

Troubleshooting Cisco Data Center Unified Fabric

Volume 2

Version 5.0

Student Guide

Text Part Number: 97-3214-01



Americas Headquarters

Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV Amsterdam,
The Netherlands

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.

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)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 2

<u>FCoE Troubleshooting.....</u>	<u>4-1</u>
Overview	4-1
Module Objectives	4-1
<u>Troubleshooting FCoE.....</u>	<u>4-3</u>
Overview	4-3
Objectives.....	4-3
Troubleshooting FIP	4-4
Troubleshooting FCoE Performance	4-17
Summary.....	4-24
<u>Troubleshooting DCBX.....</u>	<u>4-25</u>
Overview	4-25
Objectives.....	4-25
Troubleshooting DCBX	4-26
Troubleshooting PFC	4-32
Summary.....	4-38
Module Summary.....	4-39
Module Self-Check	4-41
Module Self-Check Answer Key.....	4-43
<u>Platform-Specific Issue Troubleshooting</u>	<u>5-1</u>
Overview	5-1
Module Objectives.....	5-1
<u>Troubleshooting Cisco Nexus 7000 Series Switches</u>	<u>5-3</u>
Overview	5-3
Objectives.....	5-3
Troubleshooting Licensing.....	5-4
Troubleshooting Installs, Upgrades, and Reboots.....	5-12
Troubleshooting Cisco Fabric Services	5-22
Troubleshooting VDCs.....	5-32
Troubleshooting Routing.....	5-39
Troubleshooting Unicast Traffic	5-43
Troubleshooting Memory	5-50
Troubleshooting CPU	5-57
Troubleshooting Switch Fabric	5-63
Troubleshooting CoPP and Data Plane Rate Limiters	5-66
Summary.....	5-69
<u>Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches</u>	<u>5-71</u>
Overview	5-71
Objectives.....	5-71
Troubleshooting Licensing.....	5-72
Troubleshooting Cisco IOS ISSU	5-75
Troubleshooting Configuration Synchronization	5-78
Troubleshooting QoS.....	5-87
Troubleshooting CRC Errors	5-96
Troubleshooting CPU	5-99
Troubleshooting Unified Ports	5-101
Summary.....	5-103
<u>Troubleshooting Cisco Nexus 2000 Series Fabric Extenders.....</u>	<u>5-105</u>
Overview	5-105
Objectives.....	5-105
Troubleshooting Fabric-Extender Configuration and Management.....	5-106
Troubleshooting Fabric-Extender Queuing and Packet Drops.....	5-110
Summary.....	5-112

Troubleshooting Cisco MDS Series Switches 5-113

Overview	5-113
Objectives	5-113
Troubleshooting Licensing.....	5-114
Troubleshooting Installs, Upgrades, and Reboots	5-119
Troubleshooting Ports	5-121
Troubleshooting Cisco Fabric Services.....	5-124
Troubleshooting VSANs	5-127
Troubleshooting Zones and Zone Sets	5-133
Summary	5-136
Module Summary.....	5-137
Module Self-Check	5-139
Module Self-Check Answer Key	5-142

FCoE Troubleshooting

Overview

This module identifies common issues that relate to Fibre Channel over Ethernet (FCoE). The module also presents methods for troubleshooting these issues. Topics include issues that are related to FCoE, FCoE Initialization Protocol (FIP), and data center bridging.

Module Objectives

Upon completing this module, you will be able to identify and resolve issues that relate to FCoE in the Cisco data center architecture. This ability includes being able to meet these objectives:

- Identify and resolve issues that relate to FIP, FCoE, and FCoE performance
- Identify and resolve FCoE and FCoE performance issues that relate to incorrect or mismatched configuration

Troubleshooting FCoE

Overview

This lesson describes how to identify and resolve problems that can occur with Fibre Channel over Ethernet (FCoE) in Cisco Nexus or Cisco MDS Series switches.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that relate to FCoE Initialization Protocol (FIP), FCoE, and FCoE performance. This ability includes being able to meet these objectives:


- Explain how to troubleshoot FCoE issues that relate to FIP on a Cisco Nexus or Cisco MDS Series switch
- Explain how to troubleshoot FCoE issues that relate to QoS on a Cisco Nexus switch

Troubleshooting FIP

This topic explains how to troubleshoot FCoE issues that relate to FIP on a Cisco Nexus or Cisco MDS Series switch.

FCoE Overview

- Fibre Channel traffic requires a lossless transport.
- FCoE allows Fibre Channel traffic to be encapsulated over a physical Ethernet link.
- FCoE encapsulate a Fibre Channel frame in an Ethernet packet with a specific EtherType (0x8906 = FCoE, 0x8914 = FIP).
- The other header fields are standard Ethernet fields.



SOF = Start of Frame
EOF = End of Frame

© 2012 Cisco and/or its affiliates. All rights reserved.


DCUFT v5.0--4.4

Ethernet is a best-effort protocol: If congestion occurs, Ethernet discards packets and lets higher-level protocols provide retransmission and other reliability mechanisms if necessary. However, Fibre Channel traffic requires a lossless transport layer. As a data-storage protocol, losing even one data packet is unacceptable.

FCoE offers the capability to transport Fibre Channel payloads in addition to an Ethernet network. FCoE is implemented by encapsulating a Fibre Channel frame in an Ethernet packet with a specific EtherType: 0x8906 for FCoE or 0x8914 for FIP. The other header fields in the frame (the source and destination MAC addresses, VLAN tags, and frame markers) are all standard Ethernet fields.

FIP Overview

- FIP is a control protocol responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices over an Ethernet LAN.
- Two types
 - FIP
 - Pre-FIP



© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.6

FIP is the FCoE control protocol that is responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices that are connected to an Ethernet LAN. FIP and FCoE are both supported on the Cisco Nexus 5000 Series, 5500 Platform, and 7000 F-Series modules, and the 10-Gb/s 8-Port FCoE Module. FIP performs the device discovery, initialization, and link maintenance and is available in two modes:

- **FIP:** The Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) protocol supports T11-compliant, second-generation converged network adapters (CNAs).
- **Pre-FIP:** This protocol supports only first-generation CNAs; for example, Cisco Nexus 5500 Platform switches do not support Pre-FIP.

You can use these commands to discover the supported DCX Protocol:


```
Switch1# show system internal dcbx info interface ethernet 1/21
```

```
Interface info for if_index: 0x1a014000(Eth1/21)
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE
DCX Protocol: CEE
DCX CEE NIV extension: disabled
... <output omitted>
```

```
Switch2# show system internal dcbx info interface ethernet 1/18
```

```
Interface info for if_index: 0x1a011000(Eth1/18)
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE
DCX Protocol: CIN
DCX CEE NIV extension: disabled
... <output omitted>
```

Single-Hop FIP Virtual Link Establishment



© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.6


For single hop between ENode and Fibre Channel Forwarder (FCF), FIP aims to establish virtual Fibre Channel (vFC) links between virtual node ports (VN Ports) and virtual fabric ports (VF Ports). For multihop, the links are established between virtual expansion ports (VE ports). FIP between VN Ports and VF Ports performs the device discovery, initialization, and link maintenance and uses these protocols:

- **FIP VLAN Discovery:** Discovery of the FCoE VLAN is used by all other FIP protocols as well as by the FCoE encapsulation for Fibre Channel payloads on the established virtual link.
- **FCF Discovery:** A discovery solicitation message is sent out when an FCoE device is connected to the fabric. An FCF or switch responds to this message with a solicited advertisement that provides an FCF MAC address for subsequent logins.
- **FCoE Virtual Link Instantiation:** FIP defines the encapsulation of fabric login (FLOGI), fabric discovery (FDISC), logout (LOGO), and exchange link parameters (ELP) frames along with corresponding reply frames. FCoE devices use these messages to perform a FLOGI.
- **FCoE Virtual Link Maintenance:** Periodically, FIP sends out maintenance messages between the switch and the CNA. These messages are used to ensure that the connection is still valid.

FCoE has three Ethernet group addresses reserved for multicast operations:

- **ALL_FCoE_MACS:** 01-10-18-01-00-00, which is the group address for all FCoE devices
- **ALL_ENODE_MACS:** 01-10-18-01-00-01, which is the group address for all ENodes and is used by multicast discovery advertisements
- **ALL_FCF_MACS:** 01-10-18-01-00-02, which is the group address for all FCFs and is used by VLAN discovery request and multicast discovery solicitation

Multihop FCoE Example



Multihop FCoE is supported across VE Ports that are established between two Cisco Nexus 5500 Platform, Cisco Nexus 7000 Series, or Cisco MDS Series switches.

A VE Port is a port that emulates an E Port over a non-Fibre Channel link. This port is supported over point-to-point links between FCFs. These links can be Ethernet interfaces or members of an Ethernet port channel interface. For each FCF-connected Ethernet interface, you must create and bind a vFC interface to the Ethernet interface.

VE Ports have these guidelines:

- Auto mode on the vFC interface is unsupported.
- VE-Port trunking is supported over FCoE-enabled VLANs.
- VE-Port interface binding to MAC addresses is unsupported.
- A VE Port is enabled for trunk mode by default.

The Spanning Tree Protocol (STP) is disabled on the FCoE VLANs on any interface to which a vFC is bound, including the interfaces to which the VE Ports are bound.

The FIP virtual-link establishment process for multihop FCoE is similar to the single-hop process. After FCF discovery, there is a process of ELP between two VE Ports. Fibre Channel commands follow as part of the FCoE protocol.

vFC Down Because of FIP Failure

Problem	Host is incapable 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.
Solution	Check for correct FIP supporting firmware and drivers on the CNA and FIP supporting adapters.



```
N5548-3# show platform software fcoe_mgr info interface vfc 3
vfc3(0x83e5384), if_index: 0x1e000002, VFC RID vfc3
FSM current state: FCOE_MGR_VFC_ST_PHY_UP
<...>
FIP Capable ? : TRUE
<...>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.8

Problem

The host cannot support FIP-related type, length, values (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 DCBX exchanges over the bound interface the three basic TLVs that are required for FIP, 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 by using these commands:

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

```
N5548-3# sh platform software fcoe_mgr info interface vfc 3
vfc3(0x83e5384), if_index: 0x1e000002, VFC RID vfc3
FSM current state: FCOE_MGR_VFC_ST_PHY_UP
PSS Runtime Config:-
Type: 3
Bound IF: Eth1/3
Disable FKA: 0
PSS Runtime Data:-
IOD: 0x00000000, WWN: 20:02:54:7f:ee:3e:66:3f
Created at: Sat Jul 7 14:34:22 2012
```

```
FC Admin State: up
Oper State: up, Reason: down
Eth IF Index: Eth1/3
Port Vsan: 11
Port Mode: F port
Config Vsan: 11
Oper Vsan: 11
Solicits on vsan: 11
Isolated Vsan:
FIP Capable ? : TRUE
```

UP using DCBX ? : FALSE
Pinned Border Port : fc1/31

In the output from the commands, look for these items:


- FIP capable is TRUE.
- Triggered event is [FCOE_MGR_VFC_EV_FIP_VLAN_DISCOVERY].

The state of the vFC never progresses further to solicitation.

Solution

Make sure that you check for correct FIP-supporting firmware and drivers on the CNA and FIP-supporting adapters by using the compatibility matrix in your device release notes.

vFC Down Because of FIP Solicitation Failure



© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.0

Problem

When FIP solicitation fails, the vFC goes down.

Possible Cause

After 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 received solicitation, the vFC does not come up. The host continues trying to solicit, but never succeeds.

These reasons might be the cause for no response or advertisement:

- No active fabric-provided MAC address exists; for example possible wrong fc-map.
- Fabric is unavailable for FLOGI.
- The MAC address descriptor might be incorrect. (The CNA uses this address as the destination MAC [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.

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.

Solution

You need to ensure that the VSAN is active, the memberships are correct, and the fabric is available. Also while in N-Port Virtualization (NPV) mode, confirm that an active proxy N port (NP port) is available.


Check the output from the **ethalyzer local interface inbound-hi** command for any received type 0x8914 frames; for example:

```
N5548-3# ethalyzer local interface inbound-hi
Capturing on eth4
2012-07-07 07:41:37.710588 54:7f:ee:3e:66:26 -> 01:10:18:01:00:01 0x8914 PRI:
3 CFI: 0 ID: 1
2012-07-07 07:41:37.710738 54:7f:ee:3e:66:26 -> 01:10:18:01:00:01 0x8914 PRI:
3 CFI: 0 ID: 1
```

If no FIP frames are shown in the output, then packets might be dropped in hardware. The next step is to check for any packet drops. In Cisco Nexus 7000 Series switches, use the **attach mod number** and the **show hardware internal statistics pktflow dropped** commands to check for drops.

vFC Down Because VLAN Response Not Received by CNA

Problem	CNA does not receive VLAN response sent by the switch.
Possible Cause	If native VLAN matches the FCoE VLAN, the VLAN response sent out will be untagged.
Solution	Check the configuration on the bound Ethernet trunk interface. FIP protocol must be enabled for vFC to come up, and FCoE VLAN must be configured at CNA level or FIP will negotiate FCoE VLAN.



© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-10

Problem

Though the switch sends out a VLAN response, the CNA does not receive the response, indicating 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 that is sent out is untagged. However, the FIP adapters expect tagged frames. Therefore, the native VLAN on the trunk interface should be a non-FCoE VLAN.

Solution

Use the **show interface ethernet port trunk** command to check the configuration on the bound Ethernet trunk interface and ensure that it is a non-FCoE native VLAN.

FIP must be enabled for vFC to come up, and the FCoE VLAN must be configured at the CNA level or FIP will negotiate the FCoE VLAN.

vFC Down Because of No Active STP Port State on the Bound Ethernet Interface

Problem	No active STP port state is on the bound Ethernet interface.
Possible Cause	The bound interface should be in an STP-forwarding state for both the native VLAN and the member FCoE VLAN mapped to the active VSAN.
Solution	Check the STP port state on the bound Ethernet trunk interface.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-4.11

Problem

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

Possible Cause

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

Solution

Use the **show interface ethernet port trunk** command, **show span vlan native_vlan** command, and **show span vlan fcoe member vlan** command to check the STP port state on the bound Ethernet trunk interface for both the non-FCoE native VLAN and FCoE member VLAN. Fix the STP-port state and move it to forwarding if it is in a blocked inconsistent state or error-disable state.

In this example, all states are forwarding. VLAN 1 is the native VLAN and VLAN 1011 is an FCoE member VLAN:

```
N5548-3# sh span vlan 1
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID    Priority    32769
            Address     547f.ee3e.6641
            This bridge is the root
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
            Address     547f.ee3e.6641
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Eth1/3	Desg	FWD	2	128.131	P2p
Eth1/9	Desg	FWD	2	128.137	P2p

Eth1/10 Desg **FWD** 4 128.138 P2p

N5548-3# **sh span vlan 1011**

VLAN1011

Spanning tree enabled protocol rstp

Root ID Priority 33779

Address 547f.ee3e.6641

This bridge is the root

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 33779 (priority 32768 sys-id-ext 1011)

Address 547f.ee3e.6641

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----------	------	-----	------	----------	------

Eth1/9	Desg	FWD	2	128.137	P2p
--------	------	------------	---	---------	-----

vFC Down Because of FIP Keepalive 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>**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-4-12

Problem

The vFC goes down because of FIP keepalive misses.

Possible Cause

When FIP keepalives (FKAs) are missed for approximately 22 seconds, approximately three FKAs are not being continuously received from the host. Missed FKAs can occur for many reasons, including congestion or link issues.

An FKA timeout is equal to $2.4 * \text{FKA_adv_period}$. The FKA_adv_period is exchanged and agreed upon with the host, as in the FIP advertisement when responding to a solicitation.

FKA failure can also occur because of failures with the FIP multicast advertisement from the peer FCF.

Solution

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

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

Observe the output from these 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**

Troubleshooting Multihop FIP

- Some issues specific to multihop FIP:
 - FIP ELP issues
 - VE port-state transition errors
- Solution:
 - **show flogi database** to check the FLOGI database
 - **show run int vfc <x>** to check the virtual interface configuration
 - **show run int eth <x/y>** to check the physical interface configuration
 - **show interface ethernet <port> trunk** to check STP state forwarding in the vFC-mapped VLAN
 - **show vlan fcoe** to determine whether VLAN FCoE is enabled
 - **show platform software fcoe_mgr info interface vfc <x>** to check the vFC state machine on the virtual interface
 - **show platform software fcoe_mgr info interface Ethernet <x/y>** to check the vFC state machine on the physical interface

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-13

These are some issues that are specific to multihop FIP:

- FIP ELP issues
- VE Port-state transition errors

Solution

- Use the **show flogi database** command to check the FLOGI database.
- Use the **show run int vfc x** command and **show run int eth x/y** command to check the interface configuration.
- Use the **show interface ethernet port trunk** command to check STP state forwarding in the vFC mapped VLAN.
- Use the **show vlan fcoe** command to check if VLAN FCoE is enabled.
- Use the **show platform software fcoe_mgr info interface vfc x** or **show platform software fcoe_mgr info interface Ethernet x/y** command to check the vFC state machine.

Troubleshooting FCoE Performance

This topic explains how to troubleshoot FCoE issues that relate to quality of service (QoS) on a Cisco Nexus switch.

Verifying an MQC Object

Command	Purpose
show class-map [<i>type qos</i>] [<i>class-map-name</i> <i>conform-color-in</i> <i>conform-color-out</i> <i>exceed-color-in</i> <i>exceed-color-out</i>]	Displays information about all configured class maps or a selected class map of type QoS
show class-map [<i>type queuing</i>] [<i>class-queuing-name</i>]	Displays information about all configured class maps or a selected class map of type queuing
show table-map [<i>table-map-name</i> <i>cir-markdown-map</i> <i>pir-markdown-map</i>]	Displays information about all configured table maps or a selected table map
show policy-map [<i>type qos</i>] [<i>policy-map-name</i> <i>qos-dynamic</i>]	Displays information about all configured policy maps or a selected policy map of type QoS
show policy-map [<i>type queuing</i>] [<i>policy-map-name</i> <i>qos-dynamic</i>]	Displays information about all configured policy maps or a selected policy map of type queuing

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4-15

To display Cisco Modular QoS CLI (MQC) object configuration information on Cisco Nexus 5500 Platform or Nexus 7000 F-1 Series switches, perform one of the tasks that the figure shows. Also look at these output examples:

```
N5548-2# show class-map type qos
```

```
Type qos class-maps
```

```
=====
```

```
class-map type qos match-any class-fcoe
  match cos 3
```

```
class-map type qos match-any class-default
  match any
```

```
class-map type qos match-any class-all-flood
  match all flood
```

```
class-map type qos match-any class-ip-multicast
  match ip multicast
```

```
N5548-2# show policy-map type queuing
```

```
Type queuing policy-maps
```

```
=====
```

```
policy-map type queuing default-in-policy
  class type queuing class-default
    bandwidth percent 100
```

```
policy-map type queuing default-out-policy
  class type queuing class-default
    bandwidth percent 100
```

```
policy-map type queuing fcoe-default-in-policy
  class type queuing class-fcoe
    bandwidth percent 50
  class type queuing class-default
    bandwidth percent 50
policy-map type queuing fcoe-default-out-policy
  class type queuing class-fcoe
    bandwidth percent 50
  class type queuing class-default
    bandwidth percent 50
```

Types of Policy Maps

Policy Type	Function	Attach Point
QoS	Define traffic classification	- System QoS - Ingress interface
Queuing	- Strict priority queue - Deficit weighted round robin	- System QoS - Egress interface - Ingress interface
Network-QoS	- Define flow control mechanism (pause or tail drop) - MTU per CoS - Queue size - Marking	System QoS

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-4-16

Three steps are necessary to configure the QoS that is based on the Cisco MQC model:

Step 1 Define the class map.

