



Cisco Nexus 6000 Series Troubleshooting Guide

First Published: January 29th, 2014

Cisco Systems, Inc.

www.cisco.com

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

Text Part Number: OL-27923-01

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

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

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

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

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

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

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

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

Cisco Nexus 6000 Series Troubleshooting Guide © 2014 Cisco Systems, Inc. All rights reserved.



Troubleshooting Overview 1-1

Troubleshooting Basics 1-1 Fabric Manager Tools and CLI Commands 1-4 Failover 1-14

Troubleshooting FCoE Issues 1-1

Data Center Bridging1-1FIP1-7CNA1-9PFC1-11Registers and Counters1-15

Troubleshooting Layer 2 Switching Issues 1-1

MAC Address Table 1-1 Spanning Tree Protocol 1-3 Multicast 1-5 VLANs 1-6 Registers and Counters 1-10

Troubleshooting QoS Issues 1-1

Policy Maps 1-1 Improper Configurations 1-2 PFC 1-8 Registers and Counters 1-10

Troubleshooting SAN Switching Issues 1-1

Overview 1-1 NPV 1-2 Zoning 1-7 SAN Port Channels 1-13 FC Services 1-16 Cisco Fabric Services 1-31 VSANs 1-40

Registers and Counters 1-48

Troubleshooting Security Issues 1-1

Roles 1-1 AAA 1-4

Troubleshooting System Management Issues 1-1

SNMP 1-1 Logging 1-3 Traps 1-4 DNS 1-4

Troubleshooting Virtual Port Channel Issues 1-1

Improper Configurations 1-1

Troubleshooting Config-Sync Issues 1-1

Commit Failure 1-1 Import Failure 1-3 Merge Failure 1-5 Switch-profile Deletion Failure 1-6 Verify Failure 1-8

I



This preface describes the audience, organization, and conventions of the *Book Title*, and how to obtain related documentation.

This chapter includes the following topics:

- Audience, page -1
- Organization, page -1
- Document Conventions, page -2
- Related Documentation, page -2

Audience

This publication is for experienced users who configure and maintain Cisco Nexus 6000 Series switch.

Organization

This reference is organized as follows:

Chapter	Description
Troubleshooting Overview	Introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that may occur when configuring and using the Cisco Nexus 6000 Series switch.
Troubleshooting FCoE Issues	Describes how to identify and resolve problems that can occur with FCoE in the Cisco Nexus 6000 Series switch.
Troubleshooting Layer 2 Switching Issues	Describes how to identify and resolve problems that can occur with Layer 2 switching in the Cisco Nexus 6000 Series switch.
Troubleshooting QoS Issues	Describes how to identify and resolve problems that can occur with QoS in the Cisco Nexus 6000 Series switch.
Troubleshooting SAN Switching Issues	Describes how to identify and resolve problems that can occur with SAN switching and the Cisco Nexus 6000 Series switch.
Troubleshooting Security Issues	Describes how to identify and resolve problems that can occur with security in the Cisco Nexus 6000 Series switch.
Troubleshooting System Management Issues	Describes how to identify and resolve problems that can occur with system management and the Cisco Nexus 6000 Series switch.

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 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.



Means the following information will help you solve a problem.

Related Documentation

Documentation for the Cisco Nexus 6000 Series Switch is available at the following URL: http://www.cisco.com/en/US/products/ps12806/tsd_products_support_series_home.html The documentation set is divided into the following categories:

Release Notes

The release notes are available at the following URL: http://www.cisco.com/en/US/products/ps12806/prod_release_notes_list.html

Installation and Upgrade Guides

The installation and upgrade guides are available at the following URL: http://www.cisco.com/en/US/products/ps12806/prod_installation_guides_list.html

Command References

The command references are available at the following URL: http://www.cisco.com/en/US/products/ps12806/prod_command_reference_list.html

Technical References

The technical references are available at the following URL:

http://www.cisco.com/en/US/products/ps12806/prod_technical_reference_list.html

Configuration Guides

The configuration guides are available at the following URL:

 $http://www.cisco.com/en/US/products/ps12806/products_installation_and_configuration_guides_list.html$

Error and System Messages

The system message reference guide is available at the following URL: http://www.cisco.com/en/US/products/ps12806/products_system_message_guides_list.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus6k-docfeedback@cisco.com. We appreciate your feedback.

Г



Troubleshooting Overview

This chapter introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that may occur when configuring and using the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Troubleshooting Basics
- Fabric Manager Tools and CLI Commands
- Failover

Troubleshooting Basics

The following are the basic steps for troubleshooting:

- **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.

To identify the possible problems, you need to use a variety of tools and understand the overall configuration. The following chapters in this guide describe many approaches and specific solutions to potential problems.

Troubleshooting a Switch Crash

When a switch crashes, the cause might be from the failure of a process, and results in a reload of the switch.

A crash is usually recorded with a core file on the switch and includes the reason for the crash, such as a failed process. The following can help you determine the cause of the crash:

• Use the **show version** or **show system reset-reason** commands to display the reason for the crash.

Г

```
Service: clis hap reset
Version: 4.2(1)N2(1)
2) At 841383 usecs after Sat Nov 6 14:56:25 2010
Reason: Reset triggered due to HA policy of Reset
Service: clis hap reset
Version: 4.2(1)N2(1)
```

• Use the **show cores** command to determine if a core file was recorded. You also can use the **show process log** command to display the processes and if a core was created.

switch#show proce	ess log						
Process	PID	Normal-exit	Stack	Core	Log-creat	ce-time	
clis	4023	N	Y	Y	Sat Nov	6 15:14:53	2010
clis	4155	N	Y	Ν	Sat Nov	6 14:56:18	2010

• Use the **show processes log details** command to provide useful information about the reason for the crash:

```
switch# show processes log details
Service: clis
Description: CLI Server
Started at Sat Nov 6 14:59:10 2010 (882984 us)
Stopped at Sat Nov 6 15:14:53 2010 (614588 us)
Uptime: 15 minutes 43 seconds
Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Last heartbeat 9.35 secs ago
RLIMIT_AS: 687474867
System image name: n5000-uk9.4.2.1.N2.1.bin
System image version: 4.2(1)N2(1) S0
PID: 4023
Exit code: signal 11 (core dumped)
```

Threads: 4026 4024 4025

• Note the module-number and the PID number in the output of the **show cores** command for the process that crashed. (Usually the module number is 1 for a Nexus 5000 switch.)

switch#show cores					
Module-num	Instance-num	Process-name	PID	Core-create-time	
1	1	clis	4023	Nov 6 15:20	

- Use the **copy core**://module-id/**PID ftp**: command to export the file and contact the TAC to obtain an analysis of the file.
- Obtain the timestamp of the crash with the **show version**, **show system reset-reason**, or **show cores** commands. With the **show logging** command, review the events that happened just before the crash.

```
switch# show logging
[snip]
2010 Nov 6 08:00:50 TTPSW-5020SF1 %$ VDC-1 %$ %STP-2-BLOCK_BPDUGUARD: Received BPDU
on port Ethernet103/1/1 with BPDU Guard enabled. Disabling port.
2010 Nov 6 08:00:51 TTPSW-5020SF1 %$ VDC-1 %$ %ETHPORT-2-IF_DOWN_ERROR_DISABLED:
Interface Ethernet103/1/1 is down (Error disabled. Reason:BPDUGuard)
2010 Nov 6 14:56:18 TTPSW-5020SF1 %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service
"clis" (PID 4155) hasn't caught signal 11 (core will be saved).
```

Best Practices

Best practices are the recommended steps you should take to ensure the proper operation of your switch.

- Maintain a consistent Cisco NX-OS release across all your Cisco Nexus 5000 switches.
- Refer to the release notes for your Cisco SAN-OS release for the latest features, limitations, and caveats.
- Enable system message logging.
- Troubleshoot any new configuration changes after implementing the change.
- Use the Device Manager to manage your configuration and detect possible problems before they become critical.

Common Terms

Term	Description
DCBX	Data Center Bridging Exchange
RSTP+	Rapid Spanning Tree Protocol
FCoE	Fibre Channel over Ethernet
FCF	Fibre Channel Forwarder
FIP	FCoE Initialization Protocol
PFC	Priority Flow Control
ETS	Enhanced Transmission Selection
LLDP	Link Layer Discovery Protocol
CEE	Converged Enhanced Ethernet
VNTag	Virtual Network Tag
Lossless Ethernet	No-Drop Ethernet
CNA	Consolidated Network Adapter
HBA	Host Bus Adapter
NPV/NPIV	N Port Virtualizer
VN-Link	Virtual Network Link
FEX	Fabric Extender
PAA	Port Analyzer Adapter
RCF	Reconfigure Fabric
RSCN	Request State Change Notification
Menlo	Cisco FCoE MUX ASIC
FCP	Fibre Channel Protocol
FSPF	Fabric Shortest Path First

Fabric Manager Tools and CLI Commands

This section highlights the tools and CLI commands that are commonly used to troubleshoot problems. These tools and commands are a portion of what you may use to troubleshoot your specific problem.

The following chapters in this guide may describe additional tools and commands specific to the symptoms and possible problems covered in that chapter.

NX-OS Tips

Displaying what is required from the configuration

```
switch# show running-config interface
version 4.0(1a)N2(1)
interface vfc29
  no shutdown
  bind interface Ethernet1/29
interface fc2/3
  no shutdown
  switchport speed 1000
  switchport mode SD
interface fc2/4
interface Ethernet1/1
  speed 1000
```

Displaying within Config Mode

With NX-OS, you can display required data from within the configuration mode, so there is no need to back out to the switch prompt.

```
switch(config)# show run
switch(config)# show interface brief
```

Pipe command

```
switch# show logging |
 egrep Egrep
 grep
         Grep
 head
       Stream Editor
 last Display last lines
 less
         Stream Editor
 no-more Turn-off pagination for command output
 sed
         Stream Editor
         Count words, lines, characters
 WC
 begin
         Begin with the line that matches
 count
         Count number of lines
 exclude Exclude lines that match
 include Include lines that match
```

Using the pipe command to only display required keyword

```
switch# show running-config | include switchport
system default switchport
switchport mode trunk
switchport trunk allowed vlan 1,18
switchport mode fex-fabric
switchport mode fex-fabric
switchport speed 1000
switchport mode SD
no system default switchport shutdown
```

Copy command

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
startup-config	Copy startup configuration to destination
system:	Select source filesystem
tftp:	Select source filesystem
volatile:	Select source filesystem

Redirecting output

NX-OS allows you to redirect outputs to files and flash areas in the switch.

```
switch# show tech-support aaa > bootflash:ciscolive09
switch# dir
 103557265
              Apr 01 17:39:22 2009 .tmp-system
     12451
            Apr 10 16:36:37 2009 ciscolive09
     49152
             Apr 01 17:39:22 2009 lost+found/
             Oct 21 13:10:44 2008 n5000-uk9-kickstart.4.0.0.N1.2.bin
  20058112
  20193280
             Apr 01 17:36:37 2009 n5000-uk9-kickstart.4.0.1a.N2.1.bin
             Oct 21 13:11:33 2008 n5000-uk9.4.0.0.N1.2.bin
  76930262
 103557265
             Apr 01 17:37:30 2009 n5000-uk9.4.0.1a.N2.1.bin
      4096
              Jan 01 00:03:26 2005 routing-sw/
```

Redirecting output of the show tech-support details command

Use the **tac-pac** *filename* command to redirect the output of the **show tech-support details** command to a file and then gzip the file.

Book Title

The file is stored on bootflash://filename provided that there is enough memory available. If you do not specify a filename, NX-OS creates the file as volatile:show_tech_out.gz. Copy the file from the device using the procedure in the copy command section.

```
switch# tac-pac
switch# dir volatile:
374382 Aug 16 17:15:55 2010 show_tech_out.gz
```

From volatile, copy the file to the bootflash, FTP, or TFTP server.

```
switch# copy volatile:show_tech_out.gz ?
bootflash: Select destination filesystem
debug: Select destination filesystem
ftp: Select destination filesystem
nodflash: Select destination filesystem
nvram: Select destination filesystem
running-config Copy from source to running configuration
scp: Select destination filesystem
startup-config Copy from source to startup configuration
system: Select destination filesystem
tftp: Select destination filesystem
volatile: Select destination filesystem
```

NX-OS command listing

```
switch# show cli list | include ?
    -i Ignore case difference when comparing strings
    -x Print only lines where the match is a whole line
    WORD Search for the expression
switch# show cli list | include debug | include interface
```

Narrowing scope of keywords

You can use many commands like grep and include to narrow the scope of a keyword.

```
switch(config-if)# show interface | grep fc
fc2/1 is trunking
fc2/2 is trunking
fc2/3 is up
fc2/4 is down (Administratively down)
vfc29 is up
```

Logging

You can use logging through the CLI or Device Manager. In the following examples, the **logging** command and the Device Manager display severity information:

Viewing Severity Information with the CLI

```
switch(config)# show logging
```

REVIEW DRAFT-CISCO CONFIDENTIAL

Logging	console:	enabled	(Severity:	critical)
Logging	monitor:	enabled	(Severity:	notifications)
Logging	linecard:	enabled	(Severity:	notifications)
Logging	fex:	enabled	(Severity:	notifications)
Logging	timestamp:	Seconds		
Logging	server:	enabled		
{10.91.4	42.134}			
	server severity:	notifica	tions	
	server facility:	local7		
	server VRF:	manageme	nt	
Logging	logflash:	disabled	l	
Logging	logfile:	enabled		
Name - c	ciscolive09: Severity - d	lebugging	Size - 41	94304

Viewing Severity Levels in the Device Manager



Ethanalyzer and SPAN

Ethanalyzer is a tool that collects frames that are destined to, or originate from, the Nexus 5000 control plane. Node to switch or switch to switch traffic can be seen with this tool.

SPAN is a feature whereby frames that are transient to the switch are copied to a second port for analysis. Node to switch or node to node traffic can be seen via this method.

Ethanalyzer

Ethanalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark open source code. This tool is a command-line version of Wireshark that captures and decodes packets. You can use Ethanalyzer to troubleshoot your network and analyze the control-plane traffic.

99998

Command	Description
ethanalyzer local sniff-interface	Captures packets sent or received by the supervisor and provides detailed protocol information.
ethanalyzer local sniff-interface brief	Captures packets sent or received by the supervisor and provides a summary of protocol information.
ethanalyzer local sniff-interface limit-captured-frames	Limits the number of frames to capture.
ethanalyzer local sniff-interface limit-frame-size	Limits the length of the frame to capture.
ethanalyzer local sniff-interface capture-filter	Filters the types of packets to capture.
ethanalyzer local sniff-interface display-filter	Filters the types of captured packets to display.
ethanalyzer local sniff-interface decode-internal	Decodes the internal frame header for Cisco NX-OS.
	Note Do not use this option if you plan to analyze the data using Wireshark instead of NX-OS Ethanalyzer.
ethanalyzer local sniff-interface write	Saves the captured data to a file.
ethanalyzer local sniff-interface read	Opens the captured data file and analyzes it.

Examples

```
switch# ethanalyzer local sniff-interface
No matches in current mode, matching in (exec) mode
inbound-hi Inbound(high priority) interface
inbound-low Inbound(low priority) interface
mgmt Management interface
```

```
switch# ethanalyzer local sniff-interface mgmt brief
Capturing on eth0
2008-08-13 01:34:23.776519 10.116.167.244 -> 172.18.217.80 TCP 1106 > telnet [ACK] Seq=0
Ack=0 Win=64040 Len=0
2008-08-13 01:34:23.777752 172.18.217.80 -> 10.116.167.244 TELNET Telnet Data ...
2008-08-13 01:34:23.966262 00:04:dd:2f:75:10 -> 01:80:c2:00:00:00 STP Conf. Root =
32768/00:04:c1:0f:6e:c0 Cost = 57 Port = 0x801d
[snip]
```

The following example is for viewing the Spanning Tree Protocol (STP) and Fibre Channel: Using 0 in the command captures output until you press **Ctrl-C**. The FCID is a well-known name for switch domain controller.

switch# ethanalyzer local sniff-interface inbound-hi brief limit-captured-frames 0

Capturing on eth4

```
2008-08-13 01:37:16.639896 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 1 0 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 0x0 0x0 STP RST. Root = 32769/00:0d:ec:6b:cd:41 Cost = 0 Port = 0x8093 2008-08-13 01:37:18.639992 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 1 0 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 1 0 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 0x0 0x0 STP RST. Root = 32769/00:0d:ec:6b:cd:41 Cost = 0 Port = 0x8093
```

[snip]

2008-08-13 01:37:23.220253 00:0d:ec:6b:cd:40 -> fc:fc:fc:ff:ff:fd 4 0 ff.ff.fd -> ff.ff.fd 0x5384 0xffff SW ILS ELP ff.ff.fd -> 2008-08-13 01:37:23.220615 00:0d:ec:6b:cd:40 -> aa:bb:cc:dd:01:04 4 0 ff.ff.fd 0x5384 0x2b3f FC Link Ctl, ACK1 2008-08-13 01:37:23.227202 00:0d:ec:6b:cd:40 -> aa:bb:cc:dd:01:04 4 0 ff.ff.fd -> ff.ff.fd 0x5384 0x2b3f SW_ILS SW_ACC (ELP) ff.ff.fd -> 2008-08-13 01:37:23.229927 00:0d:ec:6b:cd:40 -> fc:fc:fc:ff:ff:fd 4 0 ff.ff.fd 0x5384 0x2b3f FC Link Ctl, ACK1

Detailed BPDU

```
switch# ethanalyzer local sniff-interface inbound-hi limit-captured-frames 0
Capturing on eth4
Frame 1 (57 bytes on wire, 57 bytes captured)
   Arrival Time: Aug 13, 2008 01:41:32.631969000
    [Time delta from previous captured frame: 1218591692.631969000 seconds]
    [Time delta from previous displayed frame: 1218591692.631969000 seconds]
    [Time since reference or first frame: 1218591692.631969000 seconds]
   Frame Number: 1
   Frame Length: 57 bytes
   Capture Length: 57 bytes
    [Frame is marked: False]
    [Protocols in frame: eth:vlan:llc:stp]
[snip]
   DSAP: Spanning Tree BPDU (0x42)
   IG Bit: Individual
   SSAP: Spanning Tree BPDU (0x42)
   CR Bit: Command
   Control field: U, func=UI (0x03)
       000. 00.. = Command: Unnumbered Information (0x00)
        .... ..11 = Frame type: Unnumbered frame (0x03)
[snip]
```

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, a Fibre Channel Analyzer, or other Remote Monitoring (RMON) probes.

SPAN sources refer to the interfaces from which traffic can be monitored. The Cisco Nexus 5000 Series switch supports Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, port channels, SAN port channels, VLANs, and VSANs as SPAN sources. With VLANs or VSANs, all supported interfaces in the specified VLAN or VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for Ethernet, virtual Ethernet, Fibre Channel, and virtual Fibre Channel source interfaces:

- Ingress source (Rx)—Traffic entering the switch through this source port is copied to the SPAN destination port.
- Egress source (Tx)—Traffic exiting the switch through this source port is copied to the SPAN destination port.



For the Cisco Nexus 5548 Switch, Fibre Channel ports cannot be configured as ingress source ports in a SPAN session.

Source Ports

A source port, also called a monitored port, is a switched interface that you monitor for network traffic analysis. The switch supports any number of ingress source ports (up to the maximum number of available ports on the switch) and any number of source VLANs or VSANs.

A source port has these characteristics:

- Can be of any port type: Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, port channel, SAN port channel, VLAN, and VSAN.
- Cannot be monitored in multiple SPAN sessions.
- Cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For VLAN, VSAN, port channel, and SAN port channel sources, the monitored direction can only be ingress and applies to all physical ports in the group. The rx/tx option is not available for VLAN or VSAN SPAN sessions.
- Beginning with Cisco NX-OS Release 5.0(2)N1(1). Port channel and SAN port channel interfaces can be configured as ingress or egress source ports.
- Source ports can be in the same or different VLANs or VSANs.
- For VLAN or VSAN SPAN sources, all active ports in the source VLAN or VSAN are included as source ports.
- The Cisco Nexus 5010 switch supports a maximum of two egress SPAN source ports. This limit does not apply to the Cisco Nexus 5020 Switch and the Cisco Nexus 5548 switch.

SPAN Destinations

SPAN destinations refer to the interfaces that monitors source ports. The Cisco Nexus 5000 Series switch supports Ethernet and Fibre Channel interfaces as SPAN destinations.

Source SPAN	Destination SPAN
Ethernet	Ethernet
Fibre Channel	Fibre Channel
Fibre Channel	Ethernet (FCoE)
Virtual Ethernet	Ethernet
Virtual Fibre Channel	Fibre Channel
Virtual Fibre Channel	Ethernet (FCoE)

Characteristics of Destination Ports

Each local SPAN session must have a destination port (also called a monitoring port) that receives a copy of traffic from the source ports, VLANs, or VSANs. A destination port has these characteristics:

- Can be any physical port, Ethernet, Ethernet (FCoE), or Fibre Channel. Virtual Ethernet and virtual Fibre Channel ports cannot be destination ports.
- Cannot be a source port.
- Cannot be a port channel or SAN port channel group.

- Does not participate in spanning tree while the SPAN session is active.
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session.
- Receives copies of sent and received traffic for all monitored source ports. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.

Monitor Caveats

Limitations of Nexus 5000 SPAN CoS values are not preserved at the monitor (span) destination.

- Packets coming in on the monitor source with an unknown VLAN tag are spanned ouf with a 0 VLAN tag (priority tag).
- For Ethernet destination, the monitor session is up only if the destination port is configured as switch port monitor.
- Out of 18 configurable sessions, only two are active (up state). The rest are in down state (hardware resource unavailable).

Configuration limitations: VLAN or port-channel cannot be configured as egress source

- VLAN or port channel cannot be a monitor destination.
- Only two egress sources supported.
- Only one destination port can be configured for a session.

SPAN Configuration

Example:

```
switch(config)# interface fc3/2
switch(config-if)# switchport mode sd
switch(config-if)# switchport speed 1000
switch(config-if)# no shut
switch(config-if)# exit
switch(config)# monitor session 1
switch(config-monitor)# source interface fc3/1 tx
switch(config-monitor)# source interface fc3/1 rx
switch(config-monitor)# destination interface fc3/2
```



Verifying the SPAN Session

Example:

switch#	show monitor	session	DESCRIPTION
SESSION	STATE	REASON	
 1	 up	The session is up	

or session 1
: local
: up
:
: fc3/1
: fc3/1
: fc3/1
:
:
:
:
: fc3/2

Suspending the SPAN Session

Example:

switch(config) # monitor session 1 suspend

```
switch(config)# show monitor session 1
  session 1
  _____
                : local
type
state
                : down (Session suspended)
state
source intf
rx
               :
               : fc3/1
               : fc3/1
   tx
   both
               : fc3/1
source VLANs
               :
   rx
               :
source VSANs
                :
   rx
                :
destination ports : fc3/2
```

Debugging

Command-Line Debugging

Available debugs depend on features enabled in NX-OS. There are many different options to select when turning on debugs.

Determine the destination of the output:

- Logfile—Data file in switch memory.
- Capture to direct to screen via console, Telnet, or SSH.

You must have administrator privileges to run debugs. Debugs can only be run from the CLI.

Debug Logging

Set the log file as CiscoLive_debugs, using the **debug logfile** command. Then, use the **show debug** command to see name of the debug file.

switch# debug logfile CiscoLive_debugs
switch# show debug

Display debugging to the screen with the following command:

switch# show debug logfile CiscoLive_debugs

Copy the debug file from MDS to a server with the **copy** command. When you enter the VRF, if none is specified then the default is used.

switch# copy log:CiscoLive_debugs tftp:

Enter vrf: management Enter hostname for the tftp server: 10.91.42.134 Trying to connect to tftp server..... Connection to Server Established.

To delete the debug logfile, use one of the following commands:

switch# clear debug-logfile CiscoLive_debugs

switch# undebug all

If you do not use one of these commands, the debug logfile will be cleared and overwritten when the next debug logfile is created. The system only allows one debug logfile to exist.

Debugs to the Direct Telnet Window

- Use a Telnet, SSH, or console application that captures the expected output to buffer or file.
- Undebug all or no debug of a specific debug command is required to turn trace off.
- The debugs are not persistent across reboots
- Most debugs are easy to read and understand, but some are not.

Cisco Discover Protocol

Cisco Discover Protocol (CDP) version 2 is applied to the physical Ethernet interface and only works when enabled at both ends of the link. LLDP standard is derived from CDP.

CDP is used to verify proper connectivity to correct network devices, very useful at switch deployment.

The following example shows the arguments that can be used with the show CDP command:

switch# show cd	p
all	Show interfaces that are CDP enabled
entry	Show CDP entries in database
global	Show CDP global parameters
interface	Show CDP parameters for an interface
neighbors	Show CDP neighbors
traffic	Show CDP traffic statistics
switch# show cd Global CDP info CDP enabled Sending CDP Sending a h Sending CDP Sending Dev	p global rmation: globally packets every 60 seconds oldtime value of 180 seconds v2 advertisements is enabled iceID TLV in Default Format

Г

```
Device ID:TM-6506-1
System Name:
Interface address(es):
   IPv4 Address: 11.1.1.1
Platform: cisco WS-C6506, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet1/4, Port ID (outgoing port): TenGigabitEthernet1/2 ? Verifies proper
port connections
Holdtime: 133 sec
Version:
Cisco Internetwork Operating System Software
IOS (tm) s72033_rp Software (s72033_rp-IPSERVICES_WAN-VM), Version 12.2(18)SXF11, RELEASE
SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2007 by cisco Systems, Inc.
Compiled Fri 14-Sep-07 23:09 by kellythw
Advertisement Version: 2
Native VLAN: 1
                 ? Sent on Native VLAN
Duplex: full
```

Failover

FCoE Traffic

When the Nexus 5000 experiences loss of fabric connectivity, it brings down all the affected vFC interfaces.

The following methods are used to signal the host of loss of connectivity to the FC fabric

- FIP Clear Link Virtual Link to the CNA will be signaled to indicate the 'shut' state of vFC. Throughout the 'shut' period FCF Advertisements indicate 'not available for login'.
- In case the loss of connectivity is over the FCoE network, FIP keep-alives are used by the FCF and the CNA to timeout the login sessions. The keep-alive timers are configurable.

Non-FCoE traffic

Under certain failure scenarios where the access switch has lost all uplink connectivity to the aggregation layer, the CNA needs to be signaled of the loss of LAN connectivity. This helps the CNA failover the host traffic to the standby port. Traditionally, such a failure is signaled by bringing down the host facing link. Bringing down the link achieves two purposes:

- Host is signaled of loss of connectivity.
- The access switch stops forwarding traffic to and from the host-facing link.

However, in the converged network, even though the LAN connectivity is lost at the access switch, the SAN connectivity might still be intact. Bringing down the entire host-facing link is not desirable. Instead, the loss of connectivity is signaled over protocols. Loss of SAN connectivity is signaled using the FIP Clear Virtual Link message. Loss of LAN connectivity is signaled using logical link status TLVs defined in DCBX and VIC protocols.

REVIEW DRAFT-CISCO CONFIDENTIAL

LAN Traffic

When LAN connectivity is lost for a particular VLAN on the uplinks, the VLAN is also brought down on the host-facing link.

Dedicating a VLAN solely for FCoE traffic helps with shutting down non-FCoE traffic to and from the host-facing link without disrupting FCoE traffic from the same host.

1-16



Troubleshooting FCoE Issues

Fibre Channel over Ethernet (FCoE) provides a method of transporting Fibre Channel traffic over a physical Ethernet connection. FCoE requires that the underlying Ethernet be full duplex and provides lossless behavior for Fibre Channel traffic.

This chapter describes how to identify and resolve problems that can occur with FCoE in the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Data Center Bridging
- FIP
- CNA
- PFC
- Registers and Counters

Data Center Bridging

VFC (FCoE) interface not online

This section includes the following topics:

- General troubleshooting
- Nexus 5548 Troubleshooting

General troubleshooting

Possible Cause

An FCoE-attached server has no connectivity to FC, or FCoE-attached storage, and the **show interface** command for the virtual Fibre Channel interface mapped to this server's port reveals that the VFC interface is down.

<u>Solution</u>

• Verify the configuration using the **show running-config** command.

Example:

Γ

```
<u>Note</u>
```

The default setting for VFC is shutdown, however, in the following example was changed by the setup script.

```
switch# show running-config
feature fcoe
vlan 1
vlan 100
fcoe
vsan database
vsan 100
interface vfc4
bind interface Ethernet1/4
no shutdown
vsan database
vsan 100 interface vfc4
interface fc2/1
no shutdown
interface Ethernet1/4
switchport mode trunk
switchport trunk allowed vlan 100
spanning-tree port type edge trunk
```

• Check to ensure that the LLDP Transmit and Receive are enabled on the interface. Use the **show lldp interface ethernet 1/4** command.

Example:

switch# show lldp interface ethernet 1/4

```
Interface Information:
Enable (tx/rx/dcbx): Y/Y/Y Port Mac address: 00:0d:ec:d5:a3:8b
Peer's LLDP TLVs:
Type Length Value
---- -----
001 007 0400c0dd 145486
002 007 0300c0dd 145486
003 002 0078
128 061 001b2102 020a0000 00000002 00000001 04110000 c0000001 00003232
00000000 00000206 060000c0 00080108 100000c0 00890600 1b210889
14001b21 08
000 000
```

If LLDP is disabled, the VFC will not come online. You can enable LLDP transmit and receive with the **interface ethernet 1/4** command:

```
switch(config)# interface ethernet 1/4
switch(config-if)# lldp ?
receive Enable LLDP reception on interface
transmit Enable LLDP transmission on interface
```

• Check that the peer supports LLDP.

Check if remote peers exist. Check if values exist for a peer's LLDP TLVs. Use the **show lldp interface ethernet 1/4** command.

Example:

```
switch# show lldp interface ethernet 1/4
Interface Information:
Enable (tx/rx/dcbx): Y/Y/Y Port Mac address: 00:0d:ec:d5:a3:8b
Peer's LLDP TLVs:
Type Length Value
---- -----
```

```
001 007 0400c0dd 145486

002 007 0300c0dd 145486

003 002 0078

128 061 001b2102 020a0000 00000002 00000001 04110000 c0000001 00003232

00000000 00000206 060000c0 00080108 100000c0 00890600 1b210889

14001b21 08
```

• Check the peer (CNA) to see if it supports DCBX.

```
Use the show system internal dcbx info interface ethernet 1/4 command.
(For releases earlier than 4.2(1)N1, use the "sh platform software dcbx internal info interface ethernet x/y" command.)
```

```
<u>Note</u>
```

In the example, DCBX is enabled and the peer supports CEE.

Example:

```
switch# show system internal dcbx info interface ethernet 1/4
Interface info for if_index: 0x1a003000(Eth1/4)
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE
DCX Protocol: CEE
Port MAC address: 00:0d:ec:d5:a3:8b
DCX Control FSM Variables: seq_no: 0x1, ack_no: 0x2,my_ack_no: 0x1, peer_seq_no: 0x2
oper_version: 0x0, max_version: 0x0 fast_retries 0x0
Lock Status: UNLOCKED
PORT STATE: UP
```

• In the output from the **show system internal dcbx info interface ethernet 1/4** command, check the peers LLDP values.

Make sure that the mandatory LLDP values exist.

Example:

switch# show system internal dcbx info interface ethernet $1/4\,$

```
LLDP Neighbors
Remote Peers Information on interface Eth1/4
Remote peer's MSAP: length 12 Bytes:
00 c0 dd 14 54 86 00 c0 dd 14 54 86
LLDP TLV's
LLDP TLV type: Chassis ID LLDP TLV Length: 7
04 00 c0 dd 14 54 86
Chassis type: 04 Chassis ID:00 c0 dd 14 54 86
LLDP TLV type:Port ID LLDP TLV Length: 7
03 00 c0 dd 14 54 86
Port ID subtype: 03 Port ID:00 c0 dd 14 54 86
LLDP TLV type: Time to Live LLDP TLV Length: 2
00 78
TTL = 00
LLDP TLV type:Unknown 128 LLDP TLV Length: 61
00 1b 21 02 02 0a 00 00 00 00 00 02 00 00 01 04 11 00 00 c0 00
00 01 00 00 32 32 00 00 00 00 00 00 02 06 06 00 00 c0 00 08 01 08
10 00 00 c0 00 89 06 00 1b 21 08 89 14 00 1b 21 08
LLDP TLV type: END of LLDP PDU LLDP TLV Length: 0
```

• In the output from the **show system internal dcbx info interface ethernet 1/4** command, check the peers DCBX TLVs.

Make sure that PFC and FCoE TLV were negotiated as willing and enabled, and that there are no errors.

Example:

switch# show system internal dcbx info interface ethernet 1/4

```
Peer's DCX TLV:
DCBX TLV Proto(1) type: 1(Control) DCBX TLV Length: 10 DCBX TLV Value
00 00 02 00 00 01 00 00 00
sub_type 0, error 0, willing 0, enable 0, max_version 0, oper_version 0
DCBX TLV Proto(1) type: 2(PriGrp) DCBX TLV Length: 17 DCBX TLV Value
00 00 c0 00 00 1 00 00 32 32 00 00 00 00 00 00 02
sub_type 0, error 0, willing 1, enable 1, max_version 0, oper_version 0
DCBX TLV Proto(1) type: 3(PFC) DCBX TLV Length: 6 DCBX TLV Value
00 00 c0 00 08 01
sub_type 0, error 0, willing 1, enable 1, max_version 0, oper_version 0
DCBX TLV Proto(1) type: 4(App(Fcce)) DCBX TLV Length: 16 DCBX TLV Value
00 00 c0 00 89 06 00 1b 21 08 89 14 00 1b 21 08
sub_type 0, error 0, willing 1, enable 1, max_version 0, oper_version 0
```

• Check the peer PFC and FCoE subtypes.

Use the show **system internal dcbx info interface ethernet 1/4** command. (For releases earlier than 4.2(1)N1, use the **sh platform software dcbx internal info interface ethernet x/y** command.)

Example:

switch# show system internal dcbx info interface ethernet 1/4

```
Feature type PFC (3)
feature type 3(PFC)sub_type 0
Feature State Variables: oper_version 0 error 0 local error 0 oper_mode 1
feature_seq_no 0 remote_feature_tlv_present 1 remote_tlv_aged_out 0
remote_tlv_not_present_notification_sent 0
Feature Register Params: max_version 0, enable 1, willing 0 advertise 1
disruptive_error 0 mts_addr_node 0x101 mts_addr_sap 0x179
Desired config cfg length: 2 data bytes:08 08
Operating config cfg length: 2 data bytes:08 08
Peer config cfg length: 0 data bytes:
Feature type App(Fcoe) (4) sub_type FCoE (0)
feature type 4(App(Fcoe))sub_type 0
Feature State Variables: oper_version 0 error 0 local error 0 oper_mode 1
feature_seq_no 0 remote_feature_tlv_present 1 remote_tlv_aged_out 0
remote_tlv_not_present_notification_sent 0
Feature Register Params: max_version 0, enable 1, willing 0 advertise 1
disruptive_error 0 mts_addr_node 0x101 mts_addr_sap 0x179
```

• Check the DCBX counters located at the very bottom of the output display from the **show system** internal dcbx info interface ethernet 1/4 command. Look for any errors.

Example:

```
Traffic Counters
DCBX pkt stats:
Total frames out: 15383
Total Entries aged: 97
Total frames in: 15039
DCBX frames in: 15033
Total frames received in error: 6
Total frames discarded: 6
Total TLVs unrecognized: 0
```

- Check for the same values for the FCoE Data Center Bridging and the Type-Length-Value on the host CNA software.
- Ensure that the VSAN trunk protocol has been enabled.

Use the **configuration terminal** command to enter into configuration mode and use the **trunk protocol enable** command to enable the trunking protocol.

Solution Summary

• Review every feature negotiation result.

Use the **show system internal dcbx info interface ethernet 1/4** command. (For releases earlier than 4.2(1)N1, use the **sh platform software dcbx internal info interface ethernet x/y** command.)

Example:

switch# show system internal dcbx info interface ethernet 1/4

feature type 3 sub_type 0
feature state variables: oper_version 0 error 0 oper_mode 1 feature_seq_no 0
remote_feature_tlv_present 1
remote_tlv_not_present_notification_sent 0 remote_tlv_aged_out 0
feature register params max_version 0, enable 1, willing 0 advertise 1,
disruptive_error 0 mts_addr_node
0x101mts_addr_sap 0x1e5
Desired config cfg length: 1 data bytes:08
Operating config cfg length: 1 data bytes:08

- Errors
 - Indicates negotiation error.
 - Never expected to happen when connected to CNA.
 - When two Nexus 5000 switches are connected back-to-back, and if PFC is enabled on different CoS values, then a negotiation error can occur.
- Operating configuration
 - Indicates negotiation result.
 - Absence of operating configuration indicates that the peer does not support the DCBX TLV or that there is a negotiation error.
 - The remote_feature_tlv_present message indicates whether the remote peer supports this feature TLV or not.
- DCBX feature might not be working because:
 - Peer does not support the LLDP Protocol.
 - Peer does not support the DCBX Protocol.
 - Peer does not support some DCBX TLVs.
 - Unexpected DCBX negotiation result.
- An option exists to force PFC mode on an interface. Use the inerfacet ethernet 1/21 command and the priority-flow-control mode command to force the PFC mode.

Example:

```
switch(config)# int eth1/21
switch(config-if)# priority-flow-control mode ?
```

L

```
auto Advertise priority-flow-control capability on Turn on priority-flow-control
```

<u>Note</u>

The default setting for this command is auto. The **no** option returns the mode to auto.

Nexus 5548 Troubleshooting

Possible Cause

The type of Converged Network Adapter might not be supported.

Solution

Ensure that the type of adapter is supported. The FCoE interface only supports a Generation-2 Converged Network Adapter.

Possible Cause

The FCoE class-fcoe system class is not enabled in the QoS configuration.

Solution

For a Cisco Nexus 5548 switch, the FCoE class-fcoe system class is not enabled by default in the QoS configuration. Before enabling FCoE, you must include class-fcoe in each of the following policy types:

- Network-QoS
- Queuing
- QoS

The following is an example of a service policy that needs to be configured:

```
F340.24.10-5548-1
class-map type qos class-fcoe
class-map type queuing class-fcoe
match qos-group 1
class-map type queuing class-all-flood
match qos-group 2
class-map type queuing class-ip-multicast
match qos-group 2
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-all-flood
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
system qos
```

service-policy type qos input fcoe-default-in-policy service-policy type queuing input fcoe-default-in-policy service-policy type queuing output fcoe-default-out-policy service-policy type network-qos fcoe-default-nq-policy

FIP

<u>Note</u>

FIP Generation-1 CNAs are not supported on the Nexus 2232 FEX. Only FIP Generation-2 CNAs are supported on the Nexus 2232 FEX.

VFC down due to FIP failure

Host is not capable of supporting FIP-related TLVs.

Possible Cause

When the connected host does not support FIP, the first step of VLAN-discovery fails based on which VFC is brought up. Use show commands to verify that the three basic TLVs required for FIP are exchanged by DCBX over the bound interface, and that FCOE-MGR is enabled for FIP. The three TLVs are FCoE TLV, PriGrp TLV, and PFC TLV. These three TLVs should be checked for both local and peer values.

Verify the TLVs with the following commands:

- show system internal dcbx info interface <bound-ethernet-interface-id>
- show platform software fcoe_mgr info interface vfc<id>

In the output from the commands:

- Check for FIP capable is TRUE.
- Check for triggered event: [FCOE_MGR_VFC_EV_FIP_VLAN_DISCOVERY].

The state of the VFC never progresses further to solicitation.

<u>Solution</u>

Make sure you check for correct FIP supporting firmware and drivers on the CNA and FIP supporting adapters.

VFC down due to FIP solicitation failure

When the FIP solicitation fails, the VFC goes down.

Possible Cause

Once the first step of FIP VLAN-discovery has succeeded, the host sends FIP solicitations. The switch should respond with FIP advertisements in detail. If the response is not sent or the advertisement is not sent back to the solicitation received, the VFC does not come up. The host continues trying to solicit, but never succeeds.

The following are possible reasons for no response or advertisement:

- No active fabric-provided MAC address exists. (Possible wrong fc-map, etc.)
- Fabric is not available for FLOGI.
- MAC address descriptor may be incorrect. (This is the address the CNA uses as the DMAC when it sends responses.)

Use the **show platform software fcoe_mgr info interface vfc**<*id*> command to view the status of the FIP solicitation.

Г

In the output from the command, check for triggered event: [FCOE_MGR_VFC_EV_FIP_VLAN_DISCOVERY];

followed by triggered event: [FCOE_MGR_VFC_EV_FIP_SOLICITATION].

If the solicitation is successful, then triggered event: [FCOE_MGR_VFC_EV_FIP_FLOGI] is displayed.

If the solicitation has failed, then triggered event: [FCOE_MGR_VFC_EV_FIP_FLOGI] is not displayed and no further progress occurs.

<u>Solution</u>

Need to check and ensure that the VSAN is active, the memberships are correct, and that the fabric is available. Also while in NPV mode, check that an active border/NP port is available.

VFC down because VLAN response not received by CNA

Though the switch sends out a VLAN response, the response is not received by the CNA. This indicates that the VFC is down.

Possible Cause

A bound interface native VLAN ID should be a non-FCoE VLAN. If not, and the native VLAN matches the FCoE VLAN, the VLAN response sent out will be untagged. However, the FIP adapters expect tagged frames. This means that the native VLAN on the trunk interface should be a non-FCoE VLAN.

<u>Solution</u>

Check the configuration on the bound Ethernet trunk interface and ensure that it is a non-FCoE native VLAN.

VFC down because no active STP port-state on the bound Ethernet interface

No active STP port-state on the bound Ethernet interface causes the VFC to be down.

Possible Cause

The bound interface should be in a STP-forwarding state for both the native VLAN and the member FCOE VLAN mapped to the active VSAN. If there are no STP active ports on the VLAN, then the switch drops all FIP packets received on the VLAN over the bound interface. This means that the FIP is not initiated to bring up the VFC.

<u>Solution</u>

Check the STP port state on the bound Ethernet trunk interface for both non-FCOE native VLAN and FCOE member VLAN. Fix the STP port state and move it to forwarding, if in blocked inconsistent state or error-disable state.

VFC down due to FIP keepalive misses

The VFC goes down due to FIP keepalive misses.

Possible Cause

When FIP keepalives (FKA) are missed for a period of approximately 22 seconds, this means that approximately three FKAs are not continuously received from the host. Missed FKAs can occur for many reasons, including congestion or link issues.

1-8

FKA timeout : 2.4 * FKA_adv_period.

The FKA_adv_period is exchanged and agreed upon with the host as in the FIP advertisement when responding to a solicitation.

Observe the output from the following commands to confirm FKA misses:

- show platform software fcoe_mgr info interface vfc<id>
- show platform software fcoe_mgr event-history errors
- show platform software fcoe_mgr event-history lock
- show platform software fcoe_mgr event-history msgs
- show platform fwm info pif ethernet <bound-ethernet-interface-id>

<u>Solution</u>

Sometimes when congestion is relieved, the VFC comes back up. If the symptom persists, then additional analysis is required. The possible considerations are:

- The host stopped sending the FKA.
- The switch dropped the FKA that was received.

CNA

This section includes an overview of best practices for the topology of the Converged Network Adapter (CNA), a description of troubleshooting with host-based tools, followed by a description of common problems and their solutions.

Best practice topology for CNA

Best Practice Topology for Direct Connected CNA

- A unique dedicated VLAN must be configured at every converged access switch to carry traffic for each virtual fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If MSTP is enabled, a separate MST instance must be used for FCoE VLANs
- Unified Fabric (UF) links must be configured as trunk ports. FCoE VLAN must not be configured as a native VLAN. All FCoE VLANs must be configured as members of the UF links. This allows it to be extendible for VF_Port trunking and VSAN management for the VFC interfaces.
- UF links must be configured as spanning tree edge ports.
- FCoE VLANs must not be configured as members of Ethernet links that are not designated to carry FCoE traffic. This ensures to limit the scope of the spanning-tree protocol for FCoE VLANs to UF links only.
- If the converged access switches (in the same SAN fabric or in the other) need to be connected to
 each over Ethernet links for the purposes of LAN alternate pathing, then such links must explicitly
 be configured to exclude all FCoE VLANs from membership. This ensures to limit the scope of the
 Spanning Tree Protocol for FCoE VLANs to UF links only.
- Separate FCoE VLANs must be used for FCoE in SAN-A and SAN-B.

Г

Best Practice Topology for Remote Connected CNAs

- A unique dedicated VLAN must be configured at every converged access switch and every blade switch to carry traffic for each virtual fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1002 for VSAN 2, and so on). If MSTP is enabled, a separate MST instance must be used for FCoE VLANs.
- Unified Fabric (UF) links must be configured as trunk ports. FCoE VLAN must not be configured as a native VLAN. All FCoE VLANs must be configured as members of the UF links. This allows it to be extendible for VF_Port trunking and VSAN management for the VFCs.
- UF links between the CNAs and the blade switches must be configured as spanning tree edge ports.
- A blade switch must connect to exactly one converged access switch, preferably over an Ethernet port channel to avoid disruption due to STP reconvergence on events such as provisioning of new links or blade switches.

Troubleshooting with Host tools

You can troubleshoot the CNA with following host-based tools:

- Emulex
 - Emulex provides the OneCommand GUI tool to manage Emulex CNAs. The CEE tab of this tool displays details about DCB configurations and FIP settings within the FC interface.
- Qlogic
 - Qlogic provides the SanSurfer tool. The Data Center Bridging tab of this tool displays the DCB configuration learned from the switch alone with TLV exchange data. The DCE Statistics tab of this tool displays the ethernet statistics.
- Microsoft Windows
 - Microsoft Windows provides tools to view the configuration and registers for many CNA vendor products.

CNA not recognized by Host OS

Although the CNA is installed on the host, the Converged Network Adapter (CNA) is not recognized.

Possible Cause

The host operating system may not have the appropriate drivers to support the installed Converged Network Adapter model.

Solution

- **Step 1** 1) Obtain the following information:
 - a. Operating system of the host.
 - **b.** Specific model of installed CNA.
- Step 2 Reference the appropriate vendor support page for the CNA model and host OS.
- **Step 3** Determine if an existing driver is already installed on the host OS.
- **Step 4** Ensure that the latest driver is installed from the CNA vendor support page or the host OS support page.

PFC

This section includes an overview of how to view standard pause frames, followed by a description of common problems and their solutions.

Standard pause frames

For ports with standard, non-CNA type host connections, the Nexus 5000 supports standard pause frames. These are enabled with the interface setting, as shown in the following example:

Example:

```
switch(config)# interface ethernet 1/16
switch(config-if)# flowcontrol ?
    receive Receive pause frames
    send Send pause frames
switch(config-if)# flowcontrol receive on
switch(config-if)# flowcontrol send on
```

To view standard pause frames, use the show interface flowcontrol command.

Example:

switch(config-if)# show interface flowcontrol

Port	Send admin	FlowControl oper	Receive admin	FlowControl oper	RxPause	TxPause
Eth1/1	off	off	off	off	0	0
Eth1/2	off	off	off	off	0	0
Eth1/3	off	off	off	off	0	0
Eth1/4	off	off	off	off	0	0
Eth1/5	off	off	off	off	0	0
Eth1/6	off	off	off	off	0	0
Eth1/7	off	off	off	off	0	0
Eth1/8	off	off	off	off	0	0
Eth1/9	off	off	off	off	0	0
Eth1/10	off	off	off	off	0	0
Eth1/11	off	off	off	off	0	0
Eth1/12	off	off	off	off	0	0
Eth1/13	off	off	off	off	0	0
Eth1/14	off	off	off	off	0	0
Eth1/15	off	off	off	off	0	0
Eth1/16	on	on	on	on	0	0
Eth1/17	off	off	off	off	0	0

PFC not negotiated with FCOE-capable adapters (CNA)

Priority flow control (PFC) is not negotiated with FCOE-capable adapters (CNA).

This causes packet drop to be noticed on FCoE traffic from the servers.

Possible Causes

The CNA may not support DCBX and the PFC TLV is not negotiated.

<u>Solution</u>

Use the following information to verify DCBX support and that the PFC TLV is negotiated:

• Check the status of the PFC. Use the **show int ethx/x priority-flow-control** command. (Connected to CNA.)

Example:

```
      switch# show interface ethernet 1/13 priority-flow-control

      Port
      Mode Oper(VL bmap)

      RxPPP
      TxPPP

      Ethernet1/13
      Auto Off
      0
```

• Check for LLDP neighbor or PFC/DCBX TLV advertised by the peer. Use the **show system** internal dcbx info int ethx/x command.

Example:

switch(config-if)# show system internal dcbx info interface ethernet 1/1

```
Interface info for if_index: 0x1a00000(Eth1/1)
tx_enabled: FALSE
rx_enabled: FALSE
dcbx_enabled: TRUE
DCX Protocol: CIN
Port MAC address: 00:0d:ec:c9:c8:08
DCX Control FSM Variables: seq_no: 0x1, ack_no: 0x0,my_ack_no: 0x0, peer_seq_no:
    0x0 oper_version: 0x0, max_version: 0x0 fast_retries 0x0
Lock Status: UNLOCKED
PORT STATE: UP
LLDP Neighbors
No DCX tlvs from the remote peer
```

• If the peer does not support DCBX, configure the priority-flow-control mode setting to on to enable PFC.

Switch Interface connected to CNA receives constant pause frames (PFC)

Constant pause frames (PFC) are received when the switch interface is connected to a CNA.

Possible Cause

If the Nexus 5000 switch is connected to a CNA, then the CNA might be sending Xon PFC frames to the switch. This increments pause counters when using the **show interface ethx/x** command.

To verify this situation, perform the following:

- For a few iterations, check using the **show interface ethx/x** command and make sure the pause frame count is incrementing using the **show interface ethx/x** lgrep i pause command.
- For a few iterations, **check using the show interface ethx/x** command and ensure that the PFC frame count is incrementing **using the show interface ethx/x priority-flow-control** command.
- For a few iterations, use the **show queuing interface ethx/x** command to check the pause status. Example:

```
Per-priority-pause status
```

: Rx (Inactive), Tx (Inactive)

Book Title
If the Rx (Inactive) and pause counter increment over time (as shown with the **show interface ethx/x priority-flow-control** command), then this indicates that the issue is due to Xon frames received from the CNA.

Possible Cause

If the Nexus 5000 switch is connected to a CNA along with slow servers that are not able to handle the traffic from the switch port, then the server sends Xoff pause frames to the switch to slow it down. This increments the pause counters when using the **show interface ethx/x** command.

To verify this situation, perform the following:

- For a few iterations, check using the **show interface ethx/x lgrep i pause** command and ensure that the pause frame count is incrementing.
- For a few iterations, check using the **show interface ethx/x priority-flow-control** command and ensure that the PFC frame count is incrementing.
- For a few iterations, use the **show queuing interface ethx/x** command and check the pause status.

Example:

Per-priority-pause status

: Rx (Active), Tx (Inactive)

If the Rx (Active) and pause counter increment (as shown with the **show interface ethx/x priority-flow-control** command), this indicates that the issue is due to Xoff frames received from the server.

<u>Solution</u>

Xoff pause frames from the server pause the Nexus 5000 interface and reduces the throughput from the switch to the CNA. On the server, investigate the OS/PCI slot to ensure that they are high-speed servers. Replace the servers that can run 10gb throughput.

Check if switch is sending pause frames or getting paused

FCoE throughput on servers is very low due to pause frames from the switch. It is then necessary to check if the switch is sending pause frames or if it is getting paused.

Possible Cause

If the egress FC port is congested, the switch sends PFC frames to the servers. The PFC frames are sent to reduce its FCoE rate and avoid a drop. If the server is slow or congested, the server sends PFC frames to the switch interface.

To verify this situation, perform the following:

- For a few iterations, check using the **show interface ethx/x lgrep i pause** command and ensure that the pause frame count (Rx/TX) is incrementing.
- For a few iterations, check using the **show interface ethx/x priority-flow-control** command and ensure that the PFC frame count (RX/TX) is incrementing.
- For a few iterations, check using the **show queuing interface ethx/x** command to check the pause status.



Note

PFC frames are a MAC-level type of packet and cannot be viewed using the SPAN feature. Analyzer in-line is required to actually see the PFC frames on the wire.

Г

Example:

Per-priority-pause status

: Rx (Active), Tx (Inactive)

If the Rx (Active) and pause RX counter increment (as shown with the **show interface ethx/x priority-flow-control** command), then this indicates that this issue is due to Xoff frames received from the server.

If the Tx(Active) and pause TX counter increment (as shown with the **show interface ethx/x priority-flow-control** command), this indicates that this issue is due to Xoff frames transmitted by the switch.

Solution

Identify the source of the congestion and try to resolve it by increasing the FC bandwidth or change it to a more powerful server. If congestion is expected, then Pause is expected for FCoE traffic.

Switch ports err-disabled due to pause rate-limit

Switch ports go into error-disable state due to pause rate limit

Possible Cause

If the switch interface receives excessive Xoff pause frames from the server, ports become error-disabled due to the high rate of pause frames received. Usually the port goes into an err-disable state due to pause frames, only if the drain rate is less than 5Mbps on a 10Gb port. This means that the server is very slow and is sending a large number of pause frames to the switch ports.

To verify this situation, use the **show interface ethernet 1/14 brief** command.

Example:

switch# show int eth1/14 brief Ethernet VLAN Type Mode Status Reason Speed Port Interface Ch # Eth1/14 110 eth trunk down pauseRateLimitErrDisable 100(D) 110

- Check if the RX pause count is a large value. Use the **show interface ethx/x** command to display the pause counters.
- Check for pause error-disable logs using the **show hardware internal gatos event-history errors lgrep -i err** command.

<u>Solution</u>

Pause error-disable recovery can be enabled to get the ports out of this state, if the port is error-disabled due to transient condition as follows:

If the port is error-disabled due to transient condition listed below, then pause error-disable recovery can be enabled to move the ports out of this state.

- Error-disable recovery causes the pause rate limit.
- The error-disable recovery interval is 30.

If there is a consistent port error-disable condition due to the pause rate limit, determine if the issue is that the server is too slow. Replace the slow server.

How to enable link pause (flow control) on switch that connects DCBX capable devices

Link pause is not enabled on the switch ports that are connected to servers. It is necessary to enable link pause (flow control) on a Nexus 5000 switch that connects DCBX-capable devices.

Possible Cause

If the peer supports PFC TLV with DCBX, then configuring the flowcontrol send on and the flowcontrol receive on does not enable link pause. You have to disable PFC TLV sent by DCBX on the interface.

To verify this situation, perform one of the following:

- Check if the operating state is off using the show interface ethx/y flowcontrol command.
- Check if the operating state is on using the show interface ethx/y priority-flow-control command.

Solution

Use the following commands under the **interface ethx/y** command to enable link pause instead of PFC with DCBX capable devices:

- no priority-flow-control mode on
- flowcontrol receive on
- flowcontrol send on

How to clear PFC counters

How to clear priority flow counters.

Possible Cause

Currently there are no CLI commands to clear PFC frames (Bug ID is CSCtg08068).

<u>Solution</u>

Although there are no CLI commands to clear the PFC counters, a workaround does exist. You can the clear interface counters and then enter the **show interface ethx/x flowcontrol** command to see the PFC frame count.



Note

The PFC frame count is incremented using the **show int ethx/x flowcontrol** command. This is a known bug.

Registers and Counters

Interface level errors

To view any interface level errors, use the show interface counters errors command.

Example:

switch# show interface counters errors

Book Title

Г

Port	Align-Err	FCS-Err	Xmit-Err	Rev-Err	Undersize	OutDiscards
Eth1/1	0	0	0	0	0	0
Eth1/2	0	0	0	0	0	0
Eth1/3	0	0	0	0	0	0
Eth1/4	0	0	0	0	0	0
Eth1/5	0	0	0	0	0	0
Eth1/6	0	0	0	0	0	0
Eth1/7	0	0	0	0	0	0
Eth1/8	0	0	0	0	0	0
Eth1/9	0	0	0	0	0	0
Eth1/10	0	0	0	0	0	0
Eth1/11	0	0	0	0	0	0
Eth1/12	0	0	0	0	0	0
Eth1/13	0	0	0	0	0	0
Eth1/14	0	0	0	0	0	0
Eth1/15	0	0	0	0	0	0
Eth1/16	0	0	0	0	0	0
Eth1/17	0	0	0	0	0	0
Eth1/18	0	0	0	0	0	0
Eth1/19	0	0	0	0	0	0
Eth1/20	0	0	0	0	0	0
Eth2/1	0	0	0	0	0	0
Eth2/2	0	0	0	0	0	0
Eth2/3	0	0	0	0	0	0
Eth2/4	0	0	0	0	0	0
Po300	0	0	0	0	0	0
mgmt0						

Packet byte counts

To view packet byte counts, use the show interface counters detailed command.

Example:

show interface ethernet 1/11 counters detailed

Ethernet 1/11	
Rx Packets:	430908
Rx Unicast Packets:	129965
Rx Multicast Packets:	300932
Rx Broadcast Packets:	11
Rx Jumbo Packets:	3
Rx Bytes:	41893521
Rx Packets from 0 to 64 bytes:	47
Rx Packets from 65 to 127 bytes:	353478
Rx Packets from 128 to 255 bytes:	60265
Rx Packets from 256 to 511 bytes:	17095
Rx Packets from 512 to 1023 bytes:	16
Rx Packets from 1024 to 1518 bytes:	4
Rx Trunk Packets:	387901
Tx Packets:	172983
Tx Unicast Packets:	129959
Tx Multicast Packets:	43024
Tx Jumbo Packets:	3
Tx Bytes:	18220330
Tx Packets from 0 to 64 bytes:	7
Tx Packets from 65 to 127 bytes:	112452
Tx Packets from 128 to 255 bytes:	60461
Tx Packets from 256 to 511 bytes:	40
Tx Packets from 512 to 1023 bytes:	19

Тx	Packets	from	1024	to	1518	bytes:		1
Тx	Trunk Pa	ackets	5:				1	30019

Verification of SNMP readouts

To view the verification of SNMP readouts, use the **sh interface ethernet 1/11 counters snmp** command.

Example:

switch# show interface ethernet 1/11 counters snmp

Port	InOctets	InUcastPkts
Eth1/11	41908130	130009
Port	InMcastPkts	InBcastPkts
Eth1/11	301038	11
Port	OutOctets	OutUcastPkts
Eth1/11	18226503	130003
Port	OutMcastPkts	OutBcastPkts
Eth1/11	43039	0

Traffic rates

To view traffic rates, use the show interface ethernet 1/11 counters brief command.

Example:

switch# show interface ethernet 1/11 counters brief

Interface	Input Rate (avg)		Output Rate (avg)			
	Rate	Total	Rate	Total	Rate averaging	
	MB/s	Frames	MB/s	Frames	interval (seconds)	
Ethernet 1/11	0	0	0	0	30	
	0	0	0	0	300	

REVIEW DRAFT-CISCO CONFIDENTIAL

1-18



Troubleshooting Layer 2 Switching Issues

Layer 2 is the Data Link Layer of the Open Systems Interconnection model (OSI model) of computer networking.

This chapter describes how to identify and resolve problems that can occur with Layer 2 switching in the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- MAC Address Table
- Spanning Tree Protocol
- Multicast
- VLANs
- Registers and Counters

MAC Address Table

Data traffic is flooding

Data is not getting forwarded, instead it is being flooded to all ports of a VLAN.

Possible Cause

MAC address learning is disabled because of loop detection. A severity 2 syslog message, STM_LOOP_DETECT, should have been received.

Solution

After waiting for 180 seconds, learning is enabled automatically. A severity 2 syslog message, STM_LEARNING_RE_ENABLE, should be received.

Possible Cause

The MAC address table is full. A severity 2 syslog message, STM_LIMIT_REACHED, should have been received.

Solution

After waiting for 180 seconds, the MAC table is flushed and learning is enabled automatically. Alternatively, wait until some MAC entries are aged out, so that the total learned entries fall below 1500, or enter the **clear mac address-table dynamic [address** <*mac*>] command to clear out entries. This creates free space for new MAC entries to be learned. A severity 2 syslog message, STM_LEARNING_RE_ENABLE, should be received.

Possible Cause

MAC address learning is disabled due to a learning overload (that is, too many new addresses in a short time). A severity 4 syslog message, STM_LEARNING_OVERLOAD, should have been received.

<u>Solution</u>

After waiting for 120 seconds, learning is enabled automatically.

MAC address not learned

MAC address is not learned by the switch. This causes the MAC address not to be listed in the MAC table.

Possible Cause

MAC address learning is disabled because of loop detection. A severity 2 syslog message, STM_LOOP_DETECT, should have been received.

Solution

After waiting for 180 seconds, learning is enabled automatically. A severity 2 syslog message, STM_LEARNING_RE_ENABLE, should be received.

Possible Cause

The MAC address table is full. A severity 2 syslog message, STM_LIMIT_REACHED, should have been received.

Solution

After waiting for 180 seconds, the MAC table is flushed and learning is enabled automatically. Alternatively, wait until some MAC entries are aged out, so that the total learned entries fall below 1500, or enter the **clear mac address-table dynamic [address** <*mac*>] command to clear out entries. This creates free space for new MAC entries to be learned. A severity 2 syslog message, STM_LEARNING_RE_ENABLE, should be received.

Possible Cause

MAC address learning is disabled due to a learning overload (that is, too many new addresses in a short time). A severity 4 syslog message, STM_LEARNING_OVERLOAD, should have been received.

Solution

After waiting for 120 seconds, learning is enabled automatically.

Possible Cause

No egress path was set for the incoming data traffic. The MAC address from a data stream is not learned if there is no path for that data going out of the switch.

Solution

Configure an outgoing path for the data.

For example, the VLAN may not have been enabled on any of the interfaces other than the one on which data is coming in on. Alternatively, the outgoing interfaces may be down. If this is the case, you need to bring up those interfaces.

Traffic flooding in a VPC setup

Data is not getting forwarded, instead it is being flooded in the presence of a VPC scenario.

Possible Cause

The MAC address is learned on one switch only. Typically, this situation would be a bug regarding the synchronization of the MAC address with the VPC peer.

<u>Solution</u>

Clear the MAC address from the switch where it was learned. This triggers new learning and synchronization of the MAC addresses across the VPC switches.

Spanning Tree Protocol

HIFs go down with the BPDUGuard errDisable message

HIFs go down accompanied with the message, BPDUGuard errDisable.

Possible Cause

By default, the HIFs are in STP edge mode with the BPDU guard enabled. This means that the HIFs are supposed to be connected to hosts or non-switching devices. If they are connected to a non-host device/switch that is sending BPDUs, the HIFs become error-disabled upon receiving a BPDU.

Solution

Enable the BPDU filter on the HIF and on the peer connecting device. With the filter enabled, the HIFs do not send or receive any BPDUs. Use the following commands to confirm the details of the STP port state for the port:

- show spanning-tree interface <id> detail
- show spanning-tree interface <id>

FWM-2-STM_LOOP_DETECT detected on switch, dynamic learning disabled

When FWM-2-STM_LOOP_DETECT is detected on the switch, dynamic learning is disabled.

Possible Cause

- MAC addresses are moving because of incorrect STP-port state convergence.
- MAC addresses are moving because the source of the data being physically moved across all switches while STP states are converged and in correct states.

Use the following commands to verify the STP port state across VLANS on the switches:

- show spanning-tree
- show spanning-tree vlan <*id*>

Solution

• Check for a correct STP convergence and for STP port states across all switches in the topography. Also confirm that there are no disputes or incorrect port states.

Г

- If the source of the data frames, which are physically moving, is identified, then control the source to halt rapid and continuous moves.
- By default, dynamic learning is opened after 180 seconds. At that point, any STP disputes or inconsistencies should be resolved.

Port stuck in STP blocking state with BLK*(Type_Inc)

A port is stuck in STP blocking state with BLK*(Type_Inc).

Possible Cause

A type inconsistency might exist on an access port when it is connected to a trunk port on the other end. The port becomes BLK*(Type_Inc) to indicate that there is an incorrect configuration on the link. Use the following commands to confirm the details of the STP port state for the port:

- show spanning-tree interface <*id*> detail
- **show spanning-tree interface** *<id>*

<u>Solution</u>

Check the switch port modes configured at both ends (ports) of the link. Ensure that they are in the same mode. Both should be in access or trunk mode. Once the modes are synchronized, the port moves out of the inconsistency state.

Port stuck in STP blocking state with BLK*(PVID_Inc)

A port is stuck in STP blocking state with BLK*(PVID_Inc).

Possible Cause

A PVID inconsistency may exist when there is a native VLAN mismatch across a trunk link. When this occurs, the port state becomes BLK* (PVID_Inc). Use the following commands to confirm the details of the STP port state for the port:

- show spanning-tree interface <id> detail
- **show spanning-tree interface** *<id>*

<u>Solution</u>

Check the native VLAN configured at both ends (ports) of the link. Ensure that they have the same native VLAN. Once the native VLANs are synchronized, the port moves out of the inconsistency state.

Port stuck in STP blocking state with BLK*(Loop_Inc)

A port is stuck in STP blocking state with BLK*(Loop_Inc).

Possible Cause

This situation occurs when the loop guard is configured on the port and the port stops receiving BPDUs. This is supposed to prevent loops when unidirectional link failures occur. However, the port is put into a BLK* (Loop_Inc) state. Use the following commands to confirm the details of the STP port state for the port:

- show spanning-tree interface <*id*> detail
- **show spanning-tree interface** *<id>*

<u>Solution</u>

Check the native VLAN configured at both ends (ports) of the link. Ensure that they have the same native VLAN. Once the native VLANs are synchronized, the port moves out of the inconsistency state.

Multicast

Source MAC addresses of IGMP joins are learned

In this situation, the source MAC addresses of IGMP joins are learned. However, source MAC addresses of IGMP joins are usually not learned by the switch in order to conserve MAC address space.

Possible Cause

Receiving joins and performing an ISSU simultaneously might cause the situation.

<u>Solution</u>

The MAC addresses age out (expire) if the join stops. Alternatively, you can clear the MAC addresses specifically by using the **clear mac address-table dynamic mac** <*mac*> command.

Multicast data traffic not received by host

Host does not receive multicast data traffic.

Possible Cause

The join is not registered.

Solution

- Ensure that the host application is sending the joins.
- Check if the switch port is configured for the VLAN on which the joins are being sent using the **show vlan id** *<vlan>* command.
- Check if the relevant VLAN is active by using the show vlan id <vlan> command.
- Check if the switch port is in STP forwarding state by using the **show spanning-tree vlan** <*vlan*> command.

Multicast data traffic not received when host is registered for group

Multicast data traffic is not received when the host is registered for the group.

Possible Cause

A bug may exist in exist in the communication between the IGMP and FWM processes.

Review the output from the following commands:

- show ip igmp snooping groups vlan 1001
- show mac address-table multicast vlan 1001 igmp-snooping
- show platform fwm info vlan 1001 all_macgs verbose

Solution

Г

Perform a shut/no-shut operation on the host interface and send the join again.

Multicast traffic is being flooded in a VPC setup

In a VPC setup multicast traffic is being flooded.

Possible Cause

IGMP snooping is disabled on one of the switches.

<u>Solution</u>

Enable IGMP snooping on both switches.



Groups for link local IP addresses (that is, 224.0.0.X) are not created.

VLANs

Nexus 5000 does not have the same VLANs as switch running VTP server

VLANs for the Nexus 5000 are not the same as for the switch running the VTP server.

Possible Cause

The Nexus 5000 currently supports VTP only in transparent mode (4.2(1)N1(1) and later releases).

<u>Solution</u>

This situation indicates that VLANs must be configured locally. However a VTP client and server can both communicate through a Nexus 5000 by using the following commands:

```
switch(config)# feature vtp
switch(config)# vtp mode transparent
switch(config)# exit
switch# show vtp status
```

VLAN cannot be created

A VLAN cannot be created.

Possible Cause

An internal VLAN range is used.

<u>Solution</u>

Use a VLAN number that is not being reserved for internal use.



The VLAN range of 3968 to 4047 is reserved for internal use.

Example:

switch(config)# vlan ?

<1-3967,4048-4093> VLAN ID 1-4094 or range(s): 1-5, 10 or 2-5,7-19

Interface VLAN is down

The interface VLAN is down.

Possible Cause

The VLAN was not created.

Solution

Although VLAN <###> is not yet created, the NX-OS allows the configuration of the **interface vlan** <###>. As a result, the **interface vlan** <###> does not come up. Use the **show vlan** command to determine if VLAN <###> exists. If it does not exist, use the **vlan** <###> command to create the VLAN. After the VLAN is created, you must bounce the interface VLAN to have it come up.

Example:

switch(config)# interface vlan 600 switch(config-if)# no shutdown switch(config-if)# show interface vlan 600 brief					
Interface Secondary VLAN(Type)		Status	Reason		
Vlan600 switch(config)# show vlan id 600 VLAN 600 not found in current VLAN dat switch(config-if)# vlan 600 switch(config)# show vlan id 600	tabase	down	other		
VLAN Name	Status	Ports			
600 VLAN0600 switch(config-if)# interface vlan 600 switch(config-if)# shut switch(config-if)# no shutdown switch(config-if)# show interface vlan	active n 600 brief	Po1, Po1:	L, Po30, Po31		
Interface Secondary VLAN(Type)		Status	Reason		
Vlan600		up			

Possible Cause

VLAN was suspended by the vPC configuration on the Nexus 5000 pair.

Solution

Show that the vPC consistency parameters are global and make sure that the VLAN was not suspended. Otherwise, fix the configuration mismatch on the Nexus 5000 pair:

Example:

switch# sh vpc consistency-parameters global

Legend: Type 1 : vPC will be suspended in case of mismatch Name Type Local Value Peer Value

QoS	1	([], [3], [], [], [], [])	([], [3], [], [], [] [])
Network QoS (MTU)	1	(1538, 2240, 0, 0, 0, 0)	(1538, 2240, 0, 0, 0 0)
Network Qos (Pause)	1	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Priority)			
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	None	None
STP MST Region Name	1		
STP MST Region Revision	1	0	0
STP MST Region Instance to	1		
VLAN Mapping			
STP Loopguard	1	Disabled	Disabled
STP Bridge Assurance	1	Enabled	Enabled
STP Port Type, Edge	1	Normal, Disabled,	Normal, Disabled,
BPDUFilter, Edge BPDUGuard		Disabled	Disabled
STP MST Simulate PVST	1	Enabled	Enabled
Allowed VLANs	-	1-2	1-2
Local suspended VLANs		2 -	
switch#			

Configuring interface to access port does not allow VLAN <###> to go through

After configuring an interface to access a port for allowing VLAN <###>, the VLAN <###> does not go through.

Possible Cause

The VLAN was not created.

Solution

In NX-OS, configuring with the **switchport access vlan** <**###**> command on an interface does not automatically create VLAN <**###**>. You must specifically create VLAN <**###**> using the **vlan** <**###**> command. Use the **show vlan** command to determine if VLAN <**###**> exists. If it does not exist, then use the **vlan** <**###**> command to create the VLAN.

Cannot create VLAN

The VLAN cannot be created.

Possible Cause

All VLAN resources are exhausted.

<u>Solution</u>

For the Nexus 5000, the maximum number of active VLANs and VSANs per switch is 512 (31 reserved for VSAN; remainder reserved for VLAN). Use the **show resource vlan** command to determine the number of available VLANs.

Example:

switch(config)# show resource vlan

Resource	Min	Max	Used	Unused	Avail
vlan	16	512	25	0	487

Cannot create SVI

The SVI cannot be created.

Possible Cause

The interface-vlan feature is not enabled.

<u>Solution</u>

The interface-vlan feature must be enabled before configuring the SVI. Use the **show feature** command to determine which features are enabled.

Example:

switch(config)#	featu	ıre	interfa	ce-vlan
<pre>switch(config)#</pre>	show	fea	ature	
Feature Name		Ir	nstance	State
tacacs		1		disabled
lacp		1		enabled
interface-vlan		1		enabled
private-vlan		1		enabled
udld		1		enabled
vpc		1		enabled
fcoe		1		disabled
fex		1		enabled

Cannot create private VLAN (PVLAN)

The private VLAN (PVLAN) cannot be created.

Possible Cause

The private-vlan feature is not enabled.

Solution

The private-vlan feature must be enabled prior to PVLAN configuration, which makes the PVLAN command available. Use the **show feature** command to determine which features are enabled.

Example:

switch(config)#	feature privat	e-vlan
switch(config)#	show feature	
Feature Name	Instance	State
tacacs	1	disabled
lacp	1	enabled
interface-vlan	1	enabled
private-vlan	1	enabled
udld	1	enabled
vpc	1	enabled
fcoe	1	disabled
fex	1	enabled

Registers and Counters

Identifying drops

There are logical and physical causes for the Nexus 5000 to drop a frame. There are also situations when a frame cannot be dropped because of the cut-through nature of the switch architecture. If a drop is necessary, but the frame is being switched in a cut-through path, then the only option is to stomp the Ethernet frame check sequence (FCS). Stomping a frame involves setting the FCS to a known value that does not pass a CRC check. This causes subsequent CRC checks to fail later in the path for this frame. A downstream store-and-forward device, or a host, will be able to drop this frame.



When a frame is received on a 10 Gb/s interface, it is considered to be in the cut-through path.

The following example output shows all discards and drops seen on a given interface, except for queuing drops. The queuing drops may be expected or resulting from errors. (Drops are more common than discards.)

Example:

```
switch# show platform fwm info pif ethernet 1/1 ...
Eth1/1 pd: tx stats: bytes 19765995 frames 213263 discard 0 drop 0
Eth1/1 pd: rx stats: bytes 388957 frames 4232 discard 0 drop 126
```

For some commands, you need to know on which chip your port resides.

In the following example, the chip is called Gatos. The example shows which Gatos and which Gatos port is associated with ethernet 1/1.

```
switch# show hardware internal gatos port ethernet 1/1 | include
instance|mac
    gatos instance : 7 <- Gatos 7
    mac port : 2 <- Port 2</pre>
```

: 2

Expected/Logical drops

During normal operation, the Nexus 5000 encounters frames that cannot be forwarded based on logical conclusions.

For example, if you learn a MAC address on a given interface and receive traffic on that interface with a destination MAC address on the source interface, then you cannot forward the frame. Because it is a known address and cannot be flooded, you can never send traffic out from where it came. This is a requirement to avoid looping Layer 2 topologies.

The error counter, shown in the following example, increments when the ingress port is the only port in the VLAN.

Example (same Gatos instance as in earlier example):

fw_instance

```
switch# show platform fwm info gatos-errors 7
Printing non zero Gatos error registers:
DROP_SRC_MASK_TO_NULL 9
```

REVIEW DRAFT-CISCO CONFIDENTIAL



The show platform fwm info gatos-errors command increments 3 times for a given drop.

Other expected drops

Drop	Description
VLAN_MASK_TO_NULL	Traffic destined for the CPU interfaces; not actually a VLAN.
DROP_NO_FABRIC_SELECTED	Increments with VLAN_MASK_TO_NULL.
DROP_INGRESS_ACL	Increments for an access list matching the frame.
	If an ACL is not applied, this will increment if a large amount of CPU-bound traffic is being received, and the rate limit in the hardware is enabled to protect NX-OS from Denial of Service.

Queue is full

When a queue is full, you need to increment discards in the respective queue on the ingress interface.

Example:

```
switch# show queuing interface ethernet 1/1
Ethernet1/1 queuing information:
  TX Queuing
              sched-type oper-bandwidth
    qos-group
                WRR
                                50
        0
        1
                WRR
                                50
  RX Queuing
    qos-group 0
    q-size: 243200, HW MTU: 1600 (1500 configured)
    drop-type: drop, xon: 0, xoff: 1520
    Statistics:
        Pkts received over the port
                                               : 0
                                               : 0
        Ucast pkts sent to the cross-bar
        Mcast pkts sent to the cross-bar
                                               : 0
        Ucast pkts received from the cross-bar : 0
                                              : 0
        Pkts sent to the port
    <b>
          Pkts discarded on ingress
                                                  : 0 </b>
                                              : Rx (Inactive), Tx
        Per-priority-pause status
(Inactive)
    qos-group 1
    q-size: 76800, HW MTU: 2240 (2158 configured)
    drop-type: no-drop, xon: 128, xoff: 240
    Statistics:
        Pkts received over the port
                                               : 0
        Ucast pkts sent to the cross-bar
                                               : 0
        Mcast pkts sent to the cross-bar
                                               : 0
        Ucast pkts received from the cross-bar : 0
        Pkts sent to the port
                                               : 0
        Pkts discarded on ingress
                                               : 0 </b>
     <b>
        Per-priority-pause status
                                              : Rx (Inactive), Tx
(Inactive)
```

Γ

Total Multicast crossbar statistics: Mcast pkts received from the cross-bar : 0

MTU violation

The Nexus 5000 is a cut-through switch at 10 Gb/s. This means that an MTU can be checked, but the frame will already be transmitting before the length is known. Therefore, the frame cannot be dropped. The frame is truncated after the MTU is reached and the CRC value is stomped. The ingress interface increments an Rx Jumbo and the egress interface will increment a Tx CRC and a Tx Jumbo.

- If jumbo frames are seen with the **show interface** or the **show hardware internal gatos port ethernet 1/1 counters rx** commands, this is not an indication that the frames are being dropped. A jumbo frame is just an Ethernet frame, greater than 1500 bytes, that was received or transmitted.
- The show queuing interface <i>ex/y</i> command shows the current configured MTU (per class).
- A drop due to an MTU violation can be seen with the **show hardware internal gatos counters interrupt match mtu*** command.
- A counter that matches the Gatos number and fw_instance from the **show hardware internal gatos port ethernet 1/1** | **include instancelmac** command is the indicator that an MTU violation has taken place and that the frame has been stomped.

Handling CRC errors

When a CRC error is seen in the FCS on a cut-through port, the Rx CRC counter of the **show interface** command is incremented. However, the frame cannot be dropped because the FCS is at the end of the Ethernet frame on the wire.

The egress interface increments a Tx CRC error and it propagates through to the next device in the path.

You can use the **show hardware internal gatos counters interrupt match stomp** command to determine if the Nexus 5000 is propagating CRCs or generating them.

- If stomp values exist, they should have matching CRC values on that interface.
- If Rx CRC values exist, then you know it entered the switchport with the error already. You can move on to the connected device to trace it back.

MAC Statistics

During normal operation, a Nexus 5000 encounters frames that cannot be forwarded.

Frames are characterized as good frames or bad frames.

- A good frame is a frame that does not have a CRC error or other kind of error
- A bad frame is a frame that has a CRC error or other kind of error

All counters include MAC Control frames where applicable.

MAC TX Statistics

Counter	Description			
TX_PKT_LT64	Number of frames (good and bad frames) with transmit packet size less than 64 bytes.			
TX_PKT_64	Number of frames (good and bad frames) with transmit packet size equal to 64 bytes.			
TX_PKT_65	Number of frames (good and bad frames) with transmit packet size between 65 and 127 bytes.			
TX_PKT_128	Number of frames (good and bad frames) with transmit packet size between 128 and 255 bytes.			
TX_PKT_256	Number of frames (good and bad frames) with transmit packet size between 256 and 511 bytes.			
TX_PKT_512	Number of frames (good and bad frames) with transmit packet size between 512 and 1023 bytes.			
TX_PKT_1024	Number of frames (good and bad frames) with transmit packet size between 1024 and 1518 bytes.			
TX_PKT_1519	Number of frames (good and bad frames) with transmit packet size between 1519 and 2047 bytes.			
TX_PKT_2048	Number of frames (good and bad frames) with transmit packet size between 2048 and 4095 bytes.			
TX_PKT_4096	Number of frames (good and bad frames) with transmit packet size between 4096 and 8191 bytes.			
TX_PKT_8192	Number of frames (good and bad frames) with transmit packet size between 8192 and 9216 bytes.			
TX_PKT_GT9216	Number of frames (good and bad frames) with transmit packet size greater than 9216 bytes.			
TX_PKTTOTAL	Total number of frames (good and bad frames) transmitted.			
	This number is the sum of all transmitted packets regardless of frame length.			
TX_OCTETS	Total byte count of packet octets (good and bad frames) transmitted.			
ТХ_РКТОК	Number of good frames transmitted.			
TX_UCAST	Number of good unicast frames transmitted.			
TX_MCAST	Number of good multicast frames transmitted.			
TX_BCAST	Number of good broadcast frames transmitted.			
TX_VLAN	Number of good 802.1Q tagged VLAN frames transmitted.			
TX_PAUSE	Number of 802.3x PAUSE frames transmitted.			
TX_USER_PAUSE	Number of transmitted priority flow control frames.			
TX_FRM_ERROR	Number of frames with cl_im_tx_err set to 1 at EOP.			
TX_OCTETSOK	Total byte count of good frames.			

I

MAC RX Statistics

MAC RX Statistic	Description
RX_PKT_LT64	Number of frames (good and bad frames) with receive packet size less than 64 bytes.
RX_PKT_64	Number of frames (good and bad frames) with receive packet size equal to 64 bytes.
RX_PKT_65	Number of frames (good and bad frames) with receive packet size between 65 and 127 bytes.
RX_PKT_128	Number of frames (good and bad frames) with receive packet size between 128 and 255 bytes.
RX_PKT_256	Number of frames (good and bad frames) with receive packet size between 256 and 511 bytes.
RX_PKT_512	Number of frames (good and bad frames) with receive packet size between 512 and 1023 bytes.
RX_PKT_1024	Number of frames (good and bad frames) with receive packet size between 1024 and 1518 bytes.
RX_PKT_1519	Number of frames (good and bad frames) with receive packet size between 1519 and 2047 bytes.
RX_PKT_2048	Number of frames (good and bad frames) with receive packet size between 2048 and 4095 bytes.
RX_PKT_4096	Number of frames (good and bad frames) with receive packet size between 4096 and 8191 bytes.
RX_PKT_8192	Number of frames (good and bad frames) with receive packet size between 8192 and 9216 bytes.
RX_PKT_GT9216	Number of frames (good and bad frames) with receive packet size greater than 9216 bytes.
RX_PKTTOTAL	Total number of frames (good and bad frames) received.
	This number is the sum of all received packets regardless of frame length.
RX_OCTETS	Total byte count of packet octets (good and bad frames) received.
RX_PKTOK	Number of good frames received.
RX_UCAST	Number of good unicast frames received.
RX_MCAST	Number of good multicast frames received.
RX_BCAST	Number of good broadcast frames received.
RX_VLAN	Number of good 802.1Q tagged VLAN frames received.
RX_OVERSIZE	Number of good frames received that are greater than CFG_xg_rx_stats_max_frame_len.
RX_TOOLONG	Number of frames (good and bad frames) received that are greater than CFG_xg_rx_stats_max_frame_len.
RX_DISCARD	N/A
	The value of this counter is always 0.

MAC RX Statistic	Description	
RX_UNDERSIZE	Number of good frames received that are less than CFG_xg_rx_stats_min_frame_len.	
RX_FRAGMENT	Number of bad frames received that are less than CFG_xg_rx_stats_min_frame_len.	
RX_CRCERR_NOT_STOMPED	Number of bad frames received that are not stomped.	
RX_CRCERR_STOMPED	Number of bad frames received that are stomped.	
RX_INRANGEERR	Number of frames with a length field check error, but no CRC error.	
	Note Frame length is based on the contents of the ethertype/length field. When the ethertype/length field is less than 0x600, the contents of the field is treated as the length of the frame. This length is compared with the actual frame length. An error is raised when the lengths do not match.	
RX_JABBER	Number of bad frames received with frame length greater than CFG_xg_rx_stats_max_frame_len.	
RX_PAUSE	Number of good 802.3x MAC control frames received.	
RX_USER_PAUSE	Number of good priority flow control frames received.	
RX_OCTETSOK	Total byte count of good frames.	

REVIEW DRAFT-CISCO CONFIDENTIAL



Troubleshooting QoS Issues

The Cisco Nexus 5000 Series NX-OS quality of service (QoS) provides the most desirable flow of traffic through a network. QoS uses policies and flow control to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance.

This chapter describes how to identify and resolve problems that can occur with QoS in the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Policy Maps
- Improper Configurations
- PFC
- Registers and Counters

Policy Maps

The Nexus 5000 QoS implementation follows the Cisco Modular QoS CLI model. It takes three steps to configure the QoS:

- Define the class map.
- Create a policy map to define the action taken for each class map.
- Apply the policy-map.

The Nexus 5000 implements three different types of policy maps:

- Policy-map type qos
- Policy-map type queuing
- Policy-map type network-qos

Additionally, the Nexus 5000 introduces a new configuration context for QoS called the System QoS. The policy-map applied under the System QoS context is applied to the entire switch.

The following table summarizes the function and attach points for these three types of policy maps.

Γ

Policy Type	Function	Attach Point
QoS	Define traffic classification	System QoS
		• Ingress interface
Queuing	Strict Priority queue	System QoS
	• Deficit Weight Round Robin	• Egress interface
		• Ingress interface
Network-QoS	• Define flow control mechanism (PAUSE or tail drop)	System QoS
	• MTU per class of service	
	• Queue size	
	• Marking	

Table 1-1 Types	of	Policy	Maps
-----------------	----	--------	------

With the basic process, the incoming packets are compared to the QoS classification rules that are defined by policy-map type qos. The packets are classified into 1 of 8 qos-groups.

Next, the Network-QoS and Queuing policies are applied to the packets. The Queuing policy and the Network-QoS policy define actual QoS parameters for packets belonging to each qos-group.



- The Queuing and Network-QoS policies match the qos-group (identified by policy-map type qos) instead of the actual packet headers.
- When the same type of service policy is applied under the System QoS context and the interface level, the interface level service policy is preferred.
- The queuing policy that is applied under the ingress interface is not applied locally. The queueing policy is the bandwidth allocation for a different class of service that is exchanged with its peer using the DCBX protocol.

Improper Configurations

Cannot pass frame size larger than 2300 bytes through switch

Although the jumbo MTU has been configured for class-default, you cannot pass a frame size larger than 2300 bytes through the Nexus 5000 switch and the Nexus 2000 FEX.

Possible Cause

The CoS value may conflict with the existing MTU value.

Solution

CoS 7 is used internally for controlling traffic between the Nexus 5000 switch and the Nexus 2000 FEX. The MTU value for the traffic with CoS 7 is set to a fixed value. You must check that the incoming traffic is marked with CoS 7. Use any CoS value other than 7 to avoid this limitation.

MTU for "class-default" value is 1500 when jumbo MTU configured

When the configuration for the network-qos policy-map sets the class-default to jumbo MTU, the **show queuing interface** command indicates that the MTU for class-default is 1500.

Possible Cause

An incorrect startup configuration may exist after an upgrade.

Solution

If the switch has been upgraded to the 4.2(1)N1(1) release, make sure that you have used the **write erase** command to delete the startup configuration. You can save the configuration first to another file name.

After the Nexus 5000 switch boots up with an empty configuration, reapply the original configuration. You might lose your connectivity to the Nexus 5000 if you are using Telnet or SSH. It is recommended that you use the console for this procedure.

Traffic not queued or prioritized correctly on Nexus 2148, Nexus 2232, and Nexus 2248

After configuring all three types of policy maps (QoS, Network-QoS, and Queuing), the traffic is not queued or prioritized correctly on Nexus 2148, Nexus 2232, and Nexus 2248 switches.

Possible Cause

The Nexus 2148, Nexus 2232, and Nexus 2248 FEX can only support CoS-based traffic classification. The QoS service policy type configured under System QoS is populated from the Nexus 5000 to FEX only when all the matching criteria are match cos. If other match clauses exist, such as match dscp or match ip access-group in the QoS policy map, then the FEX does not accept the service policy. As a result, all the traffic is placed into the default queue.



Use the **show queuing interface** command to ensure that the queues have been created properly.

Solution

For the ingress traffic (from server to network) that is not marked with a CoS value, the traffic is placed into the default queue on FEX. Once the traffic is received on the Nexus 5000, it is classified based on a configured rule and are placed in the proper queue.

For the egress traffic (from Nexus 5000 to FEX, and then FEX to server), it is recommended that you mark mark the traffic with a CoS value on the Nexus 5000 so that the FEX can classify and queue the traffic properly.

The following example is a complete Nexus 5000 and Nexus 2232/Nexus 2248 configuration that classifies the traffic and configures the proper bandwidth for each type of traffic. This example applies only to the Nexus 5000 and Nexus 2248. The configuration for the Nexus 2148 is slightly different due to the fact that Nexus 2148 has only two queues for user data. The Nexus 2232/Nexus 2248 has six hardware queues for user data, which is the same as Nexus 5000.

Example:

```
//class-map for global qos policy-map, which will be used to create CoS-queue mapping.//
class-map type qos voice-global
match cos 5
class-map type qos critical-global
match cos 6
class-map type qos scavenger-global
```

L

match cos 1

REVIEW DRAFT-CISCO CONFIDENTIAL

```
class-map type qos video-signal-global
match cos 4
//This gos policy-map will be attached under "system gos". It will be downloaded to 2248
to create CoS to queue mapping.//
policy-map type qos classify-5020-global
class voice-global
set gos-group 5
class video-signal-global
set qos-group 4
class critical-global
set qos-group 3
class scavenger-global
set qos-group 2
class-map type qos Video
match dscp 34
class-map type qos Voice
match dscp 40,46
class-map type qos Control
match dscp 48,56
class-map type qos BulkData
match dscp 10
class-map type qos Scavenger
match dscp 8
class-map type qos Signalling
match dscp 24,26
class-map type gos CriticalData
match dscp 18
//This gos policy-map will be applied under all N5k and 2248 interfaces to classify all
incoming traffic based on DSCP marking. Please note that even the policy-map will be
applied under Nexus 2248 interfaces the traffic will be classified on N5k//
policy-map type qos Classify-5020
class Voice
set qos-group 5
class CriticalData
set qos-group 3
class Control
set qos-group 3
class Video
set gos-group 4
class Signalling
set qos-group 4
class Scavenger
set qos-group 2
class-map type network-qos Voice
match qos-group 5
class-map type network-gos Critical
match qos-group 3
class-map type network-qos Scavenger
match qos-group 2
class-map type network-qos Video-Signalling
match qos-group 4
//This policy-map type network-gos will be applied under "system gos" to define the MTU,
marking and queue-limit(not configured here).//
policy-map type network-qos NetworkQoS-5020
class type network-qos Voice
set cos 5
class type network-gos Video-Signalling
set cos 4
mtu 9216
class type network-gos Scavenger
```

set cos 1
mtu 9216
class type network-qos Critical
set cos 6
mtu 9216
class type network-qos class-default
mtu 9216
class-map type queuing Voice
match qos-group 5
class-map type queuing Critical
match qos-group 3
class-map type queuing Scavenger
match qos-group 2
class-map type queuing Video-Signalling
match qos-group 4

//The queuing interface will be applied under "system qos" to define the priority queue and how bandwidth is shared among non-priority queues.// policy-map type queuing Queue-5020 class type queuing Scavenger bandwidth percent 1 class type queuing Voice priority class type queuing Critical bandwidth percent 6 class type queuing Video-Signalling bandwidth percent 20 class type queuing class-fcoe bandwidth percent 0 class type queuing class-default bandwidth percent 73

//The input queuing policy determines how bandwidth are shared for FEX uplink in the direction from FEX to N5k. The output queueing policy determines the bandwidth allocation for both N5k interfaces and FEX host interfaces.// system qos service-policy type qos input classify-5020-global service-policy type network-qos NetworkQoS-5020 service-policy type queuing input Queue-5020 service-policy type queuing output Queue-5020

//Apply service-policy type qos under physical interface in order to classify traffic based on DSCP. Please note that for portchannel member the service-policy needs to be configured under interface port-channel.// interface eth1/1-40 service-policy type qos input Classify-5020 interface eth100/1/1-48 service-policy type qos input Classify-5020

The **show queuing interface** command can be used to ensure that the CoS-to-queue mapping is properly configured under the FEX interfaces. It can also be used to check the bandwidth and MTU configuration.

This same command can be used to check the QoS configuration for the Nexus 5000 interfaces.

The following is the output from the **show queuing interface** command for the Nexus 2248 interfaces when the above configrations are applied:

```
switch# sh queuing interface ethernet 100/1/1
Ethernet100/1/1 queuing information:
   Input buffer allocation:
   Qos-group: 0 2 3 4 5 (shared)
   frh: 2
   drop-type: drop
   cos: 0 1 2 3 4 5 6
```

xon	xoff	buffer-size				
21760 Queuein	26880 g:	48640				
queue	qos-group c	os	priority	bandw	idth mtu	
2	++ 0	023	+ W	+ RR	73	+ 9280
4	2	1	W	RR	1	9280
5	3	6	W	RR	6	9280
6	4	4	W	RR	20	9280
7	5	5	P	RI	0	1600
Queue l	imit: 64000	bytes				
Queue	Statistics:					
queue	rx	tx				
	-+	+				
2	1138225390	41 1				
4	0	0				
5	0	0				
6	417659797	0				
7	0	0				
Port St	atistics:					
rx dr	op r	x mcast drop	rx erro	r	tx drop	
0	0		0		0	
Prior Flow-	ity-flow-con control stat	trol enabled us:	: no			
COS	qos-group	rx pause	tx pause	masked	rx pause	
0	0	xon	xon	xon		
1	2	xon	xon	xon		
2	0	xon	xon	xon		
3	0	xon	xon	xon		
4	4	xon	xon	xon		
5	5	xon	xon	xon		
6	3	xon	xon	xon		
7	n/a	xon	xon	xon		
switch#						

The Nexus 2148 has two queues in both the ingress and egress directions. One queue is mapped to the no-drop system class and another queue is mapped to the drop system class. For the ingress direction, the two queues are scheduled using WRR (Weight Round Robin). For the egress direction, the queue for the no-drop system class is the priority queue.

In order to separate traffic for the two queues, the user has to create a no-drop system class. All no-drop system classes created on the Nexus 5000 are mapped to the no-drop queue on the Nexus 2148.

The **pause no-drop** command is added to the Network-QoS in order for the Nexus 2148 to place voice in the priority queue at the FEX egress direction.

Example:

```
policy-map type network-qos NetworkQoS-5020
class type network-qos Voice
  set cos 5
  pause no-drop
class type network-qos Video-Signalling
  set cos 4
  mtu 9216
class type network-qos Scavenger
  set cos 1
  mtu 9216
class type network-qos Critical
```

REVIEW DRAFT-CISCO CONFIDENTIAL

```
set cos 6
mtu 9216
class type network-qos class-default
mtu 9216
```

The configuration classifies the incoming voice traffic based on DSCP and marks the voice traffic to CoS 5. At the Nexus 2148 egress direction, the FEX assigns voice traffic to the priority queue.

The following is example output from the **show queuing interface** command for the Nexus 2148 with the above configuration.

Example:

```
switch# sh queuing interface ethernet 199/1/1
Ethernet199/1/1 queuing information:
 Input buffer allocation:
 Qos-group: 0 2 3 4 (shared)
 frh: 3
 drop-type: drop
 cos: 0 1 2 3 4 6 7
 xon xoff buffer-size
 16640 33280 56320
 Qos-group: 5
 frh: 2
 drop-type: no-drop
 cos: 5
 xon
       xoff
              buffer-size
 8960
       19200
              34560
 Queueing:
 queue qos-group cos
                         priority bandwidth mtu
   0 2 3 4 0 1 2 3 4 6 WRR 100 9280
3
      5
                          PRI
                                0 1600
               5
 2
 Buffer threshold: 271360 bytes
 Queue limit: Disabled
 Queue Statistics:
 queue rx
 3 241439087
 2
     0
 Port Statistics:
 tx queue drop
 _____
 0
 Priority-flow-control enabled: no
 Flow-control status:
 cos qos-group rx pause tx pause masked rx pause
 0
          0
              xon
                     xon
                            xon
 1
           2
              xon
                     xon
                            xon
 2
           0
              xon
                     xon
                            xon
 3
             xon
           0
                     xon
                            xon
 4
           4 xon
                     xon
                            xon
 5
          5 xon
                     xon
                            xon
          3 xon
 6
                     xon
                            xon
 7
          n/a xon
                     xon
                            xon
```

L

switch#

PFC

Link pause (flow control) not enabled on back to back Nexus 5000 switch links

When link pause (flow control) is not enabled on back-to-back Nexus 5000 switch links, packets are dropped while sending traffic on a no-drop class.

Possible Cause

If the peer Nexus 5000 switch supports PFC TLV with DCBX, then configuring **flowcontrol send on** and **flowcontrol receive on** will not enable the link pause. You have to disable the PFC TLV sent by DCBX on that interface.

Use one of the following commands to verify:

- Use the show interface ethx/y flowcontrol command and check to see if the operating state is off.
- Use the **show interface ethx/y priority-flow-control** command and check to see if the operating state is on.

Solution

Configure the following commands under **interface ethx/y** to enable link pause instead of PFC on back-to-back switch links.

- no priority-flow-control mode on
- flowcontrol receive on
- flowcontrol send on

Cannot enable "pause no-drop" on more than one ethernet class

Cannot enable pause no-drop on more than one Ethernet class.

CLI commands fail with the following error when trying to enable pause no-drop.

```
ERROR: Module 1 returned status "Not enough buffer space available. Please change your configuration and re-apply"
```

Possible Cause

Nexus 5000 supports a maximum of three no drop classes (including FCoE). If five Ethernet classes are created, then there will be insufficient buffers to enable twoof the five Ethernet no-drop classes.

You will get an error if not enough buffers exist to enable the no-drop.

Example:

```
class type network-qos s4
pause no-drop
ERROR: Module 1 returned status "Not enough buffer space available. Please change your
configuration and re-apply"
```

<u>Solution</u>

1-8

Book Title

If you create five **ethernet** classes, then there will be an insufficient number of buffers to configure two of the five Ethernet no-drop classes. If you delete two Ethernet classes and configure the remaining three Ethernet classes (including class-default), then no-drop can be enabled on two of the Ethernet classes.

Changing no-drop configuration causes VPC peer-link to go down and FEX to go offline

Changing the QoS no-drop configuration causes the VPC MCT peer-link to go down and FEX to go offline.

Possible Cause

The network QoS policy parameters, such as MTU and pause, are treated as type1 parameters and should match between the VPC primary and secondary nodes. If a mismatch exists between the VPC primary and secondary nodes, then the VPC peer-link does not come up and FEX goes offline. Only CoS based class no-drop/MTU parameters are considered as type 1 consistency checked for VPC. If you configure an ACL based class, then it is not treated as a vtype 1 parameter for VPC.

Use one of the following commands to verify:

- show vpc brief
- show vpc consistency-parameters global

<u>Solution</u>

Configure the similar no-drop class configuration between the VPC primary and secondary nodes. Any mismatch of no-drop policy on nqos CoS-based class parameters causes a type1 inconsistency.

Pause enabled on all cos values when no-drop enabled on class-ip-multicast

Priority flow control enables pause on all CoS values when no-drop is enabled on the class-ip-multicast class.

Possible Cause

When you create a class-ip-multicast class and no-drop is enabled, then pause is enabled on all of the CoS values.

Use the **show interface ethx/y priority-flow-control** command and check that the VL bitmap is enabled for all CoS values (ff).

Solution

Use the following commands to enable PFC on CoS 4 only, instead of on all CoS values under the class-ip-multicast class.

- Policy-map type network-qos system
- Class type network-qos class-ip-multicast
- Pause no-drop pfc-cos 4

Г

No drop class not created on N2K-C2148T/N2K-C2248TP-1GE based FEX with default QoS configuration

The no-drop class is not created on the N2K-C2148T/N2K-C2248TP-1GE based FEX with the default QoS configuration.

The show queuing interface is different for the switchport and HIF port on N2K-C2248TP and N2K-C2148T.

Possible Cause

FCoE is not supported on the N2K-C2148T and N2K-C2248TP-1GE based FEX and the no drop class is not created with the default QoS configuration.

Use the following command to verify (check for no-drop class):

show queuing interface eth100/1/1

<u>Solution</u>

If you want an ethernet no-drop class on a N2K-C2148T/N2K-C2248TP-1GE FEX, then you have to create an ethernet no-drop class with the following:

- Policy-map type network-qos no-drop
- Class type network-qos class-0
- Pause no-drop

How to enable link pause (flow control) on Nexus 5000 interface

Configuring "lowcontrol send on and flowcontrol receive on does not enable flowcontrol on on Nexus 5000 switch port links when connected to another Nexus 5000 interface.

Possible Cause

By default, the DCBX runs on the Nexus 5000 interface. If the peer does not run DCBX, then the interface is configured for tail-drop.

Use one of the following commands to verify:

- Use the show interface ethx/y flowcontrol command and check to see if the operating state is off.
- Use the **show interface ethx/y priority-flow-control** command and check to see the if operating state is off.

Solution

Use the following commands under interface ethx/y to enable link pause:

- flowcontrol receive on
- flowcontrol send on

Registers and Counters

The following are the commands to access various registers and counters:

REVIEW DRAFT-CISCO CONFIDENTIAL

Nexus 5000 10G PFC

Use the following command:

show hard in gatos asic <gatos_num> registers match mm_CFG_pause\$

Nexus 5000 1G storm control

Use the following commands:

```
show plat fwm info lif eth1/1
show plat fwm info pif eth1/1
debug hardware internal gatos asic 0 dump-mem 0x3b9000 20
```

Nexus 5000 10G storm control

Use the following commands:

show plat fwm info lif eth1/5 show plat fwm info pif eth1/5 debug hardware internal gatos asic 1 dump-mem 0x3b9000 20

Nexus 5000 storm control counter

Use the following command:

show hardware internal gatos asic 1 counters rx_db 2 | grep storm

afm-related CLI commands and tools

Commands	Purpose
show platform afm in att br	Shows which features or groups are attached to which interface.
show platform afm in att global	Shows the IDs of policies including QoS Policies (printed as NP Policies) attached on the global interface.
show platform afm in att interface ethernet x/y	Shows the IDs of policies including QoS Policies for an interface or PC.
show platform afm in group id X asic Y	Shows the TCAM entries for a particular group on a particular ASIC/GATOS.
show platform afm in map-tbls	Shows the internal mapping tables, such as the ext-cos to qos-group, qos-group to int-cos, and int-cos to class_id maps.

FEX qosctrl debug commands

Command	Purpose
show platform	Displays the PI information for every port.
software qosctrl port 0 0 nif <0-48> [sat switch]	(Useful if port level configuration exists.)
show platform	Displays the PI information for every port.
software qosctrl port 0 0 hif <0-48> [sat switch]	(Useful if port level configuration exists.)
show platform software qosctrl policy hif	Displays the global network-qos and queueing configurations.
show platform software qosctrl global	Global PI level configurations.
show platform software qosctrl pss	Stores PSS information.
<pre>show platform software qosctrl asic <mod> <asic></asic></mod></pre>	Displays per asic level port details.
<pre>show platform software qosctrl default port <mod> <asic></asic></mod></pre>	Displays default port settings on FEX ports.
<pre>show platform software qosctrl port <mod> <asic> <port-type> <port></port></port-type></asic></mod></pre>	Displays per-port level PI and PD data structures.

N2K-C2148T FEX counters



Use the following commands (in the FEX shell) in preparation to display the statistics of MAC level traffic and pause statistics:

- show plat soft fex info satport < fex-interface-id> (for mapping except in the case of NIF in RW6)
- show plat soft redwood sts
- show plat soft redwood ss

Command	Purpose
<pre>show platform software qosctrl port 0 6 hif 1 counters</pre>	Displays counters.
show plat soft redwood rmon 6 nif0	Displays statistics of MAC level traffic and pause statistics of NIF of eth103/1/37.
show plat soft redwood rmon 6 hif5	Displays statistics of MAC level traffic and pause statistics of iHIF for eth103/1/37.

Command	Purpose
show plat soft redwood rmon 4 nif1	Displays statistics of MAC level traffic and pause statistics of iNIF for eth103/1/37.
show plat soft redwood rmon 4 hif5	Displays statistics of MAC level traffic and pause statistics of HIF for eth103/1/37.
show plat soft redwood ss	Displays mapping of HIF/NIF to SS.
show plat soft redwood ss 4 3	Displays statistics of RW4 SS3 - Host Receive from HIF4-7 to NIF0-3
show plat soft redwood ss 4 2	Displays statistics of RW4 SS2 - Host Receive from HIF0-3 to NIF0-3
show plat soft redwood rate	Displays overall statistics for non-zero traffic.
show plat soft redwood rmon 6 cif0	Helps debug traffic going from CIF to CPU.
show plat soft qosctrl port 0 6 cif 0 counters	Helps debug traffic going from CIF to CPU.

Nexus 5000 multicast-optimization

Use the following commands:

```
show plat fwm in mco-info
show plat fwm in vlan 1 all_macgs
```

Nexus 5000 FCoE classification

• For the FCoE interface, use the following commands:

show plat fwm info pif ethernet 1/1 | grep gatos debug platform hardware peek lu 7 index 5 pifTable

• For the FC interface, use the following commands. (The first command is used to get the gatos number and the fc number.)

```
show platform fwm info pif fc <id>
debug peek lu <gatos> index <fc num> pifTable
```

Nexus 5000 MTU programming

Use the following command:

show hardware internal gatos asic 0 registers match bm_port_CFG.*_max

Nexus 5000 interrupt

Use the following commands:

debug hardware internal gatos asic 0 clear-interrupt

```
show hardware internal gatos asic 0 interrupt show hardware internal gatos event-history errors
```

Untagged COS

Use the following commands:

sh platform afm info attachment interface eth3/1 sh system internal ipqos port-node eth3/1

Buffer usage and packet drop debugging on N2K-C2232P FEX

Use the following command:

show platform software qosctrl asic 0 0


Troubleshooting SAN Switching Issues

A storage area network (SAN) is a network of storage devices that provide data storage for servers.

This chapter describes how to identify and resolve problems that can occur with a SAN and the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Overview
- NPV
- Zoning
- SAN Port Channels
- FC Services
- Cisco Fabric Services
- VSANs
- Registers and Counters

Overview

The two most common symptoms of problems in a storage network are:

- A host cannot access its allocated storage.
- An application does not respond after attempting to access the allocated storage.

By answering the questions in this section, you can determine the paths you need to follow and the components that you should investigate further. These questions are independent of host, switch, or subsystem vendor.

Consider 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 SAN, host, or subsystem, or new LUNs exported to an existing host.)
- Has the host ever been able to see its storage?
- Does the host recognize any LUNs in the subsystem?
- Are you trying to solve an existing application problem (too slow, too high latency, excessively long response time) or did the problem show up recently?
- What changed in the configuration or in the overall infrastructure immediately before the applications started to have problems?

General SAN troubleshooting steps

Step 1	Obtain information on problems in your fabric.
Step 2	Verify physical connectivity between your switches and end devices.
Step 3	Verify registration to your fabric for all SAN elements.
Step 4	Verify the configuration for your end devices (storage subsystems and servers).
Step 5	Verify end-to-end connectivity and fabric configuration.

NPV

NP Uplink ports on NPV edge switch are stuck in initializing state

NP uplink ports connected to the core NPIV switch do not come online and are stuck in an initializing state.

Possible Cause

The core switch might not have been enabled for NPIV.

Example:

```
switch(config-if)# sh int fc2/2
fc2/2 is down (Initializing)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:42:00:0d:ec:a4:3b:80
Admin port mode is NP, trunk mode is on
```

<u>Solution</u>

• Check the status of the NPV external interfaces. Check that the NPIV is enabled on the core switch.

Example:

• If NPIV is disabled, then enable NPIV on the core switch.

Example:

switch(config)# feature npiv

Server interface does not come up and "NPV upstream port not available" message appears

A server port connected to the NPV edge switch does not come online, and the **show interface** command indicates a status of NPV upstream port not available.

Possible Cause

The upstream NP_Port(s) and the downstream server F_Port(s) on the NPV edge switch may not be in the same VSAN.

Example:

```
switch# sh int fc2/7
fc2/7 is down (NPV upstream port not available)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:47:00:0d:ec:a4:3b:80
Admin port mode is F, trunk mode is off
snmp link state traps are enabled
Port vsan is 99
Receive data field Size is 2112
```

Solution

• Check the VSAN membership of the upstream port and the server port.

Example:

```
switch# show vsan membership
vsan 1 interfaces:
fc2/1 fc2/2 fc2/3 fc2/4
fc2/5 fc2/6 san-port-channel 200
vsan 99 interfaces:
fc2/7 fc2/8
```

• In the example above, notice that the upstream ports (fc2/1-2) are in VSAN 1, and the server ports (fc2/7-8) are in VSAN 99.

Move the NP ports on the NPV edge, and the F ports on the NPIV core into the same VSAN as the server ports.

Example:

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 99 interface fc2/1-2
switch(config-if)# vsan database
switch(config-vsan-db)# vsan 99 interface fc1/17-18
Traffic on fc1/17 may be impacted. Do you want to continue? (y/n) y
Traffic on fc1/18 may be impacted. Do you want to continue? (y/n) y
```

Note

Alternatively, if the NPIV core and NPV edge switch are F_Port Trunking capable switches, then that would be the recommended configuration.

Uneven load balancing on the NPV NP ports

An examination of NP upstream ports that are members in the same VSAN reveals that uneven load balancing is occurring.

Possible Cause

This may be normal and a direct result of the default SID/DID load balancing that is done before the Nexus 5000 4.2(1)N1 release.

Solution

If the upstream switch is an MDS switch that is running 4.1(3) code or above, and it is a NPV F_Port Trunking capable switch, the preferred configuration would be to run the F_Port Trunking Port Channeling feature.

Example (NPIV core):

```
pod3-9222i(config)# feature npiv
pod3-9222i(config)# feature fport-channel-trunk
pod3-9222i(config)# interface port-channel 1
pod3-9222i(config-if)# switchport mode f
pod3-9222i(config-if)# switchport trunk mode on
pod3-9222i(config-if)# channel mode active
pod3-9222i(config-if)# interface fc2/13, fc2/19
pod3-9222i(config-if)# switchport mode f
pod3-9222i(config-if)# switchport rate-mode dedicated
pod3-9222i(config-if)# switchport trunk mode on
pod3-9222i(config-if)# switchport trunk mode on
pod3-9222i(config-if)# channel-group 1 force
```

In this example, fc2/13 and fc2/19 are added to port channel 100 and are disabled. Do the same operation on the switch at the other end of the port channel, then do no shutdown at both ends to bring them up.

Example:

pod3-9222i(config-if) # no shut

Example (NPV Edge):

```
pod7-5020-51(config)# interface san-port-channel 1
pod7-5020-51(config-if)# switchport mode np
pod7-5020-51(config-if)# switchport trunk mode on
pod7-5020-51(config-if)# interface fc2/1-2
pod7-5020-51(config-if)# switchport mode np
pod7-5020-51(config-if)# switchport trunk mode on
pod7-5020-51(config-if)# channel-group 1
```

In this example, fc2/1 and fc2/2 are added to port channel 1 and are disabled. Do the same operation on the switch at the other end of the port channel, then do no shutdown at both ends to bring them up

Example:

pod7-5020-51(config-if)# no shut

Server on downstream NPV edge switch does not login to the fabric

The server connected to the downstream NPV edge switch does not log in to the fabric.

Possible Cause

The server on the downstream NPV edge switch does not log in to the fabric, and/or you see a "waiting for FLOGI" message.

Example:

switch# show npv status
npiv is enabled

1-4

Solution

- Verify the configuration of both the NPV edge and core switches. If you are not running the F_Port trunking feature, then verify that there are no VSAN mismatches and that the server ports, NPV NP ports, NPIV Core F_Ports, and storage ports are all in the same VSAN and all are online.
- If the configuration is correct and you can determine where the problem might be, you can collect an Ethanalyzer trace and verify that the Fabric Login (FLOGI) frame is being received and sent to the NPIV core as a Fabric Discovery (FDISC) command.

Example Ethanalyzer trace:

```
switch# ethanalyzer local sniff-interface inbound-hi display-filter "!llc && !stp"
limit-captured-frames 0 write bootflash:npv-trace
Capturing on eth4
```

• Recreate the problem by flapping the NPV-attached server port. The trace will be written to bootflash and can be copied off the switch by using the following:

copy bootflash: ftp:

• After the trace has been copied, you can now open and verify the flow using Wireshark.

Example normal NPV login flow:

Server> FLOGI> NPV Edge Switch	Fabric Login frame = FLOGI
NPV Edge Switch> FDISC> NPIV Core Switch	Fabric DISCovery frame maps parameters from Server FLOGI
NPV Core Switch> Accept> NPV Edge Switch	NPIV Core assigns an FCID with the Accept to the FDISC from NPV Edge Switch
NPV Edge Switch> Accept> Server	Accept to original Server FLOGI with FCID assigned from NPIV Core Switch

Locating exact port that server is physically attached to

NPIV switches lose visibility into the physical port that a downstream NPV-connected server is attached to. The following process can be used to identify that physical port.

Possible Cause

When you have an NPIV core switch that has several downstream NPV edge switches attached, you might want to locate the exact port that a server is physically attached to.

Γ

Solution

• Identify the PWWN of the server and the corresponding switch that it is attached to.

Example:

NPIV-Core(config-if) # show flogi database

```
fc1/16 100 0xee00e4 21:00:00:04:cf:17:66:b7 20:00:00:04:cf:17:66:b7
fc1/16 100 0xee00e8 21:00:00:04:cf:17:66:0e 20:00:00:04:cf:17:66:0e
fc1/25 100 0xee0100 20:41:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41
fc1/26 100 0xee0200 20:42:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41
fc1/26 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:00:c0:dd:12:04:f3
```

In the example, the server is identified by this address:

fc1/26 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:c0:dd:12:04:f3

and the switch is identified by this address

fc1/26 100 0xee0200 20:42:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41

Identify the IP address of the NPV edge switch.

Example:

```
NPIV-Core(config-if)# sh fcns database npv VSAN 100:
```

```
20:64:00:0d:ec:a3:da:41 172.18.217.51 fc2/1 20:00:00:0d:ec:51:0c:00 fc1/25 20:64:00:0d:ec:a3:da:41 172.18.217.51 fc2/2 20:00:00:0d:ec:51:0c:00 fc1/26
```

• Telnet to the NPV edge switch.

Example:

NPIV-Core(config-if) # telnet 172.18.217.51

• Identify the PWWN of the server.

Example:

switch-NPV-Edge# show npv flogi-table

vfc3 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:00:c0:dd:12:04:f3 fc2/2

• If the interface is a FCoE (VFC) interface as shown in the previous example, use the show interface vfc3 command to see which port that the VFC is physically bound to.

VSANs stuck in initializing state after configuring the 4.2(1)N1 F_Port trunking feature

Using the **show interface** command of the F_Port trunking port channel or trunking member of the port channel indicates that certain VSANs are in an initializing state and do not come online.

Possible Cause

After configuring the 4.2(1)N1 F_Port Trunking feature, VSANs on the trunk ports appear to be stuck in an initializing state.

Example:

```
switch(config-if) # sh int fc2/1
fc2/1 is trunking
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:41:00:0d:ec:a4:3b:80
Admin port mode is NP, trunk mode is on
snmp link state traps are enabled
Port mode is TNP
Port vsan is 1
Speed is 4 Gbps
Transmit B2B Credit is 16
Receive B2B Credit is 16
Receive data field Size is 2112
Beacon is turned off
Belongs to san-port-channel 200
Trunk vsans (admin allowed and active) (1,99,200)
Trunk vsans (up) (1,99)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (200)
```

Under the Trunk Failures tab of Fabric Manager, you might also see the trunk VSAN listed. However, this may be normal. If no downstream devices are logged in for a particular VSAN, that VSAN stays in initializing state.

<u>Solution</u>

For the VSANs that you are working with, verify by using the following command:

Example:

switch# show npv flogi-table

fc2/7 99 0xba0002 10:00:00:00:00:00:00 10:00:00:00:00:00:02:00 Spo200 fc2/8 99 0xba0003 10:00:00:00:00:00:00 10:00:00:00:00:00:01:00 Spo200 Total number of flogi = 2.

In this example, no devices are logged into VSAN 200.

Zoning

Cannot activate zoneset and cannot configure zoning in enhanced zoning mode

The zone set cannot be activated and zoning cannot be configured in enhanced zoning mode. The error message "Zoning database update in progress, command rejected" might be received.

Possible Cause

Another user on the same switch or on a different switch is holding the enhanced zoning configuration lock.

Solution

Release the zoning lock with the following:

- **Step 1** Determine which switch (domain/ip address) has the lock.
- **Step 2** Determine which user has the lock on that switch.

Step 3 Clear the lock for that user on that switch.

 On the same switch, enter the show zone status vsan <vsan-id> command to determine which user holds the lock.

Example:

```
switch1# show zone status vsan 200
VSAN: 200 default-zone: deny distribute: active only Interop: default
mode: enhanced merge-control: allow
session: remote [dom: 121][ip: 171.165.98.20] <<==</pre>
```

In this example the remote switch with the IP address of 171.165.98.20 has the lock.

• Connect to the remote switch and enter the **show zone status vsan** command.

Example:

switch2# show zone status vsan 200

```
VSAN: 200 default-zone: deny distribute: active only Interop: default
mode: enhanced merge-control: allow
session: cli [remi] <<==
```

In the example, user Remi is holding the enhanced zoning lock.

- On the remote switch (N5K2 in the example), release the lock with the **no zone commit vsan** <*vsan-id*> command.
- To confirm that the lock had been cleared, enter the **show zone status vsan** *<vsan-id>* command. At this point, the session parameter should appear as none.
- If the lock still persists, remove the lock from the switch that holds the lock with the **clear zone lock** command.
- If the lock continues to persist, use the following commands to collect information to aid further analysis:

```
show zone internal vsan <vsan-id>
show zone status vsan <vsan-id>
show fcdomain domain-list vsan <vsan-id>
show users
show tech-support zone
show tech-support device-alias
show logging
```

Host cannot communicate with storage

In initial SAN deployments or after topology changes in the SAN, some hosts might not be able to communicate with storage. The initiator cannot access the LUNs that were allocated for them in the storage array.

Possible Cause

If the host and storage are connected to two different switches, the ISL link, (the xE port connecting both switches) might be isolated.

The xE port might be isolated in a specific VSAN for possible reasons:

• Misconfigured fabric timers

- Misconfigured port parameters
- Mismatched zoning

<u>Solution</u>

To resolve the VSAN isolation on the TE port:

• Use the **show interface fc** <*slot/port*> command on the TE port to determine the VSAN number.

The isolated VSAN number must match the VSAN number where the host and the storage are connected to.

In the display output, you see the Trunk vsans (isolated) (Vsan <*vsan-id*>).

• Use the **show port internal info interface fc** <*slot/port*> command to determine the root cause of the VSAN isolation.

Possible Cause

Host and storage are not in the same VSAN.

Solution

- Use the **show vsan membership** command to verify that both the host and the storage are in the same VSAN.
- If the host and the storage are in different VSANs, in the configuration mode use the commands **vsan database** and **vsan** <*vsan-id*> **interface fc** <*slot/port*> to move the interface connected to the host and storage devices into the same VSAN.

Possible Cause

The host and storage are not in the same zone. The zone is not in the active zone set. There is no active zone set and default zone policy is set to deny.

Solution

• Use the command **show zone status** *<vsan-id>* to determine if the default zone policy is set to deny.

Example:

```
switch# show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
```

The state default zone policy permit means all nodes can see all other nodes. Deny means all nodes are isolated when not explicitly placed in a zone.

If you are not using zoning, you can change the default zone policy with zone default-zone permit, but this is not a best practice.

• Use the **show zone member** command for host and storage to verify that they are both in the same zone. If they are not in the same zone, use the **zone name** *<zonename <vsan-id>* command to create a zone in that VSAN.

Example:

```
switch(config)# zone name testzone vsan 100
switch(config-zone)# member pwwn 21:00:00:20:37:9e:02:3e
switch(config-zone)# member pwwn 21:00:00:c0:dd:12:04:ce
```

Use the **show zone vsan** *<vsan-id>* command to verify that host and storage are now in the same zone.

• Use the **show zoneset active vsan** <*vsan-id*> command to verify the name of the active zone set.

If the zone that has the host and storage is not in the active zoneset, use the **zoneset name** command from the configuration mode to enter the zoneset sub-mode and use the **member** command to add the zone to the active zone set.

Example:

```
switch(config # zoneset name testzoneset vsan 100
switch(config-zoneset)# member testzone
```

• Use the **zoneset activate** command to activate the zone set.

Example:

switch(config)# zoneset activate testzoneset vsan 100

Zone merge failure when two switches connect using E or TE port

A zone merge failure can occur when two switches connect using the E or TE port.

Possible log messages that can be seen in the **show logging** log are shown in the example.

Example:

```
%ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc2/1 error:
Received rjt from adjacent switch:[reason:0]
%ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc1/2 error:
Member mismatch
%ZONE-2-ZS_MERGE_ADJ_NO_RESPONSE: Adjacent switch not responding,isolating interface
%ZONE-2-ZS_MERGE_FULL_DATABASE_MISMATCH: Zone merge full database mismatch on interface
```

Possible Cause

Two switches may have the same zone set name and the same zone names, but different zone members.

When merging switch fabrics, you must ensure that the zones in both active zone sets have unique names, or that any zones with the same name have exactly the same members. If either of these conditions are not met, then the E port connecting the two fabrics will appear in an isolated state.

The process to merge switch fabrics is as follows:

- The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
- If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
- If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
 - If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise the link is isolated.
 - If the setting is allow, then the merge rules are used to perform the merge. The host and storage are not in the same zone. The zone is not in the active zoneset. There is no active zoneset and default zone policy is set to deny.

Solution

If there is a zone merge failure, the issue can be resolved by using one of the following methods:

- Modify the zone members in both zone sets to match and eliminate the conflict.
 - Use the show zoneset active vsan <vsan-id> command on both switches to compare the zones and their respective members.
 - Change the membership of one of the zones to match the other zone of the same name.

- Deactivate the zone set on one of the switches and restart the zone merge process.
 - Use the no zoneset activate name <*zonesetname*> <*vsan-id*> command to deactivate the zone set configuration from one of the switch.
 - Use the show zoneset active command to confirm that the zone set has been removed.
 - Use the shutdown command to shut down the connection to the zone to be merged, and use the no shutdown command to reactivate the connection to the zone to be merged.
 - Use the show zoneset active <vsan-id> to verify that all the members are correct and use the show interface fc <slot/port> to verify that the VSAN is not isolated.
- Explicitly import or export a zone set between the switches to synchronize them.
 - Use the zoneset import interface <interface-number> vsan <vsan-id> command or the zoneset export interface <interface-number> vsan <vsan-id> command to overwrite the active zone set on one of the switches.
 - Use the show interface fc <*slot/port>* to verify that the VSAN is not isolated after this disruptive operation.

Zone set activation failure

When a zone set activation failure occurs, the possible log messages that can be seen in the **show logging** log are shown in the example.

Example:

```
ZONE-2-ZS_CHANGE_ACTIVATION_FAILED: Activation failed.
ZONE-2-ZS_CHANGE_ACTIVATION_FAILED_RESN: Activation failed : reason
```

Possible Cause

Zone set activation can fail if a new switch joins the fabric when the size of the zone database is larger than 2048 KB.

Solution

• Use the **show zone analysis active vsan** <*vsan-id*> command to analyze the active zone set database. Verify that the formatted size does not exceed 2048 KB.

If the 2048 KB limit is exceeded, then some zones or devices within a zone must be removed.

Example:

```
switch# show zone analysis active vsan 100
Zoning database analysis vsan 100
Active zoneset: vsm_vem_v100_zs [-]
Activated at: 13:13:44 UTC May 27 2010
Activated by: Merge [ Interface san-port-channel 100 ]
Default zone policy: Deny
Number of devices zoned in vsan: 1/9 (Unzoned: 8)
Number of zone members resolved: 1/3 (Unresolved: 2)
Num zones: 1
Number of IVR zones: 0
Number of IPS zones: 0
Formatted size: 92 bytes / 2048 Kb
```

- Use the **show zone internal change event-history vsan** <vsan-id> command to determine the zone set activation problem.
- To further troubleshoot this issue, capture the output from the **show tech-support zone** command and the **show logging log** command.

Full zone database synchronization failure across two switches

A full zone database synchronization failure may occur when two switches connect using the E or TE port and have different zone set distribution policies. As a result of a fabric isolation/merge, one fabric might not have the full zone set database in the running configuration.

Possible Cause

The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

The zone distribute policy can be set differently on two switches and that could cause synchronization failure.

<u>Solution</u>

Use the **show zone status** command to verify the distribution policy on both switches.

Example:

VSAN: 100 default-zone: deny distribute: active only Interop: default mode: basic merge-control: allow

When the distribute policy is set to active only the active zone set is distributed. Also verify that the distribute policy is set to full.

To enable the full zone set and active zone set distribution to all switches on a per-VSAN basis in the configuration mode, use the **zoneset distribute full vsan** *<vsan-id>* command.

Mismatched default zone policy in switches in VSAN causes unexpected results when accessing storage

A mismatched default zone policy in all switches in the VSAN in the basic zone mode might cause unexpected results for any hosts accessing storage.

Possible Cause

If the default zone policy is set to permit and if there is no active zone set for VSAN, then all the members of the VSAN can see all the other nodes.

Solution

One approach is to migrate the zone operation mode from basic to enhanced. Enhanced zoning synchronizes the zone configuration across all switches in the VSAN. This eliminates the possibility of mismatched default zone policies.

• Use the **show zone status** command to display the status of the zone.

```
VSAN: 300 default-zone: deny distribute: active only Interop: default mode: basic merge-control: allow
```

• Use the **zone default-zone** command to set the default zone policy and use the **zone mode enhanced** <*vsan-id>* command to set the operation to enhanced zoning mode.

Another approach is the foillowing:

- Use the **show zone status** command in all switches in the VSAN to verify the operation mode and the default-zone policy.
- Use the **zone mode basic** command to change any switches that are not in basic mode.

• Use the **zone default-zone** command on each switch in the VSAN to set the same default zone policy.

SAN Port Channels

Fibre channel port is down when trying to connect switches via SAN Port Channel

When trying to connect switches using SAN port channel, the Fibre Channel port is down.

The show interface brief command produces

fc slot/port is down (Error disabled - Possible port channel misconfiguration)

Possible Cause

One of the SAN port channel compatibility parameters is misconfigured in the configuration.

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a port channel. The compatibility check is performed before a port is added to the port channel.

The check ensures that the following parameters and settings match at both ends of a port channel:

- Capability parameters: type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends.
- Administrative compatibility parameters: speed, mode, rate mode, port VSAN, allowed VSAN list, and port security..
- Operational parameters: remote switch WWN and trunking mode.

Solution

• Use **show san-port-channel compatibility-parameters** to verify which parameters need to be checked in the configuration.

Generally, if the configuration is fixed and the FC port is shut or no shut, the port recovers normally.

• If the issue persists with a different error message, debug further by running one or more of the following commands:

```
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show san-port-channel internal event-history errors
show logging log | grep fc <slot/port>
show san-port-channel internal event-history all
show tech-support detail > bootflash:showtechdet
```

Newly added Fibre Channel interface does not come online in a SAN Port Channel

When a new Fibre Channel interface is added, it does not come online in a SAN port channel.

Г

The following error message during the configuration operation may appear.

Command failed: port not compatible [reason]

Possible Cause

Port channel mode is configured as on.

If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down.

<u>Solution</u>

Explicitly enable the ports again using the no shutdown command.

Possible Cause

Interface parameters are not compatible with the existing SAN port channel.

<u>Solution</u>

Use the force option to force the physical interface to take on the parameters of the SAN port channel. In the interface sub-configuration mode, use the **channel-group** *<channel-group number* **> force** command.

Cannot configure trunking

Trunking cannot be configured under the interface configuration mode.

The following error message may appear in the CLI output:

error: invalid switchport config

Possible Cause

Trunking protocol is disabled.

<u>Solution</u>

Enable trunking by using the trunk protocol enable CLI command.

VSAN traffic does not traverse trunk

The VSAN traffic is not able to traverse the trunk.

A host cannot gain access to a target that is on the same VSAN and connected to two different switches using TE ports. The VSAN traffic is not able to traverse the trunk. Depending on the path from host to target, you may observe a performance degradation or you may not be able to access any disks.

Possible Cause

VSAN is not listed in the allowed-active VSAN list.

Solution

Add VSAN to the allowed-active list by using the switchport trunk allowed vsan command.

xE port is isolated in a specific VSAN under interface of SAN Port Channel

The xE port is isolated in a specific VSAN that is under an interface of a SAN port channel.

The following error message may appear in the logging log:

```
"%$VSAN <VSAN#>%$ Interface port-channel <channel #>, vsan <vsan #> is down (isolation due to [cause])".
```

Possible Cause

The xE port can be isolated in a specific VSAN for many reasons:

- Fabric timers might be misconfigured.
- Port parameters might be misconfigured.
- Zoning mismatch.

<u>Solution</u>

To resolve the VSAN isolation on the TE port, use the **show interface fc** *<slot/port>* command on the TE port to determine the VSAN number. The isolated VSAN number must match the VSAN number where the host and the storage are connected to.

In the output of the command, look for information such as Trunk vsans (isolated) (Vsan <number>).

Use the **show port internal info interface san-port-channel** *<number>* command to determine the cause of the VSAN isolation.

Cannot create a san-port-channel interface

A SAN port channel interface cannot be created.

The following error message may appear while in configuration mode:

```
failed to create port-channel channel-id:
```

Possible Cause

The user receives the following message:

failed to create port-channel channel-id: all port-channels have been created [max channel number reached]



You can create a maximum of four SAN port channels (including Release NX-OS 4.2(1)N1(1)). This is a software limitation.

Solution

If you need to create a SAN port channel with a specific number, but four SAN port channels were already configured, then you have to delete one of the SAN port channels that is not actively used. Use the no interface **san-port-channel** $\langle x \rangle$ command to delete one of the SAN port channels.

Possible Cause

You receive the following message:

Channel group X is already an Ethernet port channel

<u>Solution</u>

Г

You need to choose another number between 1 to 256 to configure the SAN port channel.

Use the **show port-channel usage** command to determine the numbers that were used for the existing port channels.

Example:

```
show port-channel usage
Total 3 port-channel numbers used
Used : 198 - 199 , 500
Unused: 1 - 197 , 200 - 499 , 501 - 4096
(some numbers may be in use by SAN port channels)
```

FC Services

This section includes an overview of troubleshooting Cisco Fibre Channel Services followed by a description of common problems and their solutions.

Overview

Fibre Channel fabrics provide a set of services for its clients, which are the Fibre Channel nodes. These Fibre Channel services (FC services) allow the nodes to interact with the storage network to exchange information, such as connection state, connection parameters, configuration, topology changes, and so on.

The FC services can be accessed through login into ports that hold a well known address (WKA). WKAs are port FC IDs that are reserved for internal use of the fabric, usually fabric services.

The following table describes the well-known addresses and the service associated with each: (Source: www.t11.org)

Well Known Address	Description
x'FF FC 01' to x'FF FC FE'	Reserved for Domain Controllers
x'FF FF F0'	Reserved for N_Port Controller
x'FF FF F1' to x'FF FF F3'	Reserved
x'FF FF F4'	Event Service (FC-GS-5)
x'FF FF F5'	Multicast Server (FC-PH3)
x'FF FF F6'	Clock Synchronization Server (FC-PH3)
x'FF FF F7'	Security Key Distribution Service (FC-PH3)
x'FF FF F8'	Alias Server (FC-PH2)
x'FF FF F9'	Quality of Service Facilitator-Class4 (FC-PH2)
x'FF FF FA'	Management Service (FC-GS-5)
x'FF FF FB'	Time Service (FC-GS-5)
x'FF FF FC'	Directory Service (FC-GS-5)

Well Known Address	Description
x'FF FF FD'	Fabric Controller
x'FF FF FE'	F_Port Controller
x'FF FF FF'	Broadcast Address/Server

Fibre channel port remains in initializing state

A fibre channel F type port does not come online and is stuck in an initializing state.

The **show interface fc** *<slot/port>* command displays the following message.

fc slot/port is down (Initializing)

A Fibre Channel port goes into the initialization state after a successful completion of link-level initialization. For F type ports, the next step is to complete the FLOGI (fabric login) process. The port remains in the initialization state until the FLOGI process completes.

Possible Cause

The port is up because the link partner has put itself into a bypass mode.

<u>Solution</u>

Use the **show hardware internal fc-mac** *<slot-number>* **port** *<port-number>* **statistics** command to check whether the Class-3 input counter is increasing after the successful completion of link initialization.

Example:

switch# show hardware internal fc-mac 2 port 1 statistics	
ADDRESS STAT	COUNT
0x000003c FCP_CNTR_MAC_RX_LOSS_OF_SYNC	0x1
0x000003d FCP_CNTR_MAC_RX_BAD_WORDS_FROM_DECODER	0x50
0x0000042 FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ	0x152
0x0000043 FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY	0x7c
0x0000061 FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES	0x130
0x0000062 FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES	0x22
0x0000069 FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS	0x61c98
0x000006a FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS	0xff0
0x0000065 FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES	0x52
0x0000066 FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES	0x2a
0x000006d FCP_CNTR_MAC_DATA_TX_CLASS3_WORDS	0x944c
0x000006e FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS	0xec4
0xfffffff FCP_CNTR_LINK_RESET_IN	0x1
0xfffffff FCP_CNTR_OLS_IN	0x1
0xfffffff FCP_CNTR_NOS_IN	0x1
0xfffffff FCP_CNTR_LRR_IN	0x2
0xfffffff FCP_CNTR_LINK_RESET_OUT	0x1
0xfffffff FCP_CNTR_OLS_OUT	0xa
0xfffffff FCP_CNTR_NOS_OUT	0x2
0xfffffff FCP_CNTR_LRR_OUT	0xb
0xfffffff FCP_CNTR_LINK_FAILURE	0x2

Possible Cause

The FLOGI packet was dropped somewhere in the data path, starting from FC-MAC to FLOGI server.

<u>Solution</u>

Consider the following solutions:

- Use the **show hardware internal fc-mac** *<slot-number>* **port** *<port-number>* **statistics** command to check for Class-3 packet counters.
- Analyze the output of the show flogi internal all interface fc <*slot/port>* command for a possible drop of FLOGI packets somewhere in the path.
- Check the Fport server fault-injection table for any Invalid, Drop FLOGI packets.
- Use the shut CLI command followed by the no shut command to disable and enable the FC slot/port.
- If this does not clear the problem, try moving the connection to a different port on the same or another FC module.
- If the problem continues to persist, use the following commands to collect information to aid in further analysis:

```
show tech-support flogi
show logging log | grep fc <slot/port>
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show tech-support detail > bootflash:showtechdet
show platform fwm info pif fc <slot/port> {find the gatos instance for the port}
show platform fwm info gatos-errors 13 {check for the non-zero counters for drops}
```

Capture debug Flogi with the following:

Specific VSAN traffic is not being routed through SAN fabric

Each configured VSAN needs to support a separate set of fabric services. One such service is the FSPF routing protocol, which can be independently configured per VSAN. You may see that specific VSAN traffic is not being routed if inappropriate traffic engineering capabilities are used.

Possible Cause

The FSPF hello interval is misconfigured.

The following example shows possible log messages from the **show logging** command log.

Example:

```
FSPF-3-HELLO_MISMATCH: %$VSAN <vsan-id>%$ Mismatch in Hello timer in the Hello packet on
interface san-port-channel <channel-id>
%FSPF-3-FC2_PROC_ERR: %$VSAN <vsan-id>%$ Error in processing HELLO packet on interface
san-port-channel <channel-id>, Error = Bad packet received
```

Solution

To resolve a wrong hello interval on an ISL using the NX-OS CLI, perform the following steps.

Step 1 Either use the **debug fspf all** command and look for wrong hello interval messages or check the last messages in the **show logging** command log for an error message.

The debug output generates the following messages:

fspf: Wrong hello interval for packet on interface 40000c7 in VSAN 200
fspf: Error in processing hello packet , error = Bad packet received

Step 2 Use the **undebug all** command to turn off debugging.

```
<u>P</u>
Tip
```

Open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 3 Use the **show fspf vsan** *<vsan-id>* **interface** command to view the FSPF configuration on both switches.

Example:

```
switch# show fspf vsan 200 interface port-channel 200
FSPF interface port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 40 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
   Number of packets received : LSU 3 LSA 3 Hello 136 Error packets 3
   Number of packets transmitted : LSU 3 LSA
                                               3 Hello 182 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
```

```
Number of packets received : LSU 3 LSA 3 Hello 185 Error packets 169
Number of packets transmitted : LSU 3 LSA 3 Hello 139 Retransmitted LSU 0
Number of times inactivity timer expired for the interface = 24
```



In the Example:

- The hello timer is not set to the default (20 seconds) on the first switch. Check the neighboring switch (Nexus 5000) configuration to make sure it matches.
- FSPF is not in FULL state. This indicates a problem.
- **Step 4** In the interface configuration mode, change the fspf hello-interval value to match the same values on both switches.

Example:

```
switch(config)# interface san-port-channel 200
switch(config-if)# fspf hello-interval 40 vsan 200
```

Step 5 Verify that the FSPF is in FULL state after the change.

```
switch(config-if)# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 40 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x18(24)
Neighbor Interface is san-port-channel 200 (0x000400c7)
Statistics counters :
    Number of packets received : LSU 7 LSA 7 Hello 238 Error packets 218
    Number of packets transmitted : LSU 7 LSA 7 Hello 180 Retransmitted LSU 0
    Number of times inactivity timer expired for the interface = 32
```

Possible Cause

The FSPF dead interval is misconfigured.

The following example shows possible log messages from the **show logging** command:

Example:

```
%FSPF-3-HELLO_MISMATCH: %$VSAN <vsan-id>%$ Mismatch in Dead timer in the Hello packet on
interface san-port-channel <channel-id>
N5K-2 %FSPF-3-FC2_PROC_ERR: %$VSAN <vsan-id>%$ Error in processing HELLO packet on
interface san-port-channel <channel-id>, Error = Bad packet received
```

Solution

To identify a mismatch of dead intervals on an ISL using the NX-OS CLI, perform the following steps:

```
Step 1 Either use the debug fspf all command and look for wrong dead interval messages or check the last messages in the show logging command log for an error message.
```

The debug output generates the following messages:

fspf: Wrong hello interval for packet on interface 40000c7 in VSAN 200
fspf: Error in processing hello packet , error = Bad packet received

Step 2 Use the **undebug all** command to turn off debugging.



Open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 3 Use the **show fspf vsan** *<vsan-id>* **interface** command to view the FSPF configuration on both switches.

Example:

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 120 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
    Number of packets received : LSU 4 LSA 4 Hello 27 Error packets 4
```

Number of times inactivity timer expired for the interface = 4

Number of packets transmitted : LSU 4 LSA 4 Hello 38 Retransmitted LSU 0
Number of times inactivity timer expired for the interface = 0
switch# show fspf vsan 200 interface port-channel 200
FSPF interface port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
Number of packets received : LSU 4 LSA 4 Hello 41 Error packets 35
Number of packets transmitted : LSU 4 LSA 4 Hello 29 Retransmitted LSU 0

Note

In the example:

- The dead timer is not set to the default (80 seconds) on the first switch. Check the neighboring switch (MDS) configuration to make sure it matches.
- FSPF is not in FULL state. This indicates a problem.
- **Step 4** In the interface configuration mode, change the fspf dead-interval value so that the same values match on both switches.

```
switch(config)# interface san-port-channel 200
switch(config-if)# fspf dead-interval 80 vsan 200
```

Step 5 Verify that the FSPF is in FULL state after the change. Ensure that there is a route for VSAN traffic with the show fspf internal route vsan <vsan-id> command.

Example:

```
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x18(24)
Neighbor Interface is san-port-channel 200 (0x000400c7)
Statistics counters :
    Number of packets received : LSU 8 LSA 8 Hello 47 Error packets 4
    Number of packets transmitted : LSU 8 LSA 8 Hello 70 Retransmitted LSU 0
```

Number of times inactivity timer expired for the interface = 0

Possible Cause

There is a region mismatch on the switch.

The following example shows possible log messages from the **show logging** command log:

Example:

%FSPF-3-BAD_FC2_PKT: %\$VSAN 200%\$ Received bad FC2 packet on interface san-port-channel <channel-id> : Packet received for non existant region in VSAN

Solution

To identify a region mismatch problem on a switch using the NX-OS CLI, perform the following.

Step 1 Use the **show fspf vsan** *<vsan-id>* command to display the currently configured region in a VSAN.

Example (region value is 2; default region value is 0):

```
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 2
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x22(34)
Number of LSRs = 1, Total Checksum = 0x00000c10
Protocol constants :
  LS REFRESH TIME = 30 minutes (1800 sec)
  MAX_AGE
                  = 60 minutes (3600 sec)
Statistics counters :
  Number of LSR that reached MaxAge = 0
  Number of SPF computations
                                   = 0
                                    = 0
  Number of Checksum Errors
  Number of Transmitted packets : LSU 0 LSA 0 Hello 19 Retranmsitted LSU 0
  Number of received packets : LSU 0 LSA 0 Hello 0 Error packets 18
```

Step 2 Use the **debug fspf all** command and look for nonexistent region messages.

Example:

fspf: Hello timer reached for interface san-port-channel 200 in VSAN 200
fspf: FC2 packet received for non existant region 0 in VSAN 200
fspf: FC2 packet received for non existant region 0 in VSAN 200

The neighboring switch-advertising region is 0. FSPF is in the init state for each ISL.

Example:

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
    Number of packets received : LSU 0 LSA 0 Hello 0 Error packets 0
    Number of packets transmitted : LSU 0 LSA 0 Hello 49 Retransmitted LSU 0
Number of times inactivity timer expired for the interface = 9
```

- **Step 3** Use the **undebug all** command to turn off debugging.
- **Step 4** Use the **show fspf vsan** *<vsan-id>* command to show FSPF configuration and check the autonomous region.

Example:

```
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 2
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x22(34)
Number of LSRs = 1, Total Checksum = 0x00000c10
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing administration status is mabled
```

```
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x18(24)
Number of LSRs = 2, Total Checksum = 0x00014f9f
```

Step 5 Use the **fspf config vsan** command to enter the FSPF configuration mode and use the region command to change the region. The region must match on all switches in the VSAN.

Example:

```
switch(config)# fspf config vsan 200
switch(config-(fspf-config))# region 0
```

Fibre channel port is suspended due to too many invalid FLOGIs

A Fibre Channel node that is connected to an NPV feature-enabled Cisco Nexus 5000 switch or a Cisco Nexus 5000 switch that is running in fabric mode cannot log into the SAN fabric due to a FLOGI rejection.

The following example shows possible log messages from the **show logging** command log.

Example:

```
%FLOGI-1-MSG_FLOGI_REJECT_FCID_ERROR: %$VSAN <vsan-id>%$ [VSAN <vsan-id>, Interface
fcslot/port/: mode[F]] FLOGI rejected - FCID allocation failed.
PORT-5-IF_DOWN_TOO_MANY_INVALID_FLOGIS: %$VSAN <vsan-id>%$ Interface fc slot/port is down
(Suspended due to too many invalid flogis
```

The status of the interface shows invalidFlogis.

show interface fc slot/port brief
fc slot/port <vsan-id> F -- invalidFlogis

Possible Cause

The FC ID persistency table for that VSAN might be full. If the Nexus 5000 Series switch is configured as an NPV edge switch, the FC ID persistency table of the NPV core switch might be full.

FC IDs:

When an N port logs into a Cisco Nexus 5000 Series switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following situations can occur:

- An N port logs into a Cisco Nexus 5000 Series switch. The WWN of the requesting N port and the
 assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are
 not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted.

Solution

Check for FLOGI error messages with the show flogi internal command.

Example:

show flogi internal event-history debugs

```
222) Event:E_FLOGI_DEBUG, length:309, at 989582 usecs after Thu Jun 17
09:03:01 2010
fs_print_port_stats(10049): Port Stats for fc2/1, after cleanup:
    timestamp: Wed Jun 17 07:03:01 2010
    MSG_FLOGI: 52
    MSG_FC2_LS_RJT_OUT: 51
    EXCEPTION_CANNOT_ALLOCATE_FCID: 51
    EXCEPTION_TIMEOUT: 1
    EXCEPTION_FC2_INVALID_XCHG: 1
    tot_internal_exceptions: 51, since: Thu Dec 31 17:00:00 1969
    show flogi internal errors
52) Event:E_DEBUG, length:119, at 977471 usecs after Thu Jun 17 09:03:01
2010
    [102] Interface fc2/1, nwwn 20:01:00:1b:32:af:d6:8c, pwwn
21:01:00:1b:32:af:d6:8c: flogi is valid; exchange is INVALID.
```

Use the **show fcdomain address-allocation** command to check the FC domain address allocation table for any free FC IDs: (If NPV is enabled, enter the command on the NPV core switch.)

Example:

To find the auto area-list and the persistent FCIDs, use the **show flogi auto-area-list** command and the show fcdomain fcid persistent command.

Example:

```
show flogi auto-area-list
Fcid area allocation company id info:
<...>
   00:14:5E
   00:1B:32
   00:50:2E
   00:E0:69
   00:E0:8B
show fcdomain fcid persistent {entire AREA reserved for OUI 00:E0:8B}
102
       21:01:00:1b:32:2f:7f:63
                                   0x020003
                                               SINGLE FCID
                                                                YES
                                                                       DYNAMIC
102
       21:00:00:1b:32:0f:7f:63
                                   0x020004
                                               SINGLE FCID
                                                                YES
                                                                       DYNAMIC
102
       21:00:00:e0:8b:89:a7:07
                                   0x021c00
                                               ENTIRE AREA
                                                                YES
                                                                       DYNAMIC
       21:00:00:e0:8b:88:e9:22
```

If there are not enough FCIDs, you can purge dynamic and unused FC IDs in the specified VSAN with the purge fcdomain command.

ENTIRE AREA

YES

DYNAMIC

0x024300

Example:

102

switch# purge fcdomain fcid vsan <vsan-id>

The ports will soon come up.

It is also possible that HBAs are trying to login with $S_{ID} = 0x0$.

If this is the situation and there is nothing in the persistency table for the WWN of the HBA, try to assign the S_ID used by HBA to the HBA itself.

If the S_ID is already in use or is in the wrong domain, the request is rejected by fcdomain. After a number of retries, the port is suspended.

When HBAs get into this mode, they try to log in with every FCID in the FCID space, from 0x00.00.01 up to all the 0xDD.AA.PP numbers.

This behavior can be seen in the **show flogi internal event-history msgs** command output

{HBA is trying to login with different FCIDs}

Example:

- 841) Event:E_FLOGI_LRX, length:20, at 56079 usecs after Tue Jun 22 15:40:59 2010 WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x000032
- 886) Event:E_FLOGI_RX, length:20, at 897472 usecs after Tue Jun 22 15:40:58 2010 WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x000030
- 888) Event:E_FLOGI_FAIL, length:20, at 884758 usecs after Tue Jun 22 15:40:58 2010 WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 ev_id: 21 rjt reason: 7 OPC: MTS_OPC_DM_GET_FCIDS(275)
- 903) Event:E_FLOGI_RX, length:20, at 835015 usecs after Tue Jun 22 15:40:58 2010 WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x00002f

In this case, the solution is to manually configure an entry in the persistency table for the WWN of the HBA as shown in the following example. An alternative is to power-cycle the device. This usually makes the HBA start with a normal FLOGI with S_ID=0x0.

Example:

```
switch# conf t
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan <vsan-id> wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00 area
```

If the problem continues to persist, use the following commands to collect information to aid further analysis.

```
Show tech-support flogi
Show tech-support fcdomain
Show logging log
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show tech-support detail > bootflash:showtechdet
```

Capture debug flogi and debug fcdomain via following below steps:

```
switch# debug logfile flogi_fcdomain
switch# debug flogi all
switch# debug fcdomain all
switch(config)# int fc <slot/port>
switch(config-if)# shut
switch(config-if)# no shut
switch(config-if)# undebug all
switch# dir log: {check if you have the file in log: directory}
31 Aug 03 13:45:13 2010 dmesg
55941 Aug 05 07:21:15 2010 flogi_fcdomain
switch# copy log:flogi_fcdomain ftp://x.y.z.w {or use tftp/scp/sftp}
```

Having stale FCNS entries for Fibre Channel nodes

The Fibre Channel nodes are able to be logged (FLOGI) in to the SAN fabric, but the FCNS entries for those nodes are incomplete. Serves cannot reach their targets.

As a result, **fc4-types:fc4_features** will be empty in FCNS database.

Possible Cause

The Fibre Channel nodes may not be registering their FC4 types and FC4 features in the FCNS database in a topology where Nexus 5000 Series switches are configured as NPV core (feature NPIV) and connected to legacy gateway switches. The fc4-types:fc4_features can be verified by the **show fcns database detail** command as shown in the following example:

Example:

switch# show fcns da fcid 0x621400 detail vsan 2 VSAN:2 FCID:0x621400 _____ port-wwn (vendor) :21:01:00:1b:32:a3:d7:2c [z7095ib-1 T] :20:01:00:1b:32:a3:d7:2c node-wwn class :3 node-ip-addr :0.0.0.0 :ff ff ff ff ff ff ff ff ipa fc4-types:fc4_features symbolic-port-name : symbolic-node-name : port-type :N port-ip-addr :0.0.0.0 fabric-port-wwn :20:d9:00:0d:ec:e0:0e:80

hard-addr	:0x000000
permanent-port-wwn (vendor)	:20:11:00:05:1e:06:da:ea
Connected Interface	:fc2/2
Switch Name (IP address)	:N5K (10.200.220.13)

Some legacy gateway switches might require that the area part of the FCID be the same for the switch and for all the blades logged in through that port.

However, because of an old issue with Qlogic HBAs, the Cisco Nexus 5000 domain server assigns a separate area for each Qlogic HBA that matches a certain OUI by default. Therefore, a conflict between legacy gateway requirements and the Cisco domain allocation scheme exists. Cisco still implements this set up to support old existing Qlogic HBAs in the field.

Solution

Configure **no fcid-allocation area company** *<oui>* for all used Qlogic OUIs (ensuring flat FCID allocation in the future), force all affected blades to log out of the fabric, delete the already created persistent FCID entry from the Nexus 5000 switch configuration, and allow the blade to log in again.

In the following **show flogi database** command output, all devices obtain a unique area id (x01, x08, x0c):

Example:

Fc2/1	2	0x620104	20:10:00:05:1e:5e:6a:85	10:00:00:05:1e:5e:6a:85
Fc2/1	2	0x620800	21:01:00:1b:32:a3:c0:2e	20:01:00:1b:32:a3:c0:2e
Fc2/1	2	0x620c00	21:01:00:1b:32:33:8b:8e	20:01:00:1b:32:33:8b:8e

Because of the specific area ID requirement of the legacy switch, the last two blades must also have area x01. To force the Qlogic adapters to log in again and obtain FCID in 0x6201xx range, do the following steps:

Step 1 Configure (force) the future FCID allocation scheme to be flat for all WWNs matching the OUIs that are in this situation.

switch(configure)# no fcid-allocation area company 0x001B32

Step 2 Force the FCID under reconfiguration to log out of the fabric.

Note If you shut down the Nexus 5000 interface that serves as the primary uplink for that server, it only will log in through another one. The appropriate method is to shut down the affected blade and ensure that the FLOGI for the WWN is gone.

Step 3 Delete the automatically created configuration entry for persistent FCID allocation as shown in the following example:

Example:

switch(config)# fcdomain fcid database switch(config-fcid-db)# no vsan 2 wwn 21:01:00:1b:32:a3:c0:2e fcid 0x620800 area dynamic

Step 4 Bring up the blade and ensure that it gets a proper fcid.

Example:

Fc2/1	2	0x620104	20:10:00:05:1e:5e:6a:85	10:00:00:05:1e:5e:6a:85
Fc2/1	2	0x620123	21:01:00:1b:32:a3:c0:2e	20:01:00:1b:32:a3:c0:2e

Interface is isolated because of FC Domain ID overlap

The Fibre Channel or SAN port channel interface of a Cisco Nexus 5000 switch (in fabric mode) that is connected to an FC switch with the xE port type is isolated because of a domain overlap. The following example shows possible logging messages in the **show logging** command log:

Example:

PORT-5-IF_DOWN_DOMAIN_OVERLAP_ISOLATION: Interface fc <slot/port> is down (Isolation due to domain overlap). %FCDOMAIN-2-EPORT_ISOLATED: %\$VSAN <vsan-id>%\$ Isolation of interface san-port-channel <channel-id> (reason: domain ID assignment failure) %FCDOMAIN-2-EPORT_ISOLATED: %\$VSAN <vsan-id>%\$ Isolation of interface san-port-channel <channel-id> (reason: other side Eport indicates isolation)

Possible Cause

Two switch fabrics might not merge. If two fabrics with two or more switches are connected, have at least one assigned domain ID in common, and the auto-reconfigure option is disabled (this option is disabled by default), then the E ports that are used to connect the two fabrics will be isolated due to domain ID overlap.

In a Fibre Channel network, the principal switch assigns domain IDs when a new switch is added to an existing fabric. However, when two fabrics merge, the principal switch selection process determines which one of the preexisting switches becomes the principal switch for the merged fabric.

The election of the new principal switch is characterized by the following rules:

- A switch with a populated domain ID list takes priority over a switch that has an empty domain ID list. The principal switch becomes the one in the fabric with the populated domain ID list.
- If both fabrics have a domain ID list, the priority between the two principal switches is determined by the configured switch priority. This is a parameter that can be set by the user. The lower the value of the parameter, the higher the priority.
- If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the WWNs of the two switches. The lower the value of the WWN, the higher the switch priority.

<u>Solution</u>

To resolve an FC domain ID overlap, you can change the overlapping static domain ID by manually configuring a new static domain ID for the isolated switch, or disable the static domain assignment and allow the switch to request a new domain ID after a fabric reconfiguration.

To assign a static domain ID using the NX-OS CLI

All devices attached to the switch in the VSAN get a new FC ID when a new domain ID is assigned. Some hosts or storage devices may not function as expected if the FC ID of the host or storage device changes.

To verify FC domain ID overlap and reassign a new domain ID using the CLI, perform the following steps:

Step 1 Enter the show interface fc $\langle slot/port \rangle$ command to view the isolation error message for the E port.

Example:

```
switch(config)# show int fc 2/2
fc2/2 is down (Isolation due to domain other side eport isolated)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:42:00:0d:ec:d5:fe:00
```

```
Admin port mode is E, trunk mode is off
snmp link state traps are enabled
Port vsan is 3
```

Enter the **show interface san-port-channel** <channel-id> command to view the isolation error for the specific VSAN.

Example:

```
switch(config) # show interface san-port-channel 200
san-port-channel 200 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel
    Port WWN is 24:c8:00:0d:ec:d5:a3:80
   Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 8 Gbps
   Trunk vsans (admin allowed and active) (1,200)
                                            (1)
    Trunk vsans (up)
    Trunk vsans (isolated)
                                            (200)
    Trunk vsans (initializing)
                                            ()
```

Step 2 Use the **show fcdomain domain-list vsan** *<vsan-id>* command to view which domains are currently in your fabric.

Example (switch is isolated because of a domain ID 44 overlap):

switch(config)# show fcdomain domain-list vsan 3 Number of domains: 1 Domain ID WWN _____ _____ 0x2c(44) 20:03:00:0d:ec:3f:a5:81 [Local] [Principal] switch(config)# show fcdomain domain-list vsan 3 Number of domains: 1 Domain ID WWN _____ _ _ _ _ _ _ _ _ _ _ 20:03:00:0d:ec:d5:fe:01 [Local] [Principal] 0x2c(44)

If the isolation occurred for a specific VSAN under a SAN port channel interface, you can view the error with the **show port internal info interface san-port-channel** *<channel-id>* **vsan** *<vsan-id>* as shown in the following example:

switch(config)# show port internal info interface san-port-channel 200 vsan 200

Example:

```
san-port-channel 200, Vsan 200 - state(down), state reason(Isolation due to domain other
side eport isolated), fcid(0x00000)
port init flag(0x10000), num_active_ports (2),
Lock Info: resource [san-port-channel 200, vsan 200]
type[0] p_gwrap[(nil)]
FREE @ 159645 usecs after Thu Aug 5 13:35:00 2010
type[1] p_gwrap[(nil)]
FREE @ 159964 usecs after Thu Aug 5 13:35:00 2010
type[2] p_gwrap[(nil)]
FREE @ 450507 usecs after Tue Aug 3 14:14:08 2010
0x50c8efc7
current state [TE_FSM_ST_ISOLATED_DM_ZS]
RNID info not found.
first time elp: 0
Peer ELP Revision: 3
```

- **Step 3** Use the **fcdomain domain** *<domain-id>* [**static** | **preferred**] vsan *<vsan-id>* command to change the domain ID for one of the overlapping domain IDs.
 - The **static** option tells the switch to request that particular domain ID. If it does not obtain that particular address, it will isolate itself from the fabric.
 - The **preferred** option has the switch request a specified domain ID. If that ID is unavailable, it will accept another ID.

Step 4 Use the fcdomain restart vsan command to restart Domain Manager.

While the static option can be applied to runtime after a disruptive or nondisruptive restart, the preferred option is applied to runtime only after a disruptive restart.

Note

A domain ID restart is disruptive. The Fibre Channel nodes that are logged into that domain will be logged out and logged back in. A disruptive reconfiguration might affect data traffic.

To assign a dynamic domain ID after a fabric reconfiguration

You can use fabric reconfiguration to reassign domain IDs and resolve any overlapping domain IDs. If you enable the auto-reconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. The RCF functionality automatically forces a new principal switch selection and causes new domain IDs to be assigned to the different switches.

To use fabric reconfiguration to reassign domain IDs for a particular VSAN using the NX-OS CLI, perform the following these steps:

- **Step 1** Use the **show fcdomain domain-list** command to determine if you have statically assigned domain IDs on the switches.
- **Step 2** If you have statically assigned domain IDs, use the **no fcdomain domain** command to remove the static assignments.
- **Step 3** Use the show fcdomain vsan *<vsan-id>* command to determine if you have the RCF reject option enabled.

Example:

switch# show fcdomain vsan 3 The local switch is the Principal Switch.

Local switch run time information: State: Stable 20:03:00:0d:ec:d5:fe:01 Local switch WWN: Running fabric name: 20:03:00:0d:ec:d5:fe:01 Running priority: 128 Current domain ID: 0x2c(44) Local switch configuration information: State: Enabled FCID persistence: Enabled Auto-reconfiguration: Disabled Contiguous-allocation: Disabled Configured fabric name: 20:01:00:05:30:00:28:df Optimize Mode: Disabled Configured priority: 128 Configured domain ID: 0x2c(44) (preferred)

```
Principal switch run time information:

Running priority: 128

Interface Role RCF-reject

fc2/2 Isolated Enabled
```

Step 4 If you have the rcf-reject option enabled, use the **interface** command and then the **no fcdomain rcf-reject vsan** *<vsan-id>* command in interface mode.

Example:

```
switch(config)# interface fc 2/2
switch(config-if)# no fcdomain rcf-reject vsan 3
switch(config-if)#
```

- **Step 5** Use the **fcdomain auto-reconfigure vsan** *<vsan-id>* command in the EXEC mode on both switches to enable auto-reconfiguration after a Domain Manager restart.
- **Step 6** Use the **fcdomain restart vsan** *<vsan-id>* command to restart Domain Manager.

This is a disruptive operation and disruptive reconfiguration and can affect data traffic.

Cisco Fabric Services

This section includes an overview of troubleshooting Cisco Fabric Services (CFS) followed by a description of common problems and their solutions.

Overview

Begin troubleshooting CFS issues by checking the following:

- Verify that CFS is enabled for the same applications on all affected switches.
- Verify that CFS distribution is enabled for the same applications on all affected switches.

If the CFS Regions feature is in use, verify that the application is in the same region on all the affected switches.

- Verify that there are no pending changes for an application and that a CFS commit was issued for any configuration changes in a CFS-enabled application.
- Verify that there are no unexpected CFS locked sessions.

Clear any unexpected locked sessions.

Verifying CFS using CLI

To verify CFS using the CLI, follow these steps:

Step 1 By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation. To determine the state of CFS distribution on a switch, enter the show cfs status command. Example:

```
switch(config)# show cfs status
Distribution : Enabled
Distribution over IP : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
Distribution over Ethernet : Disabled
```

switch(config)# show cfs merge status name rscn

Step 2 To verify that an application is listed and enabled, issue the **show cfs application** command to all switches.

Example:

switch# show cfs application

Application	Enabled	Scope
fwm	Yes	Physical-eth
ntp	No	Physical-fc-ip
stp	Yes	Physical-eth
fscm	Yes	Physical-fc
role	No	Physical-fc-ip
rscn	No	Logical
radius	No	Physical-fc-ip
fctimer	No	Physical-fc
syslogd	No	Physical-fc-ip
callhome	No	Physical-fc-ip
fcdomain	No	Logical
device-alias	Yes	Physical-fc
		2

Total number of entries = 12

Note

The Physical scope means that CFS applies the configuration for that application to the entire switch. The Logical scope means that CFS applies the configuration for that application to a specific VSAN.

Step 3 Verify the set of switches in which an application is registered with CFS, using the show cfs peers name application-name command for physical scope applications, and the show cfs peers name <a price application-name > vsan <vsan-id> command for logical scope applications.

Example:

switch# show cf peers name device-alias

```
      Scope
      : Physical-fc

      Switch WWN
      IP Address

      20:00:00:0d:ec:da:6e:00
      172.25.183.124

      20:00:00:0d:ec:24:5b:c0
      172.25.183.123

      20:00:00:0d:ec:50:09:00
      172.25.183.42
```

```
Total number of entries = 3
```



Book Title

The **show cfs peers name** *< application-name >* command displays the peers for all VSANs when applied to a logical application.

Example:

switch(config) # show cfs peers name rscn

Scope	: Logical [VSAN 1]		
Domain	Switch WWN	IP Address	
106 98 238	20:00:00:0d:ec:da:6e:00 20:00:00:0d:ec:24:5b:c0 20:00:00:0d:ec:50:09:00	172.25.183.124 172.25.183.123 172.25.183.42	[Local]
Total nu	umber of entries = 3		
Scope	: Logical [VSAN 10]		
Domain	Switch WWN	IP Address	
82 5 83	20:00:00:0d:ec:da:6e:00 20:00:00:0d:ec:50:09:00 20:00:00:0d:ec:24:5b:c0	172.25.183.124 172.25.183.42 172.25.183.123	[Local]
Total nu	umber of entries = 3		
Scope	: Logical [VSAN 50]		
Domain	Switch WWN	IP Address	
66 28 235	20:00:00:0d:ec:da:6e:00 20:00:00:0d:ec:24:5b:c0 20:00:00:0d:ec:50:09:00	172.25.183.124 172.25.183.123 172.25.183.42	[Local]
Total nu	umber of entries = 3		
Scope	: Logical [VSAN 100]		
Domain	Switch WWN	IP Address	
90 100 111	20:00:00:0d:ec:da:6e:00 20:00:00:0d:ec:24:5b:c0 20:00:00:0d:ec:50:09:00	172.25.183.124 172.25.183.123 172.25.183.42	[Local]
Total nu	umber of entries = 3		

Step 4 To determine if all the switches in the fabric constitute one CFS fabric, or a multitude of partitioned CFS fabrics, enter the show cfs merge status name <application-name> command and the show cfs peers name <application-name> command and compare the outputs. If the two outputs contain the same list of switches, the entire set of switches constitutes one CFS fabric. When this is the case, the merge status should always show success at all switches.

Example:

switch(config)# show cfs merge status name rscn Logical [VSAN 1] Merge Status: Success [Thu Aug 5 11:33:50 2010] Local Fabric Domain Switch WWN IP Address 98 20:00:00:0d:ec:24:5b:c0 172.25.183.123 [Merge Master] 238 20:00:00:0d:ec:50:09:00 172.25.183.42 106 20:00:00:0d:ec:da:6e:00 172.25.183.124

```
switch
Total number of switches = 3
Logical [VSAN 10] Merge Status: Success [ Thu Aug 5 11:36:43 2010 ]
Local Fabric
Domain Switch WWN
                      IP Address
_____
83
     20:00:00:0d:ec:24:5b:c0 172.25.183.123
                                              [Merge Master]
5
     20:00:00:0d:ec:50:09:00 172.25.183.42
82
    20:00:00:0d:ec:da:6e:00 172.25.183.124
                      switch
Total number of switches = 3
Logical [VSAN 50] Merge Status: Success [ Thu Aug 5 11:36:23 2010 ]
Local Fabric
_____
Domain Switch WWN
                     IP Address
    _____
28
    20:00:00:0d:ec:24:5b:c0 172.25.183.123
                                             [Merge Master]
235 20:00:00:0d:ec:50:09:00 172.25.183.42
66 20:00:00:0d:ec:da:6e:00 172.25.183.124
                      switch
Total number of switches = 3
Logical [VSAN 100] Merge Status: Success [ Thu Aug 5 11:33:50 2010 ]
Local Fabric
_____
Domain Switch WWN
                     IP Address
_____
100
    20:00:00:0d:ec:24:5b:c0 172.25.183.123
                                             [Merge Master]
     20:00:00:0d:ec:50:09:00 172.25.183.42
111
90
     20:00:00:0d:ec:da:6e:00 172.25.183.124
                      switch
Total number of switches = 3
```

If the list of switches in the **show cfs merge status name** command output is shorter than that of the **show cfs peers name** command output, then the fabric is partitioned into multiple CFS fabrics and the merge status may show that the merge has failed, is pending, or is waiting.

Merge failure troubleshooting

During a merge, the merge managers in the merging fabrics exchange their configuration databases with each other. The application on one of the fabrics merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

When a merge is successful, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. A merge failure indicates that the merged fabrics contain inconsistent data that could not be merged.

If a new switch is added to the fabric and the merge status for any application shows In Progress for a prolonged period of time, then there may be an active session for that application in a switch. Check the lock status for that application on all the switches by using the **show cfs lock** command. If any locks exist, the merge does not proceed. Commit the changes or clear the session lock so that the merge proceeds.

Note

Merge failures must be analyzed correctly. Exercise caution when choosing a switch for blank commit. Small configurations may wipe out the large configurations.

Recovering from a Merge Failure with the CLI

To recover from a merge failure using the CLI, perform the following steps:

Step 1 To identify a switch that shows a merge failure, enter the **show cfs merge status name** <*application-name*> command.

Example:

switch(config) # show cfs merge status name ntp Physical-fc-ip Merge Status: Success [Thu Aug 5 11:47:58 2010] Local Fabric _____ Switch WWN IP Address _____ 20:00:00:0d:ec:da:6e:00 172.25.183.124 [Merge Master] switch Total number of switches = 1 switch(config) # show cfs merge status name ntp Physical-fc-ip Merge Status: Success [Thu Aug 5 11:43:39 2010] Local Fabric _____ Switch WWN IP Address _____ 20:00:00:0d:ec:50:09:00 172.25.183.42 [Merge Master] MDS-9134 20:00:00:0d:ec:da:6e:00 172.25.183.124 Total number of switches = 2

Step 2 For a more detailed description of the merge failure, enter the **show cfs internal session-history name** <*application name* > **detail** command.

Example:

 switch(config)# show cfs internal session-history name ntp

 Time Stamp
 Source WWN

 User Name
 Session ID

 Thu Aug 5 11:45:19 2010
 20:00:00:0d:ec:da:6e:00
 LOCK_ACQUIRED

 admin
 34684

 Thu Aug 5 11:45:19 2010
 20:00:00:0d:ec:da:6e:00
 COMMIT[2]

 admin
 34689

 Thu Aug 5 11:45:20 2010
 20:00:00:0d:ec:da:6e:00
 LOCK_RELEASED

```
admin
                            34684
Step 3
      Enter configuration mode and enter the commit command for the application to restore all peers in the
      fabric to the same configuration database.
      Example:
      switch(config) # ntp commit
      switch(config)# show cfs merge status name ntp
      Physical-fc-ip Merge Status: Success [ Thu Aug 5 11:51:02 2010 ]
       Local Fabric
       _____
       Switch WWN
                          IP Address
      _____
       20:00:00:0d:ec:50:09:00 172.25.183.42
                                                            [Merge Master]
       20:00:00:0d:ec:da:6e:00 172.25.183.124
                          switch
      Total number of switches = 2
      switch(config) # show cfs merge status name ntp
      Physical-fc-ip Merge Status: Success [ Thu Aug 5 11:51:02 2010 ]
       Local Fabric
      _____
       Switch WWN
                          IP Address
           _____
                                     _____
       20:00:00:0d:ec:50:09:00 172.25.183.42
                                                            [Merge Master]
                          MDS-9134
       20:00:00:0d:ec:da:6e:00 172.25.183.124
      Total number of switches = 2
```

Lock failure troubleshooting

In order to distribute a configuration in the fabric, a lock must first be acquired on all switches in the fabric. Once accomplished, a commit can be issued which distributes the data to all switches in the fabric before releasing the lock.

When a lock has been acquired by another application peer, you cannot commit new configuration changes. This is a normal situation and you should postpone any changes to an application until the lock is released. Use the troubleshooting steps in this section only if you believe the lock has not been properly released.

A lock occurs when an administrator configures a change for a CFS-enabled application. If two administrators on the same switch attempt to configure the same application, only one administrator is given the lock. The other administrator is prevented from making changes to that application until the first administrator commits a change or discards any changes. Use the show cfs lock name command to determine the name of the administrator who holds the lock for an application. You should check with that administrator before clearing the lock.

A CFS lock can also be held by another switch in your fabric. Use the show cfs peers name command to determine all the switches that participate in the CFS distribution for the application. Then use the show cfs lock name command on each switch to determine who owns the CFS lock for that application. You should check with that administrator before clearing the lock. Use the CFS abort option to release the lock without distributing the data to the fabric.
Resolving lock failure issues using the CLI

To resolve a lock failure using the CLI, perform the following steps:

Step 1 Enter the **show cfs lock name** < *name* > command to determine the lock holder.

Example:

switch(config) # show cfs lock name ntp

Scope : Physical-f	c-ip		
Switch WWN	IP Address	User Name	User Type
20:00:00:0d:ec:50:09:0	0 172.25.183.42	admin	CLI/SNMP v3

Total number of entries = 1

Step 2 For a detailed description of the lock failure, enter the show cfs internal session-history name application <name> detail command.

Example:

switch(config)# show cfs internal session-history name ntp detail

Time Stamp User Name	þ			Source WWN Session ID	Event
Thu Aug S admin	5	11:51:02	2010	20:00:00:0d:ec:da:6e:00 35035	LOCK_REQUEST
Thu Aug S admin	5	11:51:02	2010	20:00:00:0d:ec:da:6e:00 35035	LOCK_ACQUIRED
Thu Aug S admin	5	11:51:03	2010	20:00:00:0d:ec:da:6e:00 35040	COMMIT[2]
Thu Aug S admin	5	11:51:03	2010	20:00:00:0d:ec:da:6e:00 35035	LOCK_RELEASE_REQUEST
Thu Aug S admin	5	11:51:03	2010	20:00:00:0d:ec:da:6e:00 35035	LOCK_RELEASED
Thu Aug S admin	5	12:03:18	2010	20:00:00:0d:ec:50:09:00 284072	REMOTE_LOCK_REQUEST
Thu Aug S admin	5	12:03:18	2010	20:00:00:0d:ec:50:09:00 284072	LOCK_OBTAINED

Step 3 If the lock is being held by a remote peer, eneter the application-name commit command or an application-name abort command at that switch.

Example:

An example of the *<application-name>* commit command follows:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ntp commit
switch(config)#
```

Example:

An example of the *<application-name>* **abort** command follows:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ntp abort
switch(config)#
```

Γ

System state inconsistent and locks being held

An inconsistent system state occurs for one of the following situations:

- When locks are not held on all of the switches in the fabric.
- When locks are held on all switches in the fabric, but a session does not exist with the lock holding the switch.

In either case, it is necessary to use the clear option to release the locks.

Clearing locks using the CLI

When a lock is being held on a remote peer and entering the *<application-name>* commit command or the *<application-name>* abort command does not clear the lock, issue the clear *<application-name>* session command to clear all locks in the fabric. After all locks are cleared, a new distribution must be started to restore all the switches in the fabric to the same state.

Example:

```
switch# clear ntp session
switch# config terminal
switch(config)# ntp commit
switch(config)#
```

Distribution status verification

After configuring an application and committing the changes, you may want to verify that CFS is distributing the configuration change throughout the fabric or VSAN.

Verifying distribution using the CLI

Use the **show cfs lock name** <*application-name*> command to determine if a distribution is in progress on the fabric. If the application does not show in the output, the distribution has completed.

Example:

```
switch(config)# show cfs lock name ntp
Scope : Physical-fc-ip
Switch WWN IP Address User Name User Type
20:00:00:0d:ec:50:09:00 172.25.183.42 admin CLI/SNMP v3
Total number of entries = 1
```

CFS regions troubleshooting

The following rules apply to CFS Regions:

- When using CFS Regions, an application on a given switch can only belong to one region at a time.
- CFS Regions are only applicable to applications within the physical scope. You cannot create a CFS Region in the logical scope of an application.

REVIEW DRAFT-CISCO CONFIDENTIAL

- Assigning a region to an application takes precedence in distribution over its initial physical scope.
- CFS Regions configuration is not supported for deregistered applications (conditional services) or a physical scope application that is currently locked.
- Regions 1 through 200 are available for user configuration. Regions 201 through 255 are reserved regions and are not available for user configuration.

Distribution failure

To resolve a configuration distribution failure to all switches for a CFS Region, perform the following steps:

- **Step 1** Verify that application distribution is enabled. For more information, see "Overview" section on page 1-31.
- **Step 2** Verify that the application is in the same region on all switches.

Using the CLI from each switch, enter the **show cfs** *<application>* **name** *<application-name>* command.

Example (for device-alias application):

switch(config)# show cfs lock name ntp

;	Scope	:	Physical-fc-	-ip		
	Switch WWN	[IP Address	User Name	User Type
	20:00:00:0	d:	ec:50:09:00	172.25.183.42	admin	CLI/SNMP v3

Total number of entries = 1

Example (application is capable of being merged; application is in Region Default):

switch(config)# show cfs application name device-alias

Enabled	: Yes
Timeout	: 20s
Merge Capable	: Yes
Scope	: Physical-fc
Region	: Default

Example (application is capable of being merged; application is in Region 1):

```
switch# show cfs application name device-alias
Enabled : Yes
Timeout : 20s
Merge Capable : Yes
Scope : Physical-fc
Region : 1
```

Γ

Regions for conditional service

When a conditional service goes down (deregisters with CFS), it loses its region configuration. When the conditional service is restarted, it is automatically placed into the default region. To avoid this situation, reconfigure the appropriate region information for the conditional service before starting it again.

Changing regions

If you move an application from one region to another, you might encounter a database mismatch when attempting a merge. To identify and resolve the conflicts, see "Merge failure troubleshooting" section on page 1-34.



When an application is moved from one region to another (including the default region), it loses all history.

VSANs

This section includes an overview of troubleshooting VSANs followed by a description of common problems and their solutions.

Overview

Most VSAN problems can be avoided by following the best practices for VSAN implementation.

However, if necessary, you can use the fabric analysis tool in Fabric Manager to verify different categories of problems such as VSANs, zoning, FC domain, admin issues, or switch-specific or fabric-specific issues.

Fabric Manager provides the configuration consistency check tool.

To use the Fabric Configuration option to analyze the configuration of a switch, follow these steps:

Step 1	From the Fabric Manager tools menu, choose Health > Fabric Configuration .
	The Fabric Configuration Analysis dialog box appears.
Step 2	Determine whether you want to compare the selected switch to another switch or to a policy file.
	• To compare the selected switch to another switch, select Policy Switch and then select a switch from the drop-down list of switches.
	• To make a policy file comparison, select Policy File and then click the button on the right to browse your file system and select a policy file (*.XML).
Step 3	Click Rules to set the rules to apply when running the Fabric Configuration Analysis tool.
	The Rules window appears.
Step 4	Change the existing rules as appropriate and click OK .
Step 5	Click Compare to have the system to compare the configuration.
	The results of the analysis are displayed.

- **Step 6** In the Resolve column, select the issues that you want to resolve.
- **Step 7** Click **Resolve Issues** to resolve the identified issues.
- **Step 8** Click **Clear** to remove the contents of the window.
- **Step 9** Click **Close** to complete the operation and close the window.

For more information about the configuration consistency check tool, see the *Cisco DCNM* Fundamentals Guide, Release 5.x.

Note

When suspending or deleting VSANs, make sure that you suspend and unsuspend one VSAN at a time. You should wait a minimum of 60 seconds after you enter the **vsan suspend** command before you enter any other configuration command. If you fail to wait, some Fibre Channel interfaces or member ports in a port channel might become suspended or error-disabled.

Troubleshooting a SAN problem involves gathering information about the configuration and connectivity of individual devices as well as the status of the entire SAN fabric.

VSAN Troubleshooting Activities

Common troubleshooting tools in Fabric Manager

Verify the VSAN with the following Fabric Manager procedures:

- To view the VSAN configuration in the Information pane, select Fabricxx > VSANxx.
- To view the VSAN members, select Fabricxx > VSANxx, then click the Host or Storage tab in the Information pane.
- To view the FC domain configuration in the Information pane, select Fabricxx > VSANxx > Domain Manager.

Common troubleshooting CLI commands

Use the following CLI commands to display VSAN, FC domain, and FSPF information:

```
show vsan
show vsan <vsan-id>
show vsan membership
show interface fc <slot/port> trunk <vsan-id>
show <vsan-id> membership
show vsan membership interface fc <slot/port>
```

Checklist

Check for the following:

- Verify the domain parameters for switches in the VSAN.
- Verify the physical connectivity for any problem ports or VSANs.
- Verify that both devices are in the name server.
- Verify that both end devices are in the same VSAN.
- Verify that both end devices are in the same zone.

L

Nexus 5000 trunk port does not connect to upstream SAN switch

The Nexus 5000 trunk port does not connect to the upstream SAN switch because:

- Status of the trunk port connected to the upstream switch is isolated.
- The switch port trunk mode is enabled on both sides.
- Physical cabling has been checked and verified.
- Ports are up on both switches.

By examining the interface state and querying the interface, the issue is displayed as shown in the following example.

Example:

switch(config-if)# show interface brief

Interface Vsan Admin Admin Status SFP Oper Oper Port Mode Speed Channel Mode Trunk (Gbps) Mode fc2/3 1 Е on isolated swl _ _ switch(config-if) # show interface fc 2/3 fc2/3 is down (Isolation due to no common vsans with peer on trunk) Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN) Port WWN is 20:43:00:0d:ec:da:6e:00 Admin port mode is E, trunk mode is on

Possible Cause

The VSAN allow list for both interfaces is not the same. Specifically, there is no common VSAN allowed on both interfaces.

This situation might be caused by the following:

- No common VSANs on both switches.
- The trunk allowed VSAN members that do not contain common members.

In the example, the trunk VSAN allow list on the Nexus 5000 and MDS FC interfaces do not match.

<u>Solution</u>

Determine the connected ports and resolve the allowed VSANs on the trunk for both FC interfaces.

Example:

```
switch(config-if)# show run interface fc 2/3
!Command: show running-config interface fc2/3
!Time: Wed Aug 4 16:06:04 2010
version 4.2(1)N1(1)
interface fc2/3
switchport mode E
switchport trunk allowed vsan 1
no shutdown
switch(config-if)# show run interface fc 1/1
!Command: show running-config interface fc1/1
!Time: Wed Aug 4 16:20:07 2010
```

version 5.0(1a)

```
interface fc1/1
  switchport rate-mode dedicated
  switchport mode E
  switchport trunk allowed vsan 100
 no shutdown
switch(config-if)# interface fc 2/3
switch(config-if) # switchport trunk allowed vsan
add
    al]
switch(config-if)# switchport trunk allowed vsan add 100
switch(config-if) # show run interface fc 2/3
!Command: show running-config interface fc2/3
!Time: Wed Aug 4 16:07:25 2010
version 4.2(1)N1(1)
interface fc2/3
 switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 100
 no shutdown
switch(config-if) # switchport trunk allowed vsan add 1
switch(config-if)# show run interface fc 1/1
!Command: show running-config interface fc1/1
!Time: Wed Aug 4 16:20:54 2010
version 5.0(1a)
interface fc1/1
  switchport rate-mode dedicated
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 100
  no shutdown
fc2/3 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:43:00:0d:ec:da:6e:00
    Peer port WWN is 20:01:00:0d:ec:24:5b:c0
   Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
   Port vsan is 1
    Speed is 4 Gbps
   Transmit B2B Credit is 250
    Receive B2B Credit is 16
    Receive data field Size is 2112
   Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100)
    Trunk vsans (up)
                                           (1, 100)
    Trunk vsans (isolated)
                                           ()
    Trunk vsans (initializing)
                                            ()
switch(config-if)# show interface brief
Interface Vsan Admin Admin Status
                                                 SFP
                                                        Oper Oper
                                                                     Port
                  Mode Trunk
                                                        Mode Speed Channel
```

			Mode				(Gbps)	
fc2/3	1	Е	on	trunking	swl	TE	4	

Nexus 5000 E port (non-trunking) does not connect to upstream SAN switch

The Nexus 5000 E port does not connect to the upstream SAN switch because:

- The status of the interconnected non-trunking E ports shows that the status is up. However, all Fibre Channel services are not working between the switches.
- Devices in the same VSAN do not appear in the FCNS database for both switches.
- The show topology command does not list peer switch information.
- Zones show members are not logged in.

Example:

switch(config-vsan-db)# show interface brief

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel
fc2/4	50	Е	off	up	swl	E	2	
switch(con	fig-if)	# show	interfac	e brief				
Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel
fclfc1/2	100	 E	off	up	SW	1 E		2
The FC top	ology do	bes not s	how a va	lid peer interface.				
Example:								
switch(con	fig-if)	# show	topo					
FC Topolog	y for V	SAN 100	:					
Int	erface	Peer D	omain Pe	er Interface	Peer I	P Addro	ess	
	fc1/2	0x42(66)	Port 65795 ::				
The zonese	t shows	one me	mber is	not active				
switch(config-vsan-db)# show zoneset active vsan 100 zoneset name ZoneSet_Host_Storage vsan 100 zone name Zone_Host_Storage vsan 100 pwwn 20:00:00:25:b5:00:20:0e [Host]								
* ICId 0 switch(con	fig-if)	lpwwn # show	zoneset	:81:86:/8:39:66] active vsan 100	lStora	ge]		

- zoneset name ZoneSet_Host_Storage vsan 100
 - zone name Zone_Host_Storage vsan 100
 * fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
 pwwn 50:0a:09:81:86:78:39:66 [Storage]

The storage and hosts are in the correct VSAN.

Example:

<pre>switch(config-vsan-db)# show flogi database vsan 100</pre>										
INTERFACE	VSAN	FCID	PORT NAME	NODE NAME						
fc2/2	100	0x5a0000 [Stor	50:0a:09:81:86:78:39:66 age]	50:0a:09:80:86:78:39:66						
switch(config-if)# sho	w flogi da	tabase vsan 100							
INTERFACE	VSAN	FCID	PORT NAME	NODE NAME						
fc4/2	100	0x640114 [Host	20:00:00:25:b5:00:20:0e	20:00:00:25:b5:02:02:09						

Possible Cause

The error is displayed by the **show interface brief** command and the **show vsan membership** command. They show that the E port on one switch is in the wrong VSAN.

The non-trunking E port on one switch is in the wrong VSAN. (VSAN 100 is the correct VSAN.)

Example:

switch(config-if)# show interface brief

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel
fc1fc1/2	100	Е	off	up	swl	E	2	

Solution

Move the non-trunking E-port into VSAN 100.

Example:

switch(config-vsan-db)# vsan 100 interface fc 2/4Traffic on fc2/4 may be impacted. Do you want to continue? (y/n) [n] y

The zone set is now active and the FC topology is correct.

Example:

Г

Communication problem between host and storage devices

The communication problem between host and storage devices is because:

- Zones are active.
- Both host and storage are logged into the SAN.
- The storage port is not logged into the active zone set.

Example:

```
zoneset name ZoneSet_Host_Storage vsan 100
zone name Zone_Host_Storage vsan 100
```

```
* fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
pwwn 50:0a:09:81:86:78:39:66 [Storage]
```

Possible Cause

The host or storage port are in the wrong VSAN.

Example:

switch(config)# show fcns database

VSAN 50:

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x420000	N	50:0a:09:81:86:78:39:66 [Storage]	(NetApp)	scsi-fcp:target

Solution

Move the storage port to the correct VSAN. (VSAN 100 is the correct VSAN in the example.).

Example:

switch(config)# show flogi database vsan 50

```
_____
                              PORT NAME
INTERFACE
            VSAN FCID
                                                  NODE NAME
_____
            _____
fc2/2 50 0x420000 50:0a:09:81:86:78:39:66 50:0a:09:80:86:78:39:66
                     [Storage]
Total number of flogi = 1.
switch(config)# vsan database
switch(config-vsan-db)# vsan 100 interface fc 2/2
Traffic on fc2/2 may be impacted. Do you want to continue? (y/n) [n] y
switch(config-vsan-db)# show zoneset active vsan 100
zoneset name ZoneSet_Host_Storage vsan 100
 zone name Zone_Host_Storage vsan 100
 * fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
 * fcid 0x5a0000 [pwwn 50:0a:09:81:86:78:39:66] [Storage]
```

VSAN is down between switches

The VSAN is down between switches because:

- VSAN is configured on both switches.
- Trunk allow list allows the VSAN.

- VSAN reported to be down (Initializing state).
- Zones are active.
- Both host and storage are logged into the SAN.

In this failure, the storage port is not logged into the active zone set.

After examining the interface, the error can be seen as in the following example.

Example:

```
switch(config-if)# show interface fc 2/4 trunk vsan 10
fc2/4 is trunking
    Vsan 10 is down (Isolation due to domain id assignment failure)
switch(config-if)# show port internal info interface fc 2/4 | grep Isolation
    fc2/4, Vsan 10 - state(down), state reason(Isolation due to domain id assignment
failure), fcid(0x000000)
    fc2/4, Vsan 50 - state(down), state reason(Isolation due to vsan not configured on
peer), fcid(0x000000)
```

Possible Cause

The VSANs might have the same static Domain ID configured.

Example:

switch(config-if)# show fcdomain domain-list vsan 10

Number of domains: 1 Domain ID WWN ------ 0x53(83) 20:0a:00:0d:ec:da:6e:01 [Local] [Principal] switch(config)# show fcdomain domain-list vsan 10

<u>Solution</u>

Change the Domain ID on one of the VSANSs.

Example:

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 10 suspend
switch(config-vsan-db)# no vsan 10 suspend
switch(config-vsan-db)# show interface fc 2/4
Number of domains: 1
Domain ID
                      WWN
_____
            _____
          20:0a:00:0d:ec:da:6e:01 [Local] [Principal]
 0x52(82)
switch(config-vsan-db)# sho interface fc 2/4 | begin Trunk
   Trunk vsans (admin allowed and active) (1,10,50,100)
   Trunk vsans (up)
                                         (1, 10, 50, 100)
   Trunk vsans (isolated)
                                         ()
   Trunk vsans (initializing)
                                         ()
```

L

Registers and Counters

Identifying physical layer issues

To troubleshoot physical layer issues with Fibre Channel SFP optics, use the following command:

```
switch# show interface fc x/y transceiver details
```

In the following example, you can see that the results of the command contains useful information such as supported speed, nominal bit rate, and link lengths supported for the SFP.

Example:

```
switch# show interface fc 3/1 transceiver details
fc3/1 sfp is present
    name is CISCO-FINISAR
    part number is FTLF8524P2BNL-C2
    revision is 0000
    serial number is FNS0928K161
    fc-transmitter type is short wave laser w/o OFC (SN)
    fc-transmitter supports intermediate distance link length
    media type is multi-mode, 62.5m (M6)
    Supported speed is 400 MBytes/sec
    Nominal bit rate is 4300 MBits/sec
    Link length supported for 50/125mm fiber is 150 m(s)
    Link length supported for 62.5/125mm fiber is 70 m(s)
    cisco extended id is unknown (0x0)
    no tx fault, no rx loss, in sync state, Diag mon type 104
```

The command also provides detailed SFP diagnostic information and warnings and alarms, if any. Example:

SFP Detail Diagnostics Information

			ms		Warnings				
		High		Low		High		Low	
Temperature	41.50 C	95.00	C	-25.00	C	90.00	C	-20.00	C
Voltage	3.45 V	3.90	V	2.70	V	3.70	V	2.90	V
Current	7.18 mA	17.00	mA	1.00	mA	14.00	mA	2.00	mA
Tx Power	-4.41 dBm	-2.00	dBm	-11.74	dBm	-2.00	dBm	-11.02	dBm
Rx Power	-4.40 dBm	1.00	dBm	-20.00	dBm	-1.00	dBm	-18.24	dBm
Transmit Fa	ult Count = 0								
Note: ++ h	igh-alarm; +	high-warn:	ing;	low-	-alar	rm; - 10	sw-wa	rning	

Two example outputs from the command follow. The first shows a low alarm for Rx Power. The second shows low alarms for Tx, Rx, and Current. The interface for the second example was in an Error Disabled state due to the bit error rate being too high.

		Alar	ms	Warni	ngs
		High	Low	High	Low
Temperature	35.02 C	70.00 C	0.00 C	70.00 C	0.00 C

Low Alarm for RxPower

Voltage	0.00 V	0.00 V	0.00 V	0.00 V	0.00 V
Current	7.22 mA	16.00 mA	2.00 mA	14.00 mA	2.40 mA
Tx Power	-0.57 dBm	1.00 dBm	-8.21 dBm	0.00 dBm	-7.21 dBm
Rx Power	-18.86 dBm	1.00 dBm	-16.58 dBm	0.00 dBm	-14.44 dBm

Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning

Low Alarms for Current, Tx Power and R x Power

				Alarms			Warnings				
				High		Low		High		Low	
Temperature	32.75	с С		70.00	 С	0.00	 С	70.00	с С	0.00	- С
Voltage	0.00	v		0.00	v	0.00	v	0.00	v	0.00	v
Current	0.00	mA		16.00	mA	2.00	mA	14.00	mA	2.40	mA
Tx Power	N/A			1.00	dBm	-8.21	dBm	0.00	dBm	-7.21	dBm
Rx Power	-22.22	dBm		1.00	dBm	-16.58	dBm	0.00	dBm	-14.44	dBm
Note: ++ h:			 + ł	nigh-warn:	ina;	low-	 -alar		 ow-wa	arning	-

In the following example that the command does not provide detailed transceiver information for Twinax (copper).

```
switch# sh interface ethernet 1/19 transceiver details
Ethernet1/19
  sfp is present
  name is Molex Inc.
  part number is 74752-1301
  revision is E
  serial number is 733010037
  nominal bitrate is 0 MBits/sec
  Link length supported for 50/125mm fiber is 0 m(s)
  Link length supported for 62.5/125mm fiber is 0 m(s)
  cisco id is --
  cisco extended id number is 4
```

Invalid calibration

Displaying FcoE bound Ethernet interface counters

The show interface ethernet command, has two versions, brief and detailed. Examples of each follow.

Brief Version

Example:



Ensure that the jumbo frames are incrementing, as well as RX or TX pause frames counters, if any. Tx might indicate a congestion problem.

```
switch# show interface ethernet 1/4
Ethernet1/4 is up
Hardware: 1000/10000 Ethernet, address: 000d.ecd5.a38b (bia 000d.ecd5.a38b)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
```

Г

```
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s, media type is 10g
[snip]
RX
 9507 unicast packets 918874 multicast packets 3473 broadcast packets
 931854 input packets 76225281 bytes
 7121 jumbo packets 0 storm suppression packets
 0 runts 0 giants 0 CRC 0 no buffer
 0 input error 0 short frame 0 overrun
                                          0 underrun 0 ignored
 0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
 0 input with dribble 0 input discard
 0 Rx pause
ТΧ
 3986 unicast packets 294583 multicast packets 36307 broadcast packets
 334876 output packets 46873259 bytes
 1227 jumbo packets
 0 output errors 0 collision 0 deferred 0 late collision
 0 lost carrier 0 no carrier 0 babble
 2266 Tx pause
24 interface resets
```

Detailed Version

The output of the detailed version of the **show interface ethernet** command is shown in different parts in the following example. It includes both normal traffic counters, as well as counters for physical layer and protocol errors. These counters should be monitored anytime there is a connectivity or performance issue.

Example:

switch# sh interface ethernet 1/4 counters detailed all Ethernet1/4

InMcastPkts = 918901

64 bit	counters:		
0.	rxHCTotalPkts	=	931881
1.	txHCTotalPks	=	335522
2.	rxHCUnicastPkts	=	9507
3.	txHCUnicastPkts	=	3986
4.	rxHCMulticastPkts	=	918901
5.	txHCMulticastPkts	=	295229
6.	rxHCBroadcastPkts	=	3473
7.	txHCBroadcastPkts	=	36307
8.	rxHCOctets	=	76228116
9.	txHCOctets	=	46926647
10.	rxTxHCPkts640ctets	=	1065359
11.	rxTxHCpkts65to1270ctets	=	105246
12.	rxTxHCpkts128to2550ctets	=	43798
13.	rxTxHCpkts256to5110ctets	=	13822
14.	rxTxHCpkts512to10230ctets	=	30742
15.	rxTxHCpkts1024to15180ctets	=	88
16.	rxTxHCpkts1519to15480ctets	=	0
17.	rxHCTrunkFrames	=	895722
18.	txHCTrunkFrames	=	69387
19.	rxHCDropEvents	=	0
All Port	Counters:		
0.	InPackets	=	931881
1.	InOctets	=	76228116
2.	InUcastPkts	=	9507

3.

4.	InBcastPkts	=	3473
5.	InJumboPkts	=	7121
6.	StormSuppressPkts	=	0
7.	OutPackets	=	335522
8.	OutOctets	=	46926647
9.	OutUcastPkts	=	3986
10.	OutMcastPkts	=	295229
11.	OutBcastPkts	=	36307
12.	OutJumboPkts	=	1227
13.	rxHCPkts640ctets	=	889975
14.	rxHCPkts65to1270ctets	=	26702
15.	rxHCPkts128to2550ctets	=	6072
16.	rxHCPkts256to5110ctets	=	1913
17.	rxHCpkts512to10230ctets	=	11
18.	rxHCpkts1024to15180ctets	=	87
19.	rxHCpkts1519to15480ctets	=	0
20.	txHCPkts640ctets	=	175384
21.	txHCPkts65to1270ctets	=	78544
22.	txHCPkts128to2550ctets	=	37726
23.	txHCPkts256to5110ctets	=	11909
24.	txHCpkts512to10230ctets	=	30731
25.	txHCpkts1024to15180ctets	=	1
26.	txHCpkts1519to15480ctets	=	0
27.	ShortFrames	=	0
28.	Collisions	=	0
29.	SingleCol	=	0
30.	MultiCol	=	0
31.	LateCol	=	0
32.	ExcessiveCol	=	0
33.	LostCarrier	=	0
34.	NoCarrier	=	0
35.	Runts	=	0
36.	Giants	=	0
37.	InErrors	=	0
38.	OutErrors	=	0
39.	InputDiscards	=	0
40.	BadEtypeDrops	=	0
41.	IfDownDrops	=	0
42.	InUnknownProtos	=	0
43.	txErrors	=	0
44.	rxCRC	=	0
45.	Symbol	=	0
46.	txDropped	=	0
47.	TrunkFramesTx	=	69387
48.	TrunkFramesRx	=	895722
49.	WrongEncap	=	0
50.	Babbles	=	0
51.	Watchdogs	=	0
52.	ECC	=	0
53.	Overruns	=	0
54.	Underruns	=	0
55.	Dribbles	=	0
56.	Deferred	=	0
57.	Jabbers	=	0
58.	NoBuffer	=	0
59.	Ignored	=	0
60.	bpdu0utLost	=	0
61.	cos00utLost	=	0
62.	cos10utLost	=	U
63.	cos20utLost	=	U
64.	cos30utLost	=	U
65.	cos40utLost	=	U
66.	cos50utLost	=	U
67.	cos60utLost	=	U

68.	cos70utLost	=	0
69.	RxPause	=	0
70.	TxPause	=	2266
71.	Resets	=	0
72.	SQETest	=	0
73.	InLayer3Routed	=	0
74.	InLayer3RoutedOctets	=	0
75.	OutLayer3Routed	=	0
76.	OutLayer3RoutedOctets	=	0
77.	OutLayer3Unicast	=	0
78.	OutLayer3UnicastOctets	=	0
79.	OutLayer3Multicast	=	0
80.	OutLayer3MulticastOctets	=	0
81.	InLayer3Unicast	=	0
82.	InLayer3UnicastOctets	=	0
83.	InLayer3Multicast	=	0
84.	InLayer3MulticastOctets	=	0
85.	InLayer3AverageOctets	=	0
86.	InLayer3AveragePackets	=	0
87.	OutLayer3AverageOctets	=	0
88.	OutLayer3AveragePackets	=	0

Understanding Fibre Channel interface counters

The **show interface** command is very useful when troubleshooting physical layer or performance issues with a Fibre Channel interface.

In the output of the command, observe the input/output counters and any input/output discards or errors.

When input discards increment, the FC packet does not have a valid route in the Forwarding Information Base (FIB). All packets not having a route are considered discards and sent to the supervisor. These packets are NOT dropped; however, they are policed before being sent to the supervisor. You should also check the MAC ASIC for errors.

When output discards increment, packets are timing out in egress because the egress is too slow. Check the attached device because it may be a slow draining receiver that is not responding, or not replenishing buffer credits. This causes back pressure to occur on the Nexus 5000 FC interface.

Example:

```
switch# show interface fc2/1
fc2/1 is trunking
   Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
   Port WWN is 20:41:00:0d:ec:a4:02:80
[snip]
   1 minute input rate 5048 bits/sec, 631 bytes/sec, 9 frames/sec
   1 minute output rate 6752 bits/sec, 844 bytes/sec, 9 frames/sec
      36398816 frames input, 2422447564 bytes
       0 discards, 0 errors
       0 CRC, 0 unknown class
       0 too long, 0 too short
      36368010 frames output, 3213593392 bytes
       0 discards, 0 errors
      1 input OLS, 1 LRR, 0 NOS, 0 loop inits
     1 output OLS, 2 LRR, 0 NOS, 0 loop inits
      16 receive B2B credit remaining
      250 transmit B2B credit remaining
      0 low priority transmit B2B credit remaining
    Interface last changed at Thu Jan 28 18:26:30 2010
```

Troubleshooting Fibre Channel MAC issues

The **show hardware** command is very useful when troubleshooting FC physical layer issues.

```
Show hardware internal fc-mac <x> port <y> statistics
```

The output of the command contains the following useful information:

• Physical layer information

FCP_CNTR_MAC_RX_LOSS_OF_SYNC - Loss of Sync received counter

• Performance information

FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ - Receiver Ready's Sent FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY - Receiver Ready's Received

• Class 3 normal traffic counters

FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES FCP_CNTR_MAC_DATA_TX_CLASS5_WORDS FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS

• Fibre Channel primitive sequences

FCP_CNTR_LINK_RESET_IN - Link Resets Received FCP_CNTR_OLS_OUT- Offline Sequences Sent FCP_CNTR_NOS_OUT - Not Operational Sequence Sent FCP_CNTR_LRR_OUT - Link Reset Responses Sent FCP_CNTR_LINK_FAILURE

Example:

switch# show hardware internal fc-mac 2 port 1 statistics
ADDRESS STAT

0x5	FCP_CNTR_MAC_RX_LOSS_OF_SYNC	0x000003c
0xec	FCP_CNTR_MAC_RX_BAD_WORDS_FROM_DECODER	0x000003d
0x5ec	FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ	0x0000042
0xc41	FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY	0x0000043
0x5d2	FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES	0x00000061
0x1a	FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES	0x0000062
0x140b14	FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS	0x0000069
0xdcc	FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS	0x0000006a
0xc24	FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES	0x0000065
0x1d	FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES	0x0000066
0x4b9538	FCP_CNTR_MAC_DATA_TX_CLASS3_WORDS	0x000006d
0xabc	FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS	0x000006e
0x2	FCP_CNTR_LINK_RESET_IN	0xfffffff
0x5	FCP_CNTR_OLS_OUT	0xfffffff
0x2	FCP_CNTR_NOS_OUT	0xfffffff
0x7	FCP_CNTR_LRR_OUT	0xfffffff
0x2	FCP_CNTR_LINK_FAILURE	0xfffffff

COUNT

Г

When troubleshooting FC performance problems, review the R_RDY, Link Reset, and Link Reset Response counters. These help determine buffer to buffer credit problems that could lead to performance issues.

Troubleshooting Fibre Channel forwarding issues

To troubleshoot Fibre Channel forwarding issues, it is important to know that GATOS is the MAC/Forwarding ASIC on the Nexus 5000 switch. This section describes commands specific to this ASIC.

Each Fibre Channel interface is assigned a GATOS number. To understand forwarding issues, you must find the GATOS number for the specific FC interface.

Consider the following example:

switch# sh platform fwm info pif fc2/1 dump pif info: ifindex 0x1080000 dump_all 0 verbose 1 fc2/1: slot 1 port 0 state 0x0 pi_if 0x88bbb74 fwimpd ctx 0x889d4ec fc2/1: oper_mode 0x1 rcvd_rbind: No fc2/1: iftype 0x1 encap 0x5 bound_if? N #lifs 1 fwimpd ctx 0x88bd74c fc2/1: lif_blk(pi) 0x8523da4 vif_id_alloc_bmp 0x887360c fc2/1: cfg_lif_blk_size 0 lif_blk_base(pi) 1922 lif_blk_size(pi) 1 fc2/1: cfg_lif_blk_size(pi) 0 fc2/1: if_flags 0x0 num_sub_lif_tbls 0 Num HIFs pinned 0 fc2/1 pd: lif_entries 1 if_map_idx 49 if_lid 33 if_fcoe_lid 34 fc2/1 pd: reverse ifmap lookup 'same' ifmap_idx 49 fc2/1: SAT_HIF Port?: No

In the following part of the example, notice that gatos_num 13 is the GATOS instance for Fibre Channel interface 2/1:

```
fc2/1 pd: slot 1 logical port num 4 gatos_num 13 fwm_inst 0 fc 0
fc2/1 pd: pif_type 'data fc'(2) hw_present 1 port map idx 49
fc2/1 pd: fabric a info: voq 0-1 port_id 29 connected 1 up 1
fc2/1 pd: fabric b info: voq 0-1 port_id 29 connected 1 up 1
fc2/1 pd: subported 1 primary 1 atherton 0
fc2/1 pd: sup_src_dst_if 17 lif_blk 0-0
fc2/1 pd: policer info: uc (sel 2) mc (sel 1) bc (sel 0)
fc2/1 pd: mac-addr 000d.eca4.02b4
```

In the following part of the example, notice that the command also provides forwarding drop and discard informatio:

```
fc2/1 pd: tx stats: bytes 4958178736 frames 36360131 discard 0 drop 0 fc2/1 pd: rx stats: bytes 2421909296 frames 36390924 discard 0 drop 0
```

You can also display GATOS errors for the GATOS instance that corresponds to the FC interface. In the following part of the example, notice that the command only shows non-zero counters.

```
switch# show platform fwm info gatos-errors 13
Printing non zero Gatos error registers:
DROP_FCF_SW_VSAN_IDX_MISS: res0 = 60 res1 = 0
DROP_FCF_SW_DOMAIN_IDX_MISS: res0 = 489036 res1 = 0
DROP_FCF_SW_TBL_MISS: res0 = 489036 res1 = 0
DROP_NO_FABRIC_SELECTED: res0 = 489036 res1 = 0
DROP_VLAN_MASK_TO_NULL: res0 = 489036 res1 = 0
```

In the first of the two parts, drops and discards were described. Notice that the drop and discard counters are separate for vEthernet and VFC interfaces. Review the output in the second example to help correlate the reason for the drops.

```
switch# show platform fwm info pif ethernet 1/4
dump pif info: ifindex 0x1a003000 dump_all 0 verbose 1
Eth1/4: slot 0 port 3 state 0x0 pi_if 0x876acb4 fwimpd ctx 0x876171c
Eth1/4: oper_mode 0x100000 rcvd_rbind: No
Eth1/4: iftype 0x1 encap 0x1 bound_if? Y #lifs 1 fwimpd ctx 0x879f70c
Eth1/4: lif_blk(pi) 0x87cc9a4 foo vif_id_alloc_bmp 0x88313f4
Eth1/4:
          0
Eth1/4: cfg_lif_blk_size 0 lif_blk_base(pi) 512 lif_blk_size(pi) 128
Eth1/4: cfg_lif_blk_size(pi) 0
Eth1/4: if_flags 0x0 num_sub_lif_tbls 0 Num HIFs pinned 0
Eth1/4: max_hifpc_mbrs 0, max_hif_ports 0
Eth1/4 pd: lif_entries 1 if_map_idx 8 if_lid 35 if_fcoe_lid 36
Eth1/4 pd: reverse ifmap lookup 'same' ifmap_idx 8
Eth1/4: SAT_HIF Port?: No
Eth1/4 pd: slot 0 logical port num 3 gatos_num 0 fwm_inst 0 fc 0
Eth1/4 pd: pif_type 'data eth'(1) hw_present 1 port map idx 8
Eth1/4 pd: fabric a info: voq 0-7 port_id 55 connected 1 up 1
Eth1/4 pd: fabric b info: voq 0-7 port_id 55 connected 1 up 1
Eth1/4 pd: subported 0 primary 1 atherton 0
Eth1/4 pd: sup_src_dst_if 6 lif_blk 384-511
Eth1/4 pd: policer info: uc (sel 2) mc (sel 1) bc (sel 0)
Eth1/4 pd: mac-addr 000d.ecd5.a38b
```

In the following part of the example, notice the drops in the second line:

Eth1/4 pd: tx stats: bytes 50256531 frames 336488 discard 0 drop 0 Eth1/4 pd: rx stats: bytes 6718252 frames 77220 discard 0 drop 845482

In the following part of the example, notice that the FcoE counters are separate:

Eth1/4 pd fcoe: tx stats: bytes 2927716 frames 3919 discard 0 drop 0 Eth1/4 pd fcoe: rx stats: bytes 15307492 frames 9470 discard 0 drop 0

In the following part of the example, notice that the command helps find the cause of the drops:

```
switch# show platform fwm info gatos-errors 0
Printing non zero Gatos error registers:
DROP_INGRESS_FW_PARSING_ERROR: res0 = 93 res1 = 0
DROP_SRC_VLAN_MBR: res0 = 2567226 res1 = 0
DROP_FCF_SW_DOMAIN_IDX_MISS: res0 = 2445 res1 = 0
DROP_FCF_SW_TEL_MISS: res0 = 2445 res1 = 0
DROP_NO_FABRIC_SELECTED: res0 = 2556 res1 = 0
DROP_VLAN_MASK_TO_NULL: res0 = 2526 res1 = 0
```

L

REVIEW DRAFT-CISCO CONFIDENTIAL



Troubleshooting Security Issues

The Cisco Nexus 5000 NX-OS provides security that protects your network from degradation or failure and from data loss or compromise resulting from intentional attacks or from unintended, damaging mistakes.

This chapter describes how to identify and resolve problems that can occur with security in the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Roles
- AAA

Roles

Role assignment fails when user logs in

From the perspective of RBAC, when a user logs in, role assignment fails.

Possible Cause

The AV-pair is not configured properly on TACAC+ or the RADIUS server.

Solution

To complete the role assignment follow these steps:

Step 1 Check the TACACS+ (for example, ACS) server configuration.

• Use the following menu path to access the settings:

Interface Configuration > TACACS+ (Cisco IOS)

- Select the User box for Shell (exec)
- Select the Advanced TACACS+ Features

Display a window for each service that was selected, where you can enter customized TACACS+ attributes in the Advanced Configuration Options.

• Use the following menu path to access the settings and add a string to the Shell attributes:

User Setup > Add/Edit "admin" > TACACS+ Settings

- Select the Shell and Custom attributes boxes

Γ

- Add the following string into the textbox:

cisco-av-pair=shell:roles="network-admin"

- Step 2 Check the RADIUS (for example, ACS) server configuration.
 - Use the following menu paths to access the settings:

Network Configuration > AAA > AAA Servers > svi,20.1.1.2,CiscoSecure ACS

Network Configuration > AAA > AAA Client > 20.1.1.1 20.1.1.1 RADIUS (Cisco IOS/PIX 6.0) > SharedSecret=test1234, Authenticate Using=RADIUS (Cisco IOS/PIX 6.0)

Interface Configuration > RADIUS (Cisco IOS/PIX 6.0)

- Select User for cisco-av-pair.
- Use the following menu path to access the settings and add a string to the RADIUS attributes:

User Setup > Add/Edit <username> > Cisco IOS/PIX 6.x RADIUS Attributes

- Check the attribute box.
- Enter the following string:
 - shell:roles="network-admin"
- Step 3 Check the RADIUS (for example, RADIUSD) server configuration for settings in the user account.
 - Use the following path to access the user account definition:

.../etc/raddb

- Ensure that the user account definition contains: cisco-avpair= "shell: roles = network-admin"
- Step 4 Log in the user again.
- Step 5 Check the role assignment with the show user-account command.

Rules for Role's permit/deny action do not work correctly

When a user-defined role is assigned to a user account, the role's rule policy may not seem to take effect. For example, a rule in the role's configuration is set to deny all interface configuration commands. However, you still can configure interface commands.

Possible Cause

Order of rule configurations for the role is incorrect.

Note

The RBAC parser accesses a rule from highest to lowest rule number.

Solution

After identifying the rule that is not working correctly, check to see if any rules preceding it conflict or override it.

For example, if the rule that is not working correctly has a rule ID of 10, then check all the rules that have a rule ID greater than 10 to see if they might conflict with rule 10. To illustrate this example, we can say that rule 15 is found to be overriding rule 10. To resolve this conflict, you would have to modify rule 15 or change the rule ID of rule 10 so that it has a greater rule ID than rule 15.

Role's interface or VLAN policy does not appear to work correctly

When a user-defined role is assigned to a user account and the role's interface or VLAN policy is set to deny access to a certain interface, the user account can still use **show** commands to display configuration, status, setting, or statistics on the access-denied interface or VLAN.

Possible Cause

You are checking the interface or VLAN role policy with CLI commands, such as **show interface brief** or **show vlan**.

<u>Solution</u>

RBAC does not support filtering when displaying commands. Interface or VLAN role policies only apply to configuration or operational commands.

Possible Cause

You are not assigned to the role properly.

Solution

- Check the user role assignment with the show user-account command.
- Verify the role definition with the **show role name** < *name* > command.

Assigning multiple roles to single user does not seem to work correctly

When a user account is assigned to multiple roles, the user can access commands that are denied by one of the roles that it gets assigned to. This gives the appearance that the command parser does not work with multiple roles.

Possible Cause

You might expect that multiple roles on the same user account are parsed sequentially.

Solution

The NX-OS design is to parse multiple roles in a union-to-permit function, that each command is examined and compared to all the roles.

If any of the roles permit the command, then the CLI allows the user to continue.

For example, if the role permits the **interface eth1/1** command, then the CLI allows the you to enter the interface eth1/1 configuration mode.

Each role applies their policies (that is, interface, VLAN, VSAN, and so on) separately. If a role has an interface policy that denies eth1/1 as in the example, then that role would reject the command, but other roles might have a different interface policy allowing the same interface.

Change to role configuration does not get applied

When a user account is assigned to a role and you are logged into the Nexus 5000 switch, any changes made to the role configuration does not get applied immediately.

Possible Cause

While a user account is logged in and has been assigned to role A, the administrator makes some changes to role A with the expectation that the change would immediately affect the user that is logged in. However, the user is not assigned to the role properly.

L

Solution

NX-OS does not activate role configuration changes dynamically. You needs to log in again to have the configuration changes to the new role come into effect.

CLI rejects feature-group removal

The CLI rejects the **no role feature-group name** <*group-name*> command when the administrator tries to delete a feature-group.

Possible Cause

A CLI error indicates that the feature group is in use, which means that it is included in one of the role configurations.

Solution

To address the error, perform the following steps:

- Use the **show role | egrep role: |feature-group** command to display which feature group is associated with the role or under which role.
- Detach the association with the **no rule** command within the role configuration mode, and then delete the feature group.

AAA

User cannot login through TACACS+ or RADIUS authentication

With the server group properly configured for the Nexus 5000 switch and the server group is assigned the aaa authentication login default configuration on TACACS+ or RADIUS servers, the Telnet or SSH login fails to authenticate users with the following error:

%TACACS-3-TACACS_ERROR_MESSAGE: All servers failed to respond

Possible Cause

AAA group is not configured with the correct VRF to access servers.

Solution

Perform the following steps to enable login:

- Check which AAA group is being used for authentication with the **show running-config aaa** and **show aaa authentication** commands.
- For TACACS+, check the VRF association with the AAA group with the **show tacacs-server groups** and **show running-config tacacs**+ commands.
- For RADIUS, check the VRF association with the AAA group with the **show radius-server groups** and **show running-config radius** commands.
- Correct the VRF association, then test the VRF setting with the **test aaa group** <*name*> <*username*> <*password*> command.
- If the **test aaa** command returns the error, "user has failed authentication", then the server is accessable but the credentials for the user account are incorrect. Verify that the user configuration is correct on the server.

Possible Cause

AAA server is not accessable in network.

Solution

If the problem persists after correcting the VRF association and correcting the user-account credentials, then perform the following:

- If the **test aaa** command returns the error, "error authenticating to server", the route to the server might be missing in the configuration. Use the **ping** *<server>* command, if the AAA server is associated with the default VRF. If it is associated with VRF management, use the **ping** *<server>* **vrf management** command.
- If the message "No route to host" appears, then the static route to the server is not configured properly. Reconfigure the IP route in the corresponding VRF context.
- Enter the **ping** *<server>* command again. If the command is successful, then use the **test aaa group** *<name> <username> <password>* command.
- If the **ping** <*server*> command is unsuccessful, then check the network connectivity, such as if the ARP entry of the nexthop router is displayed in the **show ip arp [vrf management]** command or if the ARP entry of the Nexus 5000 switch exists in the nexthop router's ARP table.

Unable to decode content of packets with Wireshark

AAA packets were captured from the network, but Wireshark was unable to decode the content of the packets.

Possible Cause

AAA packets are encrypted while the host key is enabled.

<u>Solution</u>

Perform the following steps to decode the content:

- Use the no tacacs-server command to delete the TACACS server configuration.
- Reconfigure the TACACS server without specifying any key.
- Reconfigure the AAA client for the Nexus 5000 switch on the Network Configuration page in ACS while removing the host key.
- Re-do the wire tapping. The captured packetsnow should not be encrypted and the data content should be decoded properly by Wireshark.
- After the packet capturing, the administrator should revert to the host key configuration for better security.

Role assignment fails when user logs in

Role assignment fails when the user logs in. (From the perspective of the Nexus 5000 switch AAA.)

Possible Cause

Assuming that the ACS or TACACS+ and RADIUS has the Cisco av pair configured correctly, then the problem might be that the internal or local VRF assignment for the user login is not working correctly.

<u>Solution</u>

Г

Perform the following steps for role assignment:

- Check which AAA group is being used for authentication with the **show running-config aaa** and **show aaa authentication** commands.
- For TACACS+, check the VRF association with the AAA group with the **show tacacs-server groups** and **show running-config tacacs+** commands.
- For RADIUS, check the VRF association with the AAA group with the **show radius-server groups** and **show running-config radius** commands.
- If the above commands show that the association is correct, then use the **debug tacacs+ all** command to enable the trace.
- Log in the user again, and collect the debug trace.

The trace should contain information for further investigation (as shown in the example). Example:

tacacs: process_aaa_tplus_request: Group t1 found. corresponding vrf is management

• Use the **no debug tacacs+ all** command to turn off debug tracing on TACACS+.

No command accounting logs on ACS server when TACACS+ accounting enabled

When TACACS+ accounting is enabled, the command accounting logs on the ACS server are not found.

Possible Cause

The ACS server configuration is wrong or incomplete.

Solution

Perform the following steps:

- In the ACS GUI in Network Configuration, go to the AAA Client Setup for any client. Check the checkbox for Log Update/Watchdog Packets from this AAA Client. Click the Submit + Apply button.
- Verify CMD Accounting with the following menu path:

Reports and Activity > TACACS+ Administration

Open the Tacacs+Administration <activelDATE>.csv file and verify the cmd and timestamp on each row of the file.

PAP authentication does not work for **RADIUS**

PAP authentication works for TACACS+ but not for RADIUS.

Possible Cause

Starting with Release 4.2(1), NX-OS only supports ASCII (PAP) authentication for TACACS+.

Solution

In NX-OS, ASCII authentication is equivalent to PAP authentication. By default, both TACACS+ and RADIUS use CHAP. You can switch to PAP authentication with the **aaa authentication login ascii-authentication** command.

If you try to configure TACACS+ along with RADIUS, syslog messages similar toto the example, as shown in the example, appear during login.

Example:

```
2010 May 19 16:12:19 mars %$ VDC-1 %$ %RADIUS-2-RADIUS_NO_AUTHEN_INFO: ASCII
authentication not supported
2010 May 19 16:12:19 mars %$ VDC-1 %$ %AUTHPRIV-3-SYSTEM_MSG: pam_aaa:Authentication
failed for user oregon-regress from 10.193.128.5 - login[5698]
```

Authentication fallback method appears inoperable

The NX-OS supported fallback method for authentication is that if all the AAA remote RADIUS or TACACS+ servers are unreachable, then the log in attempts to authenticate the SSH/Telnet user locally. However, the login to the Nexus 5000 switch might still fail with the local authentication.

Possible Cause

The local user database does not contain the user account that the user is using to login with.

Solution

Perform the following steps to check the authentication fallback method.

- As a best practice, include the **aaa authentication login error-enable** command in the configuration. When it is included in the configuration, the login session sees whether the fallback method is operating correctly. If messages, such as "Remote AAA servers unreachable; local authentication done" or "Remote AAA servers unreachable; local authentication failed", are received, then the fallback method is operating correctly.
- If the remote AAA servers are not accessible, check to see if the local user database has the user credential for local authentication. Use the **show user-account** command to display the credential.



By using the **show user-account** command, you can determiine which user-account was created through REMOTE authentication. A user account that was created with REMOTE authentication cannot be used for a local (fallback) login.

• Create local user accounts with the **username <username> password <password> role <role name>** command until the remote AAA servers become accessible.

AAA

Г



Troubleshooting System Management Issues

The system management features of the Cisco Nexus 5000 Series switch allow you to monitor and manage your network for efficient device use, role-based access control, SNMP communications, diagnostics, and logging.

This chapter describes how to identify and resolve problems that can occur with system management and the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- SNMP
- Logging
- Traps
- DNS

SNMP

SNMP memory usage continuously increasing

The show proc mem | inc snmp command shows continuously increasing SNMP memory usage.

Possible Cause

SNMP memory usage increases when SNMP requests are processed from different monitoring stations. Typically, this situation stabilizes over time. If the memory increases continuously without stabilizing, then some of the SNMP requests are causing a memory leak.

<u>Solution</u>

Review the output from the show system internal snmp mem-stats detail command.

Take example snapshots with the following commands while processing SNMP requests:

- show clock
- show system internal mem-stats detail
- show tech snmp

Γ

SNMP not responding

No response or delayed response for SNMP request.

Possible Cause

If the switch CPU utilization is high during the SNMP operations such as GET, GETNEXT and WALK, the response may be very slow or there is no response that results in a time-out.

Solution

While SNMP is not responding, check CPU utilization with the following commands:

- show proc cpu history
- show proc cpu sort

The output from this command shows which Nexus 5000 component is using the greatest amount of CPU resources.

SNMP not responding and show snmp command reports SNMP has timed out

SNMP is not responding and the **show snmp** command reports that SNMP has timed out.

Possible Cause

The SNMP process might have exited, but the process did not crash.

Solution

Use **show system internal sysmgr service name snmpd** command which should show the state to be"SRV_STATE_HANDSHAKED.

Example:

```
Service "snmpd" ("snmpd", 74):
UUID = 0x1A, PID = 4131, SAP = 28
State: SRV_STATE_HANDSHAKED (entered at time Mon Jun 14 17:12:15 2010).
Restart count: 1
Time of last restart: Mon Jun 14 17:12:14 2010.
The service never crashed since the last reboot.
Tag = N/A
Plugin ID: 0
```

Not able to perform SNMP SET operation

The following error appears when trying to perform the SNMP SET operation:

```
bash-2.05b$ snmpset -v2c -c private 10.78.25.211 .1.3.6.1.4.1.9.9.305.1.1.6.0 i 1
Error in packet.
Reason: notWritable
```

Possible Cause

The SNMP community does not have write permission.

Solution

Check the output of the **show snmp community** command to ensure that the write permission is enabled.

Example:

Community

Group / Access context acl_filter

private	network-operator
public	network-admin

SNMP on BRIDGE-MIB

The SNMP GET on BRIDGE-MIB operation does not return correct values and results in errors.

Possible Cause

The BRIDGE-MIB may not be supported.

<u>Solution</u>

Check the release notes to make sure that BRIDGE-MIB is supported on NX-OS Release 4.2(1) or later releases.

Logging

L

System is not responsive

System performance is significantly slower or non responsive.

Possible Cause

Some system resources may be over-utilized. For example, an incorrect logging level might generate many messages resulting in an impact on system resources.

<u>Solution</u>

Check the logging level on the chassis. If you have a logging level setting, such as 6 or 7, many messages are generated and performance can be impacted. Use the following commands to display the amount of resources that are being used.

- show proc cpu | inc syslogd
- show proc cpu
- show run | inc logging
- show system resource

Syslog server not getting messages from DUT

Although the syslog server is configured, the destination syslog server is not receiving messages from DUT.

Possible Cause

Syslog server might not be accessible or the logging level might not be appropriate.

Solution

Г

- Check to see if the destination syslog server is accessible from VRF management. Use the **ping** <*dest-ip*> **vrf management** command to ping the server.
- Check that the syslog configuration on the DUT has use-vrf management.

Example:

logging server 10.193.12.1 5 use-vrf management

• Check that the appropriate logging level is enabled to send logging messages. Use the **show logging info** command. If the logging level is not appropriate, then set the appropriate level using the **logging level** <*feature>* <*log-level>* command.

Traps

Traps not received

The results of traps are not received.

Possible Cause

The traps might not be enabled or the SNMP host might not be accessible.

The following are possible causes:

- Traps might not be enabled.
- The SNMP host might not be accessible.
- A firewall might be blocking access.
- An access list might be blocking UDP port 162.

Solution

Use the following commands to check whether the proper VRF is configured for the SNMP host and that the trap is enabled:

- snmp-server enable traps <trapname>
- snmp-server host <x.x.x.x> use-vrf <vrf-name>

where *x.x.x.x* is the IP address of the trap receiving device.

DNS

DNS resolution not working correctly

When specifying a host name using DNS or VRF, the host name is not resolved and an error occurs.

Possible Cause

The DNS client is not configured correctly.

<u>Solution</u>

Use the following commands to configure the DNS client:

- config t
- vrf context management
- ip host name <address1 [address2... address6]>
- ip domain-name name [use-vrf <vrf-name>]
- ip domain-list name [use-vrf <vrf-name>]
- ip name-server < server-address1 [server-address2... server-address6] >< [use-vrf vrf-name>]
- ip domain lookup
- show hosts
- copy running-config startup-config

Specified domain not removed from domain-list

When using the **no ip domain-list** *<name>* command to remove a specified domain from the domain-list, only the most recently added domain is removed.

Possible Cause

The no ip domain-list <name> command is not locating the specified domain.

Solution

There are two possible workarounds:

- To remove a domain using the **no ip domain-list** *<name>* command that is not the most recently added domain to the domain-list, you must temporarily remove every domain in the domain-list until reaching the desired domain. Then you must add back the temporarily removed domains to the domain-list.
- An alternative approach is to copy the startup-config and delete the desired domain with a text editor. Then you must load the edited startup-config back onto the device.



Troubleshooting Virtual Port Channel Issues

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 5000 Series switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide Layer 2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

This chapter describes how to identify and resolve problems that can occur with vPC in the Cisco Nexus 5000 Series switch.

This chapter includes the following section:

• Improper Configurations

Improper Configurations

vPC fails to start

This issue may have many possible causes:

- Unable to configure vPC
- vPC in blocking state
- vPC domain ids
- Connectivity issues
- Peer-link issues
- vPC Consistency parameter issues

Unable to configure vPC

Possible Cause

vPC is not enabled or is not supported in the NX-OS release of software that you are running.

Solution

Ensure that the NX-OS release supports vPC. vPC is supported in NX-OS Release 4.1 and later releases. If the NX-OS release supports vPC, then use the command feature of vPC to enable it.

Г

vPC in blocking state

Possible Cause

A bridge protocol data unit (BPDU) only sends data on a single link of a port channel. If a bridge assurance (BA) dispute is detected, then vPC moves into a blocking state.

<u>Solution</u>

Do not enable bridge assurance on the vPC link: because of the following:

- Cannot be used on a spanning tree port type network.
- Prevents you from encountering ISSU issues. Bridge assurance should only be enabled on the vPC peer link.

vPC domain ids

Possible Cause

The vPC domain IDs of two switches do not match.

<u>Solution</u>

Compare the vPC domain IDs of the two switches and ensure that they match.

Example:

```
switch1# show vpc brief
Legend:
               (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id
                             : 500
Peer status: peer link is downvPC keep-alive status: Suspended (Destination IP not reachable)
Configuration consistency status: success
vPC role
              : secondary, operational primary
Number of vPCs configured
                            : 4
                           : Disabled
Peer Gateway
Dual-active excluded VLANs
                            : -
vPC Peer-link status
id Port Status Active vlans
_ _
    _ _ _ _
          _____
    Po500 down
1
switch2# show vpc brief
Legend:
              (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id
                             : 1
Peer status
                             : peer link is down
vPC keep-alive status
                             : Suspended (Destination IP not reachable)
Configuration consistency status: success
vPC role
                           : primarv
Number of vPCs configured
                            : 4
Peer Gateway
                            : Disabled
Dual-active excluded VLANs
                            : -
vPC Peer-link status
```

id	Port	Status	Active vlans
1	Po500	down	-

The two switches in this example have different vPC domain IDs. The vPC domain IDs of these Nexus switches must be changed to match. This can be done by entering configuration commands, one per line, and ending each with Cntl + Z.

```
switch2(config)# vpc domain 500
Changing domain id will flap peer-link and vPCs. Continue (yes/no)? [no] yes
Note:
 -----:: Re-init of peer-link and vPCs started ::------
switch2# show vpc brief
Legend:
              (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id
                            : 500
Peer status
                            : peer adjacency formed ok
vPC keep-alive status
                            : peer is alive
Configuration consistency status: success
vPC role
                           : primary, operational secondary
Number of vPCs configured
                           : 4
Peer Gateway
                            : Disabled
Dual-active excluded VLANs
                           : -
vPC Peer-link status
   Port Status Active vlans
id
          _____
___
    _ _ _ _
1
    Po500 up 1,19,91,99,757
```

Connectivity issues

Possible Cause

vPC peer keepalive link and connectivity issues over mgmt0 might exist.

Solution

• Check for the peer keepalive mgmt0 reachability.

On the other Nexus 5000 switch, enter the command:

show run interface mgmt 0

Example:

```
switch2# sh run int mgmt 0
!Command: show running-config interface mgmt0
!Time: Tue Mar 8 03:20:58 2011
version 4.2(1)N2(1)
interface mgmt0
    ip address 172.18.118.162/24
```

Ensure there is reachability from switch1:

L

```
switch1# ping 172.18.118.162 vrf management
PING 172.18.118.162 (172.18.118.162): 56 data bytes
64 bytes from 172.18.118.162: icmp_seq=0 ttl=254 time=5.306 ms
64 bytes from 172.18.118.162: icmp_seq=1 ttl=254 time=3.963 ms
64 bytes from 172.18.118.162: icmp_seq=2 ttl=254 time=4.04 ms
64 bytes from 172.18.118.162: icmp_seq=3 ttl=254 time=4.077 ms
64 bytes from 172.18.118.162: icmp_seq=4 ttl=254 time=4.057 ms
```

If the ping fails, it means that the connectivity between both mgmt0 interfaces does not exist or that they are not interconnected properly.

Make sure the mgmt0 interface is unshut and that you can ping the switch mgmt0 interface.

```
switch# sh int br | grep mgmt0
mgmt0 -- down 172.16.118.62 -- 1500
```

If the status shows that it is down, it means there is no physical connection to mgmt0 or that the interface is in admin shutdown. You need to verify the physical connectivity and unshut the port:

```
switch1# config t
switch1(config)# int mgmt 0
switch1(config-if)# no shut
switch1(config-if)# show int br | grep mgmt0
mgmt0 -- up 172.16.118.62 1000 1500
```

If pinging the other switch continues to fail, then there is an interconnection issue between the two Nexus 5000 switches.

Check the networking in between the switches:

- Switch interconnecting in access VLAN mode, using the same VLAN for both Nexus switches.
- The VLAN is allowed across and between the switches.
- Check the vPC configuration and compare the mgmt0 IP addresses that are used:

```
switch1# show run int mgmt 0
!Command: show running-config interface mgmt0
!Time: Tue Mar 8 03:53:48 2011
version 4.2(1)N2(1)
interface mgmt0
    ip address 172.18.118.163/24
switch1# show run vpc
!Command: show running-config vpc
!Time: Tue Mar 8 03:53:57 2011
version 4.2(1)N2(1)
feature vpc
vpc domain 500
    peer-keepalive destination 172.18.118.162
switch2# show run int mgmt 0
!Command: show running-config interface mgmt0
!Command: show running-config interface mgmt0
```

```
!Time: Tue Mar 8 03:53:53 2011
version 4.2(1)N2(1)
interface mgmt0
  ip address 172.18.118.162/24
switch2# sh run vpc
!Command: show running-config vpc
!Time: Tue Mar 8 03:54:01 2011
version 4.2(1)N2(1)
feature vpc
vpc domain 500
  peer-keepalive destination 172.18.118.162
```

In this example, the destination IP is not correct. The correct IP is 172.18.118.163, which is the peer IP address.

Peer-link issues

Possible Cause

The peer link is not configured.

<u>Solution</u>

Configure the peer link correctly.

Example:

In this example, the problem is that the vPC peer-link does not exist.

```
Peer status : peer link not configured
vPC keep-alive status : peer is alive
Configuration consistency status: failed
Configuration consistency reason: vPC peer-link does not exists
```

You can use the **show cdp neighbor** command to determine which physical ports are connected to the other Nexus switch.

```
switch1# show cdp neighbor
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                 S - Switch, H - Host, I - IGMP, r - Repeater,
                 V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
Device-ID
                     Local Intrfce Hldtme Capability Platform
                                                                    Port ID
switch2(SSI1324033X)Eth1/25
                                 128
                                        SIS
                                                  N5K-C5020P-BF Eth1/25
switch2(SSI1324033X)Eth1/26
                                 128
                                        SIS
                                                  N5K-C5020P-BF Eth1/26
```

In this example, ports 25 and 26 connect to the other Nexus 5000 switch and should be configured as a peer link.

Run the same command on the other Nexus 5000 switch and observe the ports.

```
switch2# show cdp neighbor
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                 S - Switch, H - Host, I - IGMP, r - Repeater,
                 V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute
Device-ID
                      Local Intrfce Hldtme Capability Platform
                                                                     Port ID
switch1(SSI14150768)Eth1/25
                                  168
                                         SIS
                                                   N5K-C5020P-BF Eth1/25
switch1(SSI14150768)Eth1/26
                                  168
                                         SIS
                                                   N5K-C5020P-BF Eth1/26
switch2# show run int e1/25
!Command: show running-config interface Ethernet1/25
!Time: Tue Mar 8 04:09:17 2011
version 4.2(1)N2(1)
interface Ethernet1/25
  switchport mode trunk
  channel-group 500
switch2# show run int e1/26
!Command: show running-config interface Ethernet1/26
!Time: Tue Mar 8 04:09:20 2011
version 4.2(1)N2(1)
interface Ethernet1/26
  switchport mode trunk
  channel-group 500
```

In this example, you can see that port-channel 500 is used on the connection to switch1 on switch2.

You now need to determine how port-channel 500 is configured on switch2.

```
switch2# show run int po 500
!Command: show running-config interface port-channel500
!Time: Tue Mar 8 04:10:38 2011
version 4.2(1)N2(1)
interface port-channel500
switchport mode trunk
vpc peer-link
spanning-tree port type network
speed 10000
```

Create a port-channel 500 on switch1 and associate it to the ports connecting to e1/25 and e1/26 on switch2.

```
switch1(config)# int po 500
switch1(config-if)# int e1/25-26
```

REVIEW DRAFT-CISCO CONFIDENTIAL

```
switch1(config-if-range)# channel-group 500
switch1(config-if-range)# int po 500
switch1(config-if)# vpc peer-link
```

Notice that the spanning tree port type has changed to a network port type on the vPC peer link.

This enables spanning tree bridge assurance on the vPC peer link, provided that STP bridge assurance is not disabled. (STP bridge assurance is enabled by default.)

Check the vPC again.

vPC domain id :	500
Peer status :	peer adjacency formed ok
vPC keep-alive status :	peer is alive
Configuration consistency status:	success
vPC role :	primary
Number of vPCs configured :	4
Peer Gateway :	Disabled
Dual-active excluded VLANs :	-
vPC Peer-link status	
id Port Status Active vlans	
1 Po500 up 1,19,91,99,757	

Port channel 500 and the peer-link are now up. The vPC is successful.

vPC Consistency parameter issues

Possible Cause

vPC is not operational if type 1 consistency parameters do not match on both Nexus 5000 switches.

<u>Solution</u>

Ensure that type 1 consistency parameters match.

The possible values for type are 1, 2, or -. Items that are type 1 must match on both Nexus 5000 switches. If they do not match, then vPC is suspended. Starting with Release 5.0, a type 2 was introduced. Items that are type 2 do not have to match on both Nexus 5000 switches for the vPC to be operational.

The command in the following example displays local and peer values. Run the command on both switches to ensure that the type 1 items match.

Example:

To check for a mismatch, display the consistency parameters.

switch1# show vpc consistency-parameters global

```
Legend:

Type 1 : vPC will be suspended in case of mismatch

Name

Type Local Value

Peer Value
```

QoS	1	([], [3], [], [], [],	([], [3], [], [], [],
Network QoS (MTU)	1	[]) (1538, 2240, 0, 0, 0, 0, 0)	[]) (9216, 2240, 0, 0, 0, 0)
Network Qos (Pause)	1	(F, T, F, F, F, F)	(F, T, F, F, F, F)
Input Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Input Queuing (Absolute Priority)	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Output Queuing (Bandwidth)	1	(50, 50, 0, 0, 0, 0)	(50, 50, 0, 0, 0, 0)
Output Queuing (Absolute	1	(F, F, F, F, F, F)	(F, F, F, F, F, F)
Priority)			
STP Mode	1	Rapid-PVST	Rapid-PVST
STP Disabled	1	None	None
STP MST Region Name	1		
STP MST Region Revision	1	0	0
STP MST Region Instance to	1		
	1		
STP Loopguard	1	DISADIEG Enchlad	DISabled Enchled
SIP BILOGE ASSULATICE	1	Marmal Disabled	Marmal Disabled
STP Port Type, Edge	T	Normal, Disabled,	Normal, Disabled,
GED MGE Gimulate DVGE	1	Disabled	Disabled
STP MST SIMulate PVST	T	Enabled	Enabled
Allowed VLANS	-	1,19,91,99,120,757	1,10,19-20,91,99,400-4 01,403,420,440,442,444 -446,451-486,499,757,7 97
Local suspended VLANs	-	120	-

In this example, there are different MTU values for Network QoS. The value for the peer switch is 9216 on the peer switch (switch2) and the value for the local switch is 1538 (switch1). vPC will not be operational until the Network QoS values match on both switches.



Troubleshooting Config-Sync Issues

This chapter describes how to identify and resolve problems that can occur with config-sync in the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- Commit Failure
- Import Failure
- Merge Failure
- Switch-profile Deletion Failure
- Verify Failure

Commit Failure

Use the show switch-profile status commit command to view commit status.

Commit failure has many possible causes:

- Command Parsing Failed
- Verify failed
- Commands that failed commit
- Another session in progress

Note

When a commit fails, commands that were entered under SP are still stored in the SP buffer. Do not configure these commands under SP again. After correcting the cause of the failure, only the commit needs to be executed.

Command Parsing Failed

Possible Cause

Appropriate conditional feature(s) are not enabled.

<u>Solution</u>

Ensure that appropriate conditional feature(s) are enabled.

This error message indicates that some feature commands have not been configured. Feature commands are not allowed to be configured within SP and have to be configured on BOTH peers from conf-t.

Verify failed

Possible Cause

The commands listed failed mutual-exclusion checks. These commands have already been configured under conf-t.

Solution

If you do not want these commands synchronized, remove the commands from conf-t.

Alternatively, delete these commands from the switch-profile buffer and reissue the commit.

To delete commands from the switch-profile buffer, perform the following:

- View commands in the SP buffer using the show switch-profile buffer command.
- Delete commands indicated by the sequence numbers with the **buffer delete <range>** command.
- Use the **buffer-move <seq id> <seq id>** command to rearrange commands in the buffer. This is command is useful when commands in the buffer are not ordered correctly.

Commands that failed commit

Possible Cause

Commands failed during commit.

<u>Solution</u>

Correct the reason for the failure and re-issue the commit.

If the commit continues to fail, issue the same command from conf-t. If it succeeds from conf-t, check for any errors relating to the command using the **show system internal csm info trace** command.

For every command executed from config-sync, there is a csm_cmd_status[0x0] line in the trace log that indicates that the command was successful.

Another session in progress

Possible Cause

Conflict occurs if conf-t or config-sync has taken a lock.

Solution

Compare the vPC domain IDs of the two switches and ensure that they match.

Use the **show system internal csm global info** command to check whether conf-t or config-sync has taken a lock.

• If conf-t has taken a lock and not released it, command output, similar to the following example, is displayed.



The client type should be set to 2 as shown in the example.

Example:

```
No of sessions: 1 (Max: 32)
Total number of commands: 0 (Max: 102400)
Session Database Lock Info: Locked
Client: 2 (1: SSN, 2: CONF-T)
Ref count: zero-based ref-count
Lock acquired for cmd : some-command
```

- Identify the command that acquired the lock using the show accounting log command.
- After identifying the command, check for its SUCCESS/FAILURE status.
- If the command did not return a status, then config-sync would not release the lock on conf-t.
- Use the test csm ssn-db-lock reset conf-t command to reset the lock.
- If switch-profile has taken the lock, the client id is reported as 1 in the **show system internal csm** global info command.
 - Use the **show switch-profile status** command to determine if a merge is in progress.

A merge is indicated by pending_merge:1 /rcvd_merge:1.

- If a merge/verify/commit session is already in progress, then sp ssn-db is locked.

Wait for the current session to complete and try again.

- If the lock is not released, use the **show cfs lock** command to determine if the CFS fabric is locked.

Identify the application that locked CFS. If the application is session-manager, then the CFS lock was taken by config-sync.

Analyze the output from the **show system internal csm info trace**, **show cfs internal notification log name session-mgr**, and **show cfs** commands.

• Use the **show system internal csm info trace** command to view the events, trace, or error debug traces.

Import Failure

Use the show switch-profile status command to view import status.

Import failure has many possible causes:

- Failed to collect running-config
- Command does not exist in global-db
- Mutual exclusion check failed on peer

The following describes import options and best practices.

Туре	Description
import	Enables import mode. Manually enter configuration and then commit to move the configuration into SP.
import running-config	Imports all SP-aware configurations from the running-config into SP. Use the buffer-move and buffer-delete commands to remove commands not to be synced and then commit the configurations.
import interface <interface-range></interface-range>	Imports running configurations for specified interfaces. Used to only import configurations of interfaces and not global configurations.
import running-config exclude-phy-interface	Imports only global and logical interfaces, not physical interface configurations.

Table 1-1 Import Options

Import best practices

- The import option is used when the system is already configured and you want to bring in an existing configuration within the switch-profile and sync it with the peer.
- When VPC peer switches are already configured, the import operation is performed on both switches independently. The user must verify that the configurations are the same under SP on peer switches and then add the sync-peer command to the configuration.

If the configurations are different within the SP when the sync-peer command is added, a merge failure occurs. Use the **show switch-profile status** command and determine which configurations failed to merge.

• If one of a pair of VPC switches was previously configured and the other switch was RMA'ed, then the switch-profile is created on both switches using the sync-peer command.

This means that the configuration from the previously configured switch is imported and committed on the other switch. Also the configurations from the previously configured switch are moved from the global-db to the database of the switch-profile.

Failed to collect running-config

Possible Cause

Failure occurs if the system is too busy and the **show running** command did not complete.

<u>Solution</u>

Determine if a system resource utilization problem exists. Correct the problem and retry the operation.

Command does not exist in global-db

Possible Cause

Command is missing from the global-db.

Solution

Use the **show system internal csm info global-db cmd-tbl** command to determine if the command exists in the global_db.

- If the command exists in the global_db, it is possible that there is not enough space in the show run for the command. Ensure that there are no trailing space/tabs in show running config generation.
- If the command does not exist in the global_db, use the **show accounting log** command to determine if the command was configured and to display the status of the command.
 - If the command status was a failure, then the command should not be displayed in show running.
 - If the command is displayed, then the application should correct it.
- If the command was configured before reload/issu, add the command back. If the accounting log shows the command's retval as success, determine if the command is getting added to the global-db.
 - If the command was added correctly, copy r s, check global-db reload, and check if the command exists in the global-db.
 - If the command does not exist in the global-db, then the issue might be that the command is not showing up in show running on boot up.
 - If the command does not exist in the global_db, investigate the csm_save_global_command function. The csm_save_global_command function is where the command gets added to the global_db

Mutual exclusion check failed on peer

Possible Cause

The imported configuration is sent to the peer. However, if the configuration is already configured on the peer outside of SP, then the import fails the mutual exclusion check on the peer.

Solution

Remove the failed commands from conf-t on the peer and then retry import verify/commit.

Use the **show system internal csm info trace** command for further investigation to look at events, trace, or error messages.

Merge Failure

A merge between peers happens when a peer becomes reachable.

A merge is initiated when CFS sends a peer add for the peer or if the peer is already reachable. Configuring the sync-peer command starts the merge session.



For a merge to succeed, the configuration in the switch-profile on both peers must match exactly.

Merge failure has many possible causes:

- First time merge failure
- Merge after peers that were in sync previously
- Merge after reload



Use the show system internal csm info trace command to view events, trace, and error messages.

First time merge failure

Possible Cause

When peer switches are trying to synchronize configurations, the merge might fail when validating received configurations.

<u>Solution</u>

Use the show switch-profile status command to view which commands failed validation.

This implies that the commands on both the switches are configured differently.

Perform the following to correct the configurations:

- Remove the sync-peers destination command from the switch-profile.
- Use the **show running switch-profile** command on both peers to ensure that the configuration is exactly the same under switch-profile.
- Add back the sync-peers destination command to the switch-profile.
- Reissue the commit.

Merge after peers that were in sync previously

Possible Cause

If peers were in sync and connectivity was lost, and conflicting configuration changes were made on the switches, then the merge would fail.

Solution

Use the show switch-profile status command to view which commands failed the merge.

Correct the configurations and reissue the commit from the peer with the corrected configuration.

Merge after reload

Possible Cause

After a switch is reloaded, it sends its switch-profile configuration to the peer. If there was a configuration change done under SP for the peer that was not reloaded, then the merge fails.

Solution

Use the show switch-profile status command to view which commands failed the merge.

Correct the configurations and reissue the commit.

Switch-profile Deletion Failure

A rollback is used to delete the configurations during a switch-profile deletion.



To check for commands that failed deletion, use the **show switch-profile status commit** command to view the status. Alternatively, use the **show switch-profile session-history** command by matching the session based on the timestamp/session type.

Switch-profile deletion failure has many possible causes:

- Application failure
- Failure from dependent commands
- Application does not respond

Other known switch-profile deletion issues:

- Rollback fails with "Deletion of switch profile failed" message (CSCti97003)
- Port-channel interface not deleted on switch-profile delete (CSCtf17697)

Application failure

Possible Cause

Switch-profile deletion failure might be that the application failed the command. It is possible that the configuration is deleted out of order.

The switch-profile does not order configurations as displayed in the show run output. There might be out of sequence issues that occur during the deletion of the switch-profile.

<u>Solution</u>

Use the **resequence-database** command in the conf-sync mode to resequence the commands in SP in the order that the commands appear in show running. After resequencing the commands, reissue the delete.

Failure from dependent commands

Possible Cause

Switch-profile deletion failure results from dependent commands in conf-t mode.

If a command inside SP is referenced by another command outside of SP and the first command inside SP is deleted, then failure occurs because the command outside of SP still references it.

Solution

Correct the commands, references, and reissue the delete.

Application does not respond

Possible Cause

The deletion fails because the application does not respond due to the application owning the command.

<u>Solution</u>

Correct the commands and reissue the delete.

Verify Failure

Verify failure has many possible causes:

- Mutual exclusion check under local information
- Mutual exclusion check under peer information
- Rollback/ISSU in progress
- Global-db modification in progress
- Peer unable to accept lock request

Note

Use the show switch-profile status command to view messages about the failure.



Determine if the failure is on local/peer side by looking at whether the error is listed under local error(s)/ peer error(s) or both.



Use the show system internal csm info trace command to view events, trace, and error messages.

Mutual exclusion check under local information

Possible Cause

Command failed mutual-exclusion check under local information because the command has already been configured from conf-t.

<u>Solution</u>

Delete the command from conf-t mode and run verify from config-sync mode.

Mutual exclusion check under peer information

Possible Cause

Command failed mutual-exclusion check under peer information because the command has already been configured from conf-t on the peer.

<u>Solution</u>

Delete the command from conf-t mode on the peer and run verify from config-sync mode.

Rollback/ISSU in progress

Possible Cause

Verify cannot be performed when rollback/ISSU is in progress.

<u>Solution</u>

Stop rollback or wait for it to complete and run verify.

Global-db modification in progress

Possible Cause

Verify cannot be performed when global-db is being updated on the local/peer side.

<u>Solution</u>

Wait for the update to complete and run verify.

Peer unable to accept lock request

Possible Cause

Verify cannot be performed when the peer is unable to accept the lock request.

<u>Solution</u>

The peer is handling a transaction and cannot accept a lock request. Run verify at a later time.

Use the show switch-profile status command to determine if there is an ongoing transaction.

If the peer remains in the same state for a long time, use the **show cfs lock** command to determine if the CFS fabric has been locked.

Also check the application that has taken the CFS lock. If the application is ssnmgr, use the **show cfs internal session-history name session-mgr** command and the **show cfs internal notification log name session-mgr** command to view information about when a lock was acquired or released. It can also show the mapping to the csm transactions displayed with the **show switch-profile session-history** command.



С

Cisco Fabric Service	es 1-31
CNA 1-9	
Config-sync	
Commit failure	1-1
Import failure	1-3
Merge failure	1-5
Verify failure	1-8

D

Data Center Bridging 1-1 DNS 1-4

F

FC Services 1-16 FIP 1-7

L

Logging 1-3

Μ

MAC address table 1-1 Multicast 1-5

Ν

NPV 1-2

Ρ

PFC FCoE 1-11 QoS 1-8 Policy maps 1-1

Q

QoS Improper configuration 1-2

R

Registers and Counters FCoE 1-15 Layer 2 switching 1-10 QoS 1-10 SAN switching 1-48

S

SAN Port Channels 1-13 Security AAA 1-4 Roles 1-1 SNMP 1-1 Spanning Tree Protocol 1-3

Т

Traps 1-4

V

VLAN 1-6 vPC Improper configuration 1-1 VSAN 1-40

Ζ

Zoning 1-7