all aclcomp debug flags by entering the **debug aclcomp all** command.
- Step 4** Execute the command once again with debug traces output to the console by entering the **service-policy** command. 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 might be easier to first install it directly on an interface.

-
- Step 1** Clear the log by entering the **module vem module-number execute vemcmd dpa clear** command.
- Step 2** Redirect the output of the dpa logs to the vemlog by entering the **module vem module-number execute vemcmd dpa sfqosagent all** command.
-

- Step 3** Start the vecmd DPA by entering the **module vem *module-number* execute vemcmd dpa start** command.
- Step 4** Execute the command once again with the DPA debug traces output to vemcmd dpa by entering the **service-policy** command.
- Step 5** Stop the vemcmd DPA by entering the **module vem *module-number* execute vemcmd dpa stop** command.
- Step 6** See the logs on the console by entering the **module vem *module-number* execute vemlog show all** command.

The output looks 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.

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.

The Cisco Nexus 1000V supports two types of SPAN:

- SPAN (local SPAN) that can monitor sources within a host or Virtual Ethernet Module (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 for Microsoft Hyper-V 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 might 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 a live migration (a live migration provides the capability to move a virtual machine (VM) from one node in a Microsoft Windows server failover cluster to another node without a perceived interruption in service by applications/clients connecting to the VM), the following might occur:
 - 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 to run a 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 live migration.
 - 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 address 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 the Cisco NX-OS VEM connectivity is working correctly. 2. Force reprogramming of the session on the VEM by entering these commands: <ul style="list-style-type: none"> – shut – no shut
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.

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 module-number execute vemcmd show span	Displays the VEM source IP and SPAN configuration. See Example 17-3 on page 17-4 .
show monitor internal errors	Displays the error logs.
show monitor internal event-history msgs	Displays the event history messages.
show monitor internal info global-info	Displays the global component information.
show monitor internal mem-stats	Displays the memory allocation statistics.

Example 17-1 show monitor command

```

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

```

Example 17-2 show monitor session command

```

n1000v(config)# show monitor session 1
session 1
-----
type           : erspan-source
state          : up
ERSPAN ID      : 999
source intf    :
  rx           : Eth3/3
  tx           : Eth3/3
  both         : Eth3/3
source VLANs   :
  rx           :
  tx           :
  both         :
source port-profile :
  rx           :
  tx           :
  both         :
filter VLANs   : filter not specified
destination IP : 10.54.54.1
ERSPAN TTL     : 64
ERSPAN IP Prec. : 0
ERSPAN DSCP    : 0
ERSPAN MTU     : 1000
ERSPAN Header Type: 2

```

Example 17-3 *module vem execute vemcmd show span command*

```
n1000v# module vem 3 execute vemcmd show span
HW SSN ID   ERSpan ID   HDR VER   DST LTL/IP
      1       180         2    10.54.54.1
RX Sources :17,18,
TX Sources :305,
Source Filter RX :261,263,264,
Source Filter TX:261,263,264,
```



Multicast IGMP Snooping

This chapter describes how to identify and resolve problems that relate to multicast Internet Group Management Protocol (IGMP) snooping.

Information About Multicast IGMP Snooping

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.

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.

IGMP snooping works as follows:

- Ethernet switches, such as Catalyst 6500 series 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 by 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 by 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 checks 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 entering the **ip multicast-routing** command.

This example shows how to turn on global multicast routing, configure an SVI interface, and turn on the PIM routing protocol:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ip multicast-routing
switch(config)# end

switch# configure 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:

- Verify that IGMP snooping is enabled by entering the **show ip igmp snooping** command.
- Make sure the upstream switch has IGMP configured.
- Verify that the Cisco Nexus 1000V switch is configured correctly and is ready to forward multicast traffic by entering the **show ip igmp snooping groups** command. In the displayed output of the command, look for the letter R under the port heading. The R indicates that the Virtual Supervisor Module (VSM) has learned the uplink router port from the IGMP query that was sent by the upstream switch, and means that the Cisco Nexus 1000V is ready to forward multicast traffic.