Step 2 Create a policy map to define the action that is taken for each class map.

Step 3 Apply the policy map.

There are three types of policies:

- **Network-QoS:** Defines the characteristics of QoS properties networkwide
- **QoS:** Defines MQC objects that you can use for marking and policing
- **Queuing:** Defines MQC objects that you can use for queuing and scheduling, as well as a limited set of marking objects

In the basic process, the incoming packets are compared to the QoS classification rules that are the QoS policy-map type defines. The packets are classified into one of eight QoS groups.

Next, the Network-QoS and Queuing policies are applied to the packets. These policies define actual QoS parameters for packets that belong to each QoS group.

FCoE Bandwidth Allocation

- Associate an egress policy map with an Ethernet interface, to guarantee the bandwidth.
- The bandwidth allocation limit applies to all traffic on the interface.
- Each Ethernet interface supports as many as eight queues with these default configurations:
 - Queue zero
 - FCoE traffic
 - Standard Ethernet traffic

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-17

You can associate an egress policy map with an Ethernet interface, to guarantee the bandwidth for the specified traffic class or to configure the egress queues.

The bandwidth allocation limit applies to all traffic on the interface (including any FCoE traffic).

Each Ethernet interface supports as many as eight queues (one for each system class). The queues have this default configuration:

- Queue zero is configured as a strict priority queue. Control traffic that is destined for the CPU uses this queue.
- FCoE traffic (traffic that maps to the FCoE system class) is assigned a queue. This queue uses weighted round robin (WRR) scheduling with 50 percent of the bandwidth.
- Standard Ethernet traffic (in the default drop system class) is assigned a queue. This queue uses WRR scheduling with 50 percent of the bandwidth.

Troubleshooting FCoE Bandwidth Allocation

Problem	The traffic is incorrectly queued or prioritized.
Possible Cause	The Cisco Nexus 2148 Fabric Extender, Cisco Nexus 2232 10GE Fabric Extenders, and Nexus 2248TP Fabric Extenders can support only CoS-based traffic classification.
Solution	Mark the traffic with a CoS value on the Cisco Nexus switch.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.18

Problem

After configuring all three types of policy maps (QoS, Network-QoS, and Queuing), the traffic is not queued or prioritized correctly on Cisco Nexus 2148, Nexus 2232TM 10GE, and Nexus 2248TP fabric Extenders.

Possible Cause

The Cisco Nexus 2148, Nexus 2232TM 10GE, and Nexus 2248TP Fabric Extenders can support only class of service (CoS)-based traffic classification. The QoS service policy type that is configured under System QoS is populated from the Cisco Nexus 5000 Series, Nexus 5500 Platform, or Nexus 7000 F-1 Series switches to the fabric extender only when all the matching criteria are “match CoS”. If other match clauses, such as “match dscp” or “match ip access-group”, exist in the QoS policy map, then the fabric extender does not accept the service policy. As a result, all the traffic is placed into the default queue.

Solution

Ingress traffic (from server to network) that is not marked with a CoS value is placed into the default queue on the fabric extender. After the traffic is received on the Cisco Nexus 5000 Series, Nexus 5500 Platform, or Nexus 7000 F-1 Series switch, that traffic is classified based on a configured rule and is placed in the proper queue. For egress traffic (from one of these switches to the fabric extender, and then from the fabric extender to the server), you should mark the traffic with a CoS value on the switch so that the fabric extender can properly classify and queue the traffic.

Verifying Jumbo MTU

```
N5K-1(config)# show interface ethernet 2/1 counters detailed
Rx Packets: 1547805598
Rx Unicast Packets: 1547805596
Rx Jumbo Packets: 1301767362
Rx Bytes: 7181776513802
Rx Storm Suppression: 33690
Rx Packets from 0 to 64 bytes: 169219
Rx Packets from 65 to 127 bytes: 10657133
Rx Packets from 128 to 255 bytes: 21644488
Rx Packets from 256 to 511 bytes: 43290596
Rx Packets from 512 to 1023 bytes: 86583071
Rx Packets from 1024 to 1518 bytes: 83693729
Rx Trunk Packets: 1547805596
Tx Packets: 1186564481
Tx Unicast Packets: 1005445334
Tx Multicast Packets: 7063
Tx Jumbo Packets: 997813205
Tx Bytes: 4813632103819
Tx Packets from 0 to 64 bytes: 137912
Tx Packets from 65 to 127 bytes: 8288443
Tx Packets from 128 to 255 bytes: 16596457
Tx Packets from 256 to 511 bytes: 33177999
Tx Packets from 512 to 1023 bytes: 66363944
Tx Packets from 1024 to 1518 bytes: 64186521
Tx Trunk Packets: 1005451729
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4.19

To verify that the jumbo maximum transmission unit (MTU) is enabled, enter the **show interface ethernet x/y** command for an Ethernet interface that carries traffic with the jumbo MTU.

To display detailed jumbo MTU information for a specific interface, use the **show interface ethernet x/y counters detailed** command.

Cannot Pass Jumbo Frames Through Switch

Problem	The jumbo MTU has been configured for class-default but jumbo frame cannot pass through the Cisco Nexus switch.
Possible Cause	The CoS value might conflict with the existing MTU value.
Solution	Use any CoS value other than 7 to avoid CoS 7 fixed value.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4-20

Problem

Although the jumbo MTU has been configured for class-default, you cannot pass a frame size larger than 2300 bytes through the Cisco Nexus 5500 Series or Nexus 7000 F-1 Series switch and the Cisco Nexus 2000 Series Fabric Extender.

Possible Cause

The CoS value might conflict with the existing MTU value.

Solution

CoS 7 is used internally to control traffic between the Cisco Nexus 5000 Series, Nexus 5500 Platform, or Nexus 7000 F-1 Series switch and the Cisco Nexus 2000 Series Fabric Extender. 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.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- FIP is the FCoE control protocol responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices.
- The Cisco Nexus 5000 Series, 5500 Platform, and Cisco Nexus 7000 F-1 Series devices implement three types of policy maps: QoS, Queuing, and Network-QoS.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-21

Troubleshooting DCBX

Overview

This lesson is designed to provide you with some examples of Data Center Bridging (DCB) and priority flow control (PFC) issues and show you how to identify and resolve these issues.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that relate to Fibre Channel over Ethernet (FCoE) and FCoE performance, as it relates to incorrect configuration and configuration mismatch. This ability includes being able to meet these objectives:

- Explain how to troubleshoot FCoE issues that relate to DCBX on a Cisco Nexus switch
- Explain how to troubleshoot FCoE issues that relate to PFC on a Cisco Nexus switch

Troubleshooting DCBX

This topic explains how to troubleshoot FCoE issues that relate to Data Center Bridging Exchange (DCBX) on a Cisco Nexus switch.

Troubleshooting DCBX

- Problem: The vFC (FCoE) interface is not online.
- Possible causes
 - An FCoE-attached server has no connectivity to Fibre Channel or to FCoE-attached storage.
 - LLDP Tx and Rx are not enabled on the interface.
 - The peer does not support LLDP.
- Review every feature-negotiation result. Use the **show system internal dcbx info interface ethernet x/y** command.

```
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
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4.4

Possible Cause

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

Verify the configuration by using the **show running-config** command. The default setting for vFC is shutdown. However, in this example the setup script changed that setting:

```
switch# show running-config
<part of the output omitted>
feature fcoe
vlan 1, 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
<rest of the output omitted>
```

Use the **show lldp interface ethernet x/y** command to ensure that Link Layer Discovery Protocol (LLDP) transmit (Tx) and receive (Rx) are enabled on the interface and that the peer supports LLDP, and to check the LLDP type, length, values (TLVs) for a peer:

```
switch# show lldp interface ethernet 1/4
<part of the output omitted>
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
<rest of the output omitted>
```

If LLDP is disabled, the vFC will not come online. You can enable LLDP Tx and Rx by using the **lldp interface** subcommand:

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

If the **show lldp interface ethernet x/y** command indicates that the peer might not support LLDP, check the peer (the converged network adapter [CNA]) to determine whether it supports DCBX. Use the **show system internal dcbx info interface ethernet x/y** or **show lldp dcbx interface ethernet x/y** command. In this example, DCBX is enabled and the peer supports Converged Enhanced Ethernet (CEE):

```
switch# show system internal dcbx info interface ethernet 1/4
<part of the output omitted>
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
<rest of the output omitted>
```

Troubleshooting DCBX (Cont.)

- Errors
 - Indicates negotiation error
 - Never expected to happen when connected to CNA
 - Can occur when two Cisco Nexus 5000 Series switches are connected back-to-back and PFC is enabled on different CoS values
- Operating configuration
 - Indicates negotiation result
 - Absence indicates that peer does not support the DCBX TLV or a negotiation error
 - The `remote_feature_tlv_present` message indicates whether the remote peer supports this feature TLV

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.6

In the output from the **show system internal dcbx info interface ethernet x/y** command, check the peers LLDP values and look for any errors:

```
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
<part of the output omitted>
Traffic Counters
DCBX pkt stats:
```

```
    Total frames out: 28814
    Total Entries aged: 1
    Total frames in: 28814
    DCBX frames in: 28812
    Total frames received in error: 2
    Total frames discarded: 2
    Total TLVs unrecognized: 0
<rest of the output omitted>
```

In the output from the **show system internal dcbx info interface ethernet x/y** command, check the peer DCBX TLVs. Make sure that the PFC and FCoE TLV were negotiated as willing and enabled, and that no errors exist:

```
switch# show system internal dcbx info interface ethernet 1/4
<part of the output omitted>
Peer's DCX TLV:
DCBX TLV Proto(1) type: 1(Control) DCBX TLV Length: 10 DCBX TLV Value
00 00 02 00 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 01 00 00 32 32 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(Fcoe)) 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
<rest of the output omitted>
```

Troubleshooting DCBX (Cont.)

- DCBX feature might not be working because:
 - Peer does not support LLDP.
 - Peer does not support DCBX.
 - Peer does not support some DCBX TLVs.
 - Unexpected DCBX negotiation result occurred.
- If DCBX feature is not supported, an option exists to force PFC mode on an interface. Use the **interface ethernet x/y** command and the **priority-flow-control mode** command to force PFC mode.

```
switch(config)# int eth1/21
switch(config-if)# priority-flow-control mode ?
auto Advertise priority-flow-control capability
on Turn on priority-flow-control
```

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

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.6

Check the DCBX counters at the bottom of the output of the **show system internal dcbx info interface ethernet x/y** command. Look for any errors:

```
switch# show system internal dcbx info interface ethernet 1/4
<part of the output omitted>
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
<rest of the output omitted>
```

Check for the same values for the FCoE DCB and the TLV on the host CNA software.

Troubleshooting Cisco Nexus 5548

- 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

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.7

Possible Cause

The FCoE class-fcoe system class is not enabled in the quality of service (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 these policy types: Network-QoS, Queuing, and QoS.

This example shows a service policy that needs to be configured:

```
NN548-3#show run
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
```

Troubleshooting PFC

This topic explains how to troubleshoot FCoE issues related to PFC on a Cisco Nexus Switch.

PFC Not Negotiated with FCoE-Capable Adapters

- Possible causes
 - The CNA might not support DCBX, and the PFC TLV is not negotiated.
- Solution
 - Check the status of the PFC. Use the **show int eth <x/y> priority-flow-control** command. (Connected to CNA.)
 - Check for LLDP neighbor or PFC or DCBX TLV advertised by the peer. Use the **show system internal dcbx info int eth x/y** command.
 - If the peer does not support DCBX, configure the priority-flow-control mode setting to on to enable PFC.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.0

Problem

PFC is not negotiated with FCoE-capable adapters (CNAs). Therefore, packet drop can be noticed on FCoE traffic from the servers.

Possible Causes

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

Solution

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

- Check the status of the PFC. Use the **show interface ethernet x/y priority-flow-control** command on the interface that is connected to CNA:

```
switch# show interface ethernet 1/13 priority-flow-control
=====
Port                Mode   Oper(VL bmap)      RxPPP  TxPPP
=====
Ethernet1/13       Auto   Off                 0      0
=====
```

- Check for LLDP neighbor or the PFC or DCBX TLV that the peer advertised. Use the **show system internal dcbx info int ethernet x/y** command:

```
switch(config-if)# show system internal dcbx info interface ethernet 1/1
<part of the output omitted>
Interface info for if_index: 0x1a000000(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
<rest of the output omitted>

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

Switch Interface Connected to CNA Receives Constant Pause Frames

- If the Cisco Nexus 5000 Series switch is connected to a CNA, then the CNA might be sending Xon or Xoff PFC frames to the switch.
- This increments the pause counters when using the **show interface eth x/x** command.
- To verify this situation, perform these steps:
 - Use the **show interface eth x/x |grep - i pause** command and ensure that the pause frame count is incrementing.
 - Use the **show interface eth x/x priority-flow-control** command and ensure that the PFC frame count is incrementing.
 - Use the **show queuing interface eth x/x** command and check the pause status.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-10

Problem

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

Possible Cause

If the Cisco Nexus 5000 Series switch is connected to a CNA along with slow servers that are unable to process 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 ethernet x/y** command. To verify this situation, follow these steps:

- Step 1** For a few iterations, use the **show interface ethernet x/y |grep - i pause** command and ensure that the pause frame count is incrementing.
- Step 2** For a few iterations, use the **show interface ethernet x/y priority-flow-control** command and ensure that the PFC frame count is incrementing.
- Step 3** For a few iterations, use the **show queuing interface ethernet x/y** command and check the pause status:

Per-priority-pause status : Rx (Active), Tx (Inactive)

If the Rx (Active) and pause counts increment (as shown with the **show interface ethernet x/y priority-flow-control** command), then the issue is probably caused by Xoff frames that are received from the server.

Solution

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

Is the Switch Sending Pause Frames or Being 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 the FCoE rate and avoid a drop.
 - If the server is slow or congested, the server sends PFC frames to the switch interface.
- Solution
 - Identify the source of the congestion.
 - Try to resolve the congestion by increasing the Fibre Channel bandwidth or changing to a more powerful server.
 - If congestion is expected, then pause is expected for FCoE traffic.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-4-11

Problem

FCoE throughput on servers is low because of pause frames from the switch. You must determine whether the switch is sending pause frames or is being paused.

Possible Cause

If the egress Fibre Channel port is congested, the switch sends PFC frames to the servers. The PFC frames are sent to reduce the FCoE rate and avoid a drop. To verify this situation, perform these steps:

- Step 1** For a few iterations, use the **show interface ethernet x/y |grep -i pause** command and ensure that the pause frame count (Rx and TX) is incrementing.
- Step 2** For a few iterations, use the **show interface ethernet x/y priority-flow-control** command and ensure that the PFC frame count (RX and TX) is incrementing.
- Step 3** For a few iterations, use the **show queuing interface ethernet x/y** command to check the pause status:

Per-priority-pause status : Rx (Active), Tx (Inactive)

If the Tx (Active) and pause Tx counters increment (as shown with the **show interface ethernet x/y priority-flow-control** command), then the issue is probably caused by Xoff frames that the switch is transmitting.

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

Solution

Identify the source of the congestion and try to resolve it by increasing the Fibre Channel bandwidth. If congestion is expected, then pause is expected for FCoE traffic.

Switch Ports Error-Disabled Because of Pause Rate-Limit

- Possible cause
 - If the switch interface receives excessive Xoff pause frames from the server, ports become error-disabled because of the high rate of received pause frames.
- Solution
 - If the port is error-disabled because of one of these transient conditions, 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 because of the pause rate limit, determine whether the issue is that the server is too slow.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.12

Possible Cause

If the switch interface receives excessive Xoff pause frames from the server, then ports become error-disabled because of the high rate of received pause frames. The port usually goes into an error-disabled state because of pause frames only when the drain rate is less than 5 Mb/s on a 10-Gb port. This rate means that the server is slow and is sending many pause frames to the switch ports. To verify this situation, use the **show interface ethernet x/y brief** command:

```
switch# show interface ethernet 1/14 brief
```

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

- Determine whether the Rx pause count is a large value. Use the **show interface ethernet x/y** command to display the pause counters.
- Look for pause error-disable logs by using the **show hardware internal gatos event-history errors |grep -i err** command.

Solution

Pause error-disable recovery can be enabled to get the ports out of this state, if the port is error-disabled because of one of these transient conditions:

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

If a consistent port error-disable condition occurs because of the pause rate limit, determine whether the issue is that the server is too slow. If so, replace the slow server.

How to Enable Link Pause (Flow Control) on 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 need to disable PFC TLV sent by DCBX on the interface.
- Solution
 - Use these commands under the **interface eth x/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**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-4-13

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

Possible Cause

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

To verify this situation, perform one of these actions:

- Use the **show interface ethernet x/y flow control** command to determine whether the operating state is off.
- Use the **show interface ethernet x/y priority-flow-control** command to determine whether the operating state is on.

Solution

Use these commands under the **interface ethernet x/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**

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- DCBX is used to negotiate capabilities between the Cisco Nexus 5000 Series switch and the adapter in the server and to send configuration values to the adapter.
- Use the **show int eth x/x priority-flow-control** command to determine the status of the PFC.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--4-14

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- FIP is the FCoE control protocol responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices.
- DCBX is used to negotiate capabilities between the Cisco Nexus 5000 Series switch and the adapter in the server and to send configuration values to the adapter.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-4.1

The Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN. FCoE quality of service (QoS) is a must for proper performance. Before enabling FCoE, you must include class-fcoe in each of these policy types: Network-QoS, Queuing, and QoS.

Data Center Bridging Exchange (DCBX) protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX runs on the physical Ethernet link between the Cisco Nexus 5000 Series Switch and the converged network adapter on the server. DCBX is used to negotiate capabilities between the switch and the adapter and to send configuration values to the adapter. The capability reduces the possibility of configuration error and simplifies administration of the adapters. Priority flow control (PFC) allows you to apply pause functionality to specific classes of traffic on a link instead of to all the traffic on the link. PFC applies pause functionality that is based on the IEEE 802.1p class of service (CoS) value. When the switch enables PFC, it informs the adapter as to which CoS values the adapter should apply the pause.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) FIP aims to establish vFC links between which types of ports? (Source: Troubleshooting FCoE)
- A) N Ports and F Ports
 - B) VN Ports and VF Ports
 - C) VE Ports and VF Ports
 - D) E Ports and E Ports
- Q2) Which process defines the encapsulation of FLOGI, FDISC, LOGO, and ELP frames? (Source: Troubleshooting FCoE)
- A) FIP discovery
 - B) FCoE virtual-link maintenance
 - C) FCoE virtual-link instantiation
 - D) FCoE discovery
- Q3) Which mode should be enabled on the Ethernet interface to support the transport of both Ethernet and FCoE frames? (Source: Troubleshooting FCoE)
- A) switchport mode trunk
 - B) switchport mode access
 - C) switchport mode fcoe
 - D) fcoe mode on
- Q4) Which message is sent to an individual virtual link to reset the virtual interface? (Source: Troubleshooting FCoE)
- A) FCoE Clear Virtual Link
 - B) FCoE Virtual Link Clear
 - C) FIP Virtual Link Clear
 - D) FIP Clear Virtual Link
- Q5) Which DCB Ethernet enhancement feature supports pause on a virtual channel? (Source: Troubleshooting DCBX)
- A) enhanced transmission selection
 - B) link level flow control
 - C) Layer 2 multipathing
 - D) PFC
- Q6) For ports with standard, non-CNA type host connections, the Cisco Nexus 5000 Series switch supports standard pause frames. (Source: Troubleshooting DCBX)
- A) true
 - B) false
- Q7) Priority groups are associated with which DCB Ethernet enhancement feature? (Source: Troubleshooting DCBX)
- A) PFC
 - B) enhanced transmission selection
 - C) end-to-end congestion management
 - D) Layer 2 multipathing

- Q8) Which protocol performs discovery of DCB capabilities in a peer? (Source: Troubleshooting DCBX)
- A) DCBX
 - B) DCBX Control Protocol
 - C) Cisco Discovery Protocol
 - D) LLDP

Module Self-Check Answer Key

Q1)	B
Q2)	C
Q3)	A
Q4)	D
Q5)	D
Q6)	A
Q7)	B
Q8)	A

Platform-Specific Issue Troubleshooting

Overview

This module identifies common issues that relate to the Cisco Nexus 5000 Series, Nexus 7000 Series, and Cisco MDS Series switches, as well as to the Cisco Nexus 2000 Series Fabric Extenders. The module also presents methods for troubleshooting and resolving these issues.

Module Objectives

Upon completing this module, you will be able to identify and resolve platform-specific issues in the Cisco data center architecture. This ability includes being able to meet these objectives:

- Identify and resolve issues that relate to Cisco Nexus 7000 Series Switches
- Identify and resolve issues that are specific to Cisco Nexus 5000 Series Switches
- Identify and resolve issues that are specific to Cisco Nexus 2000 Series Fabric Extenders
- Identify and resolve issues that are specific to Cisco MDS Series switches

Troubleshooting Cisco Nexus 7000 Series Switches

Overview

This lesson is designed to provide you with some examples of common issues that relate to Cisco Nexus 7000 Series Switches and methods to resolve those issues.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that relate to Cisco Nexus 7000 Series Switches. This ability includes being able to meet these objectives:

- Explain how to troubleshoot issues that relate to licensing on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to installs, upgrades, and reboots on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to Cisco Fabric Services on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to VDCs on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to routing on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to unicast forwarding on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to memory on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to CPU on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to the switch fabric on a Cisco Nexus 7000 Series Switch
- Explain how to troubleshoot issues that relate to CoPP and rate limiting on a Cisco Nexus 7000 Series Switch

Troubleshooting Licensing

This topic explains how to troubleshoot issues that relate to licensing on a Cisco Nexus 7000 Series Switch.

About Troubleshooting Licensing Issues

- Cisco NX-OS requires licenses for selected features.
- Cisco NX-OS provides a grace period (120 days) that allows you to try out the feature before purchasing the license.
- Cisco NX-OS disables the feature at the end of the grace period.
- Licenses are created by using the serial number of the chassis on which the license file is to be installed.
- If you swap out a chassis which included licenses, you can contact Cisco TAC or go to <http://www.cisco.com/go/license> to generate a new license.
- Use the **show license usage** command to determine which features are enabled for a license package.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5.5

The Cisco Nexus Operating System (NX-OS) requires licenses for selected features. The licenses enable those selected features on your system. You must purchase a license for each system on which you want to enable the licensed features. However, there is a way to enable a feature without installing the license—Cisco NX-OS provides a grace period that allows you to try out the feature before purchasing the license.

Licenses are created by using the serial number of the chassis where the license file is to be installed. When you order a license that is based on a chassis serial number, you cannot use that license on any other system. If you swap out a chassis that included a license, you can contact Cisco Technical Assistance Center (TAC) to generate a new license. The old license was based on the chassis serial number and will not work with the new chassis.

If you use a feature that requires a license that you have not installed, you are given a 120-grace period to evaluate the feature. You must purchase and install the number of licenses that are required for that feature before the grace period ends, or Cisco NX-OS disables the feature at the end of the grace period.

License packages can contain several features. If you disable a feature during the grace period and other features in that license package are still enabled, the clock does not stop for that package. To suspend the grace period countdown for a licensed feature, you must disable every feature in that license package. Use the **show license usage** command to determine which features are enabled for a license package.

Licensing Guidelines

- Allow 60 days before the grace period expires to allow time for ordering, shipping, and installation of a new license purchase.
- Carefully determine the license (or licenses) that you require, based on the features that require a license.
- Accurately order your license.
- Back up the license file to a remote, secure place.
- Install the correct licenses on each system, using the licenses that were ordered using that system's serial number.
- Use the **show license usage** command to verify the license installation.
- Never modify a license file or attempt to use it on a system other than the one for which it was ordered.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.6

Follow these guidelines when dealing with licenses for Cisco NX-OS:

- Do not ignore the grace-period expiration warnings. Allow 60 days before the grace period expires, to allow time for ordering, shipping, and installation of a new license purchase.
- Carefully determine the license (or licenses) that you require, based on the features that require a license.
- Accurately order your license:
 - Enter the Product Authorization Key (PAK) that appears in the Proof of Purchase document that comes with your system.
 - Enter the correct chassis serial number when ordering the license. The serial number must be for the same chassis on which you plan to install the license. Use the **show license host-id** command to obtain your chassis serial number.
 - Enter serial numbers accurately. Do not use the letter "O" instead of a zero in the serial number.
 - Order the license that is specific to your chassis.
- Back up the license file to a remote, secure place. Archiving your license files ensures that you will not lose the licenses if a failure occurs on your system.
- Install the correct licenses on each system, using the licenses that were ordered with the system serial number. Licenses are serial number- and platform-specific.
- Use the **show license usage** command to verify the license installation.
- Never modify a license file or attempt to use it on a system other than the one for which it was ordered. If you return a chassis, contact your customer support representative to order a replacement license for the new chassis.

Initial Troubleshooting Checklist

- Verify the chassis serial number for all ordered licenses.
- Verify the platform or module type for all ordered licenses.
- Verify that the PAK that you used to order the licenses comes from the same chassis from which you retrieved the chassis serial number.
- Verify that you have installed all licenses on all systems that require the licenses for the features you enable.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0—5-7

Begin troubleshooting license issues by checking these issues first:

- Verify the chassis serial number for all ordered licenses:
 - **show license host-id**
- Verify the platform or module type for all ordered licenses:
 - **show module**
- Verify that the PAK that you used to order the licenses comes from the same chassis for which you retrieved the chassis serial number.
- Verify that you have installed all licenses on all systems that require the licenses for the features that you enable:
 - **show license usage**

You can also configure the Call Home feature to receive an email when there is an issue with a license.

Displaying License Information by Using the CLI

```
switch(config)# show license usage
Feature                               Ins  Lic  Status Expiry Date Comments
                                Count
-----
LAN_ADVANCED_SERVICES_PKG            No   -   In use          Grace 102D 0H
LAN_ENTERPRISE_SERVICES_PKG          No   -   In use          Grace 103D 22H
-----

switch(config)# show license usage LAN_ENTERPRISE_SERVICES_PKG
Application
-----
pbr
Tunnel
-----

switch# show license
<...>
INCREMENT LAN_ENTERPRISE_SERVICES_PKG cisco 1.0 permanent uncounted \
  VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>N7K-LAN1K9=</SKU> \
  HOSTID=VDH=TBC10412106 \
  NOTICE="<LicFileID>20071025133322456</LicFileID>LicLineID>1/LicLineID> \
  <PAK></PAK>"
<...>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.6

Use the **show license** commands to display all license information that is configured on this system.

- This example displays information about current license usage:

```
switch(config)# show license usage
```

Feature	Ins	Lic	Status	Expiry Date	Comments
		Count			
LAN_ADVANCED_SERVICES_PKG	No	-	In use	Grace 102D 0H	
LAN_ENTERPRISE_SERVICES_PKG	No	-	In use	Grace 103D 22H	

- This example displays the list of features in a specified package:

```
switch(config)# show license usage LAN_ENTERPRISE_SERVICES_PKG
```

```
Application
```

```
-----
pbr
Tunnel
-----
```

- This example displays the host ID for the license:

```
switch# show license host-id
```

```
License hostid: VDH=FOX0646S017
```

Use the entire ID that appears after the colon (:). The VHD indicates the vendor host ID.

- This example displays all installed license key files and contents.

```
switch# show license
```

```
<part of the output omitted>
```

```
INCREMENT LAN_ENTERPRISE_SERVICES_PKG cisco 1.0 permanent uncounted \  
    VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>N7K-  
    LAN1K9=</SKU> \  
    HOSTID=VDH=TBC10412106 \  
    NOTICE="<LicFileID>20071025133322456</LicFileID>LicLineID>1/Li  
    cLineID> \  
    <PAK></PAK>"
```

```
<rest of the output omitted>
```

Licensing Installation Issues

Common problems with licenses usually result from incorrectly ordering the license file, installing the license file on an incorrect system, or not ordering the correct number of licenses for your fabric.

Serial Number Issues

- Make sure that you use the correct chassis serial number when ordering your license.
- Use the **show license host-id** command to obtain the correct chassis serial number for your system.

```
switch# show license host-id  
License hostid: VDH=FOX0646S017
```

- **Error Message:** LICMGR-3-LOG_LIC_INVALID_HOSTID: Invalid license hostid VDH=[chars] for feature [chars]
- **Explanation:** The feature has a license with an invalid license host ID.
- **Recommended action:** Reinstall the correct license for the chassis on which the supervisor module is installed.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.9

Make sure that you use the correct chassis serial number when ordering your license. Use the **show license host-id** command to use the CLI to obtain the correct chassis serial number for your system. If you use a license that is meant for another chassis, you might see this system message:

- **Error Message:** LICMGR-3-LOG_LIC_INVALID_HOSTID: Invalid license hostid VDH=[chars] for feature [chars].
- **Explanation:** The feature has a license with an invalid license host ID. This issue can occur when a supervisor module with licensed features for one system is installed on another system.
- **Recommended Action:** Reinstall the correct license for the chassis on which the supervisor module is installed.

When entering the chassis serial number during the license-ordering process, do not use the letter "O" instead of any zeros in the serial number.

A license is specific to the system and chassis for which it is issued and is valid on that system only. If you need to transfer a license from one chassis to another, contact your technical support representative.

License Listed as Missing

- Symptom:
 - A license is listed as missing.
- Possible cause:
 - The supervisor module was replaced after the license was installed.
 - The license file has already been installed and deleted somehow from the switch.
- Solution:
 - Reinstall the license.
- Possible cause:
 - Sometimes, at the time of manufacturing, few sprom bits may be set, leading to this problem.
- Solution:
 - Use the **clear license sprom** command to clear the error.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-10

After a license is installed and operating properly, it might go missing if you modify your system hardware.

If the license installation does not proceed correctly, or if you are using a feature that exists in a license package that you have not installed, you will continue to get grace-period warnings.

If the license file is copied to the system but is not installed, use the **license install** command to install the license. If the license installation failed, check your logs for any system messages for a failed license installation. Use the **show license usage** command to determine which feature is in use without a license.

Cisco NX-OS gives you a 120-day grace period. This grace period starts or continues when you are evaluating a feature for which you have not installed a license. The grace period stops if you disable a feature that you are evaluating. If you enable that feature again without a valid license, the grace period countdown continues where it left off.

The grace period operates across all features in a license package. License packages can contain several features. If you disable a feature during the grace period and other features in that license package are still enabled, then the countdown does not stop for that license package. To suspend the grace period countdown for a license package, you must disable every feature in the package.

The Cisco NX-OS license counter keeps track of all licenses on a system. If you are evaluating a feature and the grace period has started, you will receive console messages, Simple Network Management Protocol (SNMP) traps, system messages, and daily Call Home messages. The frequency of these messages becomes hourly during the last seven days of the grace period. When the grace period ends, the feature is automatically disabled. You are not allowed to use the feature until you purchase a valid license. You cannot modify the frequency of the grace-period messages.

Sometimes, at the time of manufacturing, few serial PROM (SPROM) bits might be set, leading to the problem of a missing license. Use the **clear license sprom** command to clear the error.

If you try to use an unlicensed feature, you might see one of the system messages that are listed in the following table.

System Messages that Relate to Unlicensed Features

Error Message	Explanation
LICMGR-2-LOG_LIC_GRACE_EXPIRED: Grace period expired for feature [chars]	The unlicensed feature has exceeded its grace period. Applications using this license will be shut down immediately.
LICMGR-3-LOG_LICAPP_NO_LIC: Application [chars] running without [chars] license, shutdown in [dec] days	The Application [chars] has not been licensed. The application will work for a grace period of [dec] days, after which it will be shut down unless a license file for the feature is installed.
LICMGR-3-LOG_LIC_LICENSE_EXPIRED: Evaluation license expired for feature [chars]	The feature has exceeded its evaluation time period. The feature will be shut down after a grace period.
LICMGR-3-LOG_LIC_NO_LIC: No license(s) present for feature [chars]. Application(s) shutdown in [dec] days.	The feature has not been licensed. The feature will work for a grace period, after which the application (or applications) that use the feature will be shut down.
LICMGR-6-LOG_LICAPP_EXPIRY_WARNING: Application [chars] evaluation license [chars] expiry in [dec] days	The application will exceed its evaluation period in the listed number of days and will be shut down unless a permanent license for the feature is installed.

Troubleshooting Installs, Upgrades, and Reboots

This topic explains how to troubleshoot issues that relate to installs, upgrades, and reboots on a Cisco Nexus 7000 Series Switch.

Upgrades and Reboot Checklist

- Read the release notes for the release to which you are upgrading or downgrading.
- Ensure that an FTP or TFTP server is available to download the software images.
- Copy the new image onto your supervisor modules in bootflash or slot0.
- Use the **show install all impact** command to verify that the new image is healthy and to determine the impact that the new load will have on any hardware, with regards to compatibility.
- Copy the startup configuration file to a snapshot configuration in NVRAM.
- Save your running configuration to the startup configuration.
- Back up a copy of your configuration to a remote TFTP server.
- Schedule your upgrade during an appropriate maintenance window for your network.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-12

Cisco NX-OS consists of two images: the kickstart image and the system image. To bring up the system, both images should have the same image version. Upgrades and reboots are ongoing network-maintenance activities. You should try to minimize the risk of disrupting the network when performing these operations in production environments, and you should know how to recover quickly when something goes wrong. Use this checklist to prepare for an upgrade:

- Step 1** Read the release notes for the release to which you are upgrading or downgrading.
- Step 2** Ensure that an FTP or TFTP server is available to download the software images.
- Step 3** Copy the new image onto your supervisor modules in bootflash or slot0.
- Step 4** Use the **show install all impact** command to verify that the new image is healthy and determine the impact that the new load will have on any hardware, with regards to compatibility. Check for compatibility.
- Step 5** Copy the startup configuration file to a snapshot configuration in NVRAM. This step creates a backup copy of the startup configuration file.
- Step 6** Save your running configuration to the startup configuration.
- Step 7** Back up a copy of your configuration to a remote TFTP server.
- Step 8** Schedule your upgrade during an appropriate maintenance window for your network.

After you have completed the checklist, you are ready to upgrade the systems in your network. The active supervisor typically becomes the standby supervisor during an upgrade. Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM. You can view this log at any time by entering the **show logging nvram** command.

Verifying Software Upgrades

```
switch# show install all status
There is an on-going installation...
Enter Ctrl-C to go back to the prompt.
Verifying image bootflash:/b-4.0.0.104
-- SUCCESS
Verifying image bootflash:/i-4.0.0.104
-- SUCCESS
Extracting system version from image bootflash:/i-4.0.0.104.
-- SUCCESS
Extracting kickstart version from image bootflash:/b-4.0.0.104.
-- SUCCESS
Extracting loader version from image bootflash:/b-4.0.0.104.
-- SUCCESS
```

In-Progress Installation

```
switch# show install all status
This is the log of last installation.
Verifying image bootflash:/b-4.0.0.104
-- SUCCESS
Verifying image bootflash:/i-4.0.0.104
-- SUCCESS
Extracting system version from image bootflash:/i-4.0.0.104.
-- SUCCESS
Extracting kickstart version from image bootflash:/b-4.0.0.104.
-- SUCCESS
Extracting loader version from image bootflash:/b-4.0.0.104.
-- SUCCESS
```

Log of Last Install

You can use the **show install all status** command to watch the progress of your software upgrade or to view the ongoing **install all** command or the log of the last installed **install all** command from a console, Secure Shell (SSH), or Telnet session. This command shows the **install all** output on both the active and standby supervisor modules, even when you are not connected to the console terminal.

Investigate Upgrade Failures

```
Switch# install all kickstart bootflash:///n7000-s1-kickstart.6.0.1.bin
<...>
Do you want to continue with the installation (y/n)? [n] y
Install is in progress, please wait.
Notifying services about the upgrade.
>[#          ] 0% -- FAIL. Return code 0x401E0066 (request timed out).
Please issue "show install all failure-reason" to find the cause of the failure.
Install has failed. Return code 0x401E0066 (request timed out).
Please identify the cause of the failure, and try 'install all' again.
```

Prompt for failure-reason

- If a service cannot allow the upgrade to proceed at this time, then the service aborts the upgrade.
- You are prompted to enter the **show install all failure-reason** command to determine why the upgrade cannot proceed.

```
switch# show install all failure-reason
Service: "cfs" failed to respond within the given time period.
switch#
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-14

When you initiate a nondisruptive upgrade, Cisco NX-OS notifies all services that an upgrade is about to start and finds out whether or not the upgrade can proceed. If a service cannot allow the upgrade to proceed, then the service aborts the upgrade and you are prompted to enter the **show install all failure-reason** command to determine the reason why the upgrade cannot proceed.

If a failure occurs for whatever reason (such as a save runtime state failure or module upgrade failure) after the upgrade is in progress, then the device reboots disruptively because the changes cannot be rolled back. In such cases, the upgrade has failed.

If you need further assistance to determine why an upgrade is unsuccessful, you should collect the details from the **show tech-support** command output and the console output from the installation, if available, before you contact your technical support representative.

Using ROM Monitor Mode

If your device does not find a valid system image to load, the system starts in ROM monitor mode. ROM monitor mode can also be accessed by interrupting the boot sequence during startup. From ROM monitor mode, you can boot the device or perform diagnostic tests.

On most systems, you can enter ROM monitor mode by entering the **reload EXEC** command and then pressing the Break key on your keyboard or by using the Break key-combination (Ctrl-C, by default) during the first 60 seconds of startup.

Software Upgrade Ends with Error

Possible Cause	Solution
The standby supervisor module bootflash file system does not have sufficient space to accept the updated image.	Use the delete command to remove unnecessary files from the file system (delete the license from the bootflash).
The specified system and kickstart images are incompatible.	Check the output of the installation process for details on the incompatibility. Possibly update the kickstart image before updating the system image.
The install all command is entered on the standby supervisor module.	Enter the command on the active supervisor module only.
A module was inserted while the upgrade was in progress.	Restart the installation.
The system experienced a power disruption while the upgrade was in progress.	Restart the installation.
An incorrect software image path was specified.	Specify the entire, accurate path for the remote location.
Another upgrade is already in progress.	Verify the state of the system at every stage and restart the upgrade after 10 seconds. If you restart the upgrade within 10 seconds, the command is rejected. An error message displays, indicating that an upgrade is currently in progress.
Module failed to upgrade.	Restart the upgrade or use the install module command to upgrade the failed module.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-15

If the upgrade ends with an error, there are several possible causes. The figure shows these causes and their solutions.

Power-On or Switch Reboot Stops Responding

- A power-on or switch reboot stops responding for a dual-supervisor configuration.

Possible Cause	Solution
The bootflash is corrupted.	Use the Recovery for Systems with Dual Supervisor Modules procedure.
The BIOS is corrupted.	Replace this module. Contact your customer support representative to return the failed module.
The kickstart image is corrupted.	Power cycle the switch if required, and press Ctrl-C when the switch says Loading Boot Loader, to interrupt the boot process at the loader> prompt. Use the Recovery from the loader> Prompt procedure to update the kickstart image.
Boot parameters are incorrect.	Verify and correct the boot parameters and reboot.
The system image is corrupted.	Power cycle the switch if required and press Ctrl-J when the switch says "Checking all filesystems....r. done", to interrupt the boot process at the switch#boot prompt. Use the Recovery from the switch(boot)# Prompt procedure to update the system image.


© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-16

If a power-on or switch reboot stops responding for a dual-supervisor configuration, there are several possible causes. The figure shows these causes and their solutions.

Note The Recovery for Systems with Dual Supervisor Modules procedure is available at http://docwiki-rcdn-prd.cisco.com/wiki/Cisco_Nexus_7000_Series_NX-OS_Troubleshooting_Guide_-_Troubleshooting_Installs,_Upgrades,_and_Reboots#Recovery_for_Systems_with_Dual_Supervisor_Modules.

Corrupted Bootflash Recovery



All device configurations reside in the internal bootflash. If you have a corrupted internal bootflash, you can lose your configuration. Be sure to save and back up your configuration files periodically. The regular system boot goes through this sequence:

1. The BIOS loads the loader.
2. The loader loads the kickstart image into RAM and starts the kickstart image.
3. The kickstart image loads and starts the system image.
4. The system image reads the startup configuration file.

If the images on your system are corrupted and you cannot proceed (error state), you can interrupt the system boot sequence and recover the image by entering the BIOS configuration utility, as described in the following section. Access this utility only when you need to recover a corrupted internal disk.

Recovery procedures require the regular sequence to be interrupted. The internal sequence goes through four phases between the time that you turn on the system and the time that the system prompt appears on your terminal: BIOS, bootloader, kickstart, and system.

The BIOS begins the power-on self-test, memory test, and other operating system applications. While the test is in progress, press **Ctrl-C** to enter the BIOS configuration utility and use the netboot option.

The bootloader uncompresses the loaded software to boot an image, using its filename as a reference. These images are made available through bootflash. When the memory test is over, press **Esc** to enter the boot loader prompt.

When the boot loader phase is over, press **Ctrl-]** to enter the switch(boot)# prompt. Depending on your Telnet client, these keys might be reserved, and you might need to remap the keystroke. See the documentation that your Telnet client provides. If the corruption causes the console to stop at this prompt, copy the system image and reboot the system.

The system image loads the configuration file of the last-saved running configuration and returns a switch login prompt.

Recovery from the loader> Prompt

- Enter the local IP address and subnet mask for the system at the loader> prompt, and press Enter.

```
loader> set ip 172.16.1.2 255.255.255.0
```
- Specify the IP address of the default gateway.

```
loader> set gw 172.16.1.1
```
- Boot the kickstart image file from the required server.

```
loader> boot tftp://172.16.10.100/tftpboot/n7000-s1-kickstart-4.0.bin
```
- The switch(boot)# prompt indicates that you have a usable kickstart image.
- Enter the **init system** command at the switch(boot)# prompt.

```
switch(boot)# init system
```
- Be sure that you have made a backup of the configuration files before you enter this command.
- Follow the procedure specified in the Recovery from the switch(boot)# Prompt procedure.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-18

This procedure uses the **init system** command, which reformats the file system of the device. Be sure that you have made a backup of the configuration files before you begin this procedure.

The loader> prompt is different from the regular switch# prompt. The CLI command completion feature does not work at the loader> prompt and might result in undesired errors. You must type the command exactly as you want it to appear. If you boot over TFTP from the loader> prompt, you must supply the full path to the image on the remote server.

Use the **help** command at the loader> prompt to display a list of commands that are available at this prompt or to obtain more information about a specific command in that list.

To recover a corrupted kickstart image (system error state) follow the steps that are listed in the figure.

Recovery from the switch boot)# Prompt

```
switch boot)# config t
switch boot)(config)# interface mgmt0
switch boot)(config-mgmt0)# ip address 172.16.1.2 255.255.255.0
switch boot)(config-mgmt0)# ip default-gateway 172.16.1.1
switch boot)(config-mgmt0)# no shutdown
switch boot)(config-mgmt0)# end

switch boot)# init system check-filesystem
switch boot)# copy tftp://172.16.10.100/system-image1 bootflash:system-image1
switch boot)# copy tftp://172.16.10.100/kickstart-image1 bootflash:kickstart-image1

switch boot)# '' dir bootflash: ''
12456448      Jul 30 23:05:28 1980  kickstart-image1
12288         Jun 23 14:58:44 1980  lost+found/
27602159      Jul 30 23:05:16 1980  system-image1

Usage for bootflash://sup-local
135404544 bytes used
49155072 bytes free
184559616 bytes total

switch boot)# ''load bootflash:system-image1''
Uncompressing system image: bootflash:/system-image1
cccccccccccccccccccccccccccccccccccccccccccccccccccccccccccc
Would you like to enter the initial configuration mode? (yes/no):
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-6-19

To recover a system image by using the kickstart image for a system with a single supervisor module, follow these steps:

- Step 9** Change to configuration mode and configure the IP address of the mgmt0 interface.
- Step 10** Follow this step if you entered an **init system** command. Otherwise, skip to Step 3.
 - Enter the **ip address** command to configure the local IP address and the subnet mask for the system.
 - Enter the **ip default-gateway** command to configure the IP address of the default gateway.
- Step 11** Enter the **no shutdown** command to enable the mgmt0 interface on the system.
- Step 12** Enter **end** to exit to EXEC mode.
- Step 13** If you believe that there are file system problems, enter the **init system check-filesystem** command. This command checks all internal file systems and fixes any errors that are encountered. This command takes a few minutes to complete.
- Step 14** Copy the system image from the required TFTP server.
- Step 15** Copy the kickstart image from the required TFTP server.
- Step 16** Verify that the system and kickstart image files are copied to your bootflash file system.
- Step 17** Load the system image from the bootflash file system.

System or Process Resets

- A recoverable error occurred on the system or on a process in the system:
 - The system has automatically recovered from the problem.
 - Check the syslog file to determine which process restarted and why.
 - Enter the **system cores tftp:[//servername][//path]** command to configure the system to use TFTP to send the core dump to a TFTP server.
- A nonrecoverable error occurred on the system:
 - The system cannot recover automatically from the problem.
 - Use the **show system reset-reason module [number]** command to display the last four reset-reason codes for a specific module in a given slot.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-20

When a recoverable or nonrecoverable error occurs, the system or a process on the system might reset. Every process restart generates a syslog message and a Call Home event. Even if the event does not affect service, you should identify and resolve the condition immediately because future occurrences could cause a service interruption.

An unrecoverable system restart might occur in these cases:

- A critical process fails and is not restartable.
- A process restarts more times than the system configuration allows.
- A process restarts more frequently than the system configuration allows.

The effect of a process reset is determined by the policy that is configured for each process. An unrecoverable reset might cause functionality loss, an active supervisor restart, a supervisor switchover, or a system restart.

The **show system reset-reason** command displays this information:

- The last four reset-reason codes for the supervisor modules, unless either supervisor module is absent (or **show system reset-reason module number** command for the last four reset-reason codes for a specific module in a given slot, unless absent)
- The overall history of when and why expected and unexpected reloads occur
- The time stamp of when the reset or reload occurred
- The reason for the reset or reload of a module
- The service that caused the reset or reload (if available)
- The software version that was running at the time of the reset or reload

Standby Supervisor Fails to Boot

- Error message:
`SYSMGR-2-STANDBY_BOOT_FAILED: Standby supervisor failed to boot up.`
- Possible cause:
 - Boot variables are not properly set for the standby supervisor.
 - A user intentionally interrupted the boot procedure at the loader> prompt.
- Solutions:
 - If the supervisor is at the loader> prompt, try to use the **boot** command to continue the boot procedure.
 - Issue a **reload** command for the standby supervisor from a vsh session on the active supervisor, specifying the **force-dnld** option.
 - When the standby is online, fix the problem by setting the boot variables appropriately.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-21

The standby supervisor does not boot after an upgrade. You might see this system error message:

`SYSMGR-2-STANDBY_BOOT_FAILED: Standby supervisor failed to boot up.`

This message is printed if the standby supervisor doesn't complete its boot procedure (doesn't reach the login prompt on the local console) in 3 to 6 minutes after the BIOS loads the loader. This message is usually caused by boot variables that are not properly set for the standby supervisor. This message can also be caused by a user intentionally interrupting the boot procedure at the loader prompt (by pressing Esc).

Connect to the local console of the standby supervisor. If the supervisor is at the loader prompt, try to use the **boot** command to continue the boot procedure. Otherwise, issue a **reload** command for the standby supervisor from a virtual shell (VSH) session on the active supervisor, specifying the **force-dnld** option. When the standby is online, fix the problem by setting the boot variables appropriately.

Troubleshooting Cisco Fabric Services

This topic explains how to troubleshoot issues that relate to Cisco Fabric Services on a Cisco Nexus 7000 Series Switch.

Initial Troubleshooting Checklist

- Verify that Cisco Fabric Services is enabled for the same applications on all affected devices.
- Verify that Cisco Fabric Services distribution is enabled for the same applications on all affected devices.
- If you are using Cisco Fabric Services regions, verify that the application is in the same region on all the affected devices.
- Verify that there are no pending changes for an application and that a Cisco Fabric Services commit was issued for any configuration changes in a Cisco Fabric Services-enabled application.
- Verify that there are no unexpected Cisco Fabric Services locked sessions. Clear any unexpected locked sessions.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5.23

Many features in Cisco NX-OS require configuration synchronization across multiple devices in the network. Cisco Fabric Services provides a common infrastructure for automatic configuration synchronization for an application in the network. Cisco Fabric Services provides the transport function as well as a rich set of common services to the applications. Cisco Fabric Services can discover Cisco Fabric Services-capable devices in the network as well as their application capabilities. These applications that can be synchronized by using Cisco Fabric Services:

- Call Home
- Device alias
- Dynamic Port Virtual Storage Area Network (VSAN) Membership (DPVM)
- Fibre Channel domain
- Fibre Channel port security
- Fibre Channel timer
- Inter-VSAN Routing (IVR)
- Network Time Protocol (NTP)
- RADIUS
- Registered State Change Notification (RSCN)
- TACACS+
- User roles

Do not enable Cisco Fabric Services for an application that you manage by using Cisco Data Center Network Manager (DCNM). You can use Cisco Fabric Services regions to limit the Cisco Fabric Services configuration distribution to a subset of devices on the network.

Begin troubleshooting Cisco Fabric Services issues by checking these issues first:

- Verify that Cisco Fabric Services is enabled for the same applications on all affected devices.
 - **show cfs status**
- Verify that Cisco Fabric Services distribution is enabled for the same applications on all affected devices.
 - **show cfs application**
- If you are using Cisco Fabric Services regions, verify that the application is in the same region on all the affected devices.
 - **show cfs regions**
- Verify that there are no pending changes for an application and that a Cisco Fabric Services commit was issued for any configuration changes in a Cisco Fabric Services-enabled application.
- Verify that there are no unexpected Cisco Fabric Services locked sessions. Clear any unexpected locked sessions.
 - **show cfs lock**

Verifying Cisco Fabric Services by Using the CLI

- Verify that Cisco Fabric Services is globally enabled on all devices in the network or Cisco Fabric Services region.

```
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
```

- Verify that Cisco Fabric Services is enabled for the application on all devices in the network or Cisco Fabric Services region.

```
switch(config)# show cfs application
-----
Application      Enabled      Scope &
-----
ntp              No          Physical-fc-ip
stp              Yes         Physical-eth
vpc              Yes         Physical-eth
igmp             Yes         Physical-eth
l2fm             Yes         Physical-eth
role             Yes         Physical-fc-ip
radius           Yes         Physical-fc-ip
tacacs           No          Physical-fc-ip
callhome         Yes         Physical-fc-ip
Total number of entries = 9
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-24

To verify Cisco Fabric Services by using the CLI, follow these steps:

- Step 1** Verify that Cisco Fabric Services is globally enabled on all devices in the network or Cisco Fabric Services region.

```
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
```

- Step 2** Verify that Cisco Fabric Services is enabled for the application on all devices in the network or Cisco Fabric Services region.

```
switch(config)# show cfs application
```

```
-----
Application      Enabled      Scope &
-----
ntp              No          Physical-fc-ip
stp              Yes         Physical-eth
vpc              Yes         Physical-eth
igmp             Yes         Physical-eth
l2fm             Yes         Physical-eth
role             Yes         Physical-fc-ip
radius           Yes         Physical-fc-ip
tacacs           No          Physical-fc-ip
callhome         Yes         Physical-fc-ip
Total number of entries = 9
```

A Physical-fc-ip scope means that Cisco Fabric Services uses IP to apply the configuration for that application to all devices in the network or region. A Physical-eth scope means that Cisco Fabric Services uses Ethernet to apply the configuration for that application to all devices in the network or region.

Step 3 Verify that Cisco Fabric Services distribution is enabled for the application on all devices in the network or Cisco Fabric Services region.

```
switch(config)# show cfs application name radius
```

```
Enabled      : Yes
Timeout      : 20s
Merge Capable : Yes
Scope        : Physical-fc-ip
Region       : 99
```

Step 4 If you configure Cisco Fabric Services regions, verify that the application is in the same region on all applicable devices.

```
switch(config)# show cfs regions brief
```

```
-----
Region      Application  Enabled
-----
4           callhome    yes
99          radius      yes
```

Step 5 Verify the set of devices that are registered with Cisco Fabric Services for that application.

```
switch# show cfs peers name radius
```

```
Scope      : Physical-fc-ip
```

```
-----
Switch WWN                IP Address
-----
20:00:00:0e:d7:0e:bf:c0   192.0.2.51    [Local]
20:00:00:0e:d7:00:3c:9e   192.0.2.52
Total number of entries = 2
```

Step 6 Compare the output of the **show cfs merge status name application-name** command and the **show cfs peers name application-name** command to verify that the network is not partitioned.

```
switch# show cfs merge status name radius
```

```
Physical-fc-ip Merge Status: Success [ Sat May 5 11:59:36 2012 ]
Local Fabric
```

```
-----
Switch WWN                IP Address
-----
20:00:00:05:30:00:4a:de   192.0.2.51    [Merge Master]
20:00:00:0d:ec:0c:f1:40   192.0.2.204
Total number of switches = 2
```

```
switch# show cfs peers name radius
```

```
Scope      : Physical-fc-ip
```

```

-----
Switch WWN                IP Address
-----
20:00:00:0d:ec:0c:f1:40  192.0.2.51 [Local]
20:00:00:05:30:00:4a:de  192.0.2.204
Total number of entries = 2

```

If the list of switch world wide names (sWWNs) in the **show cfs merge status name** command output is shorter than the list of sWWNs in the **show cfs peers name** command output, then the network is partitioned into multiple Cisco Fabric Services fabrics and the merge status might show that the merge has failed, is pending, or is waiting.

Step 7 Verify that a distribution is not in progress in the network for the application.

```
switch# show cfs lock
```

```
Application: callhome
```

```
Scope      : Physical-fc-ip
```

```

-----
Switch WWN                IP Address                User Name    User Type
-----
20:00:00:22:55:79:a4:c1  172.28.230.85      admin        CLI/SNMP v3
                        switch
Total number of entries = 1

```

If the application does not show in the output, then the distribution has completed.

Step 8 Verify that there are no Cisco Fabric Services sessions in progress for the application.

```
switch(config)# show radius session status
```

```
Last Action Time Stamp    : Sun Jun 24 13:25:00 2012
```

```
Last Action                : Commit
```

```
Last Action Result         : Success
```

```
Last Action Failure Reason : none
```

Step 9 You can also verify the internal event-history:

```
N7K2-pod2# show cfs internal event-history errors
```

```
Tue Jul 31 05:00:29 2012 :
```

```
Recieving over TCP failed, retval -1, errno 9 [Bad file descriptor]
```

```
Tue Jul 31 05:03:45 2012 :
```

```
Recieving over TCP failed, retval -1, errno 104 [Connection reset by peer]
```

```
Tue Jul 31 05:03:45 2012 :
```

```
Recieving over TCP failed, retval -1, errno 104 [Connection reset by peer]
```

```
Tue Jul 31 05:03:45 2012 :
```

```
Recieving over TCP failed, retval -1, errno 104 [Connection reset by peer]
```

Troubleshooting Merge Failures

- To recover from a merge failure by using the CLI, follow these steps:
 - Identify a device that shows a merge failure.

```
switch# show cfs merge status
```

Application	Scope	Vsan	Status
role	Physical-fc-ip	-	Success
radius	Physical-fc-ip	-	Success
callhome	Physical-fc-ip	-	Failed

- Commit the application configuration to restore all peers in the fabric to the same configuration database.

```
switch(config)# callhome commit
```

During a merge, the merge managers in the merging networks exchange their configuration databases with each other. The application on the merge master device merges the information, decides whether the merge is successful, and informs all devices in the combined network of the merge status. When a merge is successful, the merge master distributes the database to all devices in the combined network, and the combined network remains in a consistent state. A merge failure indicates that the merged network contains inconsistent data that could not be merged.

If you add a new device to the network and the merge status for any application shows "In Progress" for a prolonged period, then there might be an active session for that application in some other device. Use the **show cfs lock** command to determine the lock status for that application on all devices. The merge will not proceed if any locks are present for that application on any device in the network or Cisco Fabric Services region. Use the **application-name commit** command to commit the changes or use the **clear application-name session** command to clear the session lock so that the merge can proceed.

To recover from a merge failure by using the CLI, follow these steps:

Step 1 Identify a device that shows a merge failure.

```
switch# show cfs merge status
```

Application	Scope	Vsan	Status
role	Physical-fc-ip	-	Success
radius	Physical-fc-ip	-	Success
callhome	Physical-fc-ip	-	Failed

Step 2 Commit the application configuration to restore all peers in the fabric to the same configuration database.

```
switch(config)# callhome commit
```

Troubleshooting Lock Failures

- Determine all the devices that participate in the Cisco Fabric Services distribution for this application.

```
switch1# show cfs peers name radius
Scope      : Physical-fc-ip
-----
Switch WWN          IP Address
-----
20:00:00:0d:ec:0c:f1:40 192.0.2.51 [Local]
20:00:00:05:30:00:4a:de 192.0.2.204
Total number of entries = 2
```

- Check for a lock for this application on all Cisco Fabric Services peer devices.

```
switch2# show cfs lock
Application: radius
Scope      : Physical-fc-ip
-----
Switch WWN          IP Address          User Name    User Type
-----
20:00:00:05:30:00:4a:de 192.0.2.204      admin        CLI/SNMP v3
switch
Total number of entries = 1
```

- Release the Cisco Fabric Services lock on the device that owns the lock.

```
switch2# radius abort
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-28

To distribute a configuration in the network, Cisco Fabric Services must first acquire a lock on all devices in the network or Cisco Fabric Services region. After Cisco Fabric Services acquires the locks, it issues a commit to distribute the configuration to all devices in the network or Cisco Fabric Services region. Under normal circumstances, Cisco Fabric Services releases the lock after the commit.

When another application peer acquires a lock, you cannot commit new configuration changes. This operation is normal and you should postpone any changes to an application until the application peer releases the lock.

An inconsistent lock state also occurs in these scenarios:

- When locks are not held on all the devices in the network or Cisco Fabric Services region
- When locks are held on all devices in the network or region, but a Cisco Fabric Services session does not exist on the device that holds the lock

Use the troubleshooting steps in this section only when you believe that the lock has not been properly released. To troubleshoot a lock failure, follow these steps:

- Step 1** Determine all the devices that participate in the Cisco Fabric Services distribution for this application.

```
switch1# show cfs peers name radius
Scope      : Physical-fc-ip
-----
Switch WWN          IP Address
-----
20:00:00:0d:ec:0c:f1:40 192.0.2.51 [Local]
20:00:00:05:30:00:4a:de 192.0.2.204
Total number of entries = 2
```

Step 2 Check for a lock for this application on all Cisco Fabric Services peer devices, to determine the name of the administrator who owns the lock for the application.

```
switch2# show cfs lock
```

```
Application: radius
```

```
Scope      : Physical-fc-ip
```

```
-----  
Switch WWN          IP Address          User Name      User Type  
-----  
20:00:00:05:30:00:4a:de 192.0.2.204      admin          CLI/SNMP v3  
switch
```

```
Total number of entries = 1
```

Step 3 Connect to the device that owns the Cisco Fabric Services lock.

Step 4 Release the Cisco Fabric Services lock on the device that owns the lock.

```
switch2# radius abort
```

Step 5 If the device does not release the lock, then clear the Cisco Fabric Services session on the device that owns the lock.

```
switch2# clear radius session
```

Troubleshooting Cisco Fabric Services Regions

- Verify the list of devices in a region for the application.

```
switch(config)# show cfs region name radius
Region-ID : 4
Application: radius
Scope : Physical-fc-ip
-----
Switch WWN          IP Address
-----
20:00:00:22:55:79:a4:c1 172.28.230.85          [Local]
                        switch
Total number of entries = 1
```

- Verify that the application distribution is enabled and is in the same region on all devices in the region.

```
switch2# show cfs application name radius
Enabled : Yes
Timeout : 20s
Merge Capable : Yes
Scope : Physical-fc-ip
Region : 1
```

Application can be merged.

Application is in Region 1.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-27

These rules apply to Cisco Fabric Services regions:

- When using Cisco Fabric Services regions, an application on a given device can belong to only one region at a time.
- An application in a Cisco Fabric Services region ignores all Cisco Fabric Services distributions in any other region (including the default region).
- All applications that you do not assign to a Cisco Fabric Services region exist in the default region.

To resolve a configuration distribution failure to all devices in a Cisco Fabric Services region, follow these steps:

Step 1 Verify the list of devices in a region for the application.

```
switch(config)# show cfs region name radius
```

```
Region-ID : 4
```

```
Application: radius
```

```
Scope : Physical-fc-ip
```

```
-----
Switch WWN          IP Address
-----
```

```
20:00:00:22:55:79:a4:c1 172.28.230.85          [Local]
                        switch
```

```
Total number of entries = 1
```

Step 2 Verify that the application distribution is enabled and is in the same region on all devices in the region.

```
switch2# show cfs application name radius
```

```
Enabled : Yes
```

```
Timeout          : 20s
Merge Capable    : Yes
Scope            : Physical-fc-ip
Region           : 1
switch2(config)# cfs region 4
switch2(config-cfs-region)# radius
```

You must reassign an application to a region whenever you disable that application. Cisco Fabric Services assigns new applications in the default region.

If you move an application from one region to another, you might encounter a database mismatch when attempting a merge. Follow the steps for troubleshooting merge failures to identify and resolve the conflicts.

When an application is moved from one region to another (including the default region), the application loses all Cisco Fabric Services history.

Troubleshooting VDCs

This topic explains how to troubleshoot issues that relate to virtual device contexts (VDCs) on a Cisco Nexus 7000 Series Switch.

Initial Troubleshooting Checklist

- Verify that you are logged in to the device as network-admin if you are creating or modifying VDCs.
- Verify that you are in the correct VDC. You must be in the default VDC to configure VDCs.
- Verify that you have installed the Advanced Services license to configure VDCs.
- The total number of all possible VDCs is 3+1 for SUP-1, 4+1 for SUP-2 and 8+1 for SUP-2E (which needs an additional license). The +1 VDC is a dedicated admin VDC and does not support Layer 2 or Layer 3 forwarding or routing.
- CPU shares are available for SUP-2 and SUP-2E.

© 2012 Cisco and/or its affiliates. All rights reserved.DCUFT v5.0--5-29

Cisco NX-OS supports VDCs, which you can use to divide the physical NX-OS device into separate virtual devices. Each VDC appears as a unique device to connected users. A VDC runs as a separate logical entity within the physical Cisco NX-OS device, maintains its own unique set of running software processes, has its own configuration, and can be managed by a separate administrator.

VDC issues might not be directly related to VDC management. For instance, if you configure a VDC template that limits the number of port channels in that VDC, you might experience problems if you try to create more port channels than the VDC template allows. VDC templates set limits on these features:

- Port channels
- Switched Port Analyzer (SPAN) sessions
- IPv4 route map memory
- VLANs
- Virtual routing and forwarding instances (VRFs)

The minimum resource value configures the guaranteed limit for that feature. The maximum resource value represents oversubscription for the feature and is available on a first-come, first-served basis.

When you allocate an interface to a VDC, Cisco NX-OS removes all configuration for that interface.

Begin troubleshooting VDC issues by checking these issues first:

- Verify that you are logged in to the device as network-admin if you are creating or modifying VDCs.
- Verify that you are in the correct VDC. You must be in the default VDC to configure VDCs.
- Verify that you have installed the Advanced Services license to configure VDCs.
- The total number of all possible VDCs is 3+1 for SUP-1, 4+1 for SUP-2 and 8+1 for SUP-2E (needs additional license). The +1 VDC is a dedicated admin VDC and does not support Layer 2 or Layer 3 forwarding or routing.
- CPU shares are available for SUP-2 and SUP-2E.

Use these commands to display VDC information:

- **show vdc membership:** Displays information about which interfaces are assigned to a VDC
- **show vdc resource:** Displays information about the assigned resources (available only in the default VDC)
- **show vdc current-vdc:** Displays the VDC that you are currently in

Issues with Creating a VDC

Symptom	Possible Cause	Solution
You cannot create a VDC.	You are not logged in as network-admin.	Log in to the device with an account that has network-admin privileges.
	You are not logged in to the default VDC.	Use the switchback command to switch to the default VDC to allocate resources.
	There are not enough resources.	Use the show vdc resources [detail] or show vdc resource template command to determine your available resources. Modify your template or create a VDC with fewer resources by using the limit-resource command in VDC configuration mode.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-30

When you have a problem creating a VDC, you might see one of these system messages:

- **Error message:** VDC_MGR-2-VDC_BAD: vdc_mgr: There has been a failure at res_mgr
 - **Explanation:** You cannot create a VDC because not enough resources are available, based on the template configuration. If no template is used, a default template is applied.
 - **Recommended action:** Verify that you have sufficient resources available to create this VDC by using the **show vdc resources [detail]** or **show vdc resource template** command. Modify the template that you are using to create the VDC, or create a new template with resource limits that are currently available.
- **Error message:** VDC_MGR-2-VDC_BAD: vdc_mgr: : There has been a failure at sys_mgr
 - **Explanation:** Some services have failed or have not come up because of insufficient system resources, other than what can be reserved by using the resource templates. These dynamic resources are based on system utilization and might not be available to support a new VDC.
 - **Recommended action:** Use the **show system internal sysmgr service running** command to determine what caused the failure.

Issues with Logging in to a Device

Symptom	Possible Cause	Solution
You cannot log in to a device.	There is no account information for the VDC.	Log in to the device as network-admin and use the switchto command to switch to the VDC and configure the password and network connectivity for this VDC.
	You are using an incorrect VDC username.	Log in to the device with the account created for that VDC.
You cannot switch to a VDC.	You are not logged in as network-admin or network-operator.	Log in to the device with an account that has the correct privileges.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-6-31

You may have a problem when logging into a device. If you cannot log into a device, consider one of these possible causes and solutions:

- There is no account information for the VDC. Log into the device as network-admin and use the **switchto** command to switch to the VDC and configure the password and network connectivity for this VDC.
- You are using an incorrect VDC username. Log into the device with the account that was created for that VDC.
- You might have a problem when you switch to another VDC. If you cannot switch to a VDC, you might not be logged in as network-admin or network-operator. Log in to the device with an account that has the correct privileges.

Issues with Deleting a VDC

- Error message: `VDC_MGR-2-VDC_UNGRACEFUL: vdc_mgr: Ungraceful cleanup request received for vdc [dec], restart count for this vdc is [dec]`
- Explanation: Vdc_mgr has begun an ungraceful cleanup for a VDC.
- Error message: `VDC_MGR-2-VDC_OFFLINE: vdc [dec] is now offline`
- Explanation: Vdc_mgr has finished deleting a VDC.

Symptom	Possible Cause	Solution
You cannot delete a VDC.	You attempted to delete the default VDC.	You cannot delete the default VDC.
	Unknown errors occurred when deleting a VDC.	Use the show tech-support VDC command to gather more information.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-9-32

When you have a problem deleting a VDC, you might see one of these system messages:

- **Error message:** `VDC_MGR-2-VDC_UNGRACEFUL: vdc_mgr: Ungraceful cleanup request received for vdc [dec], restart count for this vdc is [dec]`
 - **Explanation:** Vdc_mgr has begun an ungraceful cleanup for a VDC.
 - **Recommended action:** No action is required.
- **Error message:** `VDC_MGR-2-VDC_OFFLINE: vdc [dec] is now offline`
 - **Explanation:** Vdc_mgr has finished deleting a VDC.
 - **Recommended action:** No action is required.

If you cannot delete a VDC, there are several possible causes and solutions:

- You attempted to delete the default VDC. You cannot delete the default VDC.
- Unknown errors occurred when deleting a VDC. Use **the show tech-support VDC** command to gather more information.

Issues with Allocating an Interface to a VDC

Symptom	Possible Cause	Solution
You cannot allocate an interface to a VDC.	You are not logged in as network-admin.	Log in to the device with an account that has the correct privileges.
	You are not logged into the correct VDC.	Use the switchback command to switch to the default VDC to allocate resources.
	The interface is part of a dedicated port group.	Use the show interface capabilities command to determine whether the port is dedicated. All ports in a dedicated port group must be in the same VDC.
	The interface is on the Cisco Nexus 7000 M-1 Series 32-Port 10 Gigabit Ethernet Module (N7K-M132XP-12).	You must allocate all ports in a port group to the same VDC for this module.
	The VDC allocation has failed.	Use the show vdc membership [status] or show interface brief command to gather more information.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5-33

When you have a problem creating a VDC, you might see this system message:

```
VDC_MGR-2-VDC_BAD: vdc_mgr: There has been a failure at gim  
(port_affected_list)
```

An interface allocation has failed. Use the **show vdc membership status** or **show interface brief** command to gather more information.

The VDC Remains in a Failed State

Symptom	Possible Cause	Solution
The VDC remains in failed state.	The VDC failed and the high-availability policy was set to bring down the VDCs.	Use the show vdc detail command to verify the high-availability policy for this VDC. Use the ha-policy command in VDC configuration mode to change the high-availability policy.
	A supervisor switchover has occurred and the switchover policy was set to bring down the VDCs.	Use the no vdc command to delete the failed VDC. Recreate the VDC with a different switchover policy using the sw-policy keyword.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-34

You might have a problem when a VDC fails. You configure switchover and high availability policies for a VDC when you create the VDC. These policies determine what happens when the VDC fails or when a Stateful Switchover occurs to the standby supervisor.

Other Issues with VDC

Symptom	Possible Cause	Solution
The VDC does not reflect a resource template change.	The template has not been reapplied to a VDC after a template change.	Use the show vdc resource template command to verify the template. Use the template command in VDC configuration mode to reapply the template to the VDC.
You cannot copy the running configuration file to the startup configuration file in a VDC.	The resource allocation was not saved in the default VDC.	Log in to the default VDC and use the copy running-config startup-config command to save the resource allocation. Log in to the nondefault VDC and save the configuration or use the copy running-config startup-config vdc-all command in the default VDC to save the configuration in all VDCs.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-35

You might have a problem when updating a resource template or when trying to save the configuration in a VDC. The figure shows the possible causes and solutions.

Troubleshooting Routing

This topic explains how to troubleshoot issues that relate to routing on a Cisco Nexus 7000 Series Switch.

Initial Troubleshooting Checklist

- Verify that the routing protocol is enabled.
- Verify that the address family is configured if necessary.
- Verify that you have configured the correct VRF for your routing protocol.
- Use these commands to display routing information:
 - **show ip arp**
 - **show ip traffic**
 - **show ip fib**
 - **show ip process**
 - **show routing**
 - **show vrf**
 - **show vrf interface**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.37

Layer 3 routing involves determining optimal routing paths and packet switching. You can use routing algorithms to calculate the optimal path from the router to a destination. This calculation depends on the algorithm that is selected, route metrics, and other considerations, such as load balancing and alternate path discovery.

Cisco NX-OS uses the VDCs to provide separate management domains per VDC and software fault isolation. Each VDC supports multiple VRFs and multiple routing information bases (RIBs) to support multiple address domains. Each VRF is associated with an RIB and this information is collected by the forwarding information base (FIB).

Troubleshooting Routing

- Verify that the routing protocol is enabled.

```
switch(config)# show ospf
                    ^
% invalid command detected at '^' marker.
```

- Verify the configuration for this routing protocol.

```
switch# show running-configuration eigrp

!Command: show running-config eigrp
!Time: Thu Jul 19 10:08:01 2012

version 6.0(4)
feature eigrp

router eigrp 10
 vrf OTV

interface Ethernet1/5
 ip router eigrp 10

interface Ethernet1/9
 ip router eigrp 10
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-38

To troubleshoot basic routing issues, follow these steps:

- Step 1** Use the **show feature | include protocol** command to verify that the routing protocol is enabled.

If the feature is not enabled, Cisco NX-OS reports that the command is invalid. Use the **feature** command to enable the routing protocol.

- Step 2** Verify the VRF configuration for this routing protocol. In addition to the command that is shown in the figure, you can use the **show running-configuration eigrp all** command:

```
switch# sh run eigrp all
```

```
!Command: show running-config eigrp all
!Time: Thu Jul 19 10:08:16 2012
```

```
version 6.0(4)
feature eigrp
```

```
router eigrp 10
  log-neighbor-warnings
  log-adjacency-changes
  graceful-restart
  timers active-time 3
  timers nsf signal 20
  timers nsf converge 120
  timers nsf route-hold 240
  distance 90 170
  metric weights 0 1 0 1 0 0
  metric rib-scale 128
```

```

metric maximum-hops 100
default-metric 1000000 100 255 1 1492
maximum-paths 8
vrf OTV
    log-neighbor-warnings
    log-adjacency-changes
    graceful-restart
    timers active-time 3
    timers nsf signal 20
    timers nsf converge 120
    timers nsf route-hold 240
    distance 90 170
    metric weights 0 1 0 1 0 0
    metric rib-scale 128
    metric maximum-hops 100
    default-metric 1000000 100 255 1 1492
    maximum-paths 8
ipv6 eigrp event-history l3vpn size small
ipv6 eigrp event-history cli size small
ipv6 eigrp event-history rib size small
ipv6 eigrp event-history packet size small
ipv6 eigrp event-history fsm size small

interface Ethernet1/5
    ip router eigrp 10
    ipv6 hold-time eigrp 10 15
    ip hold-time eigrp 10 15
    ipv6 hello-interval eigrp 10 5
    ip hello-interval eigrp 10 5
    ipv6 bandwidth-percent eigrp 10 50
    ip bandwidth-percent eigrp 10 50

interface Ethernet1/9
    ip router eigrp 10
    ipv6 hold-time eigrp 10 15
    ip hold-time eigrp 10 15
    ipv6 hello-interval eigrp 10 5
    ip hello-interval eigrp 10 5
    ipv6 bandwidth-percent eigrp 10 50
    ip bandwidth-percent eigrp 10 50

```

Troubleshooting Routing (Cont.)

- Verify that the interface is in the correct VRF.

```
switch(config)# show vrf interface ethernet 1/2
Interface          VRF-Name          VRF-ID
Ethernet1/2        default           1
```

- Verify that the routing protocol is registered with the RIB.

```
switch(config)# show routing unicast clients
CLIENT: am
index mask: 0x00000002
epid: 3908      MTS SAP: 252      MRU cache hits/misses:      2/1
Routing Instances:
  VRF: management      table: base
Messages received:
  Register      : 1      Add-route      :2      Delete-route      :1
Messages sent:
  Add-route-ack  : 2      Delete-route-ack :1
...
CLIENT: eigrp-99
index mask: 0x00002000
epid: 3148      MTS SAP: 63775     MRU cache hits/misses:      0/1
Routing Instances:
  VRF: default          table: base          notifiers: self
Messages received:
  Register      : 1      Delete-all-routes :1
...
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-39

- Step 3** Verify that the interface is in the correct VRF.
- Step 4** Verify that the routing protocol is registered with the RIB.

Troubleshooting Unicast Traffic

This topic explains how to troubleshoot issues that relate to unicast forwarding on a Cisco Nexus 7000 Series Switch. The topic provides troubleshooting for a unicast packet, from input to output and everything in between.

Packet Is Received into Interface from Wire

```
PHX2-N7K-1# show interface e1/1
Ethernet1/1 is up
<...>
Rate mode is shared
Switchport monitor is off
Last link flapped 7week(s) 4day(s)
Last clearing of "show interface" counters never
1 minute input rate 13056 bits/sec, 9 packets/sec
1 minute output rate 4608 bits/sec, 0 packets/sec
Rx
 341190251 input packets 276211313 unicast pkts 52112947 multicast pkts
12865991 broadcast packets 0 jumbo packets 0 storm suppression packets
94295027129 bytes
Tx
 462437316 output packets 85121 multicast packets
188251 broadcast packets 0 jumbo packets
648159081064 bytes
0 input error 0 short frame 0 watchdog
<...>
```

```
PHX2-N7K-1# show interface e1/1 transceiver details
Ethernet1/1
 sfp is present
 name is CISCO-AVAGO
 part number is SFBR-7700SDZ
<...>
```

If this says type is (unknown), the transceiver is not supported.

During this step, the packet is received into the Cisco Nexus 7000 Series Switch port. When troubleshooting this step, you want to ensure transceiver interoperability and determine whether you see any errors on the interface. Do so by using these commands:

- **show interface interface**
- **show interface interface transceiver**

LinkSec Decryption and Port QoS

- Use these commands to troubleshoot LinkSec and port QoS:
 - **show cts interface** [*all* | *interface*]
 - **show queuing interface** *interface*
 - **show policy-map interface** (for per-queue drop)

```
switch# show cts interface all
CTS Information for Interface Ethernet2/24:
CTS is enabled, mode: CTS_MODE_DOT1X
IFC state: CTS_IFC_ST_CTS_OPEN_STATE
Authentication Status: CTS_AUTHC_SUCCESS
Peer Identity: india1
Peer is: CTS Capable
802.1X role: CTS_ROLE_AUTH
Last Re-Authentication:
Authorization Status: CTS_AUTHZ_SUCCESS
PEER SGT: 2
Peer SGT assignment: Trusted
Global policy fallback access list:
SAP Status: CTS_SAP_SUCCESS
Configured pairwise ciphers: GCM_ENCRYPT
Replay protection: Enabled
Replay protection mode: Strict
Selected cipher: GCM_ENCRYPT
Current receive SPI: sci:1b54c1fbff0000 an:0
Current transmit SPI: sci:1b54c1fc000000 an:0
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-42

In the next step, LinkSec decryption and receive-side stage-1 quality of service (QoS) occur.

Step back and evaluate the difference between stage-1 and stage-2 QoS. The difference is that some ports on the 10 Gigabit Ethernet modules can be configured in shared mode, whereas others can be configured in dedicated mode. Therefore, 10 Gb of bandwidth can be dedicated to a port or shared among ports.

When running in shared mode, there is a chance of contention when accessing the 10-Gb bandwidth through the 4:1 multiplexer (MUX). To alleviate this risk, some QoS intelligence was passed down to the 4:1 MUX, which aggregates the ports.

In dedicated mode, no QoS is applied at the MUX. Instead, all traffic is processed in stage-2 QoS. To summarize, in shared mode, stage-1 QoS ensures fair access to the 10 Gb of port bandwidth. In both shared and dedicated mode, stage-2 QoS occurs to provide ingress queuing to the system.

For the ingress QoS, be concerned about the receive side QoS parameters in the **show queuing** command. Use the **show policy-map** command to see per-queue dropped packets.

Use these commands to troubleshoot LinkSec and port QoS:

- **show cts interface** [*all* | *interface*]
- **show queuing interface** *interface*
- **show policy-map interface** (for per-queue drop)

Layer 2 Source and Destination MAC Processing

- To validate forwarding of the Layer 2 engine, use these commands:
 - show mac address-table**
 - show mac address-table | grep macaddress**

```
PHX2-N7K-1# show mac address-table
Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC
  age - seconds since last seen, + - primary entry using vPC Peer-Link
  VLAN    MAC Address    Type    age    Secure  NTFY    Ports
-----+-----+-----+-----+-----+-----+-----
G        -    0023.ac67.dd41    static    -    False  False  sup-eth1(R)
G 5      0023.ac67.dd41    static    -    False  False  sup-eth1(R)
* 5      0000.0c07.ac01    dynamic    0    False  False  Eth1/1
* 5      000c.2943.a67e    dynamic   180    False  False  Eth1/1
* 5      000c.294b.c5ca    dynamic    0    False  False  Eth1/1
* 5      000d.ece2.0640    dynamic   180    False  False  Eth1/1
* 5      0013.5f32.aa80    dynamic    0    False  False  Eth1/1
* 5      0018.8b45.41b7    dynamic    0    False  False  Eth1/1
<->
```

In this step, the ASIC submits the packet headers to the Layer 2 engine for lookup, and the Layer 2 engine performs source and destination MAC processing.

To validate forwarding of the Layer 2 engine, you should first look at the centralized mac table that is aggregated on the supervisor to validate whether the MAC addresses are correlated as you expect them, and assigned to the ports where you expect the MACs to reside.

You can then validate the hardware programming on the ingress line card to validate that your MAC address table is properly programmed into the hardware based Layer 2 engine on the line card.

The command that is used to accomplish this, is as follows:

■ **show mac address-table**

To drill down on a specific MAC address, you can use the `grep` function with these commands to validate the MAC is associated with a particular port, and that the hardware programming reflects that.

■ **show mac address-table | grep macaddress**

When evaluating the Hardware mac table, if the Index is set to 0x00400, or the GM bit is set to “1”, that traffic will be routed. For example, you will see the index set to 0x00400 and GM bit set to 1 for traffic destined to the MAC address local to the device

Troubleshooting Layer 3 Engine Feature Processing

- ACL
 - **show access-lists** *name*
- QoS
 - **show policy-map interface** *interface*
- NetFlow
 - **show flow interface**
 - **show flow record**
 - **show flow monitor**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-44

After the Layer 2 engine is finished, it sends the header to the Layer 3 engine. The Layer 3 engine applies Layer 3 intelligent features to all packets and Layer 3 forwarding to routed packets. The Layer 3 features that are applied to all packets include access control lists (ACLs), QoS, NetFlow, and hardware intrusion prevention system (IPS).

To troubleshoot an ACL, evaluate the configuration and any relevant hit counters. You can then determine whether the hardware on the line card is programming the ACL. If you want to see per-ACL counters, you must enable statistics per-entry in the ACL.

QoS can be applied on both ingress and egress, so you should interrogate both the ingress and egress QoS.

NetFlow processing also has portions that occur in hardware. For NetFlow, you collect statistics in hardware on the line cards. You can then export those statistics via software.

Layer 3 Forwarding for Routed Traffic

```
PHX2-N7K-1# show ip adjacency
IP Adjacency Table for VRF default
Total number of entries: 1
Address      Age      MAC Address  Pref Source  Interface
86.86.86.1   00:00:37  0011.aabb.ccdd 1    Static      Vlan86

PHX2-N7K-1# show ip fib route 86.86.87.0/24 mod 1

IPv4 routes for table default/base

-----+-----+-----
Prefix      | Next-hop      | Interface
-----+-----+-----
86.86.87.0/24  86.86.86.1    Vlan86

PHX2-N7K-1# show system internal forwarding route 86.86.86.1/24 detail mod 1
RPF Flags legend:
S - Directly attached route (S_Star)
V - RPF valid
M - SMAC IP check enabled
G - SGT valid
E - RPF External table valid
86.86.87.0/24 , Vlan86
Dev: 1 , Idx: 0x19001 , RPF Flags: V , DGT: 0 , VPN: 1
RPF_Intf_5: Vlan86 (0x55 )
AdjIdx: 0x43005, LIFB: 0 , LIF: Vlan86 (0x55 ), DI: 0x0
DMAC: 0011.aabb.ccdd SMAC: 0023.ac67.dd41
```

The Layer 3 engine performs Layer 3 forwarding only for traffic that is routed through the router. This traffic has been sent to the MAC address of a valid routed interface, local to the router.

To troubleshoot the routed traffic, perform these tasks:

- Step 1** Ensure that the control plane routing is correct.
- Step 2** Ensure that the hardware forwarding entries on the ingress module have the corresponding information.

All routing of traffic is performed on the forwarding engine of the ingress module.

As you can see in the example in the figure, 86.86.87.0/24 is set to route to 86.86.86.1, out VLAN 86. This next hop is associated with MAC address 0011.aabb.ccdd.

Use these commands to accomplish this:

- **show ip route** (prefix)
- **show ip arp** (nexthop)
- **show ip adjacency**

Now you can interrogate the hardware to ensure that the hardware entries have propagated properly to the Layer 3 hardware engine. You can see that the IP FIB has properly associated 86.86.87.0/24 to the next hop of 86.86.86.1. You can also see, in the hardware entry, that this is routed out VLAN 86, and that the route entry is correctly associated with the MAC address of 0011.aabb.ccdd. This demonstrates that the routing in the forwarding plane is programmed correctly and that the forwarding will follow the information in the routing protocols.

Use these commands to accomplish this:

- **show ip fib route** *prefix module module*
- **show system internal forwarding route** *prefix detail module module*

Fabric Processing

```
switch(config)# show hardware fabric-utilization
```

Slot	Direction	Utilization
2	ingress	3%
2	egress	3%
6	ingress	1%
6	egress	1%

```
PHX2-N7K-1# show module fabric
```

Xbar	Ports	Module-Type	Model	Status
1	0	Fabric Module 1	N7K-C7010-FAB-1	ok
2	0	Fabric Module 1	N7K-C7010-FAB-1	ok
3	0	Fabric Module 1	N7K-C7010-FAB-1	ok

Xbar	Sw	Hw
1	NA	1.0
2	NA	1.0
3	NA	1.0

Xbar	MAC-Address(es)	Serial-Num
1	NA	JAF1252AHRB
2	NA	JAF1251CABF
3	NA	JAF1252AHBL

```
* this terminal session
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-48

This step occurs if the packet needs to traverse the fabric.

In this step, you need to determine whether the fabrics are functioning properly and if their utilization is at an acceptable level. You can view the fabric status and utilization by using these commands:

- **show hardware fabric-utilization**
- **show module fabric**

LinkSec Encryption and Packet Transmission

- The **show cts interface {all | interface}** is used to troubleshoot LinkSec encryption.
- The final step in the process is the transmission of the frame out of the physical egress port.
- Troubleshooting of the physical port includes these commands:
 - **show interface interface**
 - **show interface interface transceiver**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-47

Use the **show cts interface {all | interface}** command to troubleshoot LinkSec encryption.

The final step in the process is the transmission of the frame out of the physical egress port. Troubleshooting of the physical port is the same as in the first step and includes these commands:

- **show interface interface**
- **show interface interface transceiver**

Troubleshooting Memory

This topic explains how to troubleshoot issues that relate to memory on a Cisco Nexus 7000 Series Switch.

General Assessment of Platform Memory Utilization

```
N7K# show system resources
Load average: 1 minute: 0.43 5 minutes: 0.30 15 minutes: 0.28
Processes : 884 total, 1 running
CPU states : 2.0% user, 1.5% kernel, 96.5% idle
Memory usage: 4135780K total, 3423272K used, 712508K free
0K buffers, 1739356K cache
```

total: The amount of physical RAM on the platform
free: The amount of unused or available memory
used: The amount of allocated (permanent) and cached (temporary) memory

```
N7K# show processes memory
PID MemAlloc MemLimit MemUsed StackBase/Ptr Process
-----
4662 52756480 562929945 150167552 bfffd00/bfffd970 netstack
```

© 2012 Cisco and/or its affiliates. All rights reserved. DCUFT v5.0--5-49

You can assess the overall level of memory utilization on the platform by using two basic CLI commands: **show system resources** and **show processes memory**.

From these command outputs, you might be able to tell that platform utilization is higher than normal or expected, but you will not be able to tell which type of memory usage is high.

The **show system resources** command displays platform memory statistics (not per-VDC). The output is derived from the Linux memory statistics in `/proc/meminfo`.

The cache and buffers are not relevant to customer monitoring. This information provides a general representation of the platform utilization only. You need more information to troubleshoot why memory utilization is high.

The **show process memory** command displays the memory allocation per process for the current VDC (the output also contains non-VDC global processes). Although this output is more detailed, it is useful only for verifying process-level memory allocation within a specific VDC.

Detailed Assessment of Platform Memory Utilization

```
N7K# show system internal kernel meminfo
```

```
MemTotal: 4135780 kB
MemFree: 578032 kB
Buffers: 5312 kB
Cached: 1926296 kB
RAMCached: 1803020 kB
Allowed: 1033945 Pages
Free: 144508 Pages
Available: 177993 Pages
SwapCached: 0 kB
<-->
Writeback: 0 kB
Mapped: 1903768 kB
Slab: 85392 kB
<-->
```

```
N7K# show system internal kernel memory global
```

```
Total memory in system : 4129600KB
Total Free memory : 1345232KB
Total memory in use : 2784368KB
Kernel/App memory : 1759856KB
RAM FS memory : 1018616KB
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-50

Use the **show system internal kernel** or **show system internal memory-alerts-log** command for a more detailed representation of memory utilization in Cisco NX-OS.

In the output in the figure, these are the most important fields:

- **MemTotal (kB):** Total amount of memory in the system (4 GB in the Cisco Nexus 7000 Series Sup1)
- **Cached (kB):** Amount of memory that the page cache uses (including files in temporary file storage (tmpfs) mounts and data that is cached from persistent storage or bootflash)
- **RAMCached (kB):** Amount of memory that the page cache uses and that cannot be released (data that is not backed by persistent storage)
- **Available (Pages):** Amount of free memory in pages (including the space that can be made available in the page cache and free lists)
- **Mapped (Pages):** Memory that is mapped into page tables (data that nonkernel processes are using)
- **Slab (Pages):** Rough indication of kernel memory consumption

One page of memory is equivalent to 4 KB of memory. The **show system internal kernel memory global** command displays the memory usage for the page cache and kernel and process memory.

Page Cache

```
N7K# show system internal flash
```

Mount-on	1K-blocks	Used	Available	Use%	Filesystem
/	409600	43008	367616	11	/dev/root
/proc	0	0	0	0	proc
/sys	0	0	0	0	none
/isan	409600	269312	140288	66	none
/var/tmp	307200	876	306324	1	none
/var/sysmgr	1048576	999424	49152	96	none
/var/sysmgr/ftp	307200	24576	282624	8	none
/dev/shm	1048576	412672	635904	40	none
/volatile	204800	0	204800	0	none
/debug	2048	16	2032	1	none
/dev/mqueue	0	0	0	0	none
/mnt/cfg/0	76099	5674	66496	8	/dev/hda5
/mnt/cfg/1	75605	5674	66027	8	/dev/hda6
/bootflash	1796768	629784	1075712	37	/dev/hda3
/var/sysmgr/startup-cfg	409600	27536	382064	7	none
/mnt/plog	56192	3064	53128	6	
/dev/mtddblock2					
/dev/pts	0	0	0	0	devpts
/mnt/pss	38554	6682	29882	19	/dev/hda4
/slot0	2026608	4	2026604	1	/dev/hdc1
/logflash	7997912	219408	7372232	3	/dev/hde1
/bootflash_sup-remote	1767480	1121784	555912	67	
127.1.1.6:/mnt/bootflash/					
/logflash_sup-remote	7953616	554976	6994608	8	
127.1.1.6:/mnt/logflash/					

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-51

DRAM is a limited resource on all platforms and must be controlled or monitored to ensure that utilization is kept under control. Cisco NX-OS uses memory in three ways:

- **Page cache:** When you access files from persistent storage (CompactFlash), the kernel reads the data into the page cache. Therefore, when you access the data in the future, you can avoid the slow access times that are associated with disk storage. The kernel can release cached pages if other processes need the memory. Some file systems (tmpfs) exist purely in the page cache (for example, /dev/sh, /var/sysmgr, /var/tmp), so there is no persistent storage of this data, which cannot be recovered when removed from the page cache. Any tmpfs-cached files release page-cached pages only when they are deleted.
- **Kernel:** The kernel needs memory to store its own text, data, and Loadable Kernel Modules (LKMs). LKMs are pieces of code that are loaded into the kernel (as opposed to being a separate user process). An example of kernel memory usage is when an in-band port driver allocates memory to receive packets.
- **User processes:** This memory is used by Cisco NX-OS or Linux processes that are not integrated in the kernel (such as text, stack, heap, and so on).

When you are troubleshooting high memory utilization, you must first determine which type of utilization is high (process, page cache, or kernel). After you have identified the type of utilization, you can use additional troubleshooting commands to help you figure out which component is causing this behavior.

If the Cached or RAMCached utilization type is high, check the file system utilization and determine which kind of files are filling the page cache. The **show system internal flash** command displays the file system utilization. (The output is similar to `df -hT` included in the memory alerts log.)

In the example in the figure, utilization is high because /var/sysmgr (or its subfolders) is using much space. Because /var/sysmgr is a tmpfs mount, the files exist in RAM only. You need to determine which type of files are filling the partition and where they came from; for example, cores or debugs. Deleting the files will reduce utilization, but you should try to determine which type of files are taking up the space and which process left the files in tmpfs.

The **show system internal dir full directory path** command lists all the files and sizes for the specified path (hidden command).

The **filesystems delete full file path** command deletes a specific file (hidden command). Use caution when using this command. You cannot recover a deleted file.

You can also use the **show hardware internal proc-info pcacheinfo** command to determine how much space each file system is using in the page cache (Cached). The command output might help you to determine which persistent file systems are using the page cache and how much memory they are using.

User Processes

- Figure out which process is using a lot of memory.

```
N7K# show processes memory
PID MemAlloc MemLimit MemUsed StackBase/Ptr Process
-----
4662 52756480 562929945 150167552 bffdf00/bfffd970 netstack
```

- Figure out how a specific process is using memory.

```
N7K# show system internal sysmgr service pid 4727
Service "pixmap" ("pixmap", 109):
UUID = 0x133, PID = 4727, SAP = 176
State: SRV_STATE_HANDSHAKED (entered at time Fri Nov 12 01:42:01 2010).
Restart count: 1
Time of last restart: Fri Nov 12 01:41:11 2010.
The service never crashed since the last reboot.
Tag = N/A
Plugin ID: 1
```

If page cache and kernel issues have been ruled out, utilization might be high as a result of some user processes taking up too much memory or of many running processes (because of the number of VDCs or enabled features).

Cisco NX-OS defines memory limits (rlimit) for most processes. If this rlimit is exceeded, sysmgr stops the process and a core file is usually generated. Processes that are close to their rlimit might not have a large impact on platform utilization but can still become an issue if a failure occurs.

The output of the **show process memory** command might not provide a completely accurate picture of the current utilization (allocated does not mean in use). This command is useful for determining whether a process is approaching its rlimit.

To determine how much memory the processes are actually using, check the resident set size (RSS). This value will give you a rough indication of the amount of memory (in kilobytes) that the processes are consuming. The **show system internal processes memory** command displays the process information in the memory alerts log (if the event occurred).

If you see an increase in the utilization for a specific process over time, you should gather additional information about the process utilization.

If you have determined that a process is using more memory than expected, investigate how the memory is being used by the process. The **show system internal sysmgr service pid *PID_in_decimal*** command dumps the service information running the specified protocol ID (PID). The **show system internal kernel memory uuid *UUID_in_decimal*** command displays the detailed process-memory usage, including the libraries for a specific universal user ID (UUID) in the system. (Convert the UUID from the sysmgr service output.)

The **show system internal service mem-stats detail** command displays detailed memory utilization, including the libraries for a specific service. These outputs are usually requested by the Cisco customer support representative when investigating a potential memory leak in a process or its libraries.

Built-In Platform Memory Monitoring

- Memory thresholds:
 - 85% MINOR
 - 90% SEVERE
 - 95% CRITICAL
 - The thresholds are configurable.
- Memory alerts:
 - If a memory threshold is passed, the Cisco NX-OS platform manager captures a snapshot of memory utilization and logs an alert to syslog.
 - The **show system internal memory-alerts-log** command displays the memory alerts log.
 - The **show system internal memory-status** command allows you to check the current memory alert status.

```
N7K# show system internal memory-status
MemStatus: OK
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-53

Cisco NX-OS has built-in kernel monitoring of memory usage to help avoid system response disruptions, process failures, and other undesirable behavior. The platform manager periodically checks the memory utilization (relative to the total RAM present) and automatically generates an alert event if the utilization passes the configured threshold values. When an alert level is reached, the kernel attempts to free memory by releasing pages that are no longer needed; for example, the page cache of persistent files that are no longer being accessed. If critical levels are reached, the kernel stops the highest utilization process. Other Cisco NX-OS components have introduced memory alert handling such as Border Gateway Protocol (BGP) graceful low-memory handling, which allows processes to adjust their behavior to keep memory utilization under control.

Although Cisco NX-OS implements VDCs, remember that a specific VDC memory utilization is not limited. Platform memory issues affect all configured VDCs.

In Cisco NX-OS Release 4.2(4) and later, these are the memory alert thresholds:

- 85% MINOR
- 90% SEVERE
- 95% CRITICAL

This change was introduced in part because of baseline memory requirements when many features or VDCs are deployed. The thresholds are configurable, using the **system memory-thresholds minor percentage severe percentage critical percentage** command. The **show system internal memory-status** command allows you to check the current memory alert status.

If a memory threshold has been passed (from OK to MINOR, MINOR to SEVERE, or SEVERE to CRITICAL), the Cisco NX-OS platform manager captures a snapshot of memory utilization and logs an alert to syslog, as of Cisco NX-OS Release 4.2(4) for the default VDC only. This snapshot is useful in determining why memory utilization is high (process, page cache, or kernel). The log is generated in the Linux root path (/) and a copy is moved to onboard failure logging (OBFL) (/mnt/plog) if possible. This log is useful for determining whether memory utilization is high because of the memory that was consumed by the page cache, kernel, or Cisco NX-OS user processes.

Troubleshooting CPU

This topic explains how to troubleshoot issues that relate to CPU on a Cisco Nexus 7000 Series Switch.

Troubleshooting CPU

- Common reasons for high CPU:
 - Excessive CPU-bound traffic, control plane churn
 - Access-list processing, hardware programming
 - Misbehaving process
- Brief high CPU utilization not automatically a problem indication
- Suggested troubleshooting to get started:
 - **show hardware internal cpu-mac inband stat**
 - **show system internal processes cpu**
 - Ethalyzer

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-55

Brief high CPU utilization does not automatically indicate a problem:

- The Cisco Nexus 7000 Series Switch is a dual-core, Linux-based system with a robust preemptive scheduler (one functional unit for both the route processor [RP] and switch processor [SP])
- Strict control plane and data plane separation
- Scheduler ensures fair access to CPU for all processes
- Lower-level processes (drivers) run in FIFO or nonpre-emptive mode

These are common reasons for high CPU:

- Excessive CPU bound traffic or control plane turnover
- Access-list processing or hardware programming
- Misbehaving processes

To get started with troubleshooting, use these commands or tool:

- **show hardware internal cpu-mac inband stat**
- **show system internal processes cpu**
- Ethalyzer

General Health Check on Supervisor

```
N7k-3-VDC3# show system resources
Load average: 1 minute: 0.64 5 minutes: 1.08 15 minutes: 1.30
Processes : 3912 total, 2 running
CPU states : 4.5% user, 5.0% kernel, 90.5% idle
Memory usage: 4115232K total, 3434268K used, 680964K freeN7k
```

```
N7k-3-VDC3# show processes cpu history
1 2 111 11111211233 1 1 111 1 1 1 6 112 1 1 21132 1 111 123
919275058862141899918384800583739174756080779143297264026770
100
90
80
70
60
50
40
30
20
10
0...5...1...2...2...3...3...4...4...5...5...
0 5 0 5 0 5 0 5 0 5
CPU% per second (last 60 seconds)
# = average CPU%
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-58

To see how many processes were scheduled to run, use the **show system resources** command. You can also see the load in average per the entire system in the past 1, 5, and 15 minutes. With the same command, you can also see how much of the CPU cycles are used by user configured processes and kernel processes.

To see the CPU utilization 60 seconds ago, use the **show processes cpu history** command.

General Health Check on Supervisor (Cont.)

```
N7K-3-VDC3# show processes cpu | egrep "PID|--|ospf"
```

PID	Runtime(ms)	Invoked	uSecs	1Sec	Process
9337	102	72	1418	0.0%	ospfv3
22916	118	62	1905	13.1%	ospf

```
N7K-3-VDC3# show system internal sysmgr service pid 22916
```

```
Service "__inst_001_ospf" ("ospf", 58):  
  UUID = 0x41000119, PID = 22916, SAP = 320  
  State: SRV_STATE_HANDSHAKED (entered at time Thu May 3 21:53:59 2012).  
  Restart count: 1  
  Time of last restart: Thu May 3 21:53:58 2012.  
  The service never crashed since the last reboot.  
  Tag = 6467  
  Plugin ID: 1
```

In the output in the figure, these are the most important fields:

- **PID:** Process ID
- **Runtime:** Total nonidle time process that has been actively using CPU
- **Invoked:** Number of times that the process has been context-switched, both voluntarily (finished job) and involuntarily (scheduler interrupt)
- **uSecs:** Average amount of time that the process was running during a single context switch

You can see additional useful process level details by using the **show system internal sysmgr service pid number** command.

High CPU Utilization and Control Plane Instability

- Typical datacenter traffic causing high CPU utilization:
 - ARP, ND (IPv6)
 - DHCP traffic
 - Glean traffic (no ARP or ND)
 - Malicious traffic to 224.0.0.0/24 subnet
 - Fragments or malicious Layer-2 multicast or other traffic
- CPU protection via CoPP policers
- CPU protection via Layer 2 or Layer 3 hardware rate limiters
- CoPP and rate limiter default settings might need to be adjusted, based on network requirement specifics

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-58

Typical data center traffic can cause high CPU utilization:

- Address Resolution Protocol (ARP) and Neighbor Discovery (ND) in IPv6
- DHCP traffic
- Glean traffic (no ARP or ND)
- Malicious traffic to the 224.0.0.0/24 subnet
- Fragments or malicious Layer-2 multicast or other traffic

CPU protection can be implemented via Control Plane Policing (CoPP) policers or via Layer 2 or Layer 3 hardware rate limiters. CoPP and rate-limiter default settings might need to be adjusted, based on network requirement specifics. CoPP provides more granular, targeted CPU protection, whereas rate limiters work better with traffic categories in which specifics (source and destination IP) might not be known. Both CoPP and rate limiters are configured per Cisco Nexus 7000 M-1Series I/O Module, and the total RP-bound traffic that is allowed is the sum across all those I/O Modules. CoPP and rate-limiter adjustments must allow reasonable protocol convergence and CPU protection at the same time.

High CPU Utilization and Control Plane Instability (Cont.)

```

N7K-1# show hardware internal cpu-mac inband stats | egrep " Rx|
Tx|counters|Throttle|Tick|rate|total|good|XOFF p|XON p"
RMON counters
total packets          Rx          Tx
good packets          779905245    1421785114
good octets (hi)       779905245    1421650279
good octets (low)      0          0
total octets (hi)      172303021767  192965708376
total octets (low)     0          0
total octets (low)     172302724342  192974265660
XON packets            0          67627
XOFF packets           0          67208
<->
Error counters
Rx no buffers ..... 0
<->
Throttle statistics
Throttle interval ..... 2 * 100ms
Packet rate limit ..... 32000 pps
Tick counter ..... 12414130
Rx packet rate (current/max) 4993 / 20296 pps
Tx packet rate (current/max) 60 / 3474 pps
<->
MAC counters
MAC0 (R2D2)
total packets          Rx          Tx
total bytes          779905246    1421790561
XOFF packets auto-generated 5447
XOFF packets           0          7590855
XON packets           0          6731953
MAC1 (CPU)
total packets          Rx          Tx
total bytes          3996073897  504693696
XOFF packets           0          18561642
XON packets           0          18561642
<->

```

With the command that is used in the figure, you can find this information:

- Total number of frames that the CPU receives and sends
- Hard-coded maximum limit (might not be reached with larger packet size)
- How many times throttling kicked in
- CPU bound traffic (current packets per second [p/s] and maximum p/s that were reached)

Another useful output from the **show hardware internal statistics device mac qos asic-instance 0** command shows CPU-bound traffic per class of service (CoS) breakdown and tail-drops towards CPU

Another important field in the **show hardware internal cpu-mac inband stats** command output is the Rx No Buffers field, which represents how many packets were dropped toward the CPU. These are packets that already made it through CoPP and the hardware rate limiters.

The challenge is to identify the offending traffic type and its source.

High CPU Utilization and Control Plane Instability (Cont.)

```
N7K1# show hardware internal cpu-mac inband events
```

- 1) Event:TX_PPS_MAX, length:4, at 298568 usecs after Sat Jul 14 15:09:00 2012
new maximum = 1161
- 2) Event:RX_PPS_MAX, length:4, at 298567 usecs after Sat Jul 14 15:09:00 2012
new maximum = 976
- 3) Event:RX_PPS_MAX, length:4, at 298571 usecs after Sat Jul 14 15:08:59 2012
new maximum = 456
- 4) Event:TX_PPS_MAX, length:4, at 298585 usecs after Sat Jul 14 15:08:40 2012
new maximum = 601
- 5) Event:RX_PPS_MAX, length:4, at 298584 usecs after Sat Jul 14 15:08:40 2012
new maximum = 426
- 6) Event:TX_PPS_MAX, length:4, at 298583 usecs after Sat Jul 14 15:08:24 2012
new maximum = 449

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-60


The **show hardware internal cpu-mac inband events** command records a time stamp whenever the CPU hits a new maximum transmit (Tx) or receive (Rx) p/s rate. The previous example of the **show hardware internal cpu-mac inband stats** command gives the current rate only, so the **event** option allows you to see the historical events.

Troubleshooting Switch Fabric

This topic explains how to troubleshoot issues that relate to the switch fabric on a Cisco Nexus 7000 Series Switch.

Switch Fabric Utilization

- Ingress fabric interface ASIC knows all active paths through 3-stage xbar to each destination.
- First and next fragments may take different path because of missing Layer-4 information in next fragments.



Cisco Nexus 7000 Series Fabric Modules are separate modules that provide parallel fabric channels to each I/O and supervisor module slot. As many as five simultaneously active fabric modules work together to deliver up to 230 Gb/s or 550 Gb/s per slot. Through the parallel forwarding architecture, a system capacity of more than 8 terabits per second (Tb/s) is achieved with the five fabric modules.

Ingress fabric interface ASIC knows all active paths through a three-stage xbar to each destination. Unicast traffic is distributed (a 2.5-kB superframe is broken into small chunks) across all active paths to egress fabric interface ASIC. Multicast traffic selects one active path to egress fabric interface ASIC, based on the hash result calculated on Layer 2, Layer 3, and Layer 4 information. (This hash is like the EtherChannel hash but is not configurable). First and next fragments may take different path due to missing Layer 4 information in next fragments.

Note

For 1 Gb, you need one module. For 10 Gb, you need two active modules.

Switch Fabric Utilization (Cont.)

- 18 slot chassis with 5 xbar modules (10 x 23G = 230G) capacity per slot

```
N7K-3# show hardware fabric-utilization
```

Slot	Total Fabric Bandwidth	Utilization	
		Ingress %	Egress %
1	230Gbps	4.0	0.0
2	230 Gbps	0.0	2.0
3	230 Gbps	0.0	1.0
9	115 Gbps	0.0	1.0

```
<->
```



```
N7K1# show hardware fabric-utilization detail timestamp
```

Fabric Planes:
A -- Unicast fabric interface
B -- Multicast/Multidestination fabric interface

PEAK FABRIC UTILIZATION									
I/O				Ingress		Egress			
Slot	Mod	Inst	Plane	Util	Time	Util	Time	Util	Time
1	1	1	A	0%	07-19@11:02:20	0%	07-19@11:02:20	0%	07-19@11:02:20
1	1	1	B	0%	07-19@11:02:20	0%	07-19@11:02:20	0%	07-19@11:02:20
1	2	1	A	0%	07-19@11:02:20	0%	07-19@11:02:20	0%	07-19@11:02:20
1	2	1	B	0%	07-19@11:02:20	0%	07-19@11:02:20	0%	07-19@11:02:20
1	3	1	A	0%	07-19@11:02:20	0%	07-19@11:02:20	0%	07-19@11:02:20

```
<->
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-63

There are four virtual output queues (VOQs) to every egress port ASIC (every 12 1 Gigabit Ethernet ports or 4 10 Gigabit Ethernet ports in shared mode, every 1 10 Gigabit Ethernet port in dedicated mode, or 2 1 or 10 Gigabit Ethernet ports).

Unicast traffic access to fabric is arbitrated (the arbiter on the active supervisor provides access when there is enough bandwidth available to the destination VOQ). Multicast traffic access to fabric is nonarbitrated.

You can also display the peak utilization time stamp by using the **show hardware fabric-utilization detail timestamp** command.

Switch Fabric Data Collection

- For data collection, use these commands:
 - **show hardware internal xbar-driver event-history errors|msgs**
 - **show logging onboard internal xbar**
 - **show event-history xbar**

```
N7K-1-PeerA# show system internal xbar ?
```

all	Show xbar all data
dyn-mcast-info	Show xbar dynamic multicast info
dynamic-mc	Show xbar dynamic multicast table
event-history	Show internal event history
flood-mc	Show xbar flood multicast table
get-mi-slotmask	Enter the slotmask
mc	Show xbar sw multicast table
mem-stats	Show xbar allocation statistics
static-mc	Show xbar static multicast table
sw	Show xbar sw data
sync-loss-threshold	Enable setting sync-loss handling params
vqi-info	Show internal vqi-info
vqi-map	Show vqi-map information

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-64

For any packet loss-related issues, first use the **show hardware internal error module X** command. When you see any potentially related counters moving, use the **show hardware internal statistics module X device device category asic-all** command to filter out unnecessary output. (The CLI can produce long output that is difficult to read.)

To collect any VOQ-to-fabric interaction issues, use the **show hardware internal qengine asic 0|1** command.

Troubleshooting CoPP and Data Plane Rate Limiters

This topic explains how to troubleshoot issues that relate to CoPP and rate limiting on a Cisco Nexus 7000 Series Switch.

Copying the CoPP Best Practice Policy

- The CoPP best practice policy is read-only. If you want to modify its configuration, you must copy it.
- Use the **copp copy profile {strict | moderate | lenient | dense} {prefix | suffix} string** command to create a copy of the CoPP best practice policy.
 - CoPP renames all class maps and policy maps with the specified prefix or suffix.
- Use the **show copp status** command to display the CoPP status, including the last configuration operation and its status.
 - This command also lets you verify that the copied policy is not attached to the control plane.
- Use the **show running-config copp** command to display the CoPP configuration in the running configuration, including the copied policy configuration.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-66

CoPP protects the control plane and separates it from the data plane, thereby ensuring network stability, reachability, and packet delivery. Rate limits can prevent redirected packets for exceptions from overwhelming the supervisor module on a Cisco NX-OS device.

The CoPP best practice policy is read-only. If you want to modify its configuration, you must copy the policy.

CoPP and Rate Limiters CPU Protection

- Modify copp-system-acl-ospf to permit only specific IPs or subnets.
- Create copp-system-acl-224malicious access-list.
- Add copp-system-class-malicious class with zero policer.

```
N7K-1# show policy-map interface control-plane module 2 | egrep
"service-policy|critical|ospf|police cir 39600|malicious"
service-policy input: copp-system-policy
class-map copp-system-class-critical (match-any)
  match access-grp name copp-system-acl-ospf
  match access-grp name copp-system-acl-ospf6
  police cir 39600 kbps , bc 250 ms

N7K-1# show class-map type control-plane copp-system-class-critical
| egrep class|ospf
class-map type control-plane match-any copp-system-class-critical
  match access-grp name copp-system-acl-ospf
  match access-grp name copp-system-acl-ospf6

N7K-1# show ip access-lists copp-system-acl-ospf
IP access list copp-system-acl-ospf
10 permit ospf any any
```


No "malicious" class to
block malicious traffic

In the example that the figure shows, there is no "malicious" class to block any malicious traffic. What is the solution?

- Modify **copp-system-acl-ospf** to permit only specific IPs or subnets.
- Create a **copp-system-acl-224malicious** access list.
- Add a **copp-system-class-malicious** class with zero policer.
- The same approach can be used for any offending 224.0.0.0/24 traffic.

Keep in mind that CoPP is applied for all VDCs but can be modified only from the default VDC. Eventually if a unique IP scheme per VDC is used, each VDC can have different CoPP policies.

CoPP and Rate Limiters CPU Protection (Cont.)



After the steps that are shown in the figure, you can check for any offending traffic:

```
N7K-1# show policy-map interface control-plane module 2 class copp-
system-class-malicious
```

control Plane

```

service-policy input: copp-system-policy
 class-map copp-system-class-malicious (match-any)
  match access-grp name copp-system-acl-224malicious
  police cir 1 bps , bc 200 ms
 module 2 :
  conformed 0 bytes; action: drop
  violated 1799505072 bytes; action: drop
  
```

In this example, you can see that offending traffic is dropped. If you enter the same command for Module 1, you can see that offending host is only on Module 2:

```
N7K-1# show policy-map interface control-plane module 1 class copp-
system-class-malicious
```

control Plane

```

service-policy input: copp-system-policy
 class-map copp-system-class-malicious (match-any)
  match access-grp name copp-system-acl-224malicious
  police cir 1 bps , bc 200 ms
 module 1 :
  conformed 0 bytes; action: drop
  violated 0 bytes; action: drop
  
```

Depending on how routing is performed in virtual Port Channel (vPC) configuration, CoPP adjustment might be required on both vPC peers.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Use the **show license** command to display all license information configured on the system.
- Use the **show install all status** command to watch the progress of your software upgrade.
- Verify that Cisco Fabric Services is enabled for the application on all devices in the network or Cisco Fabric Services region by using the **show cfs application** command.
- Problems with VDCs usually occur from logging in to the incorrect VDC or misallocating resources for a VDC.
- Each VRF is associated with an RIB, and this information is collected by the FIB.
- To troubleshoot LinkSec and port QoS, you can use the **show cts interface [all | interface]** command.
- You can assess the overall level of memory utilization on the platform by using two basic CLI commands: **show system resources** and **show processes memory**.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-69

Summary (Cont.)

- Start troubleshooting CPU by using the **show hardware internal cpu-mac inband stat** and **show system internal processes cpu** commands.
- For switch fabric data collection, use the **show hardware internal xbar-driver event-history errors|msgs** command.
- CoPP protects the control plane by ensuring network stability, reachability, and packet delivery; rate limits can prevent redirected packets for exceptions from overwhelming the supervisor module.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-70

Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches

Overview

This lesson is designed to provide some examples of common issues that relate to Cisco Nexus 5000 Series Switches and methods to resolve them.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that are specific to Cisco Nexus 5000 Series Switches. This ability includes being able to meet these objectives:

- Explain how to troubleshoot issues that relate to licensing on a Cisco Nexus 5000 Series Switch
- Explain how to troubleshoot issues that relate to Cisco IOS ISSU on a Cisco Nexus 5000 Series Switch
- Explain how to troubleshoot issues that relate to configuration synchronization on a Cisco Nexus 5000 Series Switch
- Explain how to troubleshoot issues that relate to QoS on a Cisco Nexus 5000 Series Switch
- Explain how to troubleshoot issues that relate to CRC errors on a Cisco Nexus 5000 Series Switch
- Explain how to troubleshoot issues that relate to high CPU on a Cisco Nexus 5000 Series Switch
- Explain how to resolve issues that relate to unified ports on a Cisco Nexus 5500 Platform Switch

Troubleshooting Licensing

This topic explains how to troubleshoot issues that relate to licensing on a Cisco Nexus 5000 Series Switch.

Cisco NX-OS Licensing: Cisco Nexus 5000 Series and Nexus 5500 Platform Switches	
Package	Content
Layer 3 Base Services Package	• Static routing • Routing Information Protocol version 2 (RIPv2) • Open Shortest Path First version 2 (OSPFv2), limited to 256 dynamically learned routes • Enhanced Interior Gateway Routing Protocol (EIGRP) stub • Hot Standby Router Protocol (HSRP) • Virtual Router Redundancy Protocol (VRRP) • Internet Group Management Protocol version 2 (IGMPv2) and version 3 (IGMPv3) • Protocol Independent Multicast version 2 (PIMv2) sparse mode • Routed ACL • Unicast Reverse Path Forwarding (uRPF)
Layer 3 Enterprise Services Package	• Full EIGRP • OSPF with scalability as many as 8000 routes • Border Gateway Protocol (BGP) and Multi-VRF Customer Edge (VRF-Lite) (IP-VPN) • Maximum routes supported by Layer 3 hardware: 8000 entries
Storage Protocols Services Package	• Native Fibre Channel • FCoE • NPV • Fibre Channel port security • Fabric binding

© 2012 Cisco and/or its affiliates. All rights reserved. DCUFT v5.0—5.4


The licensing model for the Cisco Nexus Operating System (NX-OS) Software is feature-based. Feature-based licenses make features available to the entire physical device. Each license supports only the listed features.

Note	Any feature that is not included in a license package is bundled with the Cisco NX-OS Software and is provided at no extra charge.
-------------	--

Installing any license in the device is a nondisruptive process and automatically saves a copy of the permanent license to the chassis.

If you have enabled the grace period feature, then enabling a licensed feature that does not have a license key starts a counter on the grace period. You then have 120 days to install the appropriate license keys, disable the use of that feature, or disable the grace period feature. If at the end of the 120-day grace period the device does not have a valid license key for the feature, the Cisco NX-OS Software automatically disables the feature and removes the configuration from the device.

Cisco NX-OS: Obtaining the License Key File



Licenses can be obtained as factory-installed licenses for new Cisco NX-OS devices.

If you have an existing device or want to install the license yourself, you must first obtain the license key file and then install that file in the device.

The figure shows how to obtain a license key file.

To obtain the serial number for a device, use the **show license host-id** command.

To install the license, use the **install license bootflash:file_name** command.

Troubleshooting Licensing

- Alerts and error messages are the same as those used in troubleshooting licensing for Cisco Nexus 7000 Series Switches.
- To display the license configuration information, perform one of these tasks.

Command	Purpose
show license [brief]	Displays information for all installed license files
show license feature package mapping	Displays information about features available in installed license packages
show license file	Displays information for a specific license file
show license host-id	Displays the host ID for the physical device
show license usage	Displays the usage information for installed licenses

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0—5-6

Use the **show license** commands to display all license information that is configured on the system.

Troubleshooting Cisco IOS ISSU

This topic shows how to troubleshoot issues that relate to Cisco IOS In-Service Software Upgrade (ISSU) on a Cisco Nexus 5000 Series Switch.

Determining the Upgrade Impact

Command	Definition
show incompatibility system	Displays incompatible configurations on the current system that will affect the upgrade version
show install all impact	Displays information that describes the impact of the upgrade on each fabric extender, including the current and upgrade-image versions; also displays whether the upgrade is disruptive, whether the fabric extender needs to be rebooted, and why
show spanning-tree issu-impact	Displays the spanning-tree configuration and whether there are potential STP issues
show lacp issu-impact	Displays the port priority information and whether there are potential issues

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.6

The figure lists the **show** commands that identify the impact or potential problems that might occur when performing a Cisco IOS ISSU.

Monitoring the Upgrade Status

- This example shows the output from the **show install all status** command.

```
N5k-1# show install all status
There is an on-going installation...
Enter Ctrl-C to go back to the prompt.

Continuing with installation process, please wait.
The login will be disabled until the installation is completed.

Performing supervisor state verification.
SUCCESS

Supervisor non-disruptive upgrade successful.

Pre-loading modules.
SUCCESS

Module 198: Non-disruptive upgrading.
SUCCESS

Module 199: Non-disruptive upgrading.
SUCCESS

Install has been successful. (hit Ctrl-C here)
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.9

The example in the figure shows the output from the **show install all status** command.

The following example shows the output from the **show fex** command on two virtual port channel (vPC) peer switches on which fabric extenders Fex 198 and Fex 199 are upgraded:

switch-1# **show fex**

FEX Number	FEX Description	FEX State	FEX Model	FEX Serial
198	FEX0198	Hitless Upg Idle	N2K-C2248TP-1GE	JAF1342ANQP
199	FEX0199	Online	N2K-C2248TP-1GE	JAF1342ANRL

switch-2# **show fex**

FEX Number	FEX Description	FEX State	FEX Model	FEX Serial
198	FEX0198	FEX AA Upg Idle	N2K-C2248TP-1GE	JAF1342ANQP
199	FEX0199	Online	N2K-C2248TP-1GE	JAF1342ANRL

Cisco IOS ISSU Failure Conditions

- The following conditions stop a Cisco IOS ISSU process from continuing:
 - The supervisor module bootflash file system does not have sufficient space to accept the updated image.
 - The specified system and kickstart images are incompatible after an upgrade.
 - Configuration changes are made while the upgrade is in progress.
 - Hardware is installed or removed while the upgrade is in progress.
 - A power disruption occurs while the upgrade is in progress.
 - The entire path for the remote server location is not specified accurately.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-10

These conditions will stop a Cisco IOS ISSU process from continuing:

- The supervisor module bootflash file system has insufficient space to accept the updated image.
- Images are incompatible after an upgrade; for example, an I/O module image or a kickstart image might be incompatible with a system image. This condition is shown in the **show install all impact** command output in the compatibility check section of the output (under the Bootable column).
- The specified system and kickstart images are not compatible.
- Configuration changes are made while the upgrade is in progress.
- Hardware is installed or removed while the upgrade is in progress.
- A power disruption occurs while the upgrade is in progress.
- The entire path for the remote server location is not specified accurately.

The Cisco NX-OS software prevents most configuration changes while the **install all** command is in progress. However, the Cisco NX-OS software allows configuration changes from Cisco Fabric Services, and those changes might affect the Cisco IOS ISSU.

Troubleshooting Configuration Synchronization

This topic explains how to troubleshoot issues that relate to configuration synchronization on a Cisco Nexus 5000 Series Switch.

Commit Failure

- Command parsing failed:
 - Ensure that appropriate conditional feature or features are enabled.
- Verify failed:
 - Remove the commands from conf-t or delete them from the switch-profile buffer and reissue the commit.
- Commands failed during commit:
 - Correct the reason for the failure and reissue the commit.
- Another session in progress:
 - Compare the vPC domain IDs of the two switches and ensure that they match.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-12

When a commit fails, commands that were entered under the switch profile are still stored in the switch profile buffer. Do not configure these commands under the switch profile again. After correcting the cause of the failure, only the commit needs to be executed.

Use the **show switch-profile status commit** command to view commit status. Commit failure has many possible causes:

- Command parsing failed
 - **Possible cause:** Appropriate conditional feature or features are not enabled.
 - **Solution:** Ensure that appropriate conditional feature or features are enabled.
 - This error message indicates that some feature commands have not been configured. Feature commands are not allowed to be configured within the switch profile and must be configured on both peers from conf-t.
- Verify failed
 - **Possible cause:** The listed commands failed mutual-exclusion checks. These commands have already been configured under conf-t.
 - **Solution:** If you do not want the commands that are 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 these steps:

- Step 1** Use the **show switch-profile buffer** command to view commands in the switch processor (SP) buffer.
- Step 2** Use the **buffer delete range** command to delete commands that are indicated by the sequence numbers.
- Step 3** Use the **buffer-move seq id seq id** command to rearrange commands in the buffer. This command is useful when commands in the buffer are ordered incorrectly.

■ **Commands failed commit**

- **Possible cause:** Commands failed during commit.
- **Solution:** Correct the reason for the failure and reissue the commit. If the commit continues to fail, issue the same command from conf-t. If the command succeeds from conf-t, then use the **show system internal csm info trace** command to look for any errors that relate to the command. For every command that is executed from config-sync, a `csm_cmd_status[0x0]` line in the trace log 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 determine whether conf-t or config-sync has taken a lock.
 - If conf-t has taken a lock and not released it, then command output like this example is displayed. The client type should be set to 2, as the example shows.

```
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
```

Use the **show accounting log** command to identify the command that acquired the lock. After identifying the command, check for its success/failure status. If the command did not return a status, then config-sync will 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 output.

Use the **show switch-profile status** command to determine whether 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, then use the **show cfs lock** command to determine whether the Cisco Fabric Services fabric is locked. Identify the application that locked Cisco Fabric Services. If the application is session-manager, then the Cisco Fabric Services 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 configuration:
 - Determine whether a system resource-utilization problem exists. Correct the problem and retry the operation.
 - Command does not exist in global_db:
 - Use the **show system internal csm info global-db cmd-tbl** command to determine whether the command exists in global_db.
 - Mutual exclusion check failed on peer:
 - Remove the failed commands from conf-t on the peer and then retry the import verify/commit.
 - Use the **show system internal csm info trace** command for further investigation to look at events, trace, or error messages.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-13

Use the **show switch-profile status** command to view import status. Import failure has several possible causes:

- Failed to collect the running configuration
 - **Possible cause:** Failure occurs if the system is too busy and the **show running** command did not complete.
 - **Solution:** Determine whether a system resource utilization problem exists. Correct the problem and retry the operation.
- The **show switch-profile status** command does not exist in global_db
 - **Possible cause:** The command is missing from global_db.
 - **Solution:** Use the **show system internal csm info global_db cmd-tbl** command to determine whether the command exists in the global_db.
 - If the command exists in the global_db, there might not be enough space in the show run for the command. Ensure that there are no trailing spaces or tabs in show running configuration generation.
 - If the command does not exist in global_db, use the **show accounting log** command to determine whether 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 the show running. If the command is displayed, then the application should correct it.

- If the command was configured before reload or Cisco IOS ISSU, add back the command. If the accounting log shows the command retval as success, determine whether the command is getting added to global_db. If the command was added correctly, copy r s, check global_db reload, and determine whether the command exists in global_db. If the command does not exist in 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 is added to 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 the 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

- For a merge to succeed, the configuration in the switch profile on both peers must match exactly.
- Use the **show switch-profile status** command to view which commands failed the merge.
- Merge failure has many possible causes:
 - First time merge failure:
 - Remove the **sync-peers destination** command from the switch profile.
 - Ensure that the configuration is the same as under the switch profile.
 - Add back the **sync-peers destination** command to the switch profile.
 - Reissue the commit.
 - Merge after peers that were previously in sync:
 - Correct the configurations and reissue the commit from the peer with the corrected configuration.
 - Merge after reload:
 - Correct the configurations and reissue the commit.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-14

A merge between peers happens when a peer becomes reachable. A merge is initiated when Cisco Fabric Services 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 several possible causes:

■ First time merge failure

- **Possible cause:** When peer switches try to synchronize configurations, the merge might fail when validating received configurations.
- **Solution:** 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 these steps to correct the configurations:

Step 4 Remove the **sync-peers destination** command from the switch profile.

Step 5 Use the **show running switch-profile** command on both peers to ensure that the configuration is exactly the same as under the switch profile.

Step 6 Add back the **sync-peers destination** to the switch profile and reissue the commit.

■ Merge failure after peers were previously in sync

- **Possible cause:** If peers were in sync and connectivity was lost, and conflicting configuration changes were made on the switches, 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 from the peer with the corrected configuration.

- Merge failure after reload
 - **Possible cause:** After a switch is reloaded, it sends its switch-profile configuration to the peer. If a configuration change 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.
- Switch-profile deletion failure has many possible causes:
 - Application failure:
 - Use the **resequence-database** command in conf-sync mode and reissue the delete.
 - Failure from dependent commands:
 - Correct the commands and references, and reissue the delete.
 - Application does not respond:
 - Correct the commands and reissue the delete.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-15

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 that is based on the time stamp or session type. Switch-profile deletion failure has several possible causes:

- Application failure
 - **Possible cause:** Switch-profile deletion failure might occur because the application failed the command, or the configuration was 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.
 - **Solution:** Use the **resequence-database** command in 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 the switch profile is referenced by another command outside the switch profile, and the first command is deleted, then failure occurs because the second command still references the first command.
 - **Solution:** Correct the commands and references, and reissue the delete.
- Application not responding
 - **Possible cause:** The deletion fails because the application does not respond because of the application that owns the command.
 - **Solution:** Correct the commands and reissue the delete.

Verify Failure

- Mutual exclusion check under local information:
 - Delete the command from conf-t mode and run verify from config-sync mode.
- Mutual exclusion check under peer information:
 - Delete the command from conf-t mode on the peer and run verify from config-sync mode.
- Rollback or Cisco IOS ISSU in progress:
 - Stop rollback or wait for it to complete, and then run verify.
- Global_db modification in progress:
 - Wait for the update to complete, and then run verify.
- Peer unable to accept lock request:
 - The peer is processing a transaction and cannot accept a lock request; run verify later.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5.18

Use the **show switch-profile status** command to view messages about the failure. Determine whether the failure is on the local or 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. Verify failure has several possible causes:

- Mutual exclusion check under local information
 - **Possible cause:** The command failed the mutual-exclusion check under local information because the command has already been configured from conf-t.
 - **Solution:** Delete the command from conf-t mode and run verify from config-sync mode.
- Mutual exclusion check under peer information
 - **Possible cause:** The command failed the mutual-exclusion check under peer information because the command has already been configured from conf-t on the peer.
 - **Solution:** Delete the command from conf-t mode on the peer and run verify from config-sync mode.
- Rollback or Cisco IOS ISSU in progress
 - **Possible cause:** Verify cannot be performed when rollback or Cisco IOS ISSU is in progress.
 - **Solution:** Stop rollback or wait for it to complete, and then run verify.
- Global_db modification in progress
 - **Possible cause:** Verify cannot be performed when global_db is being updated on the local or peer side.
 - **Solution:** Wait for the update to complete and then run verify.

■ Peer unable to accept lock request

- **Possible cause:** Verify cannot be performed when the peer is unable to accept the lock request.
- **Solution:** The peer is processing a transaction and cannot accept a lock request. Run verify at a later time. Use the **show switch-profile status** command to determine whether 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 whether the Cisco Fabric Services fabric has been locked.

Also check the application that has taken the Cisco Fabric Services lock. If the application is ssnmgr, use the **show cfs internal session-history name session-mgr** and **show cfs internal notification log name session-mgr** commands to view information about when a lock was acquired or released. These commands can also show the mapping to the Content Switching Module (CSM) transactions that are displayed by using the **show switch-profile session-history** command.

Troubleshooting QoS

This topic explains how to troubleshoot issues that relate to quality of service (QoS) on a Cisco Nexus 5000 Series Switch.

Improper Configurations

- Cannot pass frame size larger than 2300 bytes through switch:
 - The MTU value for the traffic with CoS 7 is set to a fixed value.
 - Use any CoS value other than 7 to avoid this limitation.
- Traffic not queued or prioritized correctly on Cisco Nexus 2148, Nexus 2232, and Nexus 2248TP Fabric Extenders
 - The Cisco Nexus 2148, Nexus 2232, and Nexus 2248TP Fabric Extenders can support only CoS-based traffic classification.

© 2012 Cisco and/or its affiliates. All rights reserved.

DOC#T 50.0-5-18

The Cisco Nexus 5000 Series Switch Cisco NX-OS 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. Several events are caused by improper configurations:

- Cannot pass frame size larger than 2300 bytes through switch

Although the jumbo maximum transmission unit (MTU) has been configured for class-default, you cannot pass a frame size larger than 2300 bytes through the Cisco Nexus 5000 Series Switch and the Cisco Nexus 2000 Fabric Extender.

- **Possible cause:** The class of service (CoS) value might conflict with the existing MTU value.
- **Solution:** CoS 7 is used internally to control traffic between the Cisco Nexus 5000 Series Switch and the Cisco Nexus 2000 Series Fabric Extender. 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.

- Traffic queued or prioritized incorrectly on Cisco Nexus 2148, Nexus 2232, and Nexus 2248TP Fabric Extenders

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

- **Possible cause:** The Cisco Nexus 2148, Nexus 2232, and Nexus 2248TP Fabric Extenders can support CoS-based only traffic classification. The QoS service policy type that is configured under system QoS is populated from the Cisco Nexus 5000 Series Switch to the fabric extender only when all the matching criteria use the match cos match clause. If other match clauses exist, such as match dscp or match ip access-group in the QoS policy map, then the fabric extender 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:** Ingress traffic (from server to network) that is not marked with a CoS value is placed into the default queue on the fabric extender. After the traffic is received on the Cisco Nexus 5000 Series Switch, the traffic is classified based on a configured rule and is placed in the proper queue. Egress traffic (switch to fabric extender, and then from fabric extender to server) should be marked with a CoS value on the Cisco Nexus 5000 Series Switch so that the fabric extender can properly classify and queue the traffic.

PFC Issues

- Link pause (flow control) not enabled on back-to-back Cisco Nexus 5000 Series Switch links:
 - Determine whether both switches have FCoE enabled.
 - Determine whether PFC TLV and DCBX are enabled.
 - Enable link pause instead of PFC on back- to-back switch links.
- Cannot enable pause no-drop on more than one ethernet class:
 - Nexus 5000 Series Switch supports a maximum of three no-drop classes (including FCoE).
- Changing no-drop configuration causes vPC peer link to go down and Fabric extender to go offline:
 - 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.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-6-19

There are several possible issues when working with priority flow control (PFC):

- Link pause (flow control) not enabled on back to back Cisco Nexus 5000 Series Switch links

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

— **Possible cause:** If the peer Cisco Nexus 5000 Series Switch supports the PFC type, length, value (TLV) with Data Center Bridging Exchange (DCBX), then configuring the **flowcontrol send on** and **flowcontrol receive on** commands will not enable the link pause. You must disable the PFC TLV that is sent by DCBX on that interface. Use one of these commands to verify:

- Use the **show interface ethernet x/y flowcontrol** command and determine whether the operating state is off.
- Use the **show interface ethernet x/y priority-flow-control** command and determine whether the operating state is on.

— **Solution:** Configure these commands under **interface ethernet x/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

CLI commands fail with this error when you try 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 Fibre Channel over Ethernet [FCoE]). If five Ethernet classes are created, then there will be insufficient buffers to enable two of the five Ethernet no-drop classes. You will get an error if not enough buffers exist to enable the no-drop; for 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"
```

- **Solution:** 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 the vPC peer link to go down and the fabric extender 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 these nodes, then the vPC peer link does not come up and the fabric extender goes offline. Only CoS-based class no-drop and MTU parameters are considered as a type 1 consistency that is checked for vPC. If you configure an access control list (ACL)-based class, then it is not treated as a type 1 parameter for vPC.

Use one of these commands for verification:

- **show vpc brief**
- **show vpc consistency-parameters global**
- **Solution:** 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.

PFC Issues (Cont.)

- Pause enabled on all CoS values when no-drop enabled on class-ip-multicast
 - Enable PFC on a specific CoS only, instead of on all CoS values under the class-ip-multicast class.
- No-drop class not created on N2K-C2148T/N2K-C2248TP-1GE-based fabric extender with default QoS configuration
 - If you want an Ethernet no-drop class on these devices, you must create an Ethernet no-drop class.
- Enabling link pause (flow control) on Cisco Nexus 5000 Series interface
 - Use these commands under **interface ethx/y** to enable link pause:
 - **flowcontrol receive on**
 - **flowcontrol send on**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-20

- 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, pause is enabled on all of the CoS values. Use the **show interface ethernet x/y priority-flow-control** command and check that the virtual lane bitmap is enabled for all CoS values (ff).
 - **Solution:** Use these 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**
- No-drop class not created on N2K-C2148T/N2K-C2248TP-1GE-based fabric extender with default QoS configuration
The **show queuing interface** command is different for the switch port and host interface (HIF) port on N2K-C2248TP and N2K-C2148T.
 - **Possible cause:** FCoE is not supported on the N2K-C2148T and N2K-C2248TP-1GE-based fabric extender and the no-drop class is not created with the default QoS configuration. Use this command to verify (check for no-drop class):
show queuing interface ethernet 100/1/1
 - **Solution:** If you want an Ethernet no-drop class on a N2K-C2148T/N2K-C2248TP-1GE fabric extender, then you must create an Ethernet no-drop class by using these commands:
 - **policy-map type network-qos no-drop**
 - **class type network-qos class-0**

- **pause no-drop**

- Enabling link pause (flow control) on Cisco Nexus 5000 Series interface

Configuring the **flowcontrol send on** and **flowcontrol receive on** commands does not enable flow control on Cisco Nexus 5000 Series Switch port links when connected to another Cisco Nexus 5000 Series interface.

— **Possible cause:** By default, DCBX runs on the Cisco Nexus 5000 Series interface. If the peer does not run DCBX, then the interface is configured for tail-drop. Use one of these commands to verify:

- Use the **show interface ethernet x/y flowcontrol** command and determine whether the operating state is off.
- Use the **show interface ethernet x/y priority-flow-control** command and determine whether the operating state is off.

— **Solution:** Use these commands under **interface ethernet x/y** to enable link pause:

- **flowcontrol receive on**
- **flowcontrol send on**

Registers and Counters

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 software qosctrl port 0 0 nif <0-48> [sat switch]	Displays the PI information for every port
show platform software qosctrl port 0 0 hif <0-48> [sat switch]	Displays the PI information for every port
show platform software qosctrl policy hif	Displays the global Network-QoS and Queuing configurations
show platform software qosctrl port 0 6 hif 1 counters	Displays counters
show platform software redