Cisco Nexus 1000VE for VMware vSphere Troubleshooting Guide, Release 5.2(1)SV5(1.1)

May 11, 2020

Cisco Systems, Inc.
www.cisco.com

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

Text Part Number:

CONTENTS

Preface

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

This preface describes the following aspects of this document:

Audience

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

Document Conventions

Command descriptions use these conventions:

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

Screen examples use these conventions:

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

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

This document uses the following conventions:

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

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

Related Documentation

This section lists the documents used with the Cisco Nexus 1000VE and available on Cisco.com at the following URL:

https://www.cisco.com/c/en/us/support/switches/nexus-1000ve/tsd-products-support-series-home.html

General Information

Cisco Nexus 1000VE Release Notes

Cisco Nexus 1000VE Compatibility Information

Install and Upgrade

Cisco Nexus 1000VE Installation, Migration, and Upgrade Guide

Configuration Guides

Cisco Nexus 1000VE Layer 2 Switching Configuration Guide

Cisco Nexus 1000VE Security Configuration Guide

Cisco Nexus 1000VE System Management Configuration Guide

Reference Guides

Cisco Nexus 1000VE Command Reference

Troubleshooting, Password Recovery, System Messages Guides

Cisco Nexus 1000VE Troubleshooting Guide

Virtual Services Appliance Documentation

The Cisco Nexus Virtual Services Appliance (VSA) documentation is available at

https://www.cisco.com/c/en/us/support/switches/cloud-services-platform-2100/tsd-products-support-series-home.html

Virtual Security Gateway Documentation

The Cisco Virtual Security Gateway documentation is available at

http://www.cisco.com/en/US/products/ps13095/tsd_products_support_series_home.html

Obtaining Documentation and Submitting a Service Request

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

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

Overview

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

This chapter includes the following sections:

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.

Best Practices

We recommend that you do the following to ensure the proper operation of your networks:

- Maintain a consistent Cisco Nexus 1000VE release across all network devices.
- Refer to the release notes for your Cisco Nexus 1000VE release for the latest features, limitations, and caveats.
- Enable system message logging. See the "Overview of Symptoms" section on page 1-4.

- Verify and troubleshoot any new configuration changes after implementing the change.

Troubleshooting Basics

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

This section includes the following topics:

Troubleshooting Guidelines

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, 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 on problems in your system. See the "Gathering Information" section on page 1-2.

Step 2 Verify the Layer 2 connectivity. See the "Verifying Layer 2 Connectivity" section on page 1-3.

Step 3 Verify the configuration for your end devices (storage subsystems and servers).

Step 4 Verify end-to-end connectivity. See the "Verifying Layer 3 Connectivity" section on page 1-3.

Gathering Information

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

Use the following commands and examine the outputs:

- **show module**
- **show version**
- **show running-config**
- **show logging log**
- **show interface brief**
- **show vlan**
- **show accounting log**
- **show tech support svs**

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

Verifying Ports

Answer the following questions to verify ports:

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

Verifying Layer 2 Connectivity

Answer the following questions to verify Layer 2 connectivity:

- Are the necessary interfaces in the same VLANs?

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

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

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

Verifying Layer 3 Connectivity

Answer the following questions to verify Layer 3 connectivity:

- Have you configured a gateway of last resort?
- Are any IP access lists, filters, or route maps blocking route updates?

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

Overview of Symptoms

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

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

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

System Messages

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

This section contains the following topics:

System Message Text

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 switch %KERN-1-SYSTEM_MSG: stun_set_domain_id : Setting domain ID
(1024) - kernel
```

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

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

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

```
next_gen_fcs_ip_38# show dc hosts vse

Internal IPG tags:

inside-trunk 1:1-50,

inside-trunk 2:2047-2096,

-----------------------------------------

HOST NAME: 10.197.128.55

HOST IP: 10.197.128.55

VSE IP: 10.197.128.37

VSE UUID : 42193D34-FBB0-A6E9-4AAE-C4BC3043C013



-----------------------------------------
```

Explanation VSE module inserted successfully on slot 3.

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

syslog Server Implementation

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

This example demonstrates how to configure a Cisco Nexus 1000VE to use the syslog facility on a Solaris platform. Although a Solaris host is being used, the 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 emailed 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 1000VE 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 1000VE messages to: /var/adm/nxos_logs

To configure a syslog server, follow these steps:

Step 1 Configure the Cisco Nexus 1000VE.

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

Display the configuration.

```
switch# show logging server
Logging server: enabled
{192.0.2.1}
     server severity: notifications
     server facility: local1
```

Step 2 Configure the syslog server.

a. Modify /etc/syslog.conf to handle local1 messages. For Solaris, t at least one tab needs to be 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 that the syslog has 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 1000VE. 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 1000VE 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

Use the following commands to access and view logs in the Cisco Nexus 1000VE.

```
switch# show logging ?
> 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

```
switch# 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 1000VE

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 1000VE 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 Cisco Nexus 1000VE and support contract from Cisco, contact Cisco for Cisco Nexus 1000VE support. Cisco provides Layer 1, Layer 2, and Layer 3 support.

After you have collected this information, see the "Obtaining Documentation and Submitting a Service Request" section on page -xi.

For more information on the steps to take before calling Technical Support, see the "Gathering Information" section on page 1-2.

Troubleshooting Tools

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

Commands

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

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

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

```
switch# 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

switch# 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 destination.

Traceroute

Use traceroute to do the following:

- Trace the route followed by data traffic.
- Compute interswitch (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.

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

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

Monitoring Processes and CPUs

There CLI enables you to for monitor switch processes. CPU status, and utilization.

This section contains the following topics:

Identifying the Running Processes 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 “-” usually means a daemon 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 is currently not running.

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

Example 2-1 show processes Command

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

switch# show processes

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

Displaying CPU Utilization

Use the **show processes cpu** command to display CPU utilization. See Example 2-2. The command output includes the following:

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

Example 2-2 show processes cpu Command

```
switch# show processes cpu

PID    Runtime(ms)  Invoked   uSecs  1Sec   Process
-----  -----------  --------  -----  -----  -----------
```

```
    1          922  4294967295      0      0  init
    2          580    377810      1      0  migration/0
    3          889   3156260      0      0  ksoftirqd/0
    4         1648    532020      3      0  desched/0
    5          400    150060      2      0  migration/1
    6         1929   2882820      0      0  ksoftirqd/1
    7         1269    183010      6      0  desched/1
    8         2520  47589180      0      0  events/0
    9         1730   2874470      0      0  events/1
   10           64    158960      0      0  khelper
   15            0    106970      0      0  kthread
   24            0     12870      0      0  kacpid
  101           62   3737520      0      0  kblockd/0
  102           82   3806840      0      0  kblockd/1
  115            0     67290      0      0  khubd
  191            0      5810      0      0  pdflush
  192          983   4141020      0      0  pdflush
  194            0      5700      0      0  aio/0
  193            0      8890      0      0  kswapd0
  195            0      5750      0      0  aio/1
...
```

Displaying CPU and Memory Information

Use the **show system resources** command to display system-related CPU and memory statistics. See Example 2-3. The output includes the following:

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

Example 2-3 show system resources Command

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

RADIUS

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

- Authentication
- Authorization

- Accounting

Authentication refers to the authentication of users for access to a specific device. You can use RADIUS to manage user accounts for access to a Cisco Nexus 1000VE. When you try to log into a device, the Cisco Nexus 1000VE 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 then refer to the RADIUS server to determine the extent of access the user will have within the switch network.

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

The following is an example of an accounting log entries:

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

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

Syslog

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

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

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

Syslog messages are categorized into 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 logged to a local file or server.

Logging Levels

The Cisco Nexus 1000VE supports the following logging levels:

- 0—emergency
- 1—alert
- 2—critical
- 3—error

- 4—warning
- 5—notification
- 6—informational
- 7—debugging

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

Enabling Logging for Telnet or SSH

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

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

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

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

The **no logging console** command shown in Example 2-4 does the following:

- Disables console logging
- Enabled by default

Example 2-4 no logging console Command

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

The **terminal monitor** command shown in Example 2-5 does the following:

- Enables logging for Telnet or SSH
- Disabled by default

Example 2-5 terminal monitor Command

```
switch# terminal monitor
```

For more information about configuring syslog, see the *Cisco Nexus 1000VE System Management Configuration Guide*.

Installation

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

Isolating Installation Problems

This section explains how to isolate possible installation problems.

Verifying Your VMware License Version

Before you begin to troubleshoot any installation issues, you should verify that your ESX server has the VMware Enterprise Plus license that includes the Distributed Virtual Switch feature.

BEFORE YOU BEGIN

Before you begin, you must know or do the following:

- You are logged in to the vSphere web client on the ESX server.
- You are logged in to the Cisco Nexus 1000VE CLI in EXEC mode.
- This procedure verifies that your vSphere ESX server uses the VMware Enterprise Plus license. This license includes the Distributed Virtual Switch feature, which allows visibility to the Cisco Nexus 1000VE.
- If your vSphere ESX server does not have the Enterprise Plus license, then you must upgrade your license.

DETAILED STEPS

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

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

The Enterprise Plus licensed features are displayed.

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

- Enterprise Plus license
- Distributed Virtual Switch feature

Step 4 Do one of the following:

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

Host is Not Visible from the Distributed Virtual Switch

Scenario 1

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

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

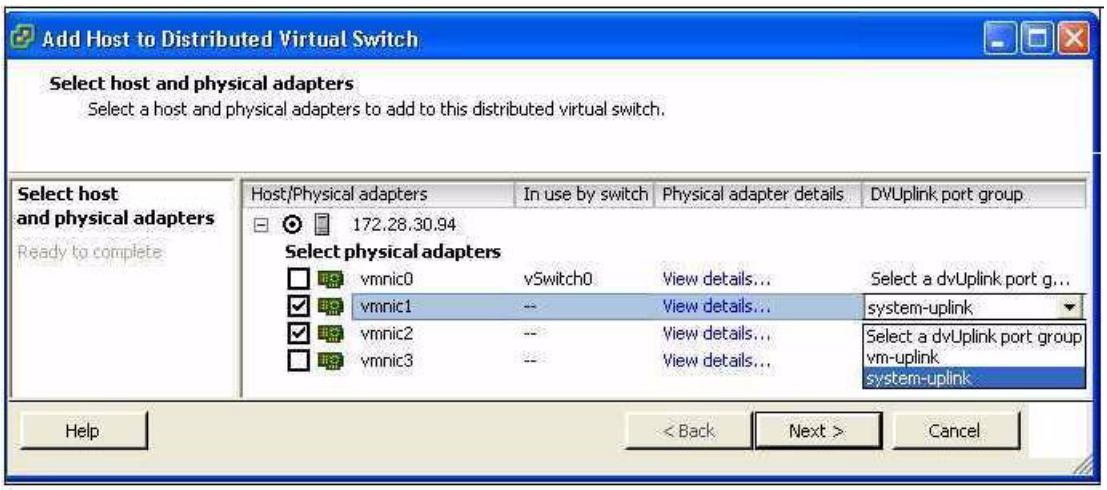

If the hosts and adapters do not appear in this dialog box, you might have the incorrect VMware license installed on your ESX server.

Use the "Verifying Your VMware License Version" procedure on page 3-1 to confirm.

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

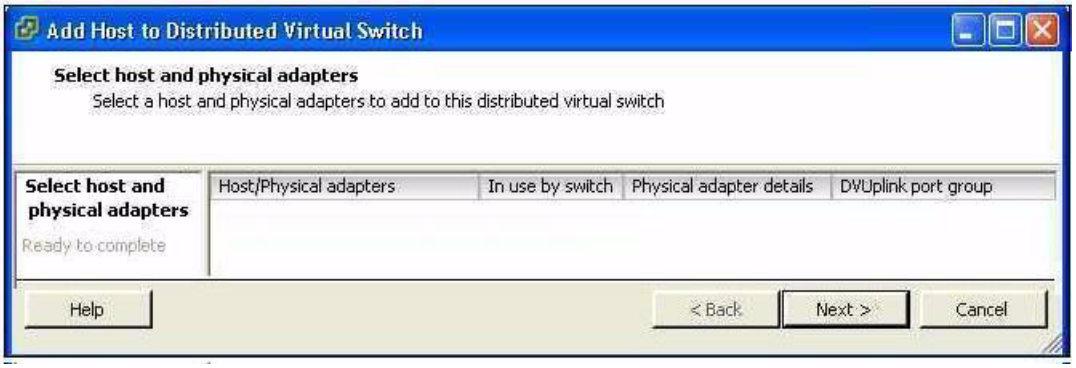

Scenario 2

If wrong map of management port group is configured during VSE deployment, host is not visible from the distributed virtual switch. Complete the following steps to resolve this issue:

Step 1 Log into the plug-in and select the correct data center and switch.

Step 2 Select the host and pNICs.

Step 3 Enter all the parameters.

Step 4 Map to the correct **`Management port group`**.

Click **Install VSE** to complete VSE installation.

Refreshing the vCenter Server Connection

You can refresh the connection between the Cisco Nexus 1000VE and vCenter Server.

Step 1 From the Cisco Nexus 1000VE Connection Configuration mode on the Virtual Supervisor Module (VSM), enter the following command sequence:

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

Step 2 You have completed this procedure.

Improving Performance on the ESX and VM

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

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

Verifying the Domain Configuration

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

Following the installation of the Cisco Nexus 1000VE, make certain that you configure a domain ID. Without a domain ID, the VSM cannot connect to the vCenter Server. Follow these guidelines:

- The domain ID should be a value within the range of 1 to 1023.
- All the control traffic between the VSM and the VSE is carried over the configured control VLAN.
- All the data traffic between the VSM and the VSE is carried over the configured packet VLAN.
- Make sure that the control VLAN and the packet VLAN are allowed on the port in the upstream switch to which the physical NIC of the host hosting the VSM and VSE VM are connected.

Verifying the Port Group Assignments for a VSM VM Virtual Interface

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

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

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

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

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

Step 1 Enter the **show mac address-table dynamic interface vlan** *control-vlan* command on the upstream switch.

Expected output: the network adapter1 MAC address of the VSM VM.

Step 2 Enter the **show mac address-table dynamic interface vlan** *mgmt-vlan* command on the upstream switch.

Expected output: the network adapter2 MAC address of the VSM VM.

Step 3 Enter the **show mac address-table dynamic interface vlan** *packet-vlan* command on the upstream switch.

Expected output: the network adapter3 MAC address of the VSM VM.

Verifying VSM and vCenter Server Connectivity

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

- Make sure that domain parameters are configured correctly.
- Make sure the Windows VM hosting the vCenter Server has the following ports open.
  - Port 80
  - Port 443
- Try reloading the VSM if after verifying the preceding steps, the connect still fails.
- Verify that the provided login and password are correct, if the **remote username** CLI was used to provide credentials to login to vCenter from the VSM.
- If the **register-plugin** CLI was used to register the VSM extension key with the vCenter, check if the VSM extension is created by the vCenter Server by pointing your web browser to *https://your-virtual-center/mob/*, and choosing **Content > Extension Manager**.

Step 1 Ensure that the Nexus 1000VE VSM VM network adapters are configured properly.

Step 2 Make sure that the Windows VM machine hosting the vCenter Server has the following ports open:

- Port 80
- Port 443

Step 3 Ping the vCenter Server from the Cisco Nexus 1000VE VSM.

Step 4 Ensure that the VMware VirtualCenter Server service is running.

Recovering the Network Administrator Password

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

Managing Extension Keys

This section includes the following topics:

Known Extension Problems and Resolutions

Use the following table to troubleshoot and resolve known problems with plug-ins and extensions.

Problem	Resolution
The extension does not show up immediately in the plugin.	Close the VI client and then open the VI client again.
You cannot delete the extension from the VI client.	If you delete the extension using Manager Object Browser (MOB), the VI client screen might not refresh and indicate that the extension was deleted. In this case, close the VI client and then open the VI client again.

Resolving a Plug-In Conflict

If you see "The specified parameter was not correct," when Creating a Nexus 1000VE plug-in on vCenter Server, you have tried to register a plug-in that is already registered.

Use the following procedure to resolve this problem.

Step 1 Make sure that you are using the correct cisco_nexus1000ve_extension.xml file.

Step 2 Make sure that you have refreshed your browser because it caches this file and unless refreshed it might cache obsolete content with the same filename.

Step 3 Follow the steps described in the "Verifying Extension Keys" section on page 3-7 to compare the extension key installed on the VSM with the plug-in installed on the vCenter Server.

Finding the Extension Key on the Cisco Nexus 1000VE

You can find the extension key on the Cisco Nexus 1000VE.

BEFORE YOU BEGIN

- Log in to the Cisco Nexus 1000VE VSM CLI in EXEC mode.
- Know that you can use the extension key in the "Unregistering the Extension Key in the vCenter Server" section on page 3-11.

DETAILED STEPS

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

show vmware vc extension-key

Example:

```
switch# show vmware vc extension-key
Extension ID: Cisco_Nexus_1000Ve_1935882621
switch#
```

Finding the Extension Key Tied to a Specific DVS

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

Step 1 From the vSphere Client, choose the DVS whose extension key you want to find.

Step 2 Click the **Summary** tab.

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

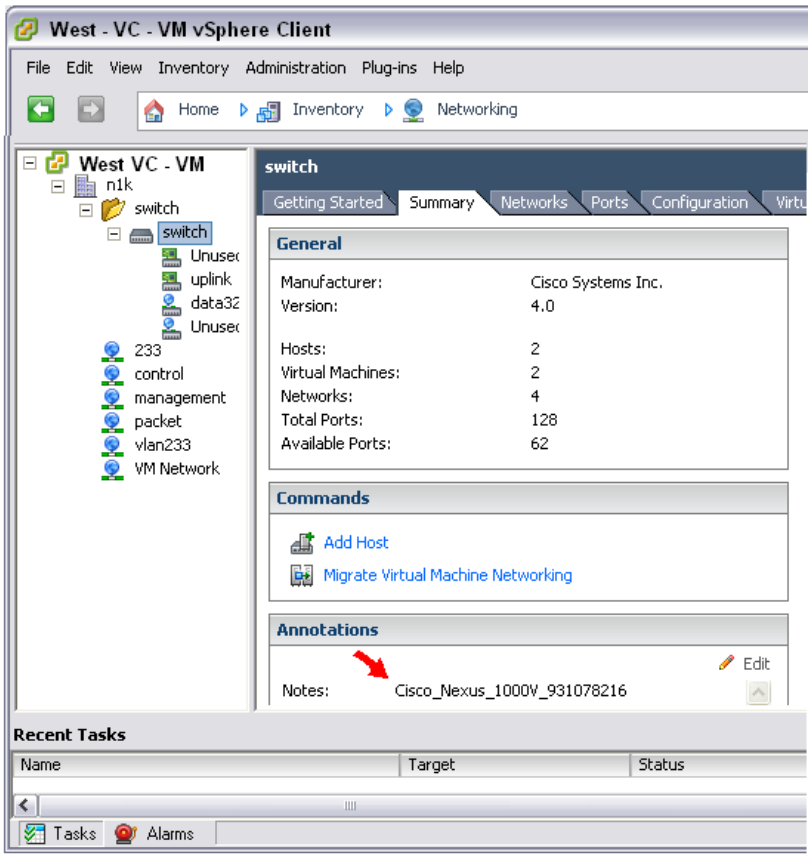

Verifying Extension Keys

You can verify that the Cisco Nexus 1000VE and vCenter Server are using the same extension key.

DETAILED STEPS

Step 1 Find the extension key used on the Cisco Nexus 1000VE using the "Finding the Extension Key on the Cisco Nexus 1000VE" section on page 3-6.

Step 2 Find the extension key used on the vCenter Server using the "Finding the Extension Key Tied to a Specific DVS" section on page 3-7.

Step 3 Verify that the two extension keys (the one found in Step 1 with that in Step 2) are the same.

Recreating the Cisco Nexus 1000VE Installation

You can re-create the complete Cisco Nexus 1000VE configuration in the event of a persistent problem that cannot be resolved using any other workaround.

Flowchart: Re-creating the Cisco Nexus 1000VE Installation

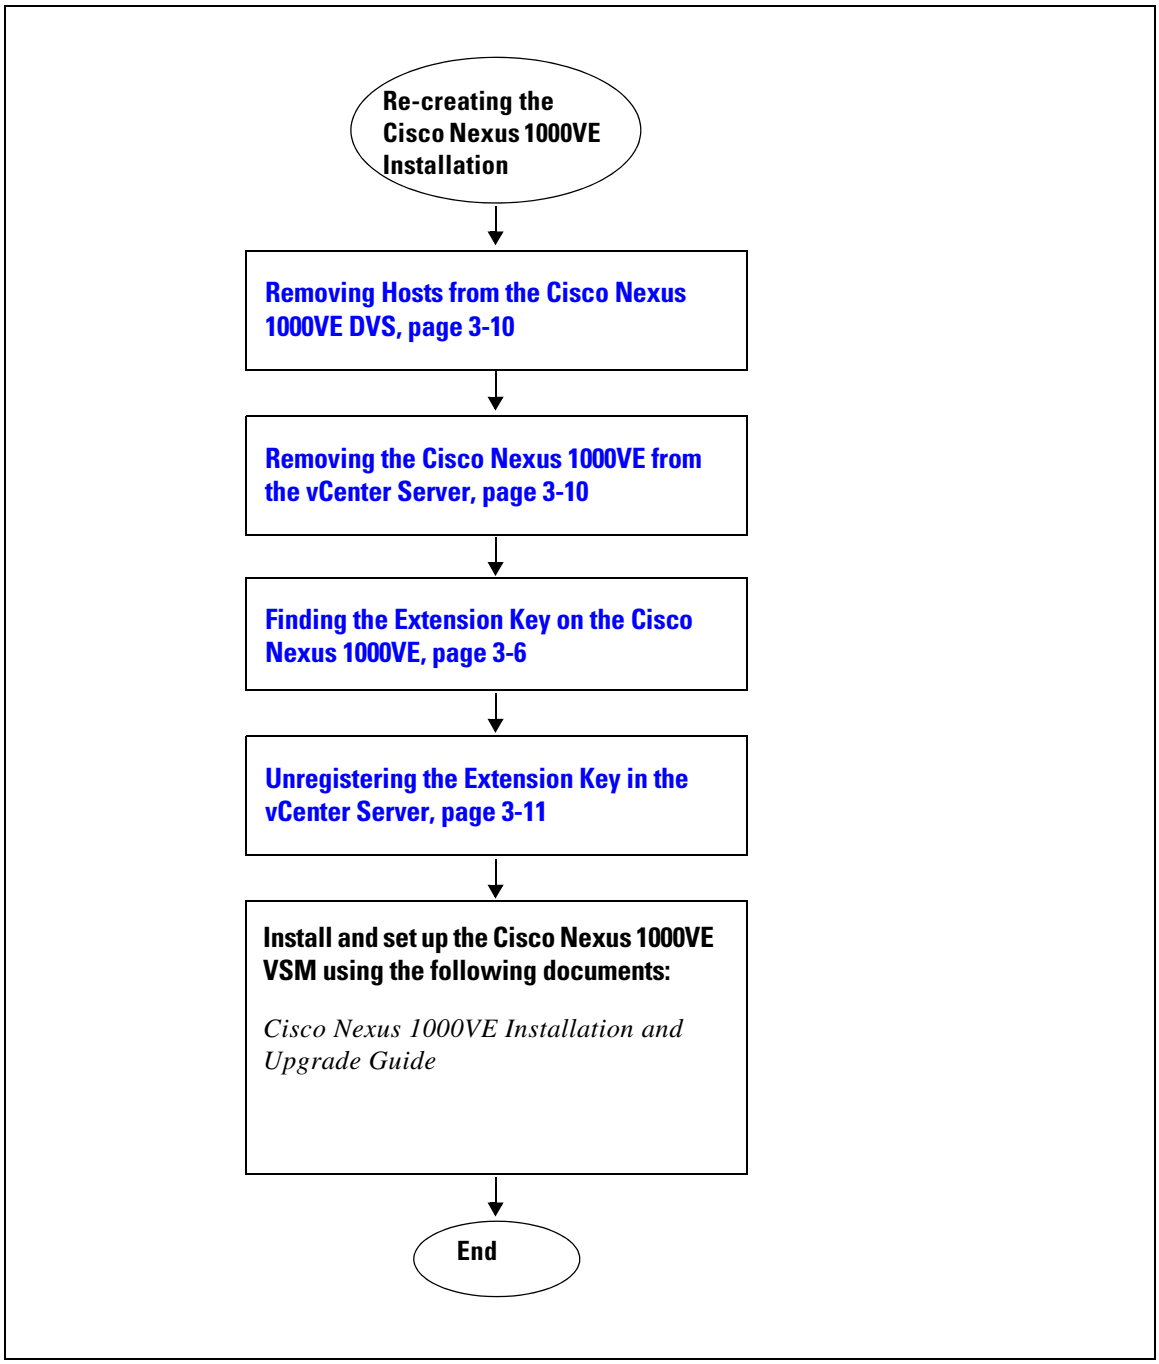

Removing Hosts from the Cisco Nexus 1000VE DVS

You can remove hosts from the Cisco Nexus 1000VE DVS.

BEFORE YOU BEGIN

- Log in to vSphere Client.
- Know the name of the Cisco Nexus 1000VE DVS to remove from vCenter Server.

DETAILED STEPS

Step 1 From vSphere Client, choose **Inventory > Networking**.

Step 2 Choose the DVS for the Cisco Nexus 1000VE and click the **Hosts** tab.

The Host tab opens.

Step 3 Right-click each host, and choose **Remove from Distributed Virtual Switch**.

The hosts are now removed from the DVS.

Removing the Cisco Nexus 1000VE from the vCenter Server

You can remove the Cisco Nexus 1000VE DVS from vCenter Server.

BEFORE YOU BEGIN

- Log in to the VSM CLI in EXEC mode.

DETAILED STEPS

Step 1 From the Cisco Nexus 1000VE VSM, use the following commands to remove the DVS from the vCenter Server.

a. **config t**

b. **svs connection vc**

c. **no vmware dvs**

```
Example:
switch# conf t
switch(config)# svs connection vc
switch(config-svs-conn)# no vmware dvs
switch(config-svs-conn)#
```

The DVS is removed from vCenter Server.

Step 2 You have completed this procedure.
Return to the "Flowchart: Re-creating the Cisco Nexus 1000VE Installation" section on page 3-9.

Unregistering the Extension Key in the vCenter Server

You can unregister the Cisco Nexus 1000VE extension key in vCenter Server.

BEFORE YOU BEGIN

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

- Open a browser window.
- Paste the extension key name into the vCenter Server MOB. You should already have the extension key found in the "Finding the Extension Key on the Cisco Nexus 1000VE" section on page 3-6.
- After unregistering the extension key in vCenter Server, you can start a new installation of the Cisco Nexus 1000VE VSM software.

DETAILED STEPS

Step 1 Point your browser to the following URL:

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

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

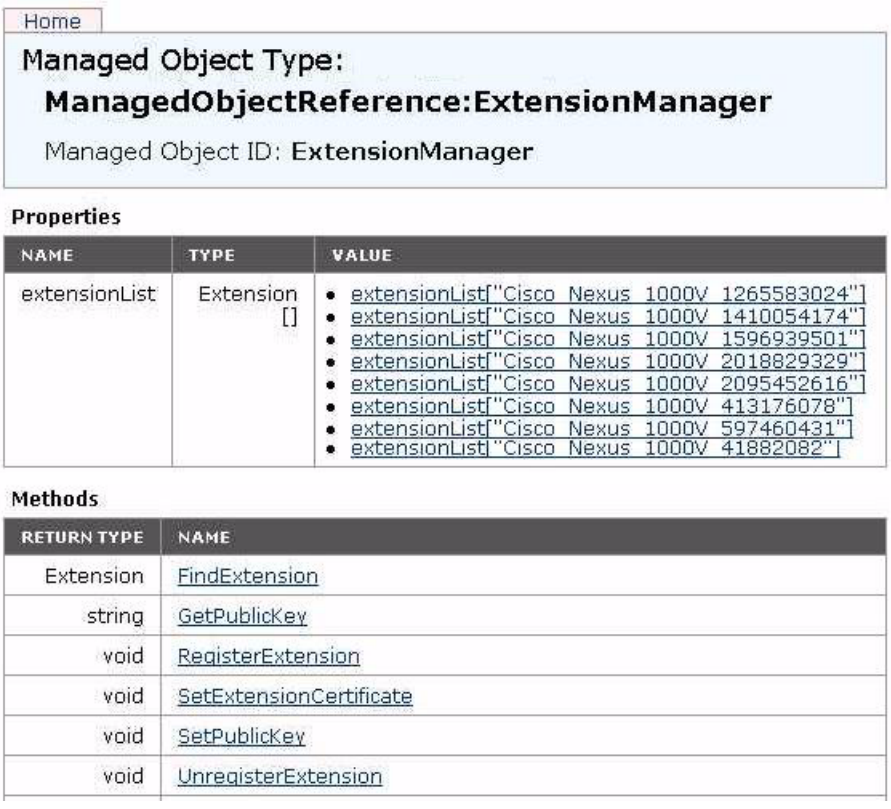

Step 2 Click **Unregister Extension**.

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

A dialog box opens to unregister the extension.

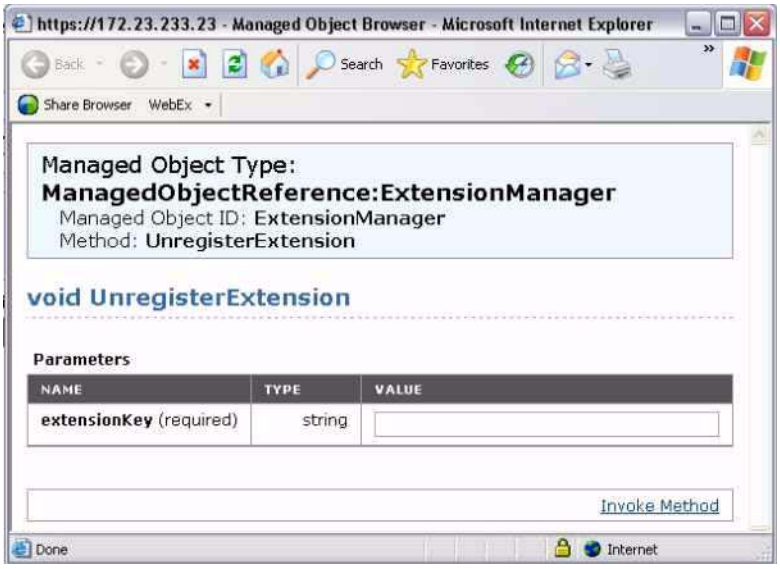

Step 3 In the value field, paste the extension key that you found in the "Finding the Extension Key on the Cisco Nexus 1000VE" section on page 3-6, and then click **Invoke Method**.

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

Step 4 You have completed this procedure.

Return to "Flowchart: Re-creating the Cisco Nexus 1000VE Installation" section on page 3-9.

Problems with the Cisco Nexus 1000VE Installation Management Center

The following are possible problems and their solutions.

.

Symptom	Problem	Recommended Action
The VSE is missing on the VSM.	• The VSE installer application finishes successfully. • The host on which the VSE is deployed is added to N1KVE DVS on vCenter, but does not display when the **show module** command is entered on the VSM.	• Verify that the VSE VM is powered on and the Nexus1000v service is active on it (systemctl status Nexus1000v) • Verify the ping connectivity between VSM and VSE. • Check the vCenter MOB for opaque data propagation.
Configuration file issue.	After loading the previously saved configuration file, the installation application does not complete.	• Check the configuration file for appropriate contents. **Note** You might need to change a few of the fields before reusing the previously saved files. • Check if a VM with the same name already exists in the DC. This can be identified by reviewing the Virtual Machine field in the configuration file.
The VSE loses its management IP after reboot and gets disconnected from the VSM.	When VSE is installed using a static IP pool from vCenter with multiple DNS server addresses, after reboot VSE gets disconnected from VSM.	• Check */var/log/messages* on VSE and check if there is any error indicating invalid DNS server address <IP1>,<IP2> • Try following workarounds: – Manually configure multiple DNS addresses in the VSE using the **nmtui** command as root. – Use only a single DNS server in the IP pool in vCenter and then re-deploy the VSE.

Cisco Nexus 1000VE Manager vCenter Plug-in

This chapter describes how to identify and resolve problems that relate to the Cisco Nexus 1000VE Manager vCenter Plug-in functionality.

This chapter includes the following topics:

About Cisco Nexus 1000VE Manager vCenter Plug-in

The Cisco Nexus 1000VE is a software-based Layer 2 switch for the virtualized server environments that are running VMware ESXi. The Cisco Nexus 1000VE provides a consistent networking experience across the physical and the virtual environments. It consists of two components: the Virtual Service Engine (VSE) that is running as a Virtual Machine Form Factor and a Virtual Supervisor Module (VSM) that manages the networking and policies for the virtual machines.

Cisco Nexus 1000VE NX-OS Release 5.2(1)SV5(1.1) is bundled with the Cisco N1KVE Manager vCenter Plug-in supported on the vSphere Web Clients.

For more information on installing the N1000VE vCenter plug-in, see Cisco Nexus 1000VE Installation, Migration, and Upgrade Guide.

The VMware vSphere Web Client enables you to connect to a VMware vCenter Server system to install a Cisco Nexus 1000VE through a browser. After the vCenter plug-in installation, the Cisco Nexus 1000VE Manager appears as a icon in the home page of vSphere web client, which displays a new tab that assists you to install the virtual service engine, uninstall and migrate from Cisco Nexus 1000V to Cisco Nexus 1000VE.

With the Cisco Nexus 1000VE Manager vCenter Plug-in, the server administrators can deploy new Virtual Service Engine (VSE) on vmware ESXi hosts or uninstall previously deployed VSEs. It also helps administrators to migrate existing Nexus 1000V networks to Cisco Nexus 1000VE.

Prerequisites for VMware vSphere Web Client

These are the prerequisites to install and configure the vCenter Plug-in functionality on the Cisco Nexus 1000VE:

- VMware vCenter Server 6.0 and/or later release.
- VMware vCenter Web Client 6.0 and/or later. The Cisco Nexus 1000VE Manager vCenter Plug-in is not supported with web client below 6.0.
- vSphere Web Client requires the Adobe Flash Player version 11.1.0 or later to be installed.
- Make sure that Cisco Nexus 1000VE Release 5.2(1)SV5(1.1) is installed and configured to a vCenter.

Troubleshooting

Problems with Cisco Nexus 1000VE vCenter Plugin Icon Display in vSphere web client

After the Cisco Nexus 1000VE vCenter plug-in is registered with vCenter, the Cisco Nexus 1000VE Manager icon should be displayed in the home screen under **Operations and Policies**.

If the Cisco Nexus 1000VE Manager icon is not displayed, do the following:

Step 1 Log out and login to the Vmware vsphere client.

Step 2 Verify if the Plugin extension has successfully registered with vCenter. To verify, navigate to **vCenter mob > content > Extension Manager > Properties**.

Under the **Properties** tab, find the plugin extension with the value **com.cisco.plugin.n1kveui and com.cisco.n1kve**. If it does not exist, then it indicates that the extension did not register with vCenter.

Step 3 Check if the http server is reachable. If it is reachable, then check for the zip file (n1kve-vcenter-plugin-1.0.1.zip) in the Http server (as specified in the server url in the N1kve plugin extension).

To check the URL, do the following:

a. Navigate to **vCenter mob > content > Extension Manager >Methods > Find Extensti**on and search for value com.cisco.plugin.n1kveui).

Note If the web server is Https then the serverThumbprint value in the N1kve plugin extension is must and matches with web sever.

Step 4 Check whether the plugin folder is downloaded to the vCenter from the Http/Https server.

To verify go to the following path in vCenter Server.

For Windows vCenter

C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity

For VMware vCenter Server Appliance (vCSA):

/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity

In this path a folder with name com.cisco.plugin.n1kveui- 1.0.1 (<extensionKey-extensionVersion>) should exists, if not present, follow these steps:

a. Download the plugin zip file. In a Web browser enter the VSM IP address in the address bar and download n1kve-vcenter-plugin-1.0.2.zip to a local location.

b. Unzip the plugin file.

c. Log into VMware vCenter and navigate to the folder /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity.

d. Copy the unzipped folder and rename it to the following naming convention com.cisco.plugin.n1kveui-1.0.2

e. Restart the Web client services.

f. Log into the vCenter, the plugin should be visible.

then log out and login from web client and confirm once again. If it continues to be unseen, then uninstall the extension keys > com.cisco.plugin.n1kveui and com.cisco.n1kve from vCenter MOB, and re-register the Cisco Nexus 1000VE manager as per the Cisco Nexus 1000VE Installation, Migration, and Upgrade Guide.

Step 5 Verify whether the version of vCenter where the plug-in is registered and the supported versions are 6.0 and later.

Problems with VSE Installation

Error: The operation is not supported on the object:

This error is displayed in the recent task list whenever the user selects the invalid VDS for the VSE installation.

Problems with Removal of VSE

Error: An internal error has occurred - Error # 1009

Tis error is displayed when you select outside VDS and try to remove the VSEs from the hosts

Reason: The outside VDS will be removed from the VCenter during the removal process if it doesn't contains any hosts. Therefore, after the VSE VMs are removed from the hosts, it tries to get the vCenter information based on the outside VDS.

Monitoring vCenter Tasks and Events

Cisco Nexus 1000VE Manager vCenter Plugin tool operations are tracked using 4 different vCenter tasks.

Table 4-1 vCenter Tasks and Events

vCenter Parent Task Name	Action
Install Cisco N1KVE VSE	VSE installation
Uninstall Cisco N1KVE VSE	VSE removal
Configuration Migration from NKV to N1KVE	Migrate supported configurations from Cisco Nexus 1000V to Cisco Nexus 1000VE.
Migration from N1KV to N1KVE	Migrate Hosts from Cisco Nexus 1000V to Cisco Nexus 1000VE. This includes all Veth ports, pnics, and vmknic migration.

For any internal errors as part of the Installation, uninstallation, or Migration process, the corresponding parent task in the vCenter will fail. This can be seen in the recent task bar with appropriate status. To get more information on the event that failed as part of the parent task go to

Vsphere webclient > Networking > N1KVE vDS name > Monitor > Tasks & Events > Tasks.

Select the Parent task name in the Task lists pane and check the related events listed below to know the exact reason for failure.

Gathering Information for Technical Support

Collecting the vsphere client virgo logs

To collect the vSphere client, do the following.

Step 1 Logon to the VCenter Server

Step 2 Find the log file named “vsphere_client_virgo.log” in the below path.

For Windows vCenter:

C:\ProgramData\VMware\vCenterServer\logs\vsphere-client\logs

For VMware vCenter Server Appliance (vCSA):

/var/log/vmware/vsphere-client/logs

Collecting Tasks and events for Plug-in operations

To collect logs, please refer section 4.3 above and capture the failed logs under Tasks & Events corresponding to the Nexus 1000VE vDS.

Generating a Log Bundle

You can collect the diagnostic information for VMware vCenter Server by collecting vSphere log files into a single location.

Step 1 Log in to the Windows server where the VMware vCenter Server is installed.

Step 2 Choose **Start > All Programs > VMware > Generate vSphere Web Client Log Bundle**.

You can use this step to generate the vSphere Web Client log bundles even when you are not able to connect to the vCenter Server using the vSphere Client. The log bundle is generated as a .zip file. See VMware documentation *Collect vSphere Log Files* for more information about collecting the log files.

Note Currently the login to the vCenter Plug-in is available through the administrator account only.

Licenses

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

Information About Licenses

The name for the Cisco Nexus 1000V license package is NEXUS1000VE_LAN_SERVICES_PKG and the version is 3.0. By default, 1024 licenses are installed with the Virtual Supervisor Module (VSM). These default licenses are valid for 60 days. You can purchase permanent licenses that do not expire.

Licensing is based on the number of CPU sockets on the ESX servers attached as Virtual Ethernet Modules (VSE) to the VSM.

A module is either licensed or unlicensed:

- Licensed module—A VSE 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 VSE 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 VSE is unlicensed, the virtual Ethernet ports correspond to the virtual machines (VMs) that are kept down and are shown as unlicensed.

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

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

Contents of the License File

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

```
switch# show license file sample.lic
sample.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT NEXUS1000VE_LAN_SERVICES_PKG cisco 1.0 permanent 16 \
        HOSTID=VDH=8449368321243879080 \
        NOTICE="<LicFileID>sample.lic</LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=34FCB2B24AE8
```

The host ID that appears in the license file must match that shown on the VSM. To verify the match, use the **show license host-id** command. See Example 5-3 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 information in this checklist:

- Make sure that the name of the license file has fewer than 32 characters by using the **show license usage** command. See Example 5-1 on page 5-5.
- Make sure that no other license file with the same name is installed on the VSM by using the **show license usage** command. See Example 5-1 on page 5-5. If there is a license file with the same name, rename your new license file to something else.
- Do not edit the contents of the license file. If you have already done so, 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 using the **show license host-id** command and the **show license file** command. See Example 5-3 on page 5-6 and Example 5-4 on page 5-6.

Problems with Licenses

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

Symptom	Possible Causes	Solution
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: `vse Unlicensed`	A license could not be obtained for the server (VSE) where the virtual machine resides.	1. Verify the license usage. **show license usage** *license_name* See Example 5-1 on page 5-5. 2. Determine the number of licenses required by viewing the sockets installed on the VSE. **show module vse license-info** See Example 5-8 on page 5-7. 3. Contact your Cisco Customer Support Account Team to acquire additional licenses.
You see the following system message: `PLATFORM-2-PFM_LIC_WARN_EXP Syslog` `2008 Dec 19 22:28:30 N1KV %PLATFORM-2-PFM_LIC_WARN_EXP: WARNING License for VSEs is about to expire in 1 days! The VSEs' 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.`	The default or evaluation license in use is about to expire. **Note** Permanent licenses do not expire.	1. Verify the license usage. **show license usage** *license_name* See Example 5-1 on page 5-5. 2. Contact your Cisco Customer Support Account Team to acquire additional licenses.

Symptom	Possible Causes	Solution
You see the following system message: `%LICMGR-2-LOG_LIC_USAGE: Feature NEXUS1000VE_LAN_SERVICES_PKG is using 17 licenses, only 16 licenses are installed.`	More licenses are being used than are installed.	1. Verify the license usage. **show license usage** *license_name* See Example 5-1 on page 5-5. 2. Contact your Cisco Customer Support Account Team to acquire additional licenses.
VSEs fails to acquire licenses even though the **show license usage** command shows there are enough licenses available. The following syslog messages are seen: `2014 Jun 7 20:15:36 vsm-demo LICMGR-3-LOG_LIC_CHECKOUT_FAIL_BAD_CLOCK: License checkout failed for feature NEXUS1000VE_LAN_SERVICES_PKG(VSE 3 - Socket 1(1.0)) because system clock has been set back. Please set the clock to the correct value.` `2014 Jun 7 20:15:36 vsm-demo VSE_MGR-2-VSE_MGR_UNLICENSED: License for VSE 3 could not be obtained. Please contact your Cisco account team or partner to purchase Licenses or downgrade to Essential Edition. To activate your purchased licenses, click on www.cisco.com/go/license.`	The clock has been changed back manually or through NTP, which has invalidated evaluation licenses. The problem is seen even if there are enough permanent licenses available to license the VSEs as long as evaluation licenses are present. You can look for the following syslog message to find the time when the clock changed: `2014 Jun 7 20:15:24 vsm-demo VSE_MGR-5-VSE_MGR_CLOCK_CHANGE : Clock setting has been changed on the system. Please be aware that, in Advanced edition, clock changes will force a recheckout of all existing VSE licenses. During this recheckout procedure, licensed VSEs which are offline will lose their licenses.`	1. Undo the clock change using the **clock set** command or uninstall all evaluation licenses using the **clear license** command. 2. Ensure there are enough permanent licenses available before uninstalling evaluation licenses. 3. Verify that the modules are licensed using the **show module VSE license-info** command.

License Troubleshooting Commands

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

Command	Purpose
show module vse license-info	Displays the VSE license information including the license type, license status, license version, and socket count. See Example 5-8 on page 5-7.
show license usage [*license_name*]	Displays information about the licenses and where they are used. If displayed for a specific license, indicates VSE and socket information. See Example 5-1 on page 5-5.

Command	Purpose
show interface veth	Displays the messages logged about port profile events within the Cisco Nexus 1000V. See Example 5-2 on page 5-6.
show license host-id	Displays the serial number for your Cisco Nexus 1000V license. See Example 5-3 on page 5-6.
show license file	Displays the contents of a named license file. See Example 5-4 on page 5-6.
svs license transfer src-vse *VSE no* **license_pool**	Transfers the licenses from a VSE to the license pool. See Example 5-5 on page 5-6.
show license brief	Displays the version and license count information for each license file. See Example 5-6 on page 5-6.
show switch edition	Displays the switch edition, advanced feature status, license expiry, module and virtual Ethernet scale. Example 5-7 on page 5-7.

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

EXAMPLES

Example 5-1 show license usage** license_name **Command

```
switch# show license usage NEXUS1000VE_LAN_SERVICES_PKG
--------------------------------------
Feature Usage Info
--------------------------------------
       Installed Licenses :    10
            Eval Licenses :   0
   Max Overdraft Licenses :    16
Installed Licenses in Use :    4
Overdraft Licenses in Use :    0
     Eval Licenses in Use :    0
       Licenses Available :   22
--------------------------------------
Application
--------------------------------------
VSE 3 - Socket 1
VSE 3 - Socket 2
VSE 4 - Socket 1
VSE 4 - Socket 2
--------------------------------------
switch#
```

Example 5-2 show interface vethernet Command

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

Example 5-3 show license host-id Command

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

Example 5-4 show license file Command

```
switch# show license file sample.lic
sample.lic:
        SERVER this_host ANY
        VENDOR cisco
        INCREMENT NEXUS1000VE_LAN_SERVICES_PKG cisco 3.0 permanent 16 \
        HOSTID=VDH=8449368321243879080 \
        NOTICE="<LicFileID>sample.lic</LicFileID><LicLineID>0</LicLineID> \
        <PAK>dummyPak</PAK>" SIGN=34FCB2B24AE8

switch#
```

Example 5-5 svs license transfer src-vse *vse no* license_pool Command

```
switch# svs license transfer src-vse 3 license_pool
switch#
```

Example 5-6 show license brief Command

```
switch# show license brief
NOTE: * is UPGRADE FILE
-----------------------------------------------------------------------------------------
File Name Feature Name Version Count Expiry
-----------------------------------------------------------------------------------------
eval.lic NEXUS1000VE_LAN_SERVICES_PKG 1.0 17 3-nov-2014
eval0715.lic NEXUS1000VE_LAN_SERVICES_PKG 3.0 17 15-jul-2015

show switch edition (purpose: Displays the switch edition, advanced feature status,
license expiry and module and veth scale)
```

Example 5-7 show switch edition Command

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

Feature Status
Name State Licensed In version
------------------------------------------------------
cts enabled Y 1.0
dhcp-snooping disabled Y 1.0
vxlan-gateway enabled Y 1.0
bgp disabled Y 3.0
bpduguard disabled Y 3.0

License Status
Edition Available In Use Expiry Date
-----------------------------------------
Advanced 17 0 03 Nov 2014

Scale Support
Edition Modules Virtual Ports
---------------------------------------
Advanced 256 12288
```

Example 5-8 show module vse license-info Command

```
n100v# show module vse license-info
Licenses are Sticky
Mod Socket Count License Usage Count License Version License Status
--- ------------ ------------------- --------------- --------------
3 2 2 3.0 licensed
4 2 2 3.0 licensed
```

High Availability

This chapter describes how to identify and resolve problems related to high availability, and includes the following sections:

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 if a failure occurs:

- Redundancy—Redundancy at every aspect of the software architecture.
- Isolation of processes—Isolation between software components to prevent a failure within one process disrupting other processes.
- Restartability—Most system functions and services are isolated so that they can be restarted independently after a failure while other services continue to run. In addition, most system services can perform stateful restarts, which allow the service to resume operations transparently to other services.
- Supervisor stateful switchover—Active/standby dual supervisor configuration. The state and configuration remain constantly synchronized between two Virtual Supervisor Modules (VSMs) to provide a seamless and statefu1 switchover if a VSM failure occurs.

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

- Virtual Ethernet Modules (VSEs) running within virtual machines (VMs) on virtualization servers. These VSEs are represented as modules within the VSM.
- A remote management component, such as VMware vCenter Server.
- One or two VSMs running within virtual machines (VMs).

System-Level High Availability

The Cisco Nexus 1000V supports redundant VSM virtual machines—a primary and a secondary—running as an HA pair. Dual VSMs operate in an active/standby capacity in which only one of the VSMs is active at any given time, while the other acts as a standby backup. The state and configuration remain constantly synchronized between the two VSMs to provide a statefu1 switchover if the active VSM fails.

Problems with High Availability

Symptom	Possible Causes	Solution
The active VSM does not see the standby VSM.	MAC addresses mismatch. • Check that the peer VSM MAC addresses that are learned by the active VSM by using the **show system redundancy status** command.	Confirm that the standby VSM MAC addresses are correctly learned by the active VSM. 1. Compare the standby VSM MAC addresses with the output MAC addresses by using the **show system redundancy status** command on the active VSM. 2. If the compared MAC addresses are different, use the **peer-sup mac-addresses clear** command to clear the stale MAC addresses that are learned by the active VSM.
The active VSM does not see the standby VSM.	Roles are not configured properly. • Check the role of the two VSMs by using the **show system redundancy status** command.	1. Confirm that the roles are the primary and secondary role, respectively. 2. If needed, use the **system redundancy role** command to correct the situation. 3. Save the configuration if roles are changed.
	Network connectivity problems. • Check that the control and management VLAN connectivity between the VSM at the upstream and virtual switches.	If network problems exist, do the following: 1. From vSphere Client, shut down the VSM, which should be in standby mode. 2. From vSphere Client, bring up the standby VSM after network connectivity is restored.

Symptom	Possible Causes	Solution
The active VSM does not complete synchronization with the standby VSM.	Version mismatch between VSMs. • Check that the primary and secondary VSMs are using the same image version by using the **show version** command.	If the active and the standby VSM software versions differ, reinstall the secondary VSM with the same version used in the primary.
	Fatal errors during gsync process. • Check the gsyncctrl log using the **show system internal log sysmgr gsyncctrl** command and look for fatal errors.	Reload the standby VSM using the **reload module** *module-number* command, where *module-number* is the module number for the standby VSM.
	• The VSM has connectivity only through the management interface. • Check the output of the **show system internal redundancy info** command and verify if the *degraded_mode* flag is set to *true*.	Check control VLAN connectivity between the primary and the secondary VSMs.
The standby VSM reboots periodically.	The VSM has connectivity only through the management interface. • Check the output of the **show system internal redundancy info** command and verify that the *degraded_mode* flag is set to true.	Check the control VLAN connectivity between the primary and the secondary VSMs.
	The VSMs have different versions. Enter the **debug system internal sysmgr all** command and look for the active_verctrl entry that indicates a version mismatch, as the following output shows: `2009 May 5 08:34:15.721920 sysmgr: active_verctrl: Stdby running diff version- force download the standby sup.`	Isolate the standby VSM and boot it. Use the **show version** command to check the software version in both VSMs. Install the image matching the active VSM on the standby.

Symptom	Possible Causes	Solution
Active-Active detected and resolved	When control and management connectivity between the active and the standby goes down for 6 seconds, the standby VSM transitions to the active state. Upon restoration of control and management connectivity, both VSMs detect an active-active condition.	1. Once the system detects active-active VSMs, one VSM is automatically reloaded based on various parameters such as VSEs attached, vCenter connectivity, last configuration time, and last active time. 2. To see any configuration changes that are performed on the rebooted VSM during the active-active condition, enter the **show system internal active-active remote accounting logs** CLI command on the active VSM.
VSM Role Collision	If another VSM is configured/provisioned with the same role (primary or secondary) in the system, the new VSM collides with the existing VSM. The **show system redundancy info** command displays the MAC addresses of the VSM(s) that collide with the working VSM.	If the problems exist, do the following: 1. Enter the **show system redundancy status** command on the VSM console. 2. Identify the VSM(s) that owns the MAC addresses that are displayed in the output of the **show system redundancy status** command. 3. Move the identified VSM(s) out of the system to stop role collision.

Symptom	Possible Causes	Solution
Both VSMs are in active mode.	Network connectivity problems. • 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: 1. From vSphere Client, shut down the VSM, which should be in standby mode. 2. From vSphere Client, bring up the standby VSM after network connectivity is restored.
	Different domain IDs in the two VSMs Check the *domain* value by using **show system internal redundancy info** command.	If needed, update the domain ID and save it to the startup configuration. • Upgrading the domain ID in a dual VSM system must be done as follows: – 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 the configuration, and power off the VSM. – Reconnect the isolated VSM and power it on.

High Availability Troubleshooting Commands

This section lists commands that can be used troubleshoot problems related to high availability.

Command	Description
attach module	See Example 6-9attach module Command, page 6-10
reload module	See Example 6-8reload module Command, page 6-10
show cores	Use to list process logs and cores. See Example 6-1show cores Command, page 6-6
show processes [pid pid]	See Example 6-2show processes log [pid pid] Command, page 6-6
show system internal active-active	See Example 6-7show system internal active-active remote accounting logs Command, page 6-10

Command	Description
show system internal redundancy info	See Example 6-4show system internal redundancy info Command, page 6-7
show system internal sysmgr state	See Example 6-5show system internal sysmgr state Command, page 6-8
show system redundancy status	See Example 6-3show system redundancy status Command, page 6-6
show system redundancy status	See Example 6-6show system redundancy status Command, page 6-9

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

Example 6-1 show cores Command

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

Example 6-2 show processes log [pid pid] Command

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

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: nexus-1000ve-mzg.4.0.4.SV1.1.bin
System image version: 4.0(4)SV1(1) S25

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

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

To check redundancy status, use the following commands:

Example 6-3 show system redundancy status Command

```
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, use the following command:

Example 6-4 show system internal redundancy info Command

```
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_unknown_pkts:       0
```

```
Redun Device 1: <-- This device maps to the mgmt interface
  name: ha1
  pdev: ad7b6860
  alarm: true
  mac: ff:ff:ff:ff:ff:ff
  tx_set_ver_req_pkts:   11589
  tx_set_ver_rsp_pkts:   0
  tx_heartbeat_req_pkts: 12
  tx_heartbeat_rsp_pkts: 0
  rx_set_ver_req_pkts:   0
  rx_set_ver_rsp_pkts:   0
  rx_heartbeat_req_pkts: 0
  rx_heartbeat_rsp_pkts: 0 <-- When communication between VSM through the control
interface is interrupted but continues through the mgmt interface, the
rx_heartbeat_rsp_pkts will increase.
  rx_drops_wrong_domain: 0
  rx_drops_wrong_slot:   0
  rx_drops_short_pkt:    0
  rx_drops_queue_full:   0
  rx_drops_inactive_cp:  0
  rx_drops_bad_src:      0
  rx_drops_not_ready:    0
  rx_unknown_pkts:       0
```

To check the system internal sysmgr state, use the following command:

Example 6-5 show system internal sysmgr state Command

```
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.

Watchdog checking is currently disabled.

Watchdog kgdb setting is currently enabled.


        Debugging info:

The trace mask is 0x00000000, the syslog priority enabled is 3.
The '-d' option is currently disabled.
The statistics generation is currently enabled.


        HA info:


slotid = 1    supid = 0
cardstate = SYSMGR_CARDSTATE_ACTIVE .
cardstate = SYSMGR_CARDSTATE_ACTIVE (hot switchover is configured enabled).
Configured to use the real platform manager.
Configured to use the real redundancy driver.
Redundancy register: this_sup = RDN_ST_AC, other_sup = RDN_ST_SB.
EOBC device name: eth0.
Remote addresses:  MTS - 0x00000201/3      IP - 127.1.1.2
```

```
MSYNC done.
Remote MSYNC not done.
Module online notification received.
Local super-state is: SYSMGR_SUPERSTATE_STABLE
Standby super-state is: SYSMGR_SUPERSTATE_STABLE
Swover Reason : SYSMGR_SUP_REMOVED_SWOVER <-- Reason for the last switchover
Total number of Switchovers: 0 <-- Number of switchovers
                              >> Duration of the switchover would be listed, if any.

        Statistics:

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

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

Example 6-6 show system redundancy status Command

```
switch# show system redundancy status
Redundancy role
---------------
administrative: secondary
operational: secondary
Redundancy mode
---------------
administrative: HA
operational: HA
This supervisor (sup-2)
-----------------------
Redundancy state: Active
Supervisor state: Active
Internal state: Active with HA standby
Other supervisor (sup-1)
------------------------
Redundancy state: Standby
Supervisor state: HA standby
Internal state: HA standby
WARNING! Conflicting sup-2(s) detected in same domain
-----------------------------------------------------
MAC Latest Collision Time
00:50:56:97:02:3b 2012-Sep-11 18:59:17
00:50:56:97:02:3c 2012-Sep-11 18:59:17
00:50:56:97:02:2f 2012-Sep-11 18:57:42
00:50:56:97:02:35 2012-Sep-11 18:57:46
00:50:56:97:02:29 2012-Sep-11 18:57:36
00:50:56:97:02:30 2012-Sep-11 18:57:42
00:50:56:97:02:36 2012-Sep-11 18:57:46
00:50:56:97:02:2a 2012-Sep-11 18:57:36

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

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

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

Example 6-7 show system internal active-active remote accounting logs Command

```
switch# show system internal active-active remote accounting logs
```

To reload a module, use the following command:

Example 6-8 reload module Command

```
switch# reload module 2
```

This command reloads the secondary VSM.

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

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

Example 6-9 attach module Command

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 VSE Modules

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

Information About Modules

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

The Cisco Nexus 1000VE implementation has two parts:

- Virtual Supervisor Module (VSM)—Control software of the Cisco Nexus 1000VE distributed virtual switch. It runs on a virtual machine (VM) and is based on NX-OS software.
- Virtual Service Engine (VSE)—Part of the Cisco Nexus 1000VE that actually switches data traffic. It runs as a VM on a VMware ESX host. Several VSEs are controlled by one VSM. All the VSEs that form a switch domain should be in the same virtual data center as defined by VMware VirtualCenter.

Troubleshooting a Module Not Coming Up on the VSM

This section includes the following topics:

Guidelines for Troubleshooting Modules

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

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

  For more information, see the *Cisco Nexus 1000VE Compatibility Information*.
- To verify network connectivity between the VSM and vCenter Server, ping the IP address of vCenter Server. If you are using a domain name service (DNS) name, use the DNS name in the ping. If a ping to vCenter Server fails, check to see if you can ping the gateway. Otherwise, check the mgmt0 interface configuration settings.
- Make sure that the firewall settings are OFF on the vCenter Server. If you want the firewall settings, and check to see if these ports are open:
  - Port 80
  - Port 443
- If you see the following error, verify that the VSM extension was created from vCenter Server:
  - ERROR: [VMware vCenter Server 4.0.0 build-150489]
    Extension key was not registered before its use

  To verity that the extension or plugin was created, see the "Finding the Extension Key on the Cisco Nexus 1000VE" section on page 3-6.

  For more information about extension keys or plugins, see the "Managing Extension Keys" section on page 3-5.
- If you see the following error, see the "Checking the vCenter Server Configuration" section on page 7-10.
  - ERROR: Datacenter not found
- For a list of terms used with the Cisco Nexus 1000VE, see the *Cisco Nexus 1000V Getting Started Guide*.

Flowchart for Troubleshooting Modules

Use the following flowchart to troubleshoot modules.

Flowchart: Troubleshooting Modules

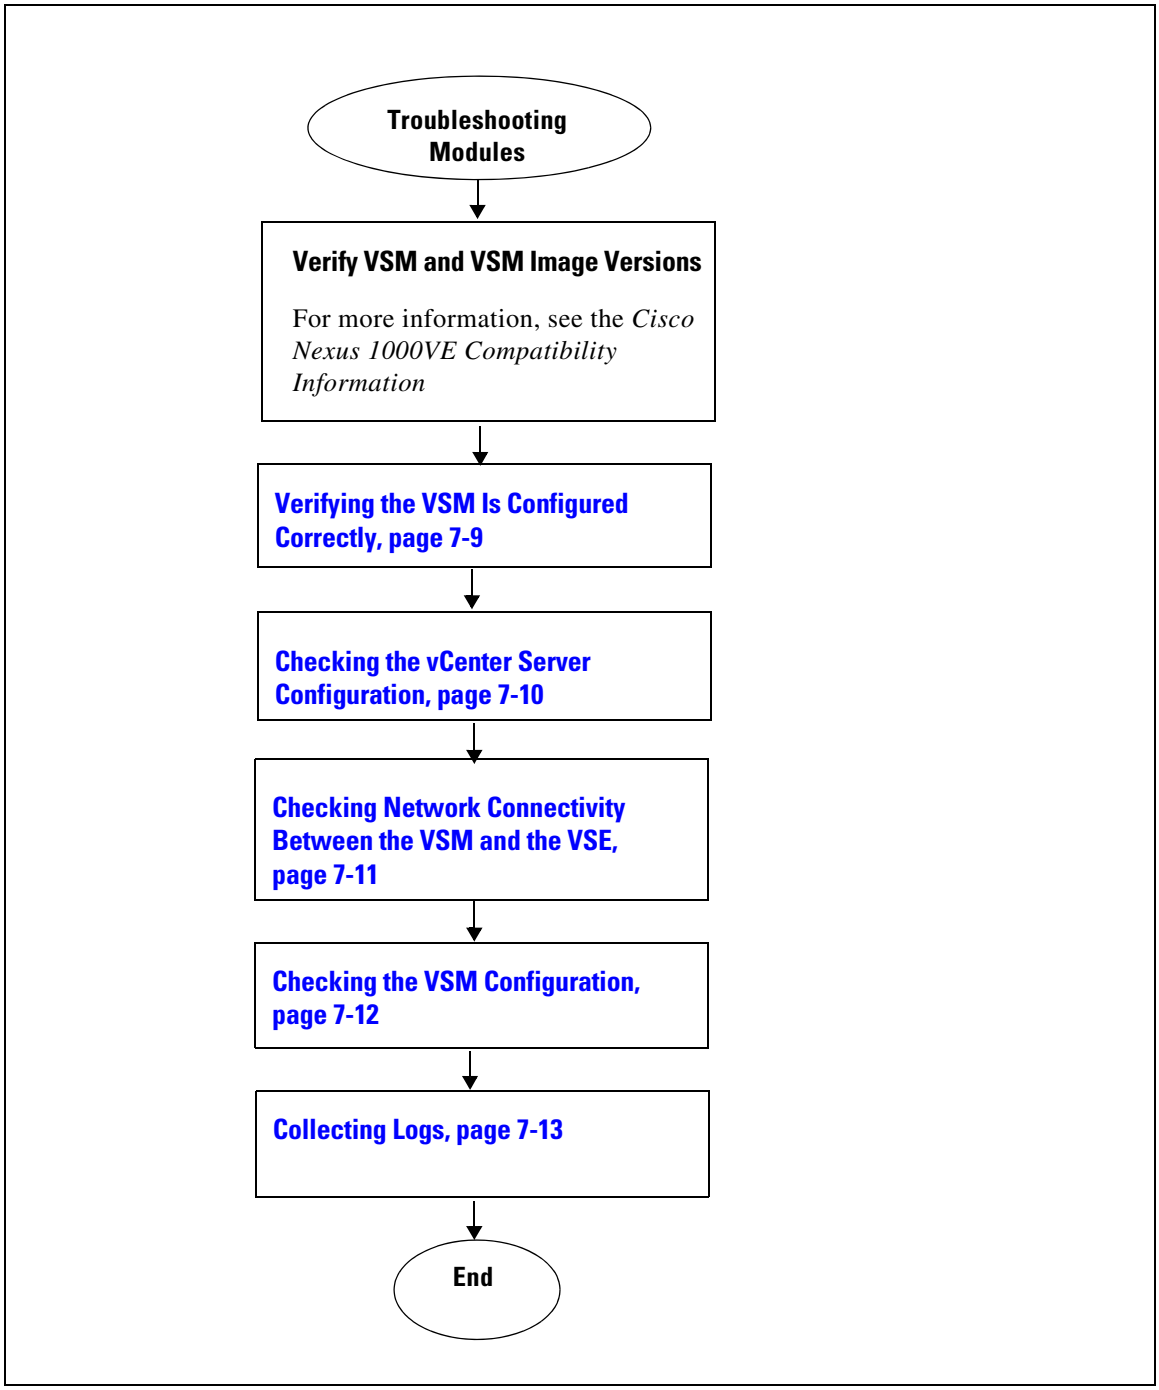

Problems with the VSM

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

Symptom	Possible Causes	Solution
You see the following error on the VSM: `ERROR: [VMware vCenter Server 4.0.0 build-150489] Extension key was not registered before its use`	A extension or plug-in was not created for the VSM.	1. Verify that the extension or plugin was created. "Finding the Extension Key Tied to a Specific DVS" procedure on page 3-7 2. If the plug-in is not found, create one using the following procedure in the *Cisco Nexus 1000V Getting Started Guide*: Creating a Cisco Nexus 1000VE Plug-In on the vCenter Server
Following a reboot of the VSM, the system stops functioning in one of the following states and does not recover on its own. Attempts to debug fail.		
After boot, VSM is in loader prompt.	Corrupt VSM kickstart image.	1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose **Option 1, Install Nexus1000ve** and bring up new image. Follow the VSM installation procedure.
	Boot variables are not set.	1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose **Option 3, Install Nexus1000v**e only if the disk unformatted and bring up 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. **reload**
After boot, VSM is in boot prompt.	Corrupt VSM system image.	1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose **Option 1, Install Nexus1000ve** and bring up new image. 3. Follow the VSM installation procedure.

Symptom	Possible Causes	Solution
After boot, VSM is reconfigured.	Startup configuration is deleted.	Do one of the following: • If you have a saved backup copy of your configuration file, restore the configuration on the VSM. **copy** *source filesystem: filename* **system:running-config** • If not, reconfigure the VSM using the following section in the *Cisco Nexus 1000V Getting Started Guide*: Setting Up the Software
After boot, VSM is stopped at "Loader Loading."	Corrupt boot menu file.	1. Boot the VSM from the CD ROM. 2. From the CD Boot menu, choose **Option 3, Install Nexus1000ve** only if the disk unformatted and bring up new image. 3. Do one of the following: • If you have a saved backup copy of your configuration file, restore the configuration on the VSM. **copy** *source filesystem: filename* **system:running-config** • If not, reconfigure the VSM using the following section in the *Cisco Nexus 1000V Getting Started Guide*: Setting Up the Software
After boot, the secondary VSM reboots continuously.	Control VLAN or control interface down	Check control connectivity between the active and the standby VSM.
	Active and standby VSMs fail to synchronize.	From the active VSM, check system manager errors to identify which application caused the failure. **show system internal sysmgr event-history errors** **show logging**

Verifying the VSM Is Connected to vCenter Server

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

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

show svs connections

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

Example:

```
switch# show svs connections
connection vc:
    hostname: -
    ip address: 172.23.43.170
    ipv6 address: -
    remote port: 80
    transport type: ipv4
    vrf: management
    protocol: vmware-vim https
    certificate: default
    datacenter name: hamilton-DC
    admin:
    max-ports: 12000
    extension key: Cisco_Nexus_1000V_342482929
    DVS uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
    dvs version: 5.0.0
    config status: Enabled
    operational status: Connected
    sync status: Complete
    version: VMware vCenter Server 6.5.0 build-4602587
    vc-uuid: d06e96b6-55e7-4cf8-9f85-511b4bdafe06
    ssl-cert: self-signed or not authenticated
switch#
```

This output indicates the ipg created

```
1. ng-vsm# sh ipg-info
--------------------------------------------------
Name VlanId PortGroupKey
--------------------------------------------------
~ipg47 47 dvportgroup-1469
~ipg3 3 dvportgroup-1425
~ipg2058 2058 dvportgroup-1484
~ipg4 4 dvportgroup-1426
~ipg2059 2059 dvportgroup-1485
~ipg5 5 dvportgroup-1427
~ipg6 6 dvportgroup-1428
~ipg7 7 dvportgroup-1429
~ipg43 43 dvportgroup-1465
~ipg8 8 dvportgroup-1430........


****IPG Creation is complete****
```

The output indicate the IPG range

```
ng-vsm# sh dc clusters
Global Internal IPG tags:
inside-trunk 1:1-50,
inside-trunk 2:2047-2096,
Cluster DefaultCluster free IPG tags:
```

```
inside-trunk 1:1-50,
inside-trunk 2:2047-2096,
Cluster DefaultCluster used IPG tags:
inside-trunk 1:,
inside-trunk 2:,
ng-vsm#
```

Step 2 Do one of the following:

- If the status is **Connected**, return to the "Flowchart: Troubleshooting Modules" section on page 7-3.
- If not, continue with the next step.

Step 3 Connect to vCenter Server.

config t

svs connection *connection_name*

connect

Example:

```
switch# conf t
switch(config)# svs connection HamiltonDC
switch(config-svs-conn)# connect
ERROR: [VMWARE-VIM] Extension-key/username-password pair needs to be configured before
connect
switch# register-plugin remote username XXX password XXX
```

Step 4 Do one of the following:

- If you see an error message about the username-password or extension key, continue with the next Step 8.
- If you see an error indicating that the DVS already exists in the show vms internal error output, go to Step 8.
- If there are no errors, go to Step 8.

Step 5 Do the following:

- Unregister the extension key using the "Unregistering the Extension Key in the vCenter Server" section on page 3-11.
- If the vCenter's login credentials were used to create N1KVE DVS, go to Step 6, else go to Step 7.

Step 6 Use the **remote username** CLI to provide vCenter credentials to the VSM and go to Step 9.

Example

```
switch# conf t
switch# svs connection HamiltonDC
switch (config-svs-conn)# remote username administrator@vsphere.local password Abcd@1234
```

Step 7 Use the **register-plugin** CLI to register VSM extension key with the vCenter and go to Step 9.

```
switch# conf t
switch# svs connection HamiltonDC
switch (config-svs-conn)# register-plugin remote username administrator@vsphere.local
password Abcd@1234
```

Step 8 Use the **vmware dvs** CLI to provide the UUID and datacenter name of the existing DVS. The DVS UUID can be obtained by pointing your web browser to https://your-vCenter-IP/mob and then navigating to the VmwareDistributedVirtualSwitch object corresponding to your DVS.

```
switch# conf t
switch# svs connection HamiltonDC
switch (config-svs-conn)# vmware dvs uuid "50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c
2e" datacenter-name hamilton-dc
```

Step 9 Verify the connection between the VSM and vCenter Server.

show svs connections

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

```
Example:
switch# show svs connections
connection vc:
    hostname: -
    ip address: 172.23.43.170
    ipv6 address: -
    remote port: 80
    transport type: ipv4
    vrf: management
    protocol: vmware-vim https
    certificate: default
    datacenter name: hamilton-DC
    admin:
    max-ports: 12000
    extension key: Cisco_Nexus_1000V_342482929
    DVS uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
    dvs version: 5.0.0
    config status: Enabled
    operational status: Connected
    sync status: Complete
    version: VMware vCenter Server 6.5.0 build-4602587
    vc-uuid: d06e96b6-55e7-4cf8-9f85-511b4bdafe06
    ssl-cert: self-signed or not authenticated
switch#
```

Step 10 Do one of the following:

- If the status is **Connected**, you have completed this procedure.
- If not, return to the "Flowchart: Troubleshooting Modules" section on page 7-3.

Verifying Internal Port Group (IPG) Information

You can verify the information about the internal group (IGP).

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.

Step 1 On the VSM, verify the IPG configuration.

This command helps in checking if the IPG was created on the VC.

show ipg-info

```
switch# sh ipg-info
--------------------------------------------------
Name VlanId PortGroupKey
--------------------------------------------------
~ipg47 47 dvportgroup-1469
~ipg3 3 dvportgroup-1425
~ipg2058 2058 dvportgroup-1484
~ipg4 4 dvportgroup-1426
```

```
~ipg2059 2059 dvportgroup-1485
~ipg5 5 dvportgroup-1427
~ipg6 6 dvportgroup-1428
~ipg7 7 dvportgroup-1429 ......
****IPG Creation is complete****
```

Verifying the VSM Is Configured Correctly

This section includes the following topics:

Verifying the Domain Configuration

You can verify the domain configuration.

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.
- Verify that the output of the show **svs domain** command indicates the following:
  - The presence of a control VLAN and a packet VLAN.
  - The domain configuration was successfully pushed to VC.

Step 1 On the VSM, verify the domain configuration.

show svs domain

Example:

```
switch# show svs domain
    SVS domain config:
    Domain id: 888
    Control vlan: NA
    Packet vlan: NA
    L2/L3 Control mode: L3
    Switch guid: 3a452e9b-a777-4a1f-ab28-c7312399a9a8
    L3 control interface: mgmt0
    Status: Config push to Management Server successful.
    Control type multicast: No
    L3Sec Status: Enabled

Note: Control VLAN and Packet VLAN are not used in L3 mode
switch#
```

Verifying the System Port Profile Configuration

You can verify the port profile configuration.

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.
- Verify that the output of the **show port-profile name** command indicates the following:
  - The control and packet VLANs are assigned.
  - The port profile is enabled.
  - If you have configured a non-default system MTU setting, check that it is the correct size.

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

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

Example:

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

Checking the vCenter Server Configuration

You can verify the configuration on vCenter Server.

Step 1 Confirm that the host is added to the data center and the Cisco Nexus 1000VE DVS in that data center.

Step 2 Confirm that at least one pnic of the server that hosts the VSE is assigned to uplink portgroup on external vDS.

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

Checking Network Connectivity Between the VSM and the VSE

You can verify the network connectivity between the VSM and the VSE.

Step 1 On the VSM, find its MAC address.

show svs neighbors

The VSM MAC address displays as the AIPC Interface MAC.

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

```
Example:
switch# show svs neighbors

Active Domain ID: 1030

AIPC Interface MAC: 0050-568e-58b7
inband/outband Interface MAC: 0050-568e-2a39

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

0002-3d44-0602     VSM      1024         0302     261058.59
```

Step 2 Do one of the following:

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

Step 3 On the upstream switch, display the MAC address table to verify the network configuration.

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

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

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

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

Step 4 Do one of the following:

- If the output from Step 3 does not display the MAC address of the VSM, then there is a problem with connectivity between the server hosting the VSM and the upstream switch. Recheck the VSM configuration and vCenter Server configuration.

- Otherwise, continue with the next step.

Step 5 Do one of the following:

- If the MAC address of the VSM does not appear in the output of Step 3, check the VSM configuration as explained in "Checking the VSM Configuration" section on page 7-12.
- Otherwise, you have completed this procedure.

Checking the VSM Configuration

You can verify that the ESX host received the VSM configuration and setup.

Step 1 On the VSE, On the VSE, verify that the vemfwd and the vssdpa processes are running.

systemctl status nexus1000v

Example:

```
cisco-vse:~# systemctl status nexus1000v
? nexus1000v.service - Cisco Nexus 1000V VSE
Loaded: loaded (/usr/lib/systemd/system/nexus1000v.service; enabled; vendor preset:
disabled)
Active: active (running) since Sat 2018-06-30 00:41:11 PDT; 1 weeks 2 days ago
Process: 25378 ExecStop=/opt/cisco/n1kv/scripts/n1kv stop (code=exited, status=0/SUCCESS)
Main PID: 25411 (n1kv)
CGroup: /system.slice/nexus1000v.service
+-25411 /bin/bash /opt/cisco/n1kv/scripts/n1kv start
+-25475 /sbin/vemfwd -c 1 -- -p 7
--config="(0,0,0),(0,1,0),(1,0,0),(1,1,0),(2,0,0),(2,1,0)" --txq-per-port=4
+-25576 /sbin/vssdpa-proto.bin -f
```

Step 2 Verify that the domain ID and the other parameters are configured correctly on the VSE.

vemcmd show card

Example:

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

Step 3 Verify that the ports of the host added to the DVS are listed and that the ports are correctly configured as access or trunk on the host.

vemcmd show port

Example:

```
~ # vemcmd show port
LTL    IfIndex   Vlan    Bndl  SG_ID Pinned_SGID  Type  Admin State  CBL Mode   Name
8           0   3969       0      2           2  VIRT     UP    UP    1 Access l20
9           0   3969       0      2           2  VIRT     UP    UP    1 Access l21
10          0   3002       0      2           2  VIRT     UP    UP    1 Access l22
11          0   3968       0      2           2  VIRT     UP    UP    1 Access l23
12          0   3003       0      2           2  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      2           2  PHYS     UP    UP    1 Trunk vmnic1
```

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

Step 4 Verify that the vmnic port that is supposed to carry the control VLAN and packet VLAN is present.

vemcmd show bd *control_vlan*
vemcmd show bd *packet_vlan*

Example:

```
~ # vemcmd show bd 3002
BD 3002, vdc 1, vlan 3002, 2 ports
Portlist:
     10  l22
     16  vmnic1
~ # vemcmd show bd 3003
BD 3003, vdc 1, vlan 3003, 2 ports
Portlist:
     12  l24
     16  vmnic1
```

Collecting Logs

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

Step 1 On the VSM, verify its UUID.

vemcmd show card info

Example:

```
~ # vemcmd show card info
Card UUID type  0: 4908a717-7d86-d28b-7d69-001a64635d18
Card name: sfish-srvr-7
Switch name: switch
Switch uuid: 50 84 06 50 81 36 4c 22-9b 4e c5 3e 1f 67 e5 ff
Card domain: 11
Card slot: 12
Control VLAN MAC: 00:02:3d:10:0b:0c
```

```
inband/outband MAC: 00:02:3d:20:0b:0c
SPAN MAC: 00:02:3d:30:0b:0c
USER DPA MAC: 00:02:3d:40:0b:0c
Management IP address: 172.28.30.56
Max physical ports: 16
Max virtual ports: 32
Card control VLAN: 3002
Card packet VLAN: 3003
```

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

show module vse mapping

Example:

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

Step 3 Using the module number from Step 2, collect the output of the following commands:

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

Note If you need to contact Cisco TAC for assistance in resolving an issue, you will need the output of the commands listed in Step 3.

VSM and VSE Troubleshooting Commands

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

Command	Description
show svs neighbors	Displays all neighbors. See Example 7-1 on page 7-16.
show svs connections	Displays the Cisco Nexus 1000VE connections. See Example 7-2 on page 7-16.
show svs domain	Displays the domain configuration. See Example 7-3 on page 7-16.
show port-profile name *name*	Displays the configuration for a named port profile. See Example 7-4 on page 7-17.

Command	Description
show running-config vlan *vlanID*	Displays the VLAN information in the running configuration. See Example 7-5 on page 7-17.
show mac address-table interface	Displays the MAC address table on an upstream switch to verify the network configuration.
module vse *module_number*	Displays the VLAN configuration on the VSM to verify that the VSM MAC appears in the control and packet VLANs.
vemcmd show card	Displays information about cards on the VSM 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 port [*port-LTL-number*]	Displays information about ports on the VSM 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 7-8 on page 7-18.
vemcmd show bd [*control_vlan_id* | *packet_vlan_id*]	Displays configured information on the VSM to verify that the VM NIC port that is supposed to carry the control VLAN and packet VLAN is present. See Example 7-10 on page 7-18.
vemcmd show trunk	Displays configured information on the VSM to verify that the DV port groups are successfully pushed from vCenter Server to the host and that the correct physical trunk port VM NIC is used. See Example 7-11 on page 7-19.
show module vse mapping	Displays information about the VSM that a VSM maps to, including the VSM module number, status, UUID, and license status. See Example 7-12 on page 7-19.
show system internal vem_mgr event-history module 13 *module-number*	Displays module FSM event information.
show module internal event-history module *module-number*	Displays the event log for a module.
show system internal im event-history module *module-number*	Displays the module IM event logs for the system.
show system internal vmm event-history module *module-number*	Displays the module VMM event logs for the system.
show system internal ethpm event-history module *module-number*	Displays the module Ethernet event logs for the system.
show system internal ethpm event-history int *type slot*	Displays the Ethernet interface logs for the system.

Example 7-1 show svs neighbors Command

```
switch# show svs neighbors

Active Domain ID: 113

AIPC Interface MAC: 0050-56b6-2bd3
inband/outband Interface MAC: 0050-56b6-4f2d

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

0002-3d40-7102     VSM       113         0302     71441.12
0002-3d40-7103     VSM       113         0402       390.77

switch#
```

Example 7-2 show svs connections Command

```
switch# show svs connections
connection vc:
    hostname: -
    ip address: 172.23.43.170
    ipv6 address: -
    remote port: 80
    transport type: ipv4
    vrf: management
    protocol: vmware-vim https
    certificate: default
    datacenter name: hamilton-DC
    admin:
    max-ports: 12000
    extension key: Cisco_Nexus_1000V_342482929
    DVS uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
    dvs version: 5.0.0
    config status: Enabled
    operational status: Connected
    sync status: Complete
    version: VMware vCenter Server 6.5.0 build-4602587
    vc-uuid: d06e96b6-55e7-4cf8-9f85-511b4bdafe06
    ssl-cert: self-signed or not authenticated
switch#
```

Example 7-3 show svs domain Command

```
switch# show svs domain
    SVS domain config:
    Domain id: 888
    Control vlan: NA
    Packet vlan: NA
    L2/L3 Control mode: L3
    Switch guid: 3a452e9b-a777-4a1f-ab28-c7312399a9a8
    L3 control interface: mgmt0
    Status: Config push to Management Server successful.
    Control type multicast: No
    L3Sec Status: Enabled

Note: Control VLAN and Packet VLAN are not used in L3 mode
switch#
```

Example 7-4 show port-profile Command

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

Example 7-5 show running-configuration vlan Command

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

switch#
```

Example 7-6 VSM-health check Command

```
switch# VSM-health check
switch# show running-config vlan 220-221

version 5.2(1)SV5(1.1)
vlan 220-221
vlan 220
name mgmt
vlan 221
name system

switch#
```

Example 7-7 vemcmd show card Command

```
switch# vemcmd show card
Card UUID type  2: 58f8afd7-e1e3-3c51-85e2-6e6f2819a7b8
Card name: sfish-srvr-1
Switch name: switch
Switch alias: DvsPortset-0
Switch uuid: 56 e0 36 50 91 1c 32 7a-e9 9f 31 59 88 0c 7f 76
Card domain: 1024
Card slot: 4
VSM Control (Control VLAN) MAC: 00:02:3d:14:00:03
```

```
VSM Packet (inband/outband) MAC: 00:02:3d:24:00:03
VSM Control Agent (DPA) MAC: 00:02:3d:44:00:03
VSM SPAN MAC: 00:02:3d:34:00:03
Management IP address: 172.23.232.102
Max physical ports: 32
Max virtual ports: 216
Card control VLAN: 3002
Card packet VLAN: 3003
       Processors: 4
  Processor Cores: 4
Processor Sockets: 2
  Physical Memory: 4290351104
```

Example 7-8 vemcmd show port Command

```
switch# vemcmd show port
switch# vemcmd show port
LTL VSM Port Admin Link State PC-LTL SGID Vem Port Type ORG svcpath Owner
21 Eth3/1 UP UP F/B* 0 eth1 0 0 dpdk-outside
53 Veth2 UP UP FWD 0 test-vm1.eth1 0 0 test-vm1
54 Veth1 UP UP FWD 0 test-vm2.eth1 0 0 test-vm2

* F/B: Port is BLOCKED on some of the vlans.
One or more vlans are either not created or
not in the list of allowed vlans for this port.
Please run "vemcmd show port vlans" to see the details.
switch#:~$
```

Example 7-9 vemcmd show port vlans Command

```
switch# vemcmd show port vlans
switch# vemcmd show port vlans
Native VLAN Allowed
LTL VSM Port Mode VLAN State* Vlans
21 Eth3/1 T 1 FWD 220-229
53 Veth2 A 222 FWD 222
54 Veth1 A 223 FWD 223

* VLAN State: VLAN State represents the state of allowed vlans.
switch #
```

Note The output `*F/B The port is blocked on some of the VLANs` means that the trunk is not forwarding all VLANs. This might be a normal situation depending on the port profile allowed VLAN list. Compare the output of the **vemcmd show port vlans** command against the port profile trunk allowed VLANs. If the lists match, all of the expected VLANs are forwarding and the Cisco Nexus 1000VE is blocking nonallowed VLANs.

Example 7-10 vemcmd show vlan Command

```
switch# vemcmd show vlan
switch vemcmd show vlan 222
VLAN 222, vdc 1, swbd 222, hwbd 6, 2 ports

Portlist:
21 eth1
53 test-vm1.eth1

switch#
```

Example 7-11 vemcmd show trunk Command

```
switch# vemcmd show trunk
switch# vemcmd show trunk
Trunk port 6 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 16 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 21 native_vlan 1 CBL 0
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
switch#
```

Example 7-12 show module vse mapping Command

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

Example 7-13 show ipg-info

```
switch# show ipg-info
Desc : shows internal port-group info created on VC


------------------------------------------------
Name            VlanId     PortGroupKey
------------------------------------------------
~ipg47          47         dvportgroup-1469
~ipg3           3          dvportgroup-1425
~ipg2058        2058       dvportgroup-1484
~ipg4           4          dvportgroup-1426
~ipg2059        2059       dvportgroup-1485
~ipg5           5          dvportgroup-1427
~ipg6           6          dvportgroup-1428
~ipg7           7          dvportgroup-1429
~ipg43          43         dvportgroup-1465
~ipg8           8          dvportgroup-1430
~ipg2070        2070       dvportgroup-1496
~ipg1           1          dvportgroup-1423
~ipg2           2          dvportgroup-1424
~ipg46          46         dvportgroup-1468
~ipg9           9          dvportgroup-1431
~ipg2071        2071       dvportgroup-1497
~ipg10          10         dvportgroup-1432
~ipg2068        2068       dvportgroup-1494
~ipg11          11         dvportgroup-1433
~ipg2069        2069       dvportgroup-1495
~ipg2084        2084       dvportgroup-1510
~ipg12          12         dvportgroup-1434
~ipg2066        2066       dvportgroup-1492
~ipg2056        2056       dvportgroup-1482
~ipg2057        2057       dvportgroup-1483
```

```
~ipg2065        2065       dvportgroup-1491
~ipg15          15         dvportgroup-1437
~ipg2067        2067       dvportgroup-1493
~ipg13          13         dvportgroup-1435
~ipg2072        2072       dvportgroup-1498
~ipg50          50         dvportgroup-1472
~ipg2073        2073       dvportgroup-1499
~ipg2095        2095       dvportgroup-1521
~ipg2074        2074       dvportgroup-1500
~ipg2075        2075       dvportgroup-1501
~ipg2076        2076       dvportgroup-1502
~ipg2077        2077       dvportgroup-1503
~ipg2078        2078       dvportgroup-1504
~ipg19          19         dvportgroup-1441
~ipg2079        2079       dvportgroup-1505
~ipg18          18         dvportgroup-1440
~ipg2080        2080       dvportgroup-1506
~ipg21          21         dvportgroup-1443
~ipg2081        2081       dvportgroup-1507
~ipg20          20         dvportgroup-1442
~ipg2082        2082       dvportgroup-1508
~ipg2083        2083       dvportgroup-1509
~ipg44          44         dvportgroup-1466
~ipg45          45         dvportgroup-1467
~ipg48          48         dvportgroup-1470
~ipg49          49         dvportgroup-1471
~ipg26          26         dvportgroup-1448
~ipg27          27         dvportgroup-1449
~ipg28          28         dvportgroup-1450
~ipg29          29         dvportgroup-1451
~ipg30          30         dvportgroup-1452
~ipg31          31         dvportgroup-1453
~ipg32          32         dvportgroup-1454
~ipg14          14         dvportgroup-1436
~ipg16          16         dvportgroup-1438
~ipg17          17         dvportgroup-1439
~ipg33          33         dvportgroup-1455
~ipg34          34         dvportgroup-1456
~ipg35          35         dvportgroup-1457
~ipg36          36         dvportgroup-1458
~ipg37          37         dvportgroup-1459
~ipg38          38         dvportgroup-1460
~ipg22          22         dvportgroup-1444
~ipg23          23         dvportgroup-1445
~ipg24          24         dvportgroup-1446
~ipg25          25         dvportgroup-1447
~ipg39          39         dvportgroup-1461
~ipg40          40         dvportgroup-1462
~ipg41          41         dvportgroup-1463
~ipg2055        2055       dvportgroup-1481
~ipg2060        2060       dvportgroup-1486
~ipg2047        2047       dvportgroup-1473
~ipg2061        2061       dvportgroup-1487
~ipg2062        2062       dvportgroup-1488
~ipg2063        2063       dvportgroup-1489
~ipg2064        2064       dvportgroup-1490
~ipg42          42         dvportgroup-1464
~ipg2048        2048       dvportgroup-1474
~ipg2049        2049       dvportgroup-1475
~ipg2050        2050       dvportgroup-1476
~ipg2051        2051       dvportgroup-1477
~ipg2052        2052       dvportgroup-1478
~ipg2053        2053       dvportgroup-1479
~ipg2054        2054       dvportgroup-1480
```

```
~ipg2085        2085       dvportgroup-1511
~ipg2086        2086       dvportgroup-1512
~ipg2087        2087       dvportgroup-1513
~ipg2088        2088       dvportgroup-1514
~ipg2089        2089       dvportgroup-1515
~ipg2090        2090       dvportgroup-1516
~ipg2091        2091       dvportgroup-1517
~ipg2092        2092       dvportgroup-1518
~ipg2093        2093       dvportgroup-1519
~ipg2094        2094       dvportgroup-1520
~ipg2096        2096       dvportgroup-1522


****IPG Creation is complete****
show dc hosts vse
```

Example 7-14 show dc hosts vse

```
switch# show dc hosts vse
Desc: Show VSE IP and Host IP/Name mapping info


Internal IPG tags:
inside-trunk 1:1-50,
inside-trunk 2:2047-2096,
----------------------------------------
HOST NAME: 10.197.148.227
HOST IP: 10.197.148.227
VSE IP: 202.1.1.227
VSE UUID : 564D5218-113F-3E4E-D034-E911AAC7FF0A

----------------------------------------
```

Example 7-15 show vms internal info host-table

```
switch# show   vms internal info host-table
Desc : Display host info which are added to N1KVE-VDS


  Host Table:
---------------
 Lock acquired [0]
 Hosts added/removed flag: [0]
 Pending host filter update: [0]
 Notification of hosts info pending flag: [0]
 Hosts filter spec created flag: [1]
 Hosts filter spec reference:
[session[5206c9d7-e0d5-68e6-9586-ffcc815e2b05]527d4768-76d4-578d-b382-08773e162514]
 VC checkforupdate version: [19]
# of host entries = 256

Host Index = 2
Host entry flags = [OLD], [4]
Host info for slot no = 2
   Module Number : 3
   Host UUID: [b7371ff4-f03f-3c40-b2b7-049ee68fc8a9]
   VSE  UUID: [564D5218-113F-3E4E-D034-E911AAC7FF0A]
   IPv4 address: [10.197.148.227], IPv6 address: [fe80::56a2:74ff:fe59:9b96], Hostname:
[10.197.148.227]
   VSE IPv4 address: [202.1.1.227], VSE Hostname: [localhost.localdomain]
   ref: [host-18]
```

```
   Cluster mor: [domain-s16], Cluster name: [10.197.148.227]
```

Example 7-16 show vms internal info host-view

```
switch# show vms internal info host-view
Desc : Display all N1KVE port details  from VC point of view


------------------------------------------------------------------
       Slot #2
------------------------------------------------------------------
Host Name       : 10.197.148.227
Host Ref        : host-18
Host UUID       : b7371ff4-f03f-3c40-b2b7-049ee68fc8a9
Host VSE IP     : 202.1.1.227
Host Slot Num   : 2
Module Num      : 3
Host ID         : 1
Cluster ID      : 1
Host ID         : 1
Cluster ID      : 1
ipg_count for trunk 1 (1)
ipg_count for trunk 2 (1)
--------------------------------------------------------------------------
        List VM(s)
--------------------------------------------------------------------------
     VM Name       : App-3
     VM Ref        : vm-25
     VM UUID       : 42087cc1-8aa7-c511-81e7-6423cf483e84
     VM Inst UUID  : 50082799-4901-f00e-49f7-9cd61d9913b1
     VM ID         : 1
     VM Host Id    : 1
     VM ID         : 1
     VM Host Id    : 1
--------------------------------------------------------------------------
        Adapter List
--------------------------------------------------------------------------
     State            : IPG allocated
     MAC Address      : 00:50:56:88:44:59
     Adapter Name     : Network adapter 1
     Port Key         : Port-4000
     DvPort           : Portgroup-356
     VLAN TAG         : 1
     Port-Group Name  : vm-313
--------------------------------------------------------------------------
     State            : IPG allocated
     MAC Address      : 00:50:56:88:52:d4
     Adapter Name     : Network adapter 2
     Port Key         : Port-4001
     DvPort           : Portgroup-388
     VLAN TAG         : 2047
     Port-Group Name  : vm-314
--------------------------------------------------------------------------
```

Example 7-17 show vms internal info ipg-profile-mapping

```
switch# show vms internal info ipg-profile-mapping
Desc : Show mapping between internal port-group and port-profile

Cluster id : 1, Cluster_ref : DefaultCluster
```

```
 IpgName ~ipg1 , Profile Id 8
 IpgName ~ipg2047 , Profile Id 9
```

Example 7-18 show dc clusters

```
switch# show dc clusters
Desc : Show ipg info per cluster basis


Global Internal IPG tags:
inside-trunk 1:1-50,
inside-trunk 2:2047-2096,
Cluster DefaultCluster free IPG tags:
inside-trunk 1:2-50,
inside-trunk 2:2048-2096,
Cluster DefaultCluster used IPG tags:
inside-trunk 1:1,
inside-trunk 2:2047,
```

Example 7-19 show vms internal info cluster-view

```
switch# show vms internal info cluster-view
Desc : Show cluster info as well as show mapping between internal port-group and port-mac

------------------------------------------------------------------
       Cluster  #1
------------------------------------------------------------------
Cluster Id         : 1
Cluster Name       : DefaultCluster
Cluster sync_flag       : 1
Cluster Ref        : DefaultCluster
Print the list of used cluster bits
1,



------------------------------------------------------------------
    Printing Mac address and ipg info of each cluster
------------------------------------------------------------------


--------------------------------------------------
       Cluster Id:  1
       Cluster Ref:  DefaultCluster
--------------------------------------------------
--------------------------------------------------
~ipg                Mac-address
--------------------------------------------------
~ipg1                    00:50:56:88:44:59
~ipg2047                 00:50:56:88:52:d4
```

Example 7-20 show vms internal info host-view module 3

```
switch# show vms internal info host-view module 3
Desc : Display all N1KVE port details on particular module from VC point of view

------------------------------------------------------------------
       Slot #2
------------------------------------------------------------------
```

```
Host Name       : 10.197.148.227
Host Ref        : host-18
Host UUID       : b7371ff4-f03f-3c40-b2b7-049ee68fc8a9
Host VSE IP     : 202.1.1.227
Host Slot Num   : 2
Module Num      : 3
Host ID         : 1
Cluster ID      : 1
Host ID         : 1
Cluster ID      : 1
ipg_count for trunk 1 (1)
ipg_count for trunk 2 (1)
------------------------------------------------------------------------
        List VM(s)
------------------------------------------------------------------------
     VM Name       : App-3
     VM Ref        : vm-25
     VM UUID       : 42087cc1-8aa7-c511-81e7-6423cf483e84
     VM Inst UUID  : 50082799-4901-f00e-49f7-9cd61d9913b1
     VM ID         : 1
     VM Host Id    : 1
     VM ID         : 1
     VM Host Id    : 1
------------------------------------------------------------------------
        Adapter List
------------------------------------------------------------------------
     State            : IPG allocated
     MAC Address      : 00:50:56:88:44:59
     Adapter Name     : Network adapter 1
     Port Key         : Port-4000
     DvPort           : Portgroup-356
     VLAN TAG         : 1
     Port-Group Name  : vm-313
------------------------------------------------------------------------
     State            : IPG allocated
     MAC Address      : 00:50:56:88:52:d4
     Adapter Name     : Network adapter 2
     Port Key         : Port-4001
     DvPort           : Portgroup-388
     VLAN TAG         : 2047
     Port-Group Name  : vm-314
------------------------------------------------------------------------
```

Example 7-21 show vms internal info ipg-duplicate

```
Desc : Show any duplicate internal port-group assigned to multiple ports per cluster

Cluster DefaultCluster:
```

Ports

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

Information About Ports

This section includes the following topics:

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.

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 operation speed).

For a complete description of port modes, administrative states, and operational states, see the *Cisco Nexus 1000V Interface Configuration Guide*.

Information About Interface Counters

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

show interface *interface-name* **counters**

See Example 8-8 on page 8-14.

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

clear counters interface *interface-name*

Information About Link Flapping

When a port continually goes up and down, it is said to be flapping, or link flapping. When a port is flapping, it cycles through the following states, in this order, and then starts over again:

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

To troubleshoot link flapping, see the "Information About Link Flapping" section on page 8-2.

Port Diagnostic Checklist

Use the following checklist to diagnose port interface activity.

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

Table 8-1 Port Diagnostic Checklist

Checklist	Example	√
Verify that the module is active. **show module**	See Example 8-1 on page 8-9.	
Verify that the VSM is connected to vCenter Server. **show svs connections**	See Example 8-3 on page 8-9.	
Verify if the internal port-group information is created on VC. **show ipg-info**		
Verify if the VSE IP to Host IP mapping is done **show dc hosts vse**		
Verify if the module is online or not. **show module**		

Table 8-1 Port Diagnostic Checklist (continued)

Checklist (continued)	Example	√
Verify that the ports have been created. **show interface brief**	See Example 8-6 on page 8-13.	
Verify the state of the interface. **show interface interface-name**	See Example 8-7 on page 8-13.	
Verify if the host and cluster MOB and the uuid info is fetched from the VC. **show vms internal info host-table**		
Verify if there are any error in the vms event-history error during the port creation. **show vms internal event-history errors**		
Verify if the VC port(s) are moved to the internal port-group from Nexus 1000VE pro-profile(s). **show vms internal info host-view**		
Verify if the VSM IPG moves event(s) received from the VC on particular port(s) on VM. **show vms internal info host-view**		
Verify if the IPG to port-profile mapping is done. **show vms internal info ipg-profile-mapping**		

Problems with Ports

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

- Cannot Enable an Interface, page 8-4
- Port Link Failure or Port Not Connected, page 8-4
- Link Flapping, page 8-4
- Port ErrDisabled, page 8-5
- Port State is Blocked on a VSE, page 8-7

Cannot Enable an Interface

Possible Cause	Solution
A Layer 2 port is not associated with an access VLAN or the VLAN is suspended.	1. Verify that the interface is configured in a VLAN. **show interface brief** 2. If not already, associate the interface with an access VLAN. 3. Determine the VLAN status. **show vlan brief** 4. If not already active, configure the VLAN as active. **config t** **vlan** *vlan-id* **state active**

Port Link Failure or Port Not Connected

Possible Cause	Solution
The port connection is bad.	1. Verify the port state. **show system internal ethpm info** 2. Disable and then enable the port. **shut** **no shut** 3. Move the connection to a different port on the same module or a different module.
The link is stuck in initialization state or the link is in a point-to-point state.	1. Check for a link failure system message. `Link Failure, Not Connected` **show logging** 2. Disable and then enable the port. **shut** **no shut** 3. Move the connection to a different port on the same module or a different module.

Link Flapping

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

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

- For a definition of link flapping, see the "Link Flapping" section on page 8-4.

Possible Cause	Solution
The bit rate exceeds the threshold and puts the port into an error-disabled state.	Disable and then enable the port. **shut** **no shut** The port should return to the normal state.
A hardware failure or intermittent hardware error causes a packet drop in the switch. A software error causes a packet drop. A control frame is erroneously sent to the device.	An external device might choose to initialize the link again when encountering the error. If so, the exact method of link initialization varies by device. 1. Determine the reason for the link flap as indicated by the MAC driver. 2. Use the debug facilities on the end device to troubleshoot the problem.
ESX errors, or link flapping, occurs on the upstream switch.	Use the troubleshooting guidelines in the documentation for your ESX or upstream switch.

Port ErrDisabled

Possible Cause	Solution
The cable is defective or damaged.	1. Verify the physical cabling. 2. Replace or repair defective cables. 3. Reenable the port. **shut** **no shut**

Possible Cause	Solution
You attempted to add a port to a port channel that was not configured identically, and the port is then errdisabled.	1. Display the switch log file and identify the exact configuration error in the list of port state changes. **show logging logfile** 2. Correct the error in the configuration and add the port to the port channel. 3. Re-enable the port. **shut** **no shut**
A VSM application error has occurred.	1. Identify the component that had an error while you were bringing up the port. **show logging logfile | grep** *interface_number* See Example 8-5 on page 8-13. 2. Identify the error transition. **show system internal ethpm event-history interface** *interface_number* 3. Open a support case and submit the output of the above commands. For more information see the "Contacting Cisco Customer Support" section on page 1-7.

Port State is Blocked on a VSE

Possible Cause	Solution
The VLAN is not created on the VSM.	1. Verify the status and of the vEthernet interface. It should be up and not inactive. **show interface vethernet** *number* 2. Verify that the VLAN on the VSM is created. **show vlan** *vlan-id* On the VSE module, do the following: 1. Verify that the VLAN is programmed. **vemcmd show vlan** *vlan-id* 2. Verify that the VLAN is allowed on the ports. **vemcmd show port vlan** 3. Create the VLAN on the VSM. **vlan** *vlan-id*
The VSE modules are unlicensed.	1. Verify that all the modules are in licensed state. **show module** 2. Verify the status of the vEthernet interface. It should be up and not "VSE Unlicensed." **show interface vethernet** *number* 3. Verify the license status of VSE modules. **show module vse license-info** On the VSE module, do the following: 1. Verify that card details show Licensed: Yes. **vemcmd show card** 2. Install the necessary licenses or move the switch to essential mode. **svs switch edition essential**

Port Troubleshooting Commands

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

Command	Purpose
show module *module-number*	Displays the state of a module. See Example 8-1 on page 8-9.
show svs domain	Displays the domain configuration. See Example 8-2 on page 8-9.
show svs connections	Displays the Cisco Nexus 1000V connections. See Example 8-3 on page 8-9.
show logging logfile	Displays logged system messages. See Example 8-4 on page 8-10.
show logging logfile | grep *interface_number*	Displays logged system messages for a specified interface. See Example 8-5 on page 8-13.
show interface brief	Displays a table of interface states. See Example 8-6 on page 8-13.
show interface *interface-name*	Displays the configuration for a named Ethernet interface, including the following: • Administrative state • Speed • Trunk VLAN status • Number of frames sent and received • Transmission errors, including discards, errors, CRCs, and invalid frames See Example 8-7 on page 8-13.
show interface *interface-name* **counters**	Displays port counters for identifying synchronization problems. For information about counters, see the "Information About Interface Counters" section on page 8-2. See Example 8-8 on page 8-14.
show interface vethernet	Displays the vEthernet interface configuration. See Example 8-9 on page 8-14.
show interface status	Displays the status of the named interface.

Command	Purpose
show interface capabilities	Displays a tabular view of all configured port profiles. See Example 8-10 on page 8-14.
show interface virtual port mapping	Displays the virtual port mapping for all vEthernet interfaces. See Example 8-11 on page 8-16.

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

EXAMPLES

Example 8-1 show module Command

```
switch# show module 3
Mod Ports Module-Type Model Status
--- ----- -------------------------------- ------------------ ------------
3 1022 Virtual Service Engine NA ok

Mod Sw Hw
--- ------------------ ------------------------------------------------
3 5.2(1)SV5(1.1) NA

Mod Server-IP Server-UUID Server-Name
--- --------------- ------------------------------------ --------------------
3 172.23.231.209 4212E360-F498-594E-219C-9040BDB93408 sfish-231-209.cisco.com

Mod VSE-IP Host-IP
--- --------------- ----------------------------------
3 172.23.231.209 172.23.233.17
switch#
```

Example 8-2 show svs domain Command

```
switch# show svs domain
SVS domain config:
  Domain id:    559
  Control vlan: 3002
  Packet vlan:  3003
  L2/L3 Aipc mode: L2
  L2/L3 Aipc interface: management interface0
  Status: Config push to VC successful.
switch#
```

Example 8-3 show svs connections Command

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

```
    operational status: Connected
switch#
```

Example 8-4 show logging logfile Command

```
switch# show logging logfile
2018 Jul 10 08:57:54 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: CDM main SAP(423)
registered
2018 Jul 10 08:57:55 switch %USER-2-SYSTEM_MSG: CLIS: loading cmd files begin  - clis
2018 Jul 10 08:57:55 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Vem_mgr SAP(744)
registered
2018 Jul 10 08:57:56 switch vdc_mgr[2427]: %VDC_MGR-5-VDC_STATE_CHANGE: vdc 1 state
changed to create pending
2018 Jul 10 08:57:56 switch platform[2301]: %PLATFORM-5-MOD_STATUS: Module 1
current-status is MOD_STATUS_ONLINE/OK
2018 Jul 10 08:57:56 switch module[2437]: %MODULE-5-ACTIVE_SUP_OK: Supervisor 1 is active
(serial: T505692DFA1)
2018 Jul 10 08:57:56 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Fwm SAP(602) registered
2018 Jul 10 08:57:56 switch fwm[2438]: %FWM-3-L3VM_SDB_OPEN: Error opening
volatile:/dev/shm/l3vm_global_sdb, errno: 0x411a000f (no such sdb exists or is destroyed)
in l3vm_open_one_sdb()
2018 Jul 10 08:57:56 switch fwm[2438]: %FWM-0-SYSLOG_SL_MSG_EMERG: l3vm_open_one_sdb
Backtrace: 0xb79accea 0x8073315 0x806c44c 0x414735c5
2018 Jul 10 08:57:56 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Aclmgr SAP(351) registered
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-3-SYSTEM_MSG: sd 0:0:0:0: [sda] Assuming
drive cache: write through - kernel
2018 Jul 10 08:57:57 switch last message repeated 1 time
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-3-SYSTEM_MSG: CMOS: Module initialized -
kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG: calling
register_stun_set_domain_id() - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG:
register_stun_set_domain_id() - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG:
stun_init_peer_mac_info_from_cmos:ha0_mac from cmos:(00:50:56:92:dd:2c) - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG:
stun_init_peer_mac_info_from_cmos: ha1_mac from cmos:(00:50:56:92:13:01) - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG: Successfully registered
SNAP client for SNAP=0x00000c0132  0xeda8b0e0 - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-1-SYSTEM_MSG: STUN : Successfully
created Socket - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-3-SYSTEM_MSG: redun_platform_ioctl :
Heartbeat interval is set to 15 - kernel
2018 Jul 10 08:57:57 switch Jul 10 08:57:56 %KERN-3-SYSTEM_MSG: redun_platform_ioctl :
Host name is set switch - kernel
2018 Jul 10 08:58:00 switch %USER-2-SYSTEM_MSG: CLIS: loading cmd files end  - clis
2018 Jul 10 08:58:00 switch %USER-2-SYSTEM_MSG: CLIS: init begin  - clis
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Acllog SAP(425) registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Pltfm_config SAP(424)
registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Qosmgr SAP(377) registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Eth PCM SAP(378)
registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: ETh SPAN SAP(174)
registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Vlan_mgr SAP(167)
registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: IGMP process MTS
queue(312) registered
2018 Jul 10 08:58:15 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Ethpm SAP(175) registered
2018 Jul 10 08:58:18 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Eth_port_sec SAP(191)
registered
```

```
2018 Jul 10 08:58:19 switch Jul 10 08:58:19 %KERN-3-SYSTEM_MSG: isec_ioctl: Aegis context
initialized - kernel
2018 Jul 10 08:58:19 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Msp SAP(444) registered
2018 Jul 10 08:58:22 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: Vns_agent SAP(753)
registered
2018 Jul 10 08:58:22 switch cdm[2340]: %CDM-5-CDM_APP_REGISTER: VIM main SAP(403)
registered
2018 Jul 10 08:58:23 switch vdc_mgr[2427]: %VDC_MGR-5-VDC_STATE_CHANGE: vdc 1 state
changed to create in progress
2018 Jul 10 08:58:23 switch ifmgr[2455]: %IM-5-IM_INTF_STATE: mgmt0 is DOWN in vdc 1
2018 Jul 10 08:58:23 switch ifmgr[2455]: %IM-5-IM_INTF_STATE: mgmt0 is UP in vdc 1
2018 Jul 10 08:58:23 switch ifmgr[2455]: %IM-5-IM_INTF_STATE: control0 is DOWN in vdc 1
2018 Jul 10 08:58:23 switch ifmgr[2455]: %IM-5-IM_INTF_STATE: control0 is UP in vdc 1
2018 Jul 10 08:58:23 switch vdc_mgr[2427]: %VDC_MGR-5-VDC_STATE_CHANGE: vdc 1 state
changed to active
2018 Jul 10 08:58:23 switch vdc_mgr[2427]: %VDC_MGR-2-VDC_ONLINE: vdc 1 has come online
2018 Jul 10 08:58:23 switch vdc_mgr[2427]: %VDC_MGR-5-VDC_HOSTNAME_CHANGE: vdc 1 hostname
changed to switch
2018 Jul 10 08:58:28 switch last message repeated 1 time
2018 Jul 10 08:58:28 switch vms[2885]: %VMS-5-CONN_CONNECT: Connection 'vc' connected to
the vCenter Server.
2018 Jul 10 08:58:31 switch bootvar[2442]: %BOOTVAR-5-NEIGHBOR_UPDATE_AUTOCOPY: auto-copy
supported by neighbor supervisor, starting...
2018 Jul 10 08:58:33 switch msp[2882]: %MSP-5-DOMAIN_CFG_SYNC_DONE: Domain config
successfully pushed to the management server.
2018 Jul 10 08:58:33 switch vshd[3564]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3564
2018 Jul 10 08:58:33 switch last message repeated 1 time
2018 Jul 10 08:58:33 switch vshd[3576]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3576
2018 Jul 10 08:58:33 switch vshd[3564]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3564
2018 Jul 10 08:58:33 switch vshd[3576]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3576
2018 Jul 10 08:58:34 switch vshd[3550]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3550
2018 Jul 10 08:58:34 switch vshd[3576]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3576
2018 Jul 10 08:58:35 switch vem_mgr[2420]: %VEM_MGR-2-VEM_MGR_DETECTED: Host
sfish-231-209.cisco.com detected as module 3
2018 Jul 10 08:58:35 switch vns_agent[2889]: %VNS_AGENT-2-VNSA_LIC_NO_ADVANCED_LIC: VSM
does not have Advanced licenses. May not be able to use VSG services. Please install
Advanced licenses.
2018 Jul 10 08:58:35 switch vem_mgr[2420]: %VEM_MGR-2-MOD_ONLINE: Module 3 is online
2018 Jul 10 08:58:30 switch %VEM_MGR-SLOT3-5-VEM_SYSLOG_NOTICE: VETH_IPG_MAPPING :
Vethernet2 is mapped to ipg1
2018 Jul 10 08:58:35 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Vethernet2 is
attached to Net Adapter 2 (test-vm1) on port 1 of module 3 with dvport id 0
2018 Jul 10 08:58:30 switch %VEM_MGR-SLOT3-5-VEM_SYSLOG_NOTICE: VETH_IPG_MAPPING :
Vethernet1 is mapped to ipg2047
2018 Jul 10 08:58:35 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Vethernet1 is
attached to Net Adapter 2 (test-vm2) on port 2 of module 3 with dvport id 0
2018 Jul 10 08:58:35 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Vethernet2 is up in
mode access
2018 Jul 10 08:58:35 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Vethernet1 is up in
mode access
2018 Jul 10 08:58:35 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Ethernet3/1 is
attached to eth1 on module 3
2018 Jul 10 08:58:35 switch ethpm[2833]: %ETHPORT-5-SPEED: Interface Ethernet3/1,
operational speed changed to 10 Gbps
2018 Jul 10 08:58:35 switch ethpm[2833]: %ETHPORT-5-IF_DUPLEX: Interface Ethernet3/1,
operational duplex mode changed to Full
2018 Jul 10 08:58:35 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Ethernet3/1 is up in
mode trunk
```

```
2018 Jul 10 08:58:36 switch vem_mgr[2420]: %VEM_MGR-2-VEM_MGR_DETECTED: Host
sfish-231-161.cisco.com detected as module 4
2018 Jul 10 08:58:36 switch vns_agent[2889]: %VNS_AGENT-2-VNSA_LIC_NO_ADVANCED_LIC: VSM
does not have Advanced licenses. May not be able to use VSG services. Please install
Advanced licenses.
2018 Jul 10 08:58:36 switch vem_mgr[2420]: %VEM_MGR-2-MOD_ONLINE: Module 4 is online
2018 Jul 10 09:11:49 switch %VEM_MGR-SLOT4-5-VEM_SYSLOG_NOTICE: VETH_IPG_MAPPING :
Vethernet3 is mapped to ipg2
2018 Jul 10 08:58:36 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Vethernet3 is
attached to Net Adapter 2 (test-vm3) on port 1 of module 4 with dvport id 0
2018 Jul 10 08:58:36 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Vethernet3 is up in
mode access
2018 Jul 10 09:11:49 switch %VEM_MGR-SLOT4-5-VEM_SYSLOG_NOTICE: VETH_IPG_MAPPING :
Vethernet4 is mapped to ipg2048
2018 Jul 10 08:58:36 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Vethernet4 is
attached to Net Adapter 2 (test-vm4) on port 2 of module 4 with dvport id 0
2018 Jul 10 08:58:36 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Vethernet4 is up in
mode access
2018 Jul 10 08:58:36 switch vim[2890]: %VIM-5-IF_ATTACHED: Interface Ethernet4/1 is
attached to eth1 on module 4
2018 Jul 10 08:58:36 switch ethpm[2833]: %ETHPORT-5-SPEED: Interface Ethernet4/1,
operational speed changed to 10 Gbps
2018 Jul 10 08:58:36 switch ethpm[2833]: %ETHPORT-5-IF_DUPLEX: Interface Ethernet4/1,
operational duplex mode changed to Full
2018 Jul 10 08:58:36 switch ethpm[2833]: %ETHPORT-5-IF_UP: Interface Ethernet4/1 is up in
mode trunk
2018 Jul 10 08:58:41 switch vms[2885]: %VMS-5-DVS_NAME_CHANGE: Changed dvswitch name to
'switch' on the vCenter Server.
2018 Jul 10 08:58:45 switch msp[2882]: %MSP-5-DOMAIN_CFG_SYNC_DONE: Domain config
successfully pushed to the management server.
2018 Jul 10 08:58:45 switch vshd[3748]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3748
2018 Jul 10 08:58:45 switch last message repeated 1 time
2018 Jul 10 08:58:45 switch vshd[3759]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3759
2018 Jul 10 08:58:45 switch vshd[3748]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3748
2018 Jul 10 08:58:46 switch vshd[3759]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3759
2018 Jul 10 08:58:46 switch vshd[3732]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3732
2018 Jul 10 08:58:46 switch vshd[3759]: %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty
by root on vsh.3759
2018 Jul 10 08:58:48 switch %SYSMGR-STANDBY-4-READCONF_STARTED: Configuration update
started (PID 3266).
2018 Jul 10 08:58:48 switch vms[2885]: %VMS-5-DVS_NAME_CHANGE: Changed dvswitch name to
'switch' on the vCenter Server.
2018 Jul 10 08:58:52 switch %SYSMGR-STANDBY-4-READCONF_STARTED: Configuration update
started (PID 3415).
2018 Jul 10 08:58:54 switch %SYSMGR-STANDBY-4-READCONF_STARTED: Configuration update
started (PID 3489).
2018 Jul 10 08:58:55 switch platform[2301]: %PLATFORM-2-MOD_DETECT: Module 2 detected
(Serial number T5056921301) Module-Type Virtual Supervisor Module Model Nexus1000V
2018 Jul 10 08:58:55 switch module[2437]: %MODULE-5-STANDBY_SUP_OK: Supervisor 2 is
standby
2018 Jul 10 08:58:55 switch %SYSMGR-STANDBY-5-MODULE_ONLINE: System Manager has received
notification of local module becoming online.
2018 Jul 10 08:58:57 switch vms[2885]: %VMS-5-DVPG_CREATE: created port-group
'inside-trunk1' on the vCenter Server.
2018 Jul 10 08:58:57 switch vms[2885]: %VMS-5-DVPG_CREATE: created port-group
'inside-trunk2' on the vCenter Server.
2018 Jul 10 08:58:59 switch vms[2885]: %VMS-5-VMS_PPM_SYNC_COMPLETE: Sync between
Port-Profile Manager and local vCenter Server cache complete
```

```
2018 Jul 10 08:59:03 switch %AUTHPRIV-3-SYSTEM_MSG: pam_aaa:Authentication failed for user
admin from 10.155.81.147 - dcos_sshd[4052]
switch#
```

Example 8-5 show logging logfile | grep Command

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

Example 8-6 show interface brief Command

```
switch# show int brief
--------------------------------------------------------------------------------
Port VRF Status IP Address Speed MTU
--------------------------------------------------------------------------------
mgmt0 -- up 172.23.232.163 1000 1500

--------------------------------------------------------------------------------
Ethernet VLAN Type Mode Status Reason Speed Port
Interface Ch #
--------------------------------------------------------------------------------
Eth3/1 1 eth trunk up none 10G
Eth4/1 1 eth trunk up none 10G

--------------------------------------------------------------------------------
Vethernet VLAN/ Type Mode Status Reason MTU Module
Segment
--------------------------------------------------------------------------------
Veth1 223 virt access up none 1500 3
Veth2 222 virt access up none 1500 3
Veth3 222 virt access up none 1500 4
Veth4 223 virt access up none 1500 4

--------------------------------------------------------------------------------
Port VRF Status IP Address Speed MTU
--------------------------------------------------------------------------------
control0 -- up -- 1000 1500

NOTE : * Denotes ports on modules which are currently offline on VSM
switch#
```

Example 8-7 show interface ethernet Command

```
switch# show interface eth3/1
Ethernet3/1 is up
  Hardware: Ethernet, address: 0050.5653.6345 (bia 0050.5653.6345)
  MTU 1500 bytes, BW -598629368 Kbit, DLY 10 usec,
     reliability 0/255, txload 0/255, rxload 0/255
  Encapsulation ARPA
  Port mode is trunk
  full-duplex, 1000 Mb/s
  Beacon is turned off
  Auto-Negotiation is turned off
```

```
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned on
  Switchport monitor is off
    Rx
    18775 Input Packets 10910 Unicast Packets
    862 Multicast Packets 7003 Broadcast Packets
    2165184 Bytes
    Tx
    6411 Output Packets 6188 Unicast Packets
    216 Multicast Packets 7 Broadcast Packets 58 Flood Packets
    1081277 Bytes
    1000 Input Packet Drops 0 Output Packet Drops
    1 interface resets
switch#
```

Example 8-8 show interface ethernet counters Command

```
switch# show interface eth3/2 counters
--------------------------------------------------------------------------------
Port                InOctets       InUcastPkts      InMcastPkts      InBcastPkts
--------------------------------------------------------------------------------
Eth3/2              2224326         11226           885          7191

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

Example 8-9 show interface vEthernet Command

```
switch# show interface veth1
Vethernet1 is up
    Port description is gentoo1, Network Adapter 1
    Hardware is Virtual, address is 0050.56bd.42f6
    Owner is VM "gentoo1", adapter is Network Adapter 1
    Active on module 33
    VMware DVS port 100
    Port-Profile is vlan48
    Port mode is access
    Rx
    491242 Input Packets 491180 Unicast Packets
    7 Multicast Packets 55 Broadcast Packets
    29488527 Bytes
    Tx
    504958 Output Packets 491181 Unicast Packets
    1 Multicast Packets 13776 Broadcast Packets 941 Flood Packets
    714925076 Bytes
    11 Input Packet Drops 0 Output Packet Drops
switch#
```

Example 8-10 show interface capabilities Command

```
switch# show interface capabilities
Ethernet3/1
  Model:                --
  Type (Non SFP):       --
  Speed:                10,100,1000,10000,auto
  Duplex:               half/full/auto
  Trunk encap. type:    802.1Q
  Channel:              yes
  Broadcast suppression: no
  Flowcontrol:          rx-(none),tx-(none)
```

```
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched

Ethernet4/1
  Model:                 --
  Type (Non SFP):        --
  Speed:                 10,100,1000,10000,auto
  Duplex:                half/full/auto
  Trunk encap. type:     802.1Q
  Channel:               yes
  Broadcast suppression: no
  Flowcontrol:           rx-(none),tx-(none)
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched

Vethernet1
  Model:                 --
  Type (Non SFP):        --
  Speed:                 10,100,1000,10000,auto
  Duplex:                half/full/auto
  Trunk encap. type:     802.1Q
  Channel:               yes
  Broadcast suppression: no
  Flowcontrol:           rx-(none),tx-(none)
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched

Vethernet2
  Model:                 --
  Type (Non SFP):        --
  Speed:                 10,100,1000,10000,auto
  Duplex:                half/full/auto
  Trunk encap. type:     802.1Q
  Channel:               yes
```

```
  Broadcast suppression: no
  Flowcontrol:           rx-(none),tx-(none)
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched

Vethernet3
  Model:                 --
  Type (Non SFP):        --
  Speed:                 10,100,1000,10000,auto
  Duplex:                half/full/auto
  Trunk encap. type:     802.1Q
  Channel:               yes
  Broadcast suppression: no
  Flowcontrol:           rx-(none),tx-(none)
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched

Vethernet4
  Model:                 --
  Type (Non SFP):        --
  Speed:                 10,100,1000,10000,auto
  Duplex:                half/full/auto
  Trunk encap. type:     802.1Q
  Channel:               yes
  Broadcast suppression: no
  Flowcontrol:           rx-(none),tx-(none)
  Rate mode:             none
  QOS scheduling:        rx-(none),tx-(none)
  CoS rewrite:           yes
  ToS rewrite:           yes
  SPAN:                  yes
  UDLD:                  no
  Link Debounce:         no
  Link Debounce Time:    no
  MDIX:                  yes
  TDR capable:           no
  FabricPath capable:    no
  Port mode:             Switched
```

Example 8-11 show interface virtual port-mapping Command

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

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

Port Profiles

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

Information About Port Profiles

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

In VMware vCenter Server, a port profile is represented as a port group. The ethernet interfaces are assigned in vCenter Server to a port profile for the following reasons:

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

Ethernet port profiles can be assigned by the server administrator to physical ports (a VMNIC or a PNIC). Port profiles not configured as Ethernet can be assigned to a VM virtual port.

Note While a manual interface configuration overrides that of the port profile, we do not recommend that you do so. Manual interface configuration is only used, for example, to quickly test a change or allow a port to be disabled without having to change the inherited port profile.

Note For VSG protected ports, some configurations related to vservice will be visible under interface level (**show running-config interface** command), even after removing the vservice configuration from port-profile or changing to non-protected port-profile; but it does not affect new port-profile related functionality. However, when any interface is being moved from VSG protected profile to a non-protected port-profile, it is recommended to move the port to VM Network port-profile. After moving the port, delete the vethernet interface using **no interface vethernet** command and then move the port to a new port-profile.

For more information about assigning port profiles to physical or virtual ports, see your VMware documentation.

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

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

To verify port profile inheritance, use the following command:

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

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

Note Inherited port profiles are automatically configured by the Cisco Nexus 1000VE when the ports are attached on the hosts. This action is done by matching up the VMware port group assigned by the system administrator with the port profile that created it.

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

Problems with Port Profiles

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

Symptom	Possible Causes	Solution
You do not see the port group on vCenter Server or the following message is displayed: `Warning: Operation succeeded locally but update failed on vCenter server. Please check if you are connected to vCenter Server.`	The connection to vCenter server is down.	1. Verify that the connection to vCenter Server is Enabled and Connected. **show svs connections** 2. Reconnect to vCenter server. For detailed instructions, see the *Connecting to vCenter Server* procedure in the *Cisco Nexus 1000V System Management Configuration Guide*.
	The domain configuration was not successfully pushed to vCenter server.	1. Verify that the domain configuration was successfully pushed to vCenter Server. **show svs domain** 2. Fix any problems with the domain configuration. For information about configuring the domain, see the *Cisco Nexus 1000V System Management Configuration Guide*.
	The port profile is configured incorrectly.	1. Verify that the **vmware port-group** is configured for the port profile and that the port profile is enabled. **show port profile name** *name* 2. Fix the port profile using the procedures in the *Cisco Nexus 1000V Port Profile Configuration Guide*.
A port configuration is not applied to an interface.	Management connectivity between vCenter server and the VSM has prevented the port profile assignment from being sent or received.	1. Display the port profile usage by interface. **show port-profile virtual usage** 2. Verify that the interface level configuration did not overwrite the port profile configuration. **show run** **show port-profile expand-interface** 3. If the **show** command output is incorrect, on vCenter server, reassign the port group to the interface.

Symptom	Possible Causes	Solution
An Ethernet interface or vEthernet interface is administratively down. A system message similar to the following is logged: `%VMS-3-DVPG_NICS_MOVED: '1' nics have been moved from port-group 'Access483' to 'Unused_Or_Quarantine_Veth'.`	The interface is inheriting a quarantined port profile. 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: • Unused_Or_Quarantine_Uplink for ethernet types • Unused_Or_Quarantine_Veth for Vethernet types	1. Verify the port profile-to-interface mapping. **show port-profile virtual usage** 2. Reassign the VMNIC or PNIC to a non-quarantined port group to enable the interface to be up and forwarding traffic. This requires changing the port group on vCenter Server.
After applying a port profile, an online interface is quarantined. A system message similar to the following is logged: `%PORT-PROFILE-2-INTERFACE_QUARANTINED: Interface Ethernet3/3 has been quarantined due to Cache Overrun`	The assigned port profile is incorrectly configured. The incorrect command fails when the port profile is applied to an interface. Although a specific command fails, the port profile-to-interface mapping is created.	1. Identify the command that failed. **show accounting log | grep FAILURE** 2. Verify that the interface is quarantined. **show port-profile sync-status** 3. Verify the port profile-to-interface mapping. **show port-profile virtual usage** 4. Fix the error in the port profile using the procedures in the *Cisco Nexus 1000V Port Profile Configuration Guide*. 5. Bring the interface out of quarantine. **no shutdown** The interface comes back online. 6. Return shutdown control to the port profile. **default shutdown**
After modifying a port profile, an assigned offline interface is quarantined. A system message similar to the following is logged: `%PORT-PROFILE-2-INTERFACE_QUARANTINED: Interface Ethernet4/3 has been quarantined due to Cache Overrun`	The interface has been removed from the DVS.	To bring the interface back online, see the "Recovering a Quarantined Offline Interface" section on page 9-5.
A module and all associated interfaces are offline. A system message similar to the following is logged: `2011 Mar 2 22:28:50 switch %VSE_MGR-2-VSE_MGR_REMOVE_NO_HB: Removing VSE 3 (heartbeats lost)` `2011 Mar 2 22:29:00 switch %VSE_MGR-2-MOD_OFFLINE: Module 3 is offline`	• The VSE or the underlying host was powered down. • There is a general loss of connectivity to the module.	Follow VSE troubleshooting guidelines to bring the module back online To bring the interface back online, see the "Recovering a Quarantined Offline Interface" section on page 9-5.

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 has been quarantined. The interface appears in the **show** command output.

show port-profile sync-status

Step 2 On vCenter server, add or associate the PNIC to a port profile (either the original port profile or a different port profile).

The interface comes back online.

Step 3 Verify that the interface has come back online.

show interface brief

Step 4 Verify the port profile-to-interface mapping.

show port-profile virtual usage

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

show port-profile sync-status

Step 6 Return shutdown control to the port profile.

default shutdown

Port Profile Logs

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

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

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

Port Profile Troubleshooting Commands

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

Command	Purpose
show port-profile	Displays the port profile configuration. See Example 9-1 on page 9-7.
show port-profile name *name*	Displays the configuration for a named port profile. See Example 9-2 on page 9-8.
show port-profile brief	Displays a tabular view of all configured port profiles. See Example 9-3 on page 9-8.
show port-profile expand-interface	Displays all configured port profiles expanded to include the interfaces assigned to them. See Example 9-4 on page 9-9.
show port-profile expand-interface name *name*	Displays a named port profile expanded to include the interfaces assigned to it. See Example 9-5 on page 9-9.
show port-profile-role [**name** *port-profile-role-name*]	Displays the port profile role configuration, including role names, descriptions, assigned users, and assigned groups. See Example 9-7 on page 9-11.
show running-config port-profile [*profile-name*]	Displays the port profile configuration. See Example 9-6 on page 9-10.
show port-profile-role	Displays the port profile role configuration. See Example 9-7 on page 9-11.
show port-profile-role users	Displays the available users and groups. See Example 9-8 on page 9-11.
show port-profile virtual usage [**name** *profile-name*]	Displays the port profile usage by interface. See Example 9-9 on page 9-11.
show msp internal info	Displays the port profile mappings on vCenter server and configured roles. See Example 9-10 on page 9-11.

Command	Purpose
show system internal port-profile profile-fsm	Displays the port profile activity on the Cisco Nexus 1000VE, including transitions such as inherits and configurations. If the following displays, then all inherits are processed: `Curr state: [PPM_PROFILE_ST_SIDLE]` See Example 9-11 on page 9-15.
show system internal port-profile event-history msgs	Displays the messages logged about port profile events within the Cisco Nexus 1000VE. See Example 9-12 on page 9-16.

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

EXAMPLES

Example 9-1 show port-profile Command

```
switch# show port-profile
port-profile inside-trunk1
 type: Vethernet
 description: Port-group created for Nexus 1000V internal usage. Do not use.
 status: enabled
 max-ports: 32
 min-ports: 1
 inherit:
 config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 1-50
  no shutdown
 evaluated config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 1-50
  no shutdown
 assigned interfaces:
 port-group: inside-trunk1
 system vlans: none
 capability l3control: no
 capability iscsi-multipath: no
 capability vxlan: no
 capability l3-vservice: no
 port-profile role: none
 port-binding: static

port-profile inside-trunk2
 type: Vethernet
 description: Port-group created for Nexus 1000V internal usage. Do not use.
 status: enabled
 max-ports: 32
 min-ports: 1
 inherit:
 config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 2047-2096
  no shutdown
 evaluated config attributes:
  switchport mode trunk
  switchport trunk allowed vlan 2047-2096
```

```
 no shutdown
assigned interfaces:
port-group: inside-trunk2
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vservice: no
port-profile role: none
port-binding: static
```

Example 9-2 show port-profile name Command

```
switch# show port-profile name vlan222
port-profile vlan222
type: Vethernet
description:
status: enabled
max-ports: 32
min-ports: 1
inherit:
config attributes:
switchport mode access
switchport access vlan 222
no shutdown
evaluated config attributes:
switchport mode access
switchport access vlan 222
no shutdown
assigned interfaces:
Vethernet2
Vethernet3
port-group: vlan222
system vlans: none
capability l3control: no
capability iscsi-multipath: no
capability vxlan: no
capability l3-vservice: no
port-profile role: none
port-binding: static

switch#
```

Example 9-3 show port-profile brief Command

```
switch# show port-profile brief
--------------------------------------------------------------------------------
Port Profile Profile Eval Max Assigned Child
Profile Type State Items Ports Ports Profs
--------------------------------------------------------------------------------
inside-trunk1 Vethernet 1 3 32 0 0
inside-trunk2 Vethernet 1 3 32 0 0
l3ctrl Vethernet 1 3 32 0 0
outside-trunk Ethernet 1 3 512 0 0
Unused_Or_Quarantine_Veth Vethernet 1 0 32 0 0
uplink-pp Ethernet 1 3 512 2 0
vlan222 Vethernet 1 3 32 2 0
vlan223 Vethernet 1 3 32 2 0
--------------------------------------------------------------------------------
Profile Assigned Total Sys Parent Child UsedBy
Type Intfs Prfls Prfls Prfls Prfls Prfls
```

```
--------------------------------------------------------------------------------
Vethernet 4 6 0 6 0 2
Ethernet 2 2 0 2 0 1
switch#
```

Example 9-4 show port-profile expand-interface Command

```
switch# show port-profile expand-interface
port-profile inside-trunk1

port-profile inside-trunk2

port-profile l3ctrl

port-profile outside-trunk

port-profile Unused_Or_Quarantine_Veth

port-profile uplink-pp
 Ethernet3/1
  switchport mode trunk
  switchport trunk allowed vlan 181,220-229
  no shutdown
 Ethernet4/1
  switchport mode trunk
  switchport trunk allowed vlan 181,220-229
  no shutdown

port-profile vlan222
 Vethernet2
  switchport mode access
  switchport access vlan 222
  no shutdown
 Vethernet3
  switchport mode access
  switchport access vlan 222
  no shutdown

port-profile vlan223
 Vethernet1
  switchport mode access
  switchport access vlan 223
  no shutdown
 Vethernet4
  switchport mode access
  switchport access vlan 223
  no shutdown
```

Example 9-5 show port-profile expand-interface name Command

```
switch# show port-profile expand-interface name uplink-pp
port-profile uplink-pp
Ethernet3/1
switchport mode trunk
switchport trunk allowed vlan 181,220-229
no shutdown
Ethernet4/1
switchport mode trunk
switchport trunk allowed vlan 181,220-229
no shutdown
switch#
```

Example 9-6 show running-config port-profile Command

```
switch# show running-config port-profile
version 5.2(1)SV5(1.1)
port-profile default max-ports 32
port-profile type vethernet Unused_Or_Quarantine_Veth
  shutdown
  port-binding static auto expand
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type ethernet outside-trunk
  switchport mode trunk
  switchport trunk allowed vlan 1-3967,4048-4093
  no shutdown
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type vethernet inside-trunk1
  switchport mode trunk
  switchport trunk allowed vlan 1-50
  no shutdown
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type vethernet inside-trunk2
  switchport mode trunk
  switchport trunk allowed vlan 2047-2096
  no shutdown
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type vethernet vlan222
  switchport mode access
  switchport access vlan 222
  no shutdown
  state enabled
  vmware port-group
port-profile type ethernet uplink-pp
  switchport mode trunk
  switchport trunk allowed vlan 181,220-229
  no shutdown
  state enabled
  vmware port-group
port-profile type vethernet l3ctrl
  switchport mode access
  switchport access vlan 181
  no shutdown
  state enabled
  vmware port-group
port-profile type vethernet vlan223
  switchport mode access
  switchport access vlan 223
  no shutdown
  state enabled
  vmware port-group

interface Vethernet1
  inherit port-profile vlan223

interface Vethernet2
  inherit port-profile vlan222

interface Vethernet3
```

```
  inherit port-profile vlan222

interface Vethernet4
  inherit port-profile vlan223

interface Ethernet3/1
  inherit port-profile uplink-pp

interface Ethernet4/1
  inherit port-profile uplink-pp
```

Example 9-7 show port-profile-role Command

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

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

Example 9-8 show port-profile-role users Command

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

Example 9-9 show port-profile virtual usage Command

```
switch# show port-profile virtual usage
--------------------------------------------------------------------------------
Port Profile Port Adapter Owner
--------------------------------------------------------------------------------
vlan222 Veth2 Net Adapter 2 test-vm1
Veth3 Net Adapter 2 test-vm3
uplink-pp Eth3/1 eth1 172.23.233.17
Eth4/1 eth1 172.23.181.156
vlan223 Veth1 Net Adapter 2 test-vm2
Veth4 Net Adapter 2 test-vm4
switch#
```

Example 9-10 show msp internal info Command

```
switch# show msp internal info
port-profile inside-trunk1
  id: 3
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
```

```
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 1
  marked for del: 0
  active used ports count: 3
  intf inherit count: 0
  Port-profile alias information
    pg name: inside-trunk1
    dvs:  (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
    pg id: inside-trunk1
      dvs uuid:
      type: 1
    pg id: dvportgroup-1676
      dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
      type: 2
    pg id: 37a5cc5a-81f2-44dc-94ee-76e9bf7e766e
      dvs uuid:
      type: 11
port-profile inside-trunk2
  id: 4
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 1
  marked for del: 0
  active used ports count: 3
  intf inherit count: 0
  Port-profile alias information
    pg name: inside-trunk2
    dvs:  (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
    pg id: inside-trunk2
      dvs uuid:
      type: 1
    pg id: dvportgroup-1677
      dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
      type: 2
    pg id: 1d066e56-afcd-46e3-a9ad-b643842e166c
      dvs uuid:
      type: 11
port-profile l3ctrl
  id: 7
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 1
  marked for del: 0
  active used ports count: 0
```

```
  intf inherit count: 0
  Port-profile alias information
    pg name: l3ctrl
    dvs:  (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
    pg id: l3ctrl
      dvs uuid:
      type: 1
    pg id: dvportgroup-1678
      dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
      type: 2
    pg id: 301ffcc4-a296-411b-ad9c-b598bfdcf59c
      dvs uuid:
      type: 11
port-profile outside-trunk
  id: 2
  capability: 0x1
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 512
  min ports: 1
  marked for del: 0
  active used ports count: 0
  intf inherit count: 0
  Port-profile alias information
    pg name: outside-trunk
    dvs:  (ignore)
    reserved ports: 512
  port-profile role:
  alias information:
    pg id: outside-trunk
      dvs uuid:
      type: 1
    pg id: dvportgroup-1679
      dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
      type: 2
    pg id: eb445392-b9f9-4c8b-9463-add3d1729d1d
      dvs uuid:
      type: 11
port-profile Unused_Or_Quarantine_Veth
  id: 1
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 2
  max ports: 32
  min ports: 1
  marked for del: 0
  active used ports count: 0
  intf inherit count: 0
  Port-profile alias information
    pg name: Unused_Or_Quarantine_Veth
    dvs:  (ignore)
    reserved ports: 1
  port-profile role:
```

```
    alias information:
      pg id: Unused_Or_Quarantine_Veth
        dvs uuid:
        type: 1
      pg id: dvportgroup-1680
        dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
        type: 2
      pg id: 1c176aec-02d2-4377-9fae-4d278548dfe5
        dvs uuid:
        type: 11
  port-profile uplink-pp
    id: 6
    capability: 0x1
    state: 0x1
    type: 0x1
    system vlan mode: -
    system vlans:
    port-binding: static
    bind_opts: 0
    max ports: 512
    min ports: 1
    marked for del: 0
    active used ports count: 0
    intf inherit count: 0
    Port-profile alias information
      pg name: uplink-pp
      dvs:  (ignore)
      reserved ports: 512
    port-profile role:
    alias information:
      pg id: uplink-pp
        dvs uuid:
        type: 1
      pg id: dvportgroup-1681
        dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
        type: 2
      pg id: d876121c-8688-4de3-bc9e-68ab7eed06ba
        dvs uuid:
        type: 11
  port-profile vlan222
    id: 5
    capability: 0x0
    state: 0x1
    type: 0x1
    system vlan mode: -
    system vlans:
    port-binding: static
    bind_opts: 0
    max ports: 32
    min ports: 1
    marked for del: 0
    active used ports count: 0
    intf inherit count: 0
    Port-profile alias information
      pg name: vlan222
      dvs:  (ignore)
      reserved ports: 32
    port-profile role:
    alias information:
      pg id: vlan222
        dvs uuid:
        type: 1
      pg id: dvportgroup-1682
        dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
```

```
      type: 2
    pg id: c5e59050-7ba8-48ab-bba5-65b32532ca5a
      dvs uuid:
      type: 11
port-profile vlan223
  id: 8
  capability: 0x0
  state: 0x1
  type: 0x1
  system vlan mode: -
  system vlans:
  port-binding: static
  bind_opts: 0
  max ports: 32
  min ports: 1
  marked for del: 0
  active used ports count: 0
  intf inherit count: 0
  Port-profile alias information
    pg name: vlan223
    dvs:  (ignore)
    reserved ports: 32
  port-profile role:
  alias information:
    pg id: vlan223
      dvs uuid:
      type: 1
    pg id: dvportgroup-1683
      dvs uuid: 50 12 e0 5d 1c 63 22 76-7b 77 69 b7 27 dc 0c 2e
      type: 2
    pg id: 8176c0d3-f714-4f86-91bf-f8584074b44a
      dvs uuid:
      type: 11
pending binds:
  global_inherit_ifindex_count: 0
  global_inherit_info.rt_data.restored_from_pss: 0
  global_inherit_info.rt_data.inherit_in_progress: 0
  third_party_app_conf.connection_state[VMWARE_VC] =1
  third_party_app_conf.sync_state[VMWARE_VC] = 1PPM restore_complete:TRUE
  opq_data_info.ppm_sdb_restored:1
Unable to read nsmgr_restore_state
  opq_data_info.nsm_sdb_restored:0
```

Example 9-11 show system internal port-profile profile-fsm Command

```
switch# 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]
    switch#
```

Example 9-12 show system internal port-profile event-history msgs Command

```
switch# 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
    ...
```

Layer 2 Switching

This chapter describes how to identify and resolve problems that relate to Layer 2 switching and includes the following sections:

Information About Layer 2 Ethernet Switching

The Cisco Nexus1000VE is a distributed Layer 2 virtual switch that extends across many virtualized hosts.

It consists of two components:

- The Virtual Supervisor Module (VSM), which is also known as the control plane (CP). The VSM acts as the supervisor and contains the Cisco CLI, configuration, and high-level features.
- The Virtual Services Engine (VSE), which is also known as the data plane (DP). The VSE acts as a line card and runs as a VM in each virtualized server to handle packet forwarding and other localized functions.

Port Model

This section includes the following topics:

Viewing Ports from the VSE

The Cisco Nexus1000VE differentiates between virtual and physical ports on each of the VSEs. Figure 10-1 shows how ports on the Cisco Nexus1000VE switch are bound to physical and virtual VMware ports within a VSE.

Figure 10-1 VSE View of Ports

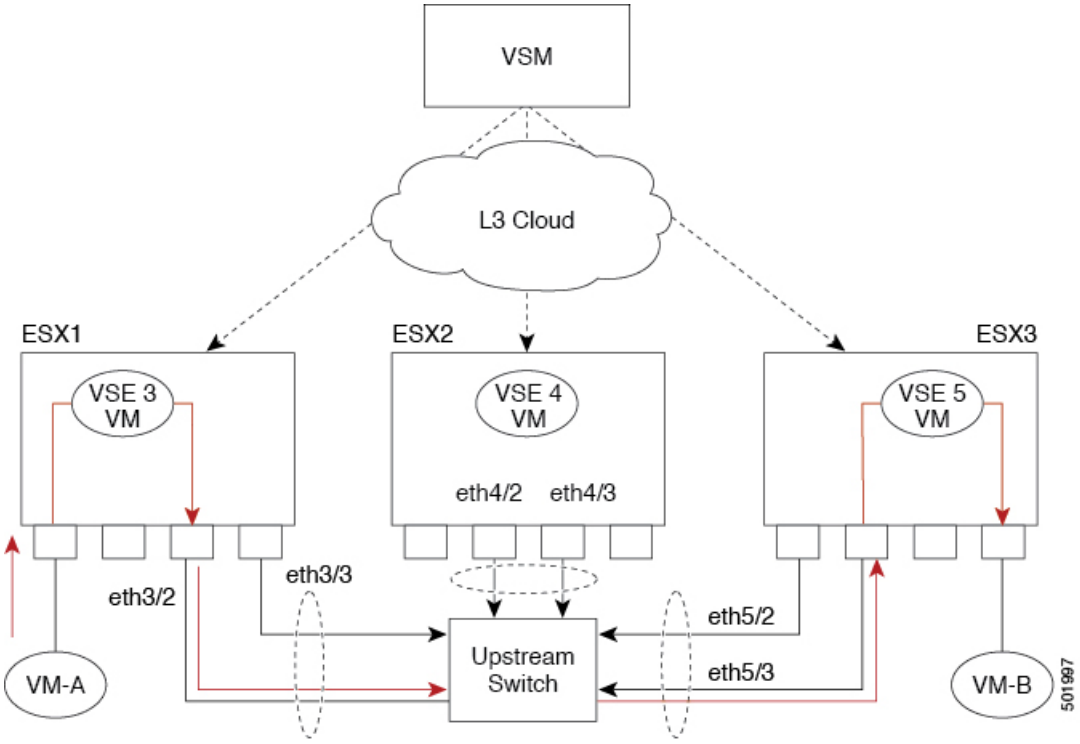

On the virtual side of the switch, three layers of ports are mapped together:

- Virtual NICs—Three types of Virtual NICs are in VMware. The virtual NIC (vnic) is part of the VM and represents the physical port of the host that is plugged into the switch. The virtual kernel NIC (VTEP) is used by the hypervisor for management, VMotion, iSCSI, network file system (NFS), and other network access needed by the kernel. This interface carries the IP address of the hypervisor itself and is also bound to a virtual Ethernet port. The vswif (not shown) appears only in CoS-based systems and is used as the VMware management port. Each type maps to a virtual Ethernet port within the Cisco Nexus1000VE.
- Virtual Ethernet Ports (VEth)—A vEth port is a port on the Cisco Nexus 1000V. 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 running the VM.

  Virtual Ethernet ports are assigned to port groups.
- Local Virtual Ethernet Ports (lveth)—Each host has a number of local vEth ports. These ports are dynamically selected for vEth ports that are needed on the host.

  These local ports do not move and are addressable by the module/port number method.

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

- Each physical NIC in VMware is represented by an interface called a vmnic. The vmnic number is allocated during VMware installation, or when a new physical NIC is installed, and remains the same for the life of the host.

Viewing Ports from the VSM

Figure 10-2 shows the VSM view ports.

Figure 10-2 VSM View of Ports

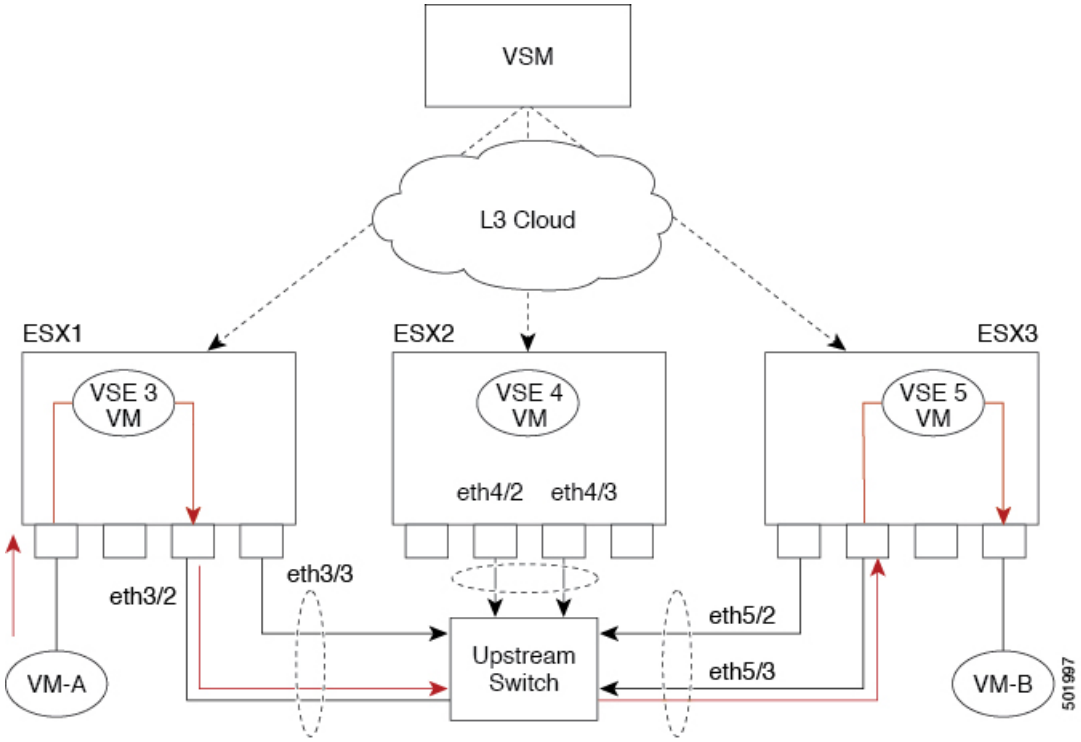

Port Types

The following types of ports are available:

- vEths can be associated with any one of the following:
  - VNICs of a Virtual Machine on the ESX host.
  - VTEPs of the ESX Host
  - VSWIFs of an ESX COS Host.
- Eths (physical Ethernet interfaces)—Correspond to the outside-trunk interface of the VSEs.

For more information about Layer 2 switching, see the *Cisco Nexus 1000VE Layer 2 Switching Configuration Guide*.

Layer 2 Switching Problems

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

Verifying a Connection Between VSE Ports

You can verify a connection between two vEth ports on a VSE.

Step 1 View the state of the VLANs associated with the port. If the VLAN associated with a port is not active, the port may be down. In this case, you must create the VLAN and activate it.

switch# **show vlan** v*lan-id*

Step 2 View the state of the ports on the VSM.

switch# **show interface brief**

Step 3 Display the ports that are present on the VSE, their local interface indices, VLAN, type (physical or virtual), port mode and port name.

switch# **module vse** *module-number* **execute vemcmd show port**

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 to identify the interface quickly in other VSE 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 port lists on a particular VSE.

switch# **module vse** *module-number* **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 VSEs

You can verify a connection between vEth ports on two separate VSEs.

Step 1 Log in to the upstream switch and make sure that the port is configured to allow the VLAN that you are looking for.

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

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

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

Isolating Traffic Interruptions

You can isolate the cause for no traffic passing across VMs on different VSEs.

Step 1 Inside the VM, verify that the Ethernet interface is up.

ifconfig –a

If not, delete that NIC from the VM, and add another NIC.

Step 2 Using any sniffer tool, verify that ARP requests and responses are received on the VM interface.

Step 3 On the upstream switch, look for the association between the IP and MAC address:

debug arp
show arp

Example:

```
switch# 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
switch#
```

Example:

```
switch# 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
```

Step 4 You have completed this procedure.

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
show mac address-table	Displays the MAC address table to verify all MAC addresses on all VSEs controlled by the VSM. See Example 10-1 on page 10-7.
show mac address-table module *module-number*	Displays all the MAC addresses on the specified VSE.
show mac address-table static *HHHH.WWWW.HHHH*	Displays the MAC address table static entries. See Example 10-2 on page 10-8.
show mac address-table address *HHHH.WWWW.HHHH*	Displays the interface on which the MAC address specified is learned or configured. • For dynamic MAC addresses, if the same MAC address appears on multiple interfaces, each of them is displayed separately. • For static MAC addresses, if the same MAC address appears on multiple interfaces, only the entry on the configured interface is displayed.
show mac address-table static | inc veth	Displays the static MAC address of vEthernet interfaces in case a VSE physical port learns a dynamic MAC addrress and the packet source is in another VSE on the same VSM. See Example 10-3 on page 10-8.
show running-config vlan *vlan-id*	Displays VLAN information in the running configuration.
show vlan [**all-ports** | **brief** | **id** *vlan-id* | **name** *name* | **dot1q tag native**]	Displays VLAN information as specified. See Example 10-4 on page 10-8.
show vlan summary	Displays a summary of VLAN information.
show interface brief	Displays a table of interface states. See Example 10-5 on page 10-9.
module vse *module-number* **execute vemcmd show port**	On the VSE, displays the port state on a particular VSE. This command can only be used from the VSE. See Example 10-6 on page 10-9.
module vse *module-number* **execute vemcmd show bd**	For the specified VSE, displays its VLANs and their port lists. See Example 10-7 on page 10-10.

Command	Purpose
module vse *module-number* **execute vemcmd show trunk**	For the specified VSE, displays the VLAN state on a trunk port**.** • If a VLAN is forwarding (active) on a port, its CBL state should be 1. • If a VLAN is blocked, its CBL state is 0. See Example 10-8 on page 10-11.
module vse *module-number* **execute vemcmd show l2** *vlan-id*	For the specified VSE, displays the VLAN forwarding table for a specified VLAN**.** See Example 10-9 on page 10-11.
show interface *interface_id* **mac**	Displays the MAC addresses and the burn-in MAC address for an interface.

Example 10-1 show mac address-table Command

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

Tip The "Module" indicates the VSE on which this MAC address is seen.

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

```
switch# show mac address-table
VLAN      MAC Address       Type    Age       Port                            Mod
---------+-----------------+-------+---------+-------------------------------+---
1         0002.3d23.7802    static  0         N1KV Internal Port              3
1         0002.3d33.7802    static  0         N1KV Internal Port              3
1         0002.3d43.7802    static  0         N1KV Internal Port              3
1         0002.3d63.7802    static  0         N1KV Internal Port              3
1         0002.3d83.7802    static  0         N1KV Internal Port              3
222       0050.56b8.7584    static  0         Veth2                           3
222       d48c.b5bc.fe01    dynamic 0         Eth3/1                          3
223       0050.56b8.0375    static  0         Veth1                           3
3968      0002.3d83.7802    static  0         N1KV Internal Port              3
3970      0002.3d83.7802    static  0         N1KV Internal Port              3
3971      0002.3d83.7802    static  0         N1KV Internal Port              3
3972      0002.3d83.7802    static  0         N1KV Internal Port              3
1         0002.3d23.7803    static  0         N1KV Internal Port              4
1         0002.3d33.7803    static  0         N1KV Internal Port              4
1         0002.3d43.7803    static  0         N1KV Internal Port              4
1         0002.3d63.7803    static  0         N1KV Internal Port              4
1         0002.3d83.7803    static  0         N1KV Internal Port              4
222       0050.56b8.8ce8    static  0         Veth3                           4
223       0050.56b8.99b6    static  0         Veth4                           4
3968      0002.3d83.7803    static  0         N1KV Internal Port              4
3970      0002.3d83.7803    static  0         N1KV Internal Port              4
3971      0002.3d83.7803    static  0         N1KV Internal Port              4
3972      0002.3d83.7803    static  0         N1KV Internal Port              4
Total MAC Addresses: 23
```

Example 10-2 show mac address-table address Command

Tip This command shows all interfaces on which a MAC is learned dynamically. In this example, the same MAC appears on Eth3/1 and Eth4/1.

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

Example 10-3 show mac address-table static | inc veth Command

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

Example 10-4 show vlan Command

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

```
switch# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Eth3/1, Eth4/1
110  VLAN0110                         active
111  VLAN0111                         active
112  VLAN0112                         active
113  VLAN0113                         active
114  VLAN0114                         active
115  VLAN0115                         active
116  VLAN0116                         active
117  VLAN0117                         active
118  VLAN0118                         active
119  VLAN0119                         active
800  VLAN0800                         active
801  VLAN0801                         active
802  VLAN0802                         active
803  VLAN0803                         active
804  VLAN0804                         active
805  VLAN0805                         active
806  VLAN0806                         active
807  VLAN0807                         active
808  VLAN0808                         active
809  VLAN0809                         active
810  VLAN0810                         active
811  VLAN0811                         active
812  VLAN0812                         active
813  VLAN0813                         active
814  VLAN0814                         active
815  VLAN0815                         active
816  VLAN0816                         active
```

```
817  VLAN0817                         active
818  VLAN0818                         active
819  VLAN0819                         active
820  VLAN0820                         active
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
---- -------------------------------- --------- -------------------------------

Remote SPAN VLANs
------------------------------------------------------------------------------

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

Example 10-5 show interface brief Command

```
switch# show interface brief
--------------------------------------------------------------------------------
Port     VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0    --           up     172.23.232.163                          1000     1500

--------------------------------------------------------------------------------
Ethernet     VLAN    Type Mode    Status  Reason                     Speed    Port
Interface                                                                    Ch #
--------------------------------------------------------------------------------
Eth3/1       1       eth  trunk   up      none                       10G
Eth4/1       1       eth  trunk   up      none                       10G

--------------------------------------------------------------------------------
Vethernet    VLAN/    Type Mode    Status  Reason                      MTU  Module
             Segment
--------------------------------------------------------------------------------
Veth1        223      virt access  up      none                        1500 3
Veth2        222      virt access  up      none                        1500 3
Veth3        222      virt access  up      none                        1500 4
Veth4        223      virt access  up      none                        1500 4

--------------------------------------------------------------------------------
Port     VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
control0 --           up     --                                      1000     1500

NOTE : * Denotes ports on modules which are currently offline on VSM
```

Example 10-6 module vse** module-number **execute vemcmd show port Command

Tip Look for the state of the port.

```
siwtch# module vse 3 execute vemcm show port
LTL   VSM Port  Admin Link  State  PC-LTL  SGID          Vem Port  Type      ORG
svcpath Owner
   21     Eth3/1     UP   UP    F/B*      0                 eth1             0
0 dpdk-outside
   53      Veth2     UP   UP    FWD       0          test-vm1.eth1           0
0 test-vm1
   54      Veth1     UP   UP    FWD       0          test-vm2.eth1           0
0 test-vm2
```

```
* F/B: Port is BLOCKED on some of the vlans.
       One or more vlans are either not created or
       not in the list of allowed vlans for this port.
 Please run "vemcmd show port vlans" to see the details.
```

Example 10-7 module vse** 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.

```
switch# module vse 3 execute vemcmd show bd
BD 1, vdc 1, vlan 1, swbd 1, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
     12  _l24

BD 2, vdc 1, vlan 3972, swbd 3972, table-id 0, 0 ports, ""
Forward type: L2
Portlist:
BD 3, vdc 1, vlan 3970, swbd 3970, table-id 0, 0 ports, ""
Forward type: L2
Portlist:
BD 4, vdc 1, vlan 3968, swbd 3968, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
     11  _l23

BD 5, vdc 1, vlan 3971, swbd 3971, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
     15  _l27

BD 6, vdc 1, vlan 222, swbd 222, table-id 0, 2 ports, ""
Forward type: L2
Portlist:
     21  eth1
     53  test-vm1.eth1

BD 7, vdc 1, vlan 220, swbd 220, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
     21  eth1

BD 8, vdc 1, vlan 221, swbd 221, table-id 0, 1 ports, ""
Forward type: L2
Portlist:
     21  eth1

BD 9, vdc 1, vlan 223, swbd 223, table-id 0, 2 ports, ""
Forward type: L2
Portlist:
     21  eth1
     54  test-vm2.eth1
```

Example 10-8 module vse** 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.

```
switch# module vse 3 execute vemcmd show trunk
Trunk port 6 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 16 native_vlan 1 CBL 1
vlan(1) cbl 1, vlan(3972) cbl 1, vlan(3970) cbl 1, vlan(3968) cbl 1, vlan(3971) cbl 1,
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
Trunk port 21 native_vlan 1 CBL 0
vlan(222) cbl 1, vlan(220) cbl 1, vlan(221) cbl 1, vlan(223) cbl 1, vlan(224) cbl 1,
vlan(225) cbl 1, vlan(226) cbl 1, vlan(227) cbl 1, vlan(228) cbl 1, vlan(229) cbl 1,
switch#
switch# module vse 3 execute vemcmd show l2
switch# module vse 3 execute vemcmd show l2 222
Bridge domain    6 brtmax 4096, brtcnt 2, timeout 300
VLAN 222, swbd 222, ""
Flags:  P - PVLAN  S - Secure  D - Drop  R - Router-mac
       Type         MAC Address   LTL   timeout   Flags    PVLAN
    Dynamic   d4:8c:b5:bc:fe:01    21         1
     Static   00:50:56:b8:75:84    53         0
```

Example 10-9 module vse** module-number **execute vemcmd show l2 Command

```
~ # module vse 5 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
```

Limitations and Restrictions

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

- PVLAN port
- Ports configured with unknown unicast flood blocking (UUFB)

Disabling Automatic Static MAC Learning on a vEthernet Interface

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

In interface configuration mode enter the following commands:

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

In port profile configuration mode enter the following commands:

```
switch(config)# port-profile type vethernet ms-nlb
```

```
switch(config-port-prof)# no mac auto-static-learn
```

Checking Status on a VSM

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

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

Step 1 In interface configuration mode, generate the VSM status.

```
switch(config-if)# show running-config int veth1
interface Vethernet1
  inherit port-profile vm59
  description Fedora117, Network Adapter 2
  no mac auto-static-learn
  vmware dvport 32 dvswitch uuid "ea 5c 3b 50 cd 00 9f 55-41 a3 2d 61 84 9e 0e c4"
```

Step 2 In port profile configuration mode, generate the VSM status.

```
switch(config-if)# show running-config port-profile ms-nlb
port-profile type vethernet ms-nlb
  vmware port-group
  switchport mode access
  switchport access vlan 59
  no mac auto-static-learn
  no shutdown
  state enabled
```

Checking the Status on a VSE

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

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

Step 1 Generate the VSE status.

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

Step 2 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
```

VLANs

This chapter describes how to identify and resolve problems that might occur when implementing VLANs and includes the following sections:

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 that you use only the following characters in a VLAN name:

- a–z or A–Z
- 0–9
- - (hyphen)
- _ (underscore)

Consider the following guidelines for VLANs:

- Keep user traffic off the management VLAN; keep the management VLAN separate from user data.

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

- IGMP 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.)
- SPAN could be configured to span both primary and secondary VLANs or to span either one if the you are is interested only in ingress or egress traffic.
- A MAC address learned in a secondary VLAN is placed in the shared table of the primary VLAN. When the secondary VLAN is associated to the primary VLAN, their MAC address tables are merged into one, shared MAC table.

Initial Troubleshooting Checklist

Troubleshooting a VLAN problem involves gathering information about the configuration and connectivity of individual devices and the entire network. In the case of VLANs, begin your troubleshooting activity as follows:

Checklist	√
Verify the physical connectivity for any problem ports or VLANs.	
Verify that both end devices are in the same VLAN.	

The following CLI commands are used to display VLAN information:

- **show system internal private-vlan info**
- **show system internal private-vlan event-history errors**
- **show system internal private-vlan event-history traces**
- **show vlan id** *vlan-id*
- **show vlan private-vlan**
- **show vlan all-ports**
- **show vlan private-vlan type**
- **show vlan internal bd-info vlan-to-bd 1**
- **show vlan internal errors**
- **show vlan internal info**
- **show vlan internal event-history errors**

Cannot Create a VLAN

Symptom	Possible Cause	Solution
Cannot create a VLAN.	Using a reserved VLAN ID	VLANs 3968 to 4047 and 4094 are reserved for internal use and cannot be changed.

Private VLANs (PVLANs)

This chapter describes how to identify and resolve problems related to private VLANs and includes the following sections:

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 Layer 2 isolated. 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 Domains

A private VLAN domain consists of one or more pairs of VLANs. The primary VLAN makes up the domain, and each VLAN pair makes up a subdomain. The VLANs in a pair are called the primary VLAN and the secondary VLAN. All VLAN pairs within a private VLAN have the same primary VLAN. The secondary VLAN ID is what differentiates one subdomain from another.

Spanning Multiple Switches

Private VLANs can span multiple switches, just like regular VLANs. Inter-switch link ports do not need to be aware of the special VLAN type and can carry frames tagged with these VLANs as like 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 along with the packet, you can maintain consistent behavior throughout the network. 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 1000VE Layer 2 Switching Configuration Guide*.

Troubleshooting Guidelines

Follow these guidelines when troubleshooting private VLAN issues:

- Use the **show vlan** *private-vlan* command to verify that a private VLAN is configured correctly.
- Use the **show interface** *switchport* command to verify the interface is up.
- Use the **module vse** *module-number* **execute vemcmd show port** command to verify the VSE is configured correctly.

Private VLAN Troubleshooting Commands

Use the commands listed in this section to troubleshoot problems related to private VLANs.

Command	Purpose
show vlan private-vlan	Displays that a private VLAN is configured correctly. See Example 12-1 on page 12-2.
show interface *name*	Displays that a physical Ethernet interface in a private VLAN trunk promiscuous mode is up. See Example 12-2 on page 12-3.
show interface *veth-name*	Displays that a virtual Ethernet interface in private VLAN host mode is up. See Example 12-3 on page 12-3.
module vse *module-number* **execute vemcmd show port**	Displays that a VSE is configured correctly. See Example 12-4 on page 12-3.

Example 12-1 show vlan private-vlan Command

```
switch# show vlan private-vlan
    Primary  Secondary  Type              Ports
    -------  ---------  ----------------  ------------------------------------------
    152      157        community
```

```
152      158        isolated
156      153        community
156      154        community
156      155        isolated
```

Example 12-2 show interface name Command

```
switch# show interface eth3/1
    Ethernet3/4 is up
      Hardware: Ethernet, address: 0050.565a.ca50 (bia 0050.565a.ca50)
      MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
         reliability 0/255, txload 0/255, rxload 0/255
      Encapsulation ARPA
      Port mode is Private-vlan trunk promiscuous
      full-duplex, 1000 Mb/s
      Beacon is turned off
      Auto-Negotiation is turned off
      Input flow-control is off, output flow-control is off
      Auto-mdix is turned on
      Switchport monitor is off
        Rx
        158776 Input Packets 75724 Unicast Packets
        76 Multicast Packets 82976 Broadcast Packets
        13861581 Bytes
        Tx
        75763 Output Packets 75709 Unicast Packets
        3 Multicast Packets 51 Broadcast Packets 0 Flood Packets
        7424670 Bytes
        5507 Input Packet Drops 0 Output Packet Drops
      2 interface resets
```

Example 12-3 show interface veth Command

```
switch# show interface vethernet3
    Vethernet3 is up
        Hardware is Virtual, address is 0050.56bb.6330
        Owner is VM "fedora9", adapter is Network Adapter 1
        Active on module 3
        VMware DVS port 10
        Port-Profile is pvlancomm153
        Port mode is Private-vlan host
        Rx
        14802 Input Packets 14539 Unicast Packets
        122 Multicast Packets 141 Broadcast Packets
        1446568 Bytes
        Tx
        15755 Output Packets 14492 Unicast Packets
        0 Multicast Packets 1263 Broadcast Packets 0 Flood Packets
        1494886 Bytes
        45 Input Packet Drops 0 Output Packet Drops
```

Example 12-4 module vse module-number execute vemcmd show port Command

```
switch# module vse 3 execute vemcmd show port-old
      LTL    IfIndex   Vlan    Bndl  SG_ID Pinned_SGID  Type  Admin State  CBL Mode   Name
        8          0   3969       0      2           2  VIRT     UP    UP    4 Access 120
        9          0   3969       0      2           2  VIRT     UP    UP    4 Access 121
       10          0    150       0      2           2  VIRT     UP    UP    4 Access 122
       11          0   3968       0      2           2  VIRT     UP    UP    4 Access 123
       12          0    151       0      2           2  VIRT     UP    UP    4 Access 124
       13          0      1       0      2           2  VIRT     UP    UP    0 Access 125
       14          0   3967       0      2           2  VIRT     UP    UP    4 Access 126
```

```
   16   1a020100      1 T     0      2           2  PHYS     UP    UP    4  Trunk
vmnic1
   18   1a020300      1 T     0      2           2  PHYS     UP    UP    4  Trunk
vmnic3
        pvlan promiscuous trunk port
            153 --> 156
            154 --> 156
            155 --> 156
            157 --> 152
            158 --> 152
   19   1a020400      1 T     0      2           2  PHYS     UP    UP    4  Trunk
vmnic4
        pvlan promiscuous trunk port
            153 --> 156
            154 --> 156
            155 --> 156
            157 --> 152
            158 --> 152
   47   1b020000    154       0      2           0  VIRT     UP    UP    4 Access
fedora9.eth0
        pvlan community 156 153
```

If additional information is required for Cisco Technical Support to troubleshoot a private VLAN issue, use the following commands:

- **show system internal private-vlan info**
- **show system internal private-vlan event-history traces**
- **show system internal private-vlan event-history errors**
- **show system internal private-vlan event-history events**

Access Control Lists (ACLs)

This chapter describes how to identify and resolve problems that relate to Access Control Lists (ACLs) and includes the following sections:

Information About Access Control Lists

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.
- IPv6—The device applies IPv6 ACLs only to IPv6 traffic

For detailed information about how ACL rules are used to configure network traffic, see the *Cisco Nexus 1000V Security Configuration Guide.*

ACL Configuration Limits

The following configuration limits apply to ACLs:

- You cannot have more that 128 rules in an ACL.

- The maximum number of ACLs is 128 (spread across all the ACLs) in one VSE.

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.
- Noninitial fragments are not subject to ACL lookup.
- You cannot have two not-equal-to (neq) operators in the same rule.

ACL Troubleshooting Commands

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

Use the following command to display configured ACLs:

- **show access-list summary**

Use following commands on the VSM to see run-time information of the ACLMGR and ACLCOMP during configuration errors and to collect ACLMGR process run-time information configuration errors:

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

Use the following commands to collect ACLCOMP process run-time information configuration errors:

- **show system internal aclcomp event-history errors**
- **show system internal aclcomp event-history msgs**
- **show system internal aclcomp pdl detailed**
- **show system internal aclcomp mem-stats** (to debug memory usage and leaks)

Displaying ACL Policies on the VSE

The commands listed in this section can be used to display configured ACL policies on the Virtual Ethernet Module (VSE).

Use the following command to list the ACLs installed on that server

```
switch(config-if)# module vse 3 execute vemcmd show acl
AclId RefCnt Type Rules StatId AclName (Stats: Permit/Deny/NoMatch)
```

```
----- ------ ---- ----- ------ ----------------------------------
1 0 IPv4 1 1 v4 (Enb: 0/0/0)
2 0 IPv6 0 2 v6 (Dis: 0/0/0)
```

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

Use the following command to list the interfaces on which ACLs have been installed

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

Debugging Policy Verification Issues

You can debug a policy verification failure.

Note This section is applicable only to VSEs that are available in older releases. The VSEs in the latest release do not have any policy verification failure issue.

Step 1 On the VSM, redirect the output to a file in bootflash.

debug logfile *filename*

Step 2 Enter the **debug aclmgr all** command.

Step 3 Enter the **debug aclcomp all** command.

For the VSEs where the policy exists, or is being applied, enter the following these steps from the VSM. The output goes to the console.

Step 4 Enter the **module vse** *module-number* **execute vemdpalog debug sfaclagent all** command.

Step 5 Enter the **module vse** *module-number* **execute vemdpalog debug sfpdlagent all** command.

Step 6 Enter the **module vse** *module-number* **execute vemlog debug sfacl all** command.

Step 7 Enter the **module vse** *module-number* **execute vemlog start** command.

Step 8 Enter the **module vse** *module-number* **execute vemlog start** command.

Step 9 Configure the policy that was causing the verify error.

Step 10 Enter the **module vse** *module-number* **execute vemdpalog show all** command.

Step 11 Enter **module vse** *module-number* **execute vemlog show all** command.

Save the Telnet or SSH session buffer to a file. Copy the logfile created in bootflash.

Troubleshooting ACL Logging

This section includes the following topics:

Using the CLI to Troubleshoot ACL Logging on a VSE

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

Viewing Current Flows

You can troubleshoot ACL logging by viewing the current flows on a VSE.

vemcmd show aclflows stats

EXAMPLE

The following example shows how to troubleshoot ACL logging:

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

Viewing Active Flows

You can display all the active flows on a VSE.

vemcmd show aclflows [permit | deny]

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

EXAMPLE

The following example shows how to display all the active flows on a VSE:

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

Flushing All ACL Flows

You can use the **vemcmd flush aclflows** command to detect any new flows that affect the VSE. Clear all the existing flows, and then you can detect new flows that match any expected traffic. Syslog messages are not sent when you do this action.

Showing Flow Debug Statistics

You can show ACL debug statistics.

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

vemcmd show aclflows dbgstats

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

vemcmd clear aclflows dbgstats

ACL Logging Troubleshooting Scenarios

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

Troubleshooting a Syslog Server Configuration

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

BEFORE YOU BEGIN

- Log in to the VSM and VSE CLI.

PROCEDURE

	Command	Description
Step 1	**show logging ip access-list status** **Example:** `switch# show logging ip access-list status` `switch #`	Verifies that the remote syslog server is configured properly.
Step 2	**vemcmd show acllog config** **Example:** `switch# module vse 3 execute vemcmd show acllog config` `switch #`	Verifies ACL logging on the VSE.
Step 3	**vemcmd show aclflows dbgstats** **Example:** `switch# module vse 3 execute vemcmd show aclflows dbgstats` `switch #`	Checks to see if any errors occurred.

Troubleshooting an ACL Rule That Does Not Have a Log Keyword

If the ACL rule does not have a **log** keyword, any flow that matches the ACL is not reported although the ACL statistics continue to advance. You can verify a **log** keyword.

BEFORE YOU BEGIN

- Log in to the VSM and VSE CLI.

PROCEDURE

	Command	Description
Step 1	**show running-config aclmg** **Example** `switch# show running-config aclmg` `switch #`	Verifies that the **log** keyword is enabled.
Step 2	**show logging ip access-list status** **Example:** `switch# show logging ip access-list status` `switch #`	Verifies that ACL logging is configured properly.
Step 3	**vemcmd show acllog config** **Example:** `switch# vemcmd show acllog config` `switch #`	Verifies ACL logging on the VSE.

Troubleshooting a Maximum Flow Limit Value That is Too Low

If the number of flows does not reach 5000 for either permit of deny flows, you can increase the maximum flows.

BEFORE YOU BEGIN

- Log in to the VSM and VSE CLI.

PROCEDURE

	Command	Description
Step 1	**show logging ip access-list status** **Example:** `switch# show logging ip access-list status` `switch #`	Verifies that ACL logging is configured properly.

	Command	Description
Step 2	**vemcmd show acllog config** **Example:** switch# vemcmd show acllog config switch #	Verifies ACL logging on the VSE.
Step 3	**logging ip access-list cache max-deny-flows** *<num>* **Example:** switch# logging ip access-list cache max-deny- flows *<num>* switch #	Increases maximum flows to the desired value.

Troubleshooting a Mismatched Configuration Between a VSM and a VSE

If syslog messages are not being sent and the flow information counters are invalid, the configuration between a VSM and a VSE might be mismatched.

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

BEFORE YOU BEGIN

- Log in to the CLI in EXEC mode.

PROCEDURE

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

SPAN

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

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 1000VE supports two types of SPAN:

- SPAN (local SPAN) that can monitor sources within a host or VSE.
- 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 1000VE 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 VMotion, 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 for a running SPAN session:
  - The limit of 64 SPAN sessions is not exceeded.
  - At least one operational source is configured.
  - At least one operational destination is configured.
  - The configured source and destination are on the same host.
  - The session is enabled with the **no shut** command.
- A session is stopped if any of the following occurs:
  - 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 VMotion.
  - The session is disabled by a **shut** command.

Problems with SPAN

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

Symptom	Possible Causes	Solution
You observe issues with VM traffic after configuring a session with Ethernet destinations.	—	Ensure that the Ethernet destination is not connected to the same uplink switch. The SPAN packets might cause problems with the IP tables, the MAC tables, or both on the uplink switch, which can cause problems with the regular traffic.
A session state is up and the packets are not received at the destination ports.	—	Verify that the correct VLANs are allowed on the trunk destination ports.
The session displays an error.	—	1. Make sure that VSM-VSE connectivity is working correctly. 2. Force reprogramming of the session on the VSE. **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 the ERSPAN ID is configured at the destination.
	An ERSPAN-enabled VMKernel NIC is not configured on the host or VSE.	Make sure that you create a VMKernel NIC on the host using a port profile configured for ERSPAN.
	The ERSPAN-enabled VMKernel NIC is not configured with a proper IP, gateway, or both.	Ping the ERSPAN IP destination from the host VMKernel NIC. **vmkping** *dest-id* Use the **vempkt** command to capture packets on the VMKernel NIC LTL and ensure ERSPAN packets are being sent. Use the **vemlog debug sfspan d** command so that the ERSPAN packets appear in the VSEpkt capture log.

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 14-1 on page 14-3.
show monitor session	Displays the current state of a SPAN session, the reason it is down, and the session configuration. See Example 14-2 on page 14-3.
module vse *module-number* **execute vemcmd show span**	Displays the VSE source IP and SPAN configuration. See Example 14-3 on page 14-4.

Additional commands:

- **show monitor internal errors**
- **show monitor internal event-history msgs**
- **show monitor internal info global-info**
- **show monitor internal mem-stats**

Example 14-1 show monitor Command

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

Example 14-2 show monitor session Command

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

Example 14-3 module vse execute vemcmd show span Command

```
switch# vemcmd show span
ERSPAN Local Encap Interface Information:
----------------------------------------
LTL: 0 cap 0 ()
SF_LTL_L3_CTRL: 0 cap 0
IP: 0.0.0.0
MAC: 0000:0000:0000
HWBD: 0
Port State: Down
----------------------------------------

VEM SOURCE IP NOT CONFIGURED.

HW SSN ID ERSPAN ID HDR VER DST LTL/IP
switch#
```

System

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

This chapter includes the following sections:

Information About the System

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

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

- Virtual Supervisor Module (VSM) – This is the control software of the Nexus 1000VE distributed virtual switch. It runs on a virtual machine (VM) and is based on NX-OS.
- Virtual Services Engine (VSE) – This is the part of Cisco Nexus 1000VE that actually switches data traffic. It runs as VM on a VMware ESX 6.0 and above host.Several VSEs are controlled by one VSM. All the VSEs that form a switch domain should be in the same virtual Datacenter as defined by VMware vCenter Server.

General Restrictions for vCenter Server

When you are troubleshooting issues related to vCenter Server, make sure that you observe the following restrictions:

- The name of a distributed virtual switch (DVS) name must be unique across Datacenters
- You create a DVS in a network folder
- A Datacenter cannot be removed unless the DVS folder or the underlying DVS is deleted.
- A DVS can be deleted only with the help of VSM using the **no vmware dvs** command in config-svs-conn mode.
- The no vmware dvs command can succeed only if there are no VMs using the DVS port-groups.
- A port group on vCenter Server can be deleted only if there are no interfaces associated with it.
- A sync operation performed in conjunction with the **connect** command helps VSM keep in sync with vCenter Server.
- Each VSM uses a unique extension key to communicate with vCenter Server and perform operations on a DVS.

Extension Key

The VSM uses the extension key when communicating with the vCenter Server. Each VSM has its own unique extension key, such as Cisco_Nexus_1000VE_32943215

Use the **show vmware vc extension-key** command to find the extension key of the VSM. It is also listed in the .xml file.

The extension key registered on the vCenter Server can be found through the MOB. For more information, see the "Finding the Extension Key Tied to a Specific DVS" procedure on page 3-7.

The same extension key cannot be used to create more than one DVS on the vCenter Server.

Recovering a DVS

You can use this procedure to recover a DVS if the VSM VM that was used to create it is lost or needs to be replaced. This section includes the following procedures:

Recovering a DVS With a Saved Copy of the VSM

You can use this procedure to recover a DVS when you have previously saved a back up copy of the VSM configuration file.

BEFORE YOU BEGIN

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

- Use this procedure if you have previously saved a back up copy of the VSM configuration file. If you have not previously saved a back up copy, the see the "Recovering a DVS Without a Saved Copy of the VSM" procedure on page 15-4.
- Make sure that the VSM VM switchname is the same as the DVS switchname on the vCenter Server. This allows the VSM configuration to synchronize with the correct DVS on the vCenter Server.

  To change the VSM switchname use the **switchname** *newname* command.

Step 1 From the MOB, find the DVS extension key.
For more information, see the "Finding the Extension Key Tied to a Specific DVS" procedure on page 3-7.

Step 2 On the VSM, add the DVS extension key found in Step 1.

The extension key allows the VSM to log in to the vCenter server.

```
Example:
switch# config t
switch(config)# vmware vc extension-key Cisco_Nexus_1000V_32943215
```

Step 3 From the MOB, unregister the extension key found in Step 1.

For more information, see the "Unregistering the Extension Key in the vCenter Server" procedure on page 3-11.

Step 4 From the VC client, register the extension (plug-in) for the VSM.

For more information see the following procedure in the *Cisco Nexus 1000V Getting Started Guide*.

- Creating a Cisco Nexus 1000VE Plug-In on the vCenter Server

Step 5 On the VSM, restore the configuration using a previously saved copy of the VSM configuration file.

copy *path/filename* **running-config**

```
Example:
switch# copy sftp://user1@172.22.36.10/backup/hamilton_cfg running-config
```

Step 6 Do one of the following:

- If the vCenter server connection is not part of the previously saved configuration, continue with the next step.
- Otherwise, go to Step 8.

Step 7 On the VSM, restore the configuration for the vCenter server connection.

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

Step 8 Connect to vCenter Server.

Example:
```
switch(config-svs-conn#) connect
```

You can now use the old DVS or remove it.

Recovering a DVS Without a Saved Copy of the VSM

You can use this procedure to recover a DVS when you have not previously saved a back up copy of the VSM configuration file.

BEFORE YOU BEGIN

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

- The folder in which the VSM resides must be:
  - At the root-level of the Data Center in which it resides. It cannot be embedded in another folder.
  - Of the same name as the VSM.

  If the folder does not meet the above criteria, the connection to vCenter server fails with the error, *the VSM already exists*.
- Use this procedure if you have not previously saved a back up copy of the VSM configuration file. If you have previously saved a back up copy, then see the "Recovering a DVS With a Saved Copy of the VSM" procedure on page 15-3.
- If you have not previously saved a back up copy of the VSM configuration file, then you may try recreating the old port profiles before connecting to the VC. This procedure has a step for recreating port profiles. If you do not recreate these before connecting to VC, then all the port groups present on the VC are removed and all ports in use are moved to the quarantine port groups.
- Make sure that the VSM VM switchname is the same as the DVS switchname on the vCenter Server. This allows the VSM configuration to synchronize with the correct DVS on the vCenter Server.

  To change the VSM switchname use the **switchname** *newname* command.

Step 1 From the MOB, find the DVS extension key.
For more information, see the "Finding the Extension Key Tied to a Specific DVS" procedure on page 3-7.

Step 2 On the VSM, add the DVS extension key found in Step 1.

The extension key allows the VSM to log in to the vCenter server.

Example:
```
switch# config t
switch(config)# vmware vc extension-key Cisco_Nexus_1000Ve_32943215
```

Step 3 From the MOB, unregister the extension key found in Step 1.

For more information, see the "Unregistering the Extension Key in the vCenter Server" procedure on page 3-11.

Step 4 From the VC client, register the extension (plug-in) for the VSM.

For more information see the following procedure in the *Cisco Nexus 1000V Getting Started Guide*.

- Creating a Cisco Nexus 1000VE Plug-In on the vCenter Server

Step 5 Manually recreate the old port profiles from your previous configuration.

For more information, see the following procedures in the *Cisco Nexus 1000V Getting Started Guide*.

- Configuring the system port profile for VSM-VSE Communication
- Configuring the uplink port profile for VM Traffic
- Configuring the data port profile for VM Traffic

Note If you do not manually recreate the port profiles, then all port groups on the vCenter Server are removed when the VSM connects.

Step 6 On the VSM, restore the configuration for the vCenter server connection.

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

Step 7 Connect to vCenter Server.

```
Example:
switch(config-svs-conn#) connect
```

You can now use the old DVS or remove it.

Problems Related to VSM and vCenter Server Connectivity

Symptom	Solution
The vCenter Server connection seems to succeed, but does not.	Make sure that the domain ID is configured correctly.
The **svs connection command** fails.	Make sure you have configured all parameters for the **svs connection** command.
	Make sure you can ping the vCenter Server IP address.
	Make sure that the proxy.xml file is correct for both the IP address and length.
	Restart the vCenter Server
The host does not show up in the Add host to DVS screen.	Make sure that the Host is installed with VMware Enterprise plus license containing the Distributed Virtual Switch feature.
The server name column of the **show module** command output shows the IP address.	The server name shows the host-name or IP address, whichever was used to add the host to the DVS on the vCenter Server.

Example 15-1 shows the **show vms internal event-history errors** command that is useful for examining VC errors in detail. It shows whether an error is caused by a VSM (client) or the server.

Example 15-1 show vms internal event-history error Command

```
switch# show vms internal event-history errors

Event:E_DEBUG, length:239, at 758116 usecs after Tue Feb  3 18:21:58 2009
    [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: "":ServerFaultCode [VMWARE-VIM]
A DVS switch with spec.name as switch already exists, cannot create DVS switch. A
specified parameter was not correct.spec.name

Event:E_DEBUG, length:142, at 824006 usecs after Tue Feb  3 18:18:30 2009
    [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: SOAP-ENV:Client [VMWARE-VIM]
Operation could not be completed due to connection failure.

Event:E_DEBUG, length:134, at 468208 usecs after Tue Feb  3 18:15:37 2009
    [102] convert_soap_fault_to_err(1179): SOAP 1.1 fault: "":ServerFaultCode [VMWARE-VIM]
Extension key was not registered before its use.
```

Setting the System MTU

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

BEFORE YOU BEGIN

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

- You are logged in to the CLI in EXEC mode.
- The system port profiles are already configured and you know the uplink profile names.

  For more information, see the *Cisco Nexus 1000V Port Profile Configuration Guide*.
- The MTU size you set for the **system mtu** on the port profile must be less than the size of the **system jumbomtu** configured on the interface.

  For more information about configuring MTU on the interface, see the *Cisco Nexus 1000V Interface Configuration Guide*.
- When you configure a system MTU on a system port profile, it takes precedence over an MTU you may have configured on the interface.
- To verify the ESX MTU settings for corresponding PNICs, use the **ESXcfg-nics -l** command.

SUMMARY STEPS

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

DETAILED STEPS

	Command	Description
Step 1	**config t** **Example:** switch# config t switch(config)#	Enters global configuration mode.
Step 2	**port-profile** *name* **Example:** switch(config)# port-profile AccessProf switch(config-port-prof)#	Enters port profile configuration mode for the named system uplink port profile.
Step 3	**system mtu** *mtu-size* **Example:** switch(config-port-prof)# system mtu 4000 switch(config-port-prof)#	Designates the MTU size. • Must be an even number between 1500 and 9000. • Must be less than the size of the **system jumbomtu** on the interface.
Step 4	**show port-profil**e [**brief** | **expand-interface** | **usage**] [**name** *profile-name*] **Example:** switch(config-port-prof)# show port-profile name AccessProf	(Optional) Displays the configuration for verification.
Step 5	**copy running-config startup-config** **Example:** switch(config-port-prof)# copy running-config startup-config	(Optional) Saves the running configuration persistently through reboots and restarts by copying it to the startup configuration.

Recovering Lost Connectivity Due to MTU Mismatch

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

BEFORE YOU BEGIN

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

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

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

SUMMARY STEPS

1. **config t**
2. **module vse** *module_number* **execute vemcmd show port** *port-LTL-number*
3. **module vse** *module_number* **execute vemcmd set mtu** *size* **ltl** *port-LTL-number*

DETAILED STEPS

	Command	Description
Step 1	**config t** **Example:** switch# config t switch(config)#	Enters global configuration mode.
Step 2	**module vse** *module_number* **execute vemcmd show port** *port-LTL-number*	Displays the port configuration including the LTL number needed for Step 3.
	Example: switch(config)# module vse 3 execute vemcmd show port 48 LTL IfIndex Vlan Bndl SG_ID Pinned_SGID Type Admin State CBL Mode Name 17 1a030100 1 T 304 1 32 PHYS UP UP 1 Trunk vmnic1 switch(config)#	
Step 3	**module vse** *module_number* **execute vemcmd set mtu** *size* **ltl** *port-LTL-number* **Example:** switch(config)# module vse 3 execute vemcmd set mtu 9000 ltl 17 switch(config)#	Designates the MTU size for the port, using the LTL number obtained in Step 2.

VSM Creation

Symptom	Possible Causes	Solution
The VSM VM is stuck at the boot prompt.	—	Make sure that you have three e1000 NICs.
The VSM VM cannot ping itself.	—	Configure the management0 interface.
The VSM VM can ping itself, but not the gateway.	—	Make sure the NIC order is correct: control, management, inband/outband.
The VSM VM can ping the gateway, but not the outside subnet.	—	Configure vrf context management.

Port Profiles

When creating a port profile, use the following commands to create the corresponding port groups on the vCenter Server:

- **vmware port-group**
- **state enabled**

Profiles that have the system VLAN configuration allow the VSE to communicate with the VSM.

Make sure that the system port-profile is defined with the right system VLANS.

Use the **show port-profile** and **show port-profile usage** commands to collect basic required information.

Problems with Port Profiles

Symptom	Possible Causes	Solution
You receive an error message "Possible failure in communication with vCenter Server."	The VSM is not connected to the vCenter Server.	Issue the **svs connection vc** command to connect to the vCenter Server.
	The port group name is not unique.	Port group names must be unique within a vCenter Server Datacenter.
Port profile or port groups do not appear on the vCenter Server.	—	Make sure you have issued the **vmware port-group** command and **state enable** command.

Problems with Hosts

Symptom	Solution
The host is visible on the vCenter Server, but not the VSM.	Issue the **vemcmd show trunk** command to verify that there is an uplink carrying the control VLAN. The profile applied to the uplink must be a system profile with a control VLAN as a system VLAN.
	Verify the control VLAN in the upstream switch port and the path to the VSM VM. Make sure that one uplink at most carries the control VLAN, or that all uplinks and upstream ports carrying the control VLAN are in port channels.
A module flap occurs.	The VSM may be overloaded. Make sure that you have 4 GB of memory and CPU shares for the VSM VM on the vCenter Server.

Problems with VM Traffic

When troubleshooting problems with intra-host VM traffic, follow these guidelines:

- Make sure that at least one of the VMware virtual NICs is on the correct DVS port group and is connected.

- If the VMware virtual NIC is down, determine if there is a conflict between the MAC address configured in the OS and the MAC address assigned by VMware. You can see the assigned MAC addresses in the vmx file.

When troubleshooting problems with inter-host VM traffic, follow these guidelines:

- Determine if there is exactly one uplink sharing a VLAN with the VMware virtual NIC. If there is more than one, they must be in a port channel.
- Ping a SVI on the upstream switch using the **show intX counters** command.

VSE Troubleshooting Commands

Use the following commands to display VSE information:

- **vemlog** – displays and controls VSE kernel logs
- **vemcmd** – displays configuration and status information
- **vem-support all** – collects support information
- **systemctl status nexus1000v**– collects status information
- **vemcmd show version**– collects version information
- **vemlog show last** *number-of-entries* – displays the circular buffer

Example 15-2 vemog show last Command

```
[root@ESX-cos1 ~]# vemlog show last 5
Timestamp                 Entry CPU  Mod Lv         Message
Oct 13 13:15:52.615416     1095   1    1  4 Warning vssnet_port_pg_data_ …
Oct 13 13:15:52.620028     1096   1    1  4 Warning vssnet_port_pg_data_ …
Oct 13 13:15:52.630377     1097   1    1  4 Warning svs_switch_state …
Oct 13 13:15:52.633201     1098   1    1  8    Info vssnet new switch …
Oct 13 13:16:24.990236     1099   1    0  0         Suspending log
```

- **vemlog show info** – displays information about entries in the log

Example 15-3 vemcmd show info Command

```
[root@ESX-cos1 ~]# vemlog show info
          Enabled: Yes
    Total Entries: 1092
  Wrapped Entries: 0
     Lost Entries: 0
  Skipped Entries: 0
Available Entries: 6898
 Stop After Entry: Not Specified
```

- **vemcmd help** – displays the type of information you can display

Example 15-4 vemcmd help Command

```
[root@ESX-cos1 ~]# vemmd help
show card              Show the card's global info
show vlan [vlan]       Show the VLAN/BD table
show bd [bd]           Show the VLAN/BD table
show l2 <bd-number>    Show the L2 table for a given BD/VLAN
show l2 all            Show the L2 table
show port [priv|vsm]   Show the port table
```

```
show pc                 Show the port channel table
show portmac            Show the port table MAC entries
show trunk [priv|vsm]   Show the trunk ports in the port table
show stats              Show port stats
```

VSE Log Commands

Use the following commands to control the VSElog:

- **VEMlog stop** – stops the log
- **VEMlog clear** – clear s the log
- **VEMog start** *number-of-entries* – starts the log and stops it after the specified number of entries
- **VEMlog stop** *number-of-entries* – stops the log after the next specified number of entries
- **VEMlog resume** – starts the log, but does not clear the stop value

Error Messages

On the vSphere Client, you can see error messages under the recent tasks tab. You can find detailed description of the error under the Tasks and Events tab. The same messages are also propagated to the VSM.

Table 15-1 lists error messages that you might see on the VSM.

Table 15-1 Error Messages on the VSM

Error	Description
ERROR: [VMWARE-VIM] Extension key was not registered before its use	This error indicates that VSM extension key is not registered.
ERROR: [VMWARE-VIM] A DVS switch with spec.name as switch already exists, cannot create DVS switch. A specified parameter was not correct. spec.name.	This error is displayed after you enter the first **connect** command, and indicates that a DVS already exists with the same name.
ERROR: [VMWARE-VIM] A DVS switch with spec.extensionKey as Cisco_Nexus_1000VE_2055343757 already exists, cannot create DVS new-switch. A specified parameter was not correct. spec.extensionKey	This error is displayed when the VSM tries to create a different DVS after changing the switch name.
ERROR: [VMWARE-VIM] A DVS switch with name as switch already exists, cannot reconfigure DVS test. A specified parameter was not correct. Spec.name	This error indicates that a DVS with the same name already exists.
Warning: Operation succeeded locally but update failed on vCenter server.[VMWARE-VIM] DVPortgroup test port 0 is in use. The resource vim.dvs.DistributedVirtualPort 0 is in use.	This warning is displayed when the VSM tries to delete the port profile if the VSM is not aware of the nics attached to the port groups.

Cisco TrustSec

This chapter describes how to identify and resolve problems that might occur when configuring Cisco TrustSec and includes the following sections:

Information About Cisco TrustSec

The Cisco TrustSec security architecture builds secure networks by establishing clouds of trusted network devices. Each device in the cloud is authenticated by its neighbors. Communication on the links between devices in the cloud is secured with a combination of encryption, message integrity checks, and data-path replay protection mechanisms.

Cisco TrustSec also uses the device and user identification information acquired during authentication for classifying, or coloring, the packets as they enter the network. This packet classification is maintained by tagging packets on ingress to the Cisco TrustSec network so that they can be properly identified for the purpose of applying security and other policy criteria along the data path. The tag, also called the security group tag (SGT), allows the network to enforce the access control policy by enabling the endpoint device to act upon the SGT to filter traffic.

See the *Cisco Nexus 1000 Virtual Edge for VMware vSphere Security Configuration Guide* for more information on the Cisco TrustSec feature on Cisco Nexus 1000VE.

Cisco TrustSec Troubleshooting Commands

This section contains the following topics:

Debugging Commands

Command	Purpose
debug cts authentication	Collects and views logs related to Cisco TrustSec authentication.
debug cts authorization	Collects and views logs related to Cisco TrustSec authorization.
debug cts errors	Collects and views logs related to Cisco TrustSec errors and warning messages.
debug cts messages	Collects and views logs related to Cisco TrustSec messages.
debug cts packets	Collects and views logs related to Cisco TrustSec packets.
debug cts relay	Collects and views logs related to Cisco TrustSec relay functionality.
debug cts sxp	Collects and views logs related to Cisco TrustSec SXP.
debug cts sap	Collects and views logs related to the Cisco TrustSec Security Association Protocol (SAP).
debug cts trace	Collects and views logs related to Cisco TrustSec trace functionality.
show cts internal debug-info	Displays Cisco TrustSec debug information.

VSE Logging Commands

You can use the commands in this section to troubleshoot commands related to VSE logging. Logging commands needs to be executed directly by login to VSE..

VSE Command	Description
echo "logfile enable" >/var/tmp/dpafifo	Enables DPA debug logging. Logs are output to the **/var/log/vemdpa.log** file.
echo "debug sfctsagent all" > /var/tmp/dpafifo	Enables TrustSec SXP agent debug logging. Logs are output to the **/var/log/vemdpa.log** file.
vemlog debug sfcts_config all	Enables the data path debug logging and captures logs for the data packets sent between the client and the server.

VSE Command	Description
vemlog debug sfipdb all	Enables the data path debug logging and captures logs corresponding to the IP database that maintains the IP addresses for all the virtual machines that are being tracked using Cisco TrustSec device tracking. To view the logs, enable Cisco TrustSec device tracking on the Cisco Nexus 1000VE.
vemcmd show learnt ip	Displays the Cisco TrustSec configuration on the Cisco Nexus 1000VE. See Example 16-1 on page 16-3
vemcmd show cts global	Displays if Cisco TrustSec is enabled on the Cisco Nexus 1000VE. See Example 16-2 on page 16-3
vemcmd show cts ipsgt	Displays the Cisco TrustSec configuration command specific to IP-to-SGT mapping on Cisco Nexus 1000VE. See Example 16-3 on page 16-4
vemcmd show cts subnet-sgt-map ip	Displays the Cisco TrustSec configuration specific to Subnet-to-SGT mapping on Cisco Nexus 1000VE. See Example 16-4 on page 16-4
vemcmd show cts access-list	Displays the Cisco TrustSec Role-Based access-list names and counters matching the ACEs in the access-list. (Permit/Deny/No-Match). See Example 16-5 on page 16-4
vemcmd show cts policy	Displays the Cisco TrustSec Role-Based policies where source sgt to destination sgt mapping with RBACL. See Example 16-6 on page 16-4.

Example

vemcmd can be executed by directly logging in to VSE or directly from VSM using **module vse #vse-module-number execute** vemcmd complete command.

Example 16-1 vemcmd show learnt ip Command

```
cisco-vse:~$ vemcmd show learnt ip
IP Address LTL VLAN BD
/SegID
10.78.1.76 49 353 7
switch#
```

Example 16-2 vemcmd show cts global Command

```
cisco-vse:~$ vemcmd show cts global
CTS Global Configuration:
CTS is: Enabled
CTS Device Tracking is: Enabled
```

```
switch#
```

Example 16-3 vemcmd show cts ipsgt Command

```
cisco-vse:~$ vemcmd show cts ipsgt
IP Address LTL VLAN BD SGT Learnt
10.78.1.76 49 353 7 6766 Device Tracking
switch#
```

Example 16-4 vemcmd show cts subnet-sgt-map ip

```
Example 16-4 vemcmd show cts subnet-sgt-map ip

cisco-vse:~$ vemcmd show cts subnet-sgt-map ip
Key (tid, ip/mask) : Data (SGT)
**** Dumping all subnet SGT entries ****
(0, 192.168.0.0/24) : 200
cisco-vse:~$
```

Example 16-5 vemcmd show cts access-list

```
cisco-vse:~$ vemcmd show cts access-list
Global RBACL List Permit/Deny/No-Match

------ ----- ---- --------------------
ise_permit_icmp 0/0/0
ise_permit_icmp_ret 4/0/0
test 0/0/0
cisco-vse:~$
```

Example 16-6 vemcmd show cts policy

```
isco-vse:~$ vemcmd show cts policy
SGT DGT RBACL
700 800 ise_permit_icmp
800 700 ise_permit_icmp_ret
cisco-vse:~$
```

show Commands

See the *Cisco Nexus 1000VE Command Reference* for more information on the **show** commands for Cisco TrustSec.

Command	Purpose
show cts	Displays the Cisco TrustSec configuration.
show cts sxp	Displays the SXP configuration for Cisco TrustSec.

Command	Purpose
show feature	Displays the features available, such as CTS, and whether they are enabled.
show running-configuration cts	Displays the running configuration information for Cisco TrustSec.
show cts device tracking	Displays the Cisco TrustSec device tracking configuration.
show cts role-based sgt-map	Displays the mapping of the IP address to SGT for Cisco TrustSec.
show cts sxp connection	Displays SXP connections for Cisco TrustSec.
show cts interface delete-hold timer	Displays the interface delete hold timer period for Cisco TrustSec.
show cts internal event-history [error |mem-stats | msgs | sxp]	Displays event logs for Cisco TrustSec.

Problems with Cisco TrustSec

This section includes symptoms, possible causes and solutions for the following problems with Cisco TrustSec.

Symptom	Possible Causes	Verification and Solution
The Cisco Nexus 1000VE is unable to form an SXP session with Cisco TrustSec.	There is no connection between the Cisco Nexus 1000VE and its peer.	Verify if the Cisco Nexus 1000VE is connected to its peer. **ping**
	The Cisco TrustSec SXP is not enabled on the Cisco Nexus 1000VE.	Verify if the Cisco TrustSec SXP is enabled on the Cisco Nexus 1000VE. **show cts sxp** If not, enable the Cisco TrustSec SXP. **cts sxp enable**
	The password configured on the Cisco Nexus 1000VE does not match the password configured on its peer.	Verify if the passwords configured on the Cisco Nexus 1000VE matches its peer. **show cts sxp**
	The default source IPv4 address is not configured on the Cisco Nexus 1000VE.	Verify if the default source IPv4 address is not configured on the Cisco Nexus 1000VE. **show cts sxp**
	The SXP peer is not configured as the listener.	Verify that the SXP peer is configured as the listener. **show cts sxp connection**

Symptom	Possible Causes	Verification and Solution
Cisco TrustSec SXP is unable to learn any IP-SGT mappings on the Cisco Nexus 1000VE.	The Cisco TrustSec device tracking is not enabled on the Cisco Nexus 1000VE.	Verify if the Cisco TrustSec device tracking is enabled on the Cisco Nexus 1000VE. **show cts device tracking** If not, enable the Cisco TrustSec device tracking. **cts sxp device tracking**

Before Contacting Technical Support

This chapter describes the steps to take before calling for technical support and includes the following sections:

Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco, contact Cisco Technical Support at this URL: http://www.cisco.com/warp/public/687/Directory/DirTAC.shtm

Cisco Support Communities

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

- Cisco Support Community for Server Networking
- Cisco Communities: Nexus 1000V

Gathering Information for Technical Support

At some point, you may need to contact your customer support representative or Cisco TAC for some additional assistance. This section outlines the steps that the you should perform prior to contacting your next level of support, so you can reduce the amount of time that you spend resolving the issue.

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

Configure your Telnet or SSH application to log the screen output to a text file. Use the **terminal length 0** CLI command and then use the **show tech-support details** CLI command.

Step 2 Capture the exact error codes you see in CLI message logs.

- **show logging log** CLI (displays the error messages)
- **show logging last** *number* (displays the last lines of the log)

Step 3 Answer the following questions before calling for technical support:

- On which switch or port is the problem occurring?
- Which Cisco Nexus 1000V software, driver versions, operating systems versions and storage device firmware are in your fabric?
- ESX and vCenter Server software that you are running?
- What is the network topology?
- Were any changes being made to the environment (VLANs, adding modules, upgrades) prior to or at the time of this event?
- Are there other similarly configured devices that could have this problem, but do not?
- Where was this problematic device connected (which switch and interface)?
- When did this problem first occur?
- When did this problem last occur?
- How often does this problem occur?
- How many devices have this problem?
- Were any traces or debug output captured during the problem time? What troubleshooting steps have you attempted? Which, if any, of the following tools were used?
  - Ethanalyzer, local, or remote SPAN
  - CLI debug commands
  - traceroute, ping

Obtaining a File of Core Memory Information

Cisco customer support engineers often use files from your system for analysis. One such file contains memory information and is referred to as a core dump. The file is sent to a TFTP server or to a flash card in slot0: of the local switch. You should set up your switch to generate this file under the instruction of your customer support representative and send it to a TFTP server so that it can be emailed to them.

To generate a file of core memory information, or a core dump, use the command in the following example.

```
switch# system cores tftp://10.91.51.200/jsmith_cores
switch# 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. An ftp/scp/tftp session will always originate from the switch and either push files to an external system or pull files from an external system.

```
File Server: 172.22.36.10
File to be copied to the switch: /etc/hosts
```

The **copy** CLI command supports four transfer protocols and 12 different sources for files.

```
switch# copy ?
    bootflash: Select source filesystem
    core: Select source filesystem
    debug: Select source filesystem
    ftp: Select source filesystem
    licenses Backup license files
    log: Select source filesystem
    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]"
```

Copy `/etc/hosts` from 172.22.36.10 using the user `user1`, where the destination would be `hosts.txt`.

```
switch# copy scp://user1@172.22.36.10/etc/hosts bootflash:hosts.txt
user1@172.22.36.10's password:
hosts 100% |*****************************| 2035 00:00
```

Back up the startup configuration to an SFTP server.

```
switch# 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:
switch#
```

Tip Back up the startup configuration to a server daily before you make any changes. You can write a short script to be run on the Cisco Nexus 1000V to perform a save and then back up the configuration. The script only needs to contain two commands: **copy running-configuration startup-configuration** and **copy startup-configuration tftp://***server/name*. To execute the script, enter the **run-script** *filename* command.