Multicast IGMP Snooping Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to multicast IGMP snooping.

Command	Purpose
show cdp neighbor	Displays Cisco Discovery Protocol (CDP) neighbors. 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 CDP on the upstream switch by entering the no cdp enable command, the show cdp neighbor command does not display any information. See Example 18-1 on page 18-3 .
show ip igmp snooping groups	Displays if IGMP snooping is enabled on the VLAN. See Example 18-2 on page 18-3 .

Command	Purpose
show ip igmp snooping groups	Displays snooping information for the group addresses.
debug ip igmp snooping vlan	Enables snooping on IGMP for events on all VLANs. See Example 18-3 on page 18-3 .

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

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
Optimised Multicast Flood (OMF) enabled
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
Link Local Groups suppression enabled
Router port detection using PIM Hellos, IGMP Queries
Number of router-ports: 0
Number of groups: 0
VLAN vPC function disabled
Active ports:
```

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 forwards 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, 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 that it is a VM and generates multicast traffic to 224.1.2.3. The traffic is 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 multicast groups, the address uses the default route, which means that the traffic is 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  vmnic3
    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
```

Problems with Multicast IGMP Snooping

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

Symptom	Solution
A VM is interested in multicast traffic but is not receiving the multicast traffic	Determine if IGMP snooping is working as expected by entering the debug ip igmp snooping vlan command. 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 Virtual Machine (VM).
	Verify that the multicast distribution table in the Virtual Ethernet Module (VEM) has the correct information by entering the module vem module-number execute vemcmd show vlan command.
	View the port table by entering the module vem module-number execute vemcmd show port command Make sure that the table has the correct information and 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 Address Resolution Protocol (ARP) Inspection (DAI)
- IP Source Guard (IPSG)

Information About DHCP Snooping

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

- Validates DHCP messages that are 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 for Microsoft Hyper-V 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 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 for Microsoft Hyper-V 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 DHCP snooping binding table.

For detailed information about configuring IP Source Guard, see the *Cisco Nexus 1000V for Microsoft Hyper-V 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 2048 DHCP entries can be snooped and learned system-wide in the distributed virtual switch (DVS). This total is for both entries learned dynamically and entries configured statically.
- Rate limits on interfaces that must be set to high values for trusted interfaces such as VSD SVM ports or vEthernet ports that connect to DHCP servers.
- If the Virtual Supervisor Module (VSM) uses the Virtual Ethernet Module (VEM) for connectivity (that is, the VSM has a VSM Asynchronous Inter-process Communication (AIPC), management, and inband ports on a particular VEM), these virtual Ethernet interfaces must be configured as trusted interfaces.
- The connecting interfaces on a device upstream from the Cisco Nexus 1000V must be configured as trusted if DHCP snooping is enabled on the device.

For detailed guidelines and limitations used in configuring these features, see the *Cisco Nexus 1000V for Microsoft Hyper-V 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, the DHCP client is not able to obtain an IP address from the server.	<p>IP address was not added to binding database.</p> <p>The connection is faulty between the DHCP server and client.</p>	<ol style="list-style-type: none"> 1. Verify the connection between the DHCP server(s) and the host connected to the client by entering the vmkping command. 2. If the connection between the 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, verify that the VLAN is allowed, and so on. – Make sure that the server itself is up and running. – Make sure that global DHCP snooping (IP DHCP snooping) is configured on the VSM
	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 by entering the show ip dhcp snooping command. 2. On the VSM, verify that the vEthernet interface attached to the server is trusted by entering the module vem module-number execute vemcmd show dhcps interfaces command.
	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 that 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.	<p>On the VSM, verify DHCP statistics by entering these commands.</p> <ul style="list-style-type: none"> • show ip dhcp snooping statistics • module vem module-number execute vemcmd show dhcps stats

Dropped ARP Response Troubleshooting

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 by entering the show ip arp inspection command.</p> <p>For detailed information about configuring DAI, see the <i>Cisco Nexus 1000V for Microsoft Hyper-V 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 by entering the show ip dhcp snooping command.</p> <p>For detailed information about enabling DHCP and configuring DAI, see the <i>Cisco Nexus 1000V for Microsoft Hyper-V 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"> From the VSM, verify the VEM DHCP snooping configuration by entering the module vem module-number execute vemcmd show dhcps vlan command. 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 for Microsoft Hyper-V 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"> On the VSM, display the binding table by entering the show ip dhcp snooping binding command. 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 for Microsoft Hyper-V Security Configuration Guide</i>.</p>
The binding that corresponds to the VM sending the ARP response is not present in the binding table.	<ol style="list-style-type: none"> On the VSM, display the binding table by entering the show ip dhcp snooping binding command. 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 for Microsoft Hyper-V Security Configuration Guide</i>.</p> <ol style="list-style-type: none"> If all configurations are correct, make sure to turn on DHCP snooping before DAI or IPSG to make sure that 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 for Microsoft Hyper-V 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.	On the VSM, verify that IP Source Guard is configured as expected by entering these commands: <ul style="list-style-type: none"> • show port-profile name <i>profile_name</i> • show running interface <i>if_ID</i> • show ip verify source For detailed information about configuring IP Source Guard, see the <i>Cisco Nexus 1000V for Microsoft Hyper-V Security Configuration Guide</i> .
	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 by entering the show ip dhcp snooping binding command. 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 by entering the show ip dhcp snooping binding command.

Collecting and Evaluating Logs

This section includes the following topics:

- [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
debug dhcp all	Enables debugging for all DHCP configuration flags.
debug dhcp errors	Enables debugging of DHCP errors.
debug dhcp mts-errors	Enables debugging of MTS errors.
debug dhcp mts-events	Enables debugging of MTS events.
debug dhcp pkt-events	Enables debugging of PKT events.
debug dhcp pss-errors	Enables debugging of PSS errors.
debug dhcp pss-events	Enables 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.

Hyper-V Host Command	Description
vemcmd.exe dpa debug sfdhcpsagent all	Enables DPA DHCP agent debug logging. Enter the vemlog.exe show all command to view the log messages.
vemlog.exe debug sfdhcps all	Enables data-path debug logging, and captures logs for the data packets sent between the client and the server.
vemlog.exe debug sfdhcps_config all	Enables data-path debug logging, and captures logs for configuration that is coming from the VSM.
vemlog.exe debug sfdhcps_binding_table all	Enables data-path debug logging, and captures logs that correspond 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
show running-config dhcp	Displays the DHCP snooping, DAI, and IP Source Guard configuration. See Example 19-1 on page 19-7 .
show ip dhcp snooping	Displays general information about DHCP snooping and whether option 82 is enabled. See Example 19-2 on page 19-7 .
show ip dhcp snooping binding	Display the contents of the DHCP snooping binding table. See Example 19-3 on page 19-7 .
show feature	Displays the features available, such as DHCP, and whether they are enabled. See Example 19-4 on page 19-7 .
show ip arp inspection	Displays the status of DAI. See Example 19-5 on page 19-8 .
show ip arp inspection interface vethernet <i>interface-number</i>	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 .

Example 19-1 show running-config dhcp Command

```
n1000v# show running-config dhcp

!Command: show running-config dhcp
!Time: Fri Feb 8 19:29:50 2013

version 5.2(1)SM1(5.1)
feature dhcp

no ip dhcp relay

n1000v#
```

Example 19-2 show ip dhcp snooping Command

```
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 Command

```
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 Command

```
n1000v# show feature
Feature Name          Instance   State
-----
dhcp-snooping         1         enabled
http-server           1         enabled
ippool                1         enabled
```

```

lacp                1          enabled
lisp                1          enabled
lispHelper          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 Command*

```

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 Command*

```

n1000v# show ip arp inspection interface vethernet 6

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

```

Example 19-7 *show ip arp inspection vlan Command*

```

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 Command*

```
n1000v# show ip verify source
```

```
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
```




System

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

Information About the System

The Microsoft System Center Virtual Machine Manager (SCVMM) is designed for the management of a large number of virtual servers that are based on the Microsoft Hyper-V model. The Microsoft SCVMM is a management application that is used to manage the windows server hosts. The Cisco Nexus 1000V is a virtual switch that provides switching for Virtual Machines (VMs) that are deployed on the servers running Microsoft Windows 2016.

VEM Troubleshooting Commands

You can use the commands in this section to troubleshoot problems related to the Virtual Ethernet Module (VEM). Enter these commands in a PowerShell window.

Command	Purpose
vemlog	Displays and controls the VEM kernel logs.
vemcmd	Displays configuration and status information.
vemcmd show version	Displays the version information. See Example 20-1 on page 20-2 .
vemlog show last <i>number-of-entries</i>	Displays the circular buffer. See Example 20-2 on page 20-2 .
vemlog show info	Displays information about entries in the log. See Example 20-3 on page 20-2 .
vemcmd help	Displays the type of information you can display. See Example 20-4 on page 20-2 .
vem-support.ps1	Navigate to the support directory under Nexus1000v and run the vem-support.ps1 script. See Example 20-5 on page 20-2 .

Example 20-1 vemcmd show version Command

```
PS C:\Program Files (x86)\Cisco\Nexus1000V> .\VemCmd.exe show version
VEM Version: 5.2.1.SM1.5.0.278-3
VSM Version: 5.2(1)SM1(5.1) [build 5.2(1)SM1(5.0.278)] [gdb]
System Version: Windows Server 2016 - Datacenter (6.2.9200, 6.30)
```

Example 20-2 vemlog show last Command

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemlog show last 5
Timestamp Entry CPU Mod Lv Message
Mar 17 14:47:30.124446 28768 0 99 4 Warning Could not get LACP Port for LTL 20
Mar 17 14:48:00.123500 28769 0 99 4 Warning Could not get LACP Port for LTL 22
Mar 17 14:48:00.123500 28770 0 99 4 Warning Could not get LACP Port for LTL 21
Mar 17 14:48:00.123500 28771 0 99 4 Warning Could not get LACP Port for LTL 20
Mar 17 14:48:00.248291 28772 6 0 0 Suspending log
```

Example 20-3 vemlog show info Command

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemlog show info
Enabled: Yes
Total Entries: 28778
Wrapped Entries: 26886
Lost Entries: 0
Skipped Entries: 0
Available Entries: 26886
Stop After Entry: Not Specified
```

Example 20-4 vemcmd help Command

```
PS C:\Program Files (x86)\cisco\Nexus1000V> .\vemcmd help
vemcmd help:

show
show version Show the VEM and VSM versions
show card Show the card's global info
show data Show switch global data
show vsm uptime Show the VSM's uptime
show acl Show ACL ids
```

Example 20-5 vem-support.ps1 Command

```
PS C:\Program Files (x86)\cisco\Nexus1000V\Support> .\vem-support.ps1

Directory: C:\Program Files (x86)\Cisco\Nexus1000V\Support

Mode LastWriteTime Length Name
----
d---- 3/17/2013 2:51 PM WIN-35-cisco-vem-2013-0317-1451
```

VEM Log Commands

Use the following commands to control the vemlog:

- **vemlog stop**—Stops the log.
- **vemlog clear**—Clears 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.



Network Segmentation Manager

This chapter describes how to identify and resolve problems with Network Segmentation Manager (NSM).

Information About Network Segmentation Manager

For information on the Network Segmentation Manager, see the *Cisco Nexus 1000V for Hyper-V Network Segmentation Manager Configuration Guide*.

Problems with Network Segmentation Manager

For more information about problems occurring with NSM see the [“Problems with Port Profiles” section on page 9-2](#).

Updating VM Fails

This problem usually occurs when you are trying to attach a vEth port to the Cisco Nexus 1000V and Microsoft SCVMM server.

-
- | | |
|---------------|---|
| Step 1 | Launch the Microsoft SCVMM UI. |
| Step 2 | Verify the system flags on the network segment and port profile on the Virtual Supervisor Module (VSM) by entering the following commands: <ul style="list-style-type: none">• show run port-profile <i>port-profile-name</i>• show nsm network segment name <i>name</i> |
| Step 3 | Repair the VM from the Microsoft SCVMM by running the <i>repair</i> option followed by <i>ignore</i> . |
| Step 4 | Once the operation is complete, refresh the VM. |
| Step 5 | Move the vEth port to the Not Connected state from the Microsoft SCVMM. |
| Step 6 | Attach the vEth port to the Microsoft SCVMM using one of the following combinations: <ul style="list-style-type: none">• Choose System Network Segment > System Port Profile• Choose System Network Segment > Non System Port Profile |

Choose **Non System Network Segment > Non System Port Profile**

Network Segment Not Visible on the Microsoft SCVMM

When creating a VM network on a Microsoft SCVMM using a network segment from a Cisco Nexus 1000V device, the network segment cannot be found. This problem can occur due to one of the following reasons:

- The network segment is in the unpublished state.
- A Switch Extension Manager refresh was not performed after creating the network segment.

-
- Step 1** Launch the Microsoft SCVMM UI.
- Step 2** From the VSM, verify the network segment configuration by entering the **show nsm network segment name** command:
- Step 3** Verify that the network segment has a publish-name.
- Step 4** After completing [Step 3](#), choose **Fabric > Switch Extension Manager/Network Service> Extension**, and choose **Refresh** to update the Microsoft SCVMM with the latest VSM configuration.
-

Network Segment Is Not Available on the Microsoft SCVMM

When you apply a VM network on a Microsoft SCVMM to a Cisco Nexus 1000V vNIC, a network error is displayed.

-
- Step 1** On the Microsoft SCVMM, identify the network segment used to create the VM network.
- Step 2** On the VSM, identify the network segment pool to which the network segment is associated.
- Step 3** Identify the uplink networks that allow the network segment pool identified in [Step 2](#).
- Step 4** On the Microsoft SCVMM, verify that the logical switch on that host is using the uplink networks identified in [Step 3](#).
- Step 5** If the links do not agree, change the appropriate uplink networks as needed.
-

Network Segmentation Manager Troubleshooting Commands

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

Command	Purpose
show nsm ip pool template name <i>name</i>	Displays the IP pool template information.
show nsm ip pool template usage network segment	Displays the network segment using an IP pool template.
show nsm logical network <i>name</i>	Displays the NSM logical network name.
show nsm network segment brief	Displays brief information about the network segment information.
show nsm network segment filter network segment pool <i>name</i>	Displays the filtered information for a network segment pool.
show nsm network segment filter vlan <i>vlan_ID</i>	Displays the network segment VLAN information.
show nsm network segment filter pvlan host <i>vlan_ID</i>	Displays the network segment PVLAN host information.
show nsm network segment filter pvlan primary <i>vlan_ID</i>	Displays the network segment PVLAN primary mode information.
show nsm network segment filter pvlan promiscuous <i>vlan_ID</i>	Displays the network segment PVLAN promiscuous mode information.
show nsm network segment filter pvlan secondary <i>vlan_ID</i>	Displays the network segment PVLAN information for a specified secondary VLAN.
show nsm network segment name <i>name</i>	Displays network segment information.
show nsm network segment pool <i>name</i>	Displays network segment pool information.
show nsm network uplink brief	Displays brief information about the network segment uplink.
show nsm network uplink filter import <i>Ethernet Port-Profile name</i>	Displays network segment uplink information filtered by the Ethernet policy port profile.
show dynamic-port-profile	Displays dynamic port profile information.
show dynamic-port-profile <i>name</i>	Displays dynamic port profile information for the specified port profile.
show dynamic-port-profile inherit <i>name</i>	Displays dynamic port profiles with inherited vEthernet policy profiles.
show dynamic-port-profile network segment <i>name</i>	Displays dynamic port profile network segment information.

For detailed information about **show** command output, see the *Cisco Nexus 1000V for Microsoft Hyper-V Command Reference Guide*.



Ethalyzer

This chapter describes how to use Ethalyzer as a Cisco NX-OS protocol analyzer tool.

Information About 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:

Command	Purpose
ethalyzer local interface <i>interface</i>	Captures packets sent or received by the supervisor and provides detailed protocol information. Note For all commands in this table, the interface is control, ha-primary, ha-secondary, inband (packet interface) or mgmt (management interface).
ethalyzer local interface <i>interface</i> limit-captured-frames	Limits the number of frames to capture.
ethalyzer local interface <i>interface</i> limit-frame-size	Limits the length of the frame to capture.
ethalyzer local interface <i>interface</i> capture-filter	Filters the types of packets to capture.
ethalyzer local interface <i>interface</i> display-filter	Filters the types of captured packets to display.
ethalyzer local interface <i>interface</i> raw	Dump the packet in HEX/ASCII with a one line summary.
ethalyzer local interface <i>interface</i> write	Saves the captured data to a file.
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 interface mgmt limit-captured-frames 4
Capturing on eth1
2016-10-01 19:15:23.794943 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=64
2016-10-01 19:15:23.796142 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=144
2016-10-01 19:15:23.796608 10.78.110.241 -> 72.163.145.51 SSH Encrypted response packet
len=144
2016-10-01 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.



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.

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 might need to contact your customer support representative or Cisco TAC for some additional assistance. This section outlines the steps that the you should perform before you contact your next level of support.



Note

Do not reload the module or the switch at least until you have completed [Step 1](#). Some logs and counters are kept in volatile storage and will not survive a reload.

Step 1

Collect the switch information and configuration before and after the issue has been resolved.

On the VSM, generate the technical support by entering the **show tech-support detail > tech-support** command. Use SCP/SFTP/FTP to get the file from the VSM.

On the VEM, generate the support directory by entering the following commands in a PowerShell window:

- **set-ExecutionPolicy Unrestricted**
- **cd c:\program files (x86)\Cisco\Nexus1000V\support**
- **vem-support.ps1**

Add the directory to a zip file to send to technical support.

- Step 2** Capture the exact error codes that you see in CLI message logs by entering one of these commands:
- **show logging log** (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?
 - Which Cisco Nexus 1000V software, driver versions, operating systems versions, and storage device firmware are in your fabric?
 - Which Microsoft Hyper-V and Microsoft SCVMM software are you 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 file that contains memory information 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.

This example shows how to generate a file of core memory information, or a core dump.

```
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 filename (indicated by jsmith_cores) must exist in the TFTP server directory.

Copying Files

You might be required to move files to or from the switch. These files might include log, configuration, or firmware files.

The Cisco Nexus 1000V always acts as a client, so that an ftp/scp/tftp session always originates from the switch and either pushes files to an external system or pulls 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]"
```

This example shows how to copy /etc/hosts from 172.22.36.10 using the user user1, where the destination would be hosts.txt:

```
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
```

This example shows how to back up the startup configuration to an SFTP server:

```
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 before you make any changes. A short script could be written to be run on the Cisco Nexus 1000V to perform a save and then a backup of the configuration. The script needs to contain two commands only: **copy running-configuration startup-configuration** and then **copy startup-configuration tftp://server/name**. To execute the script, enter the **run-script filename** command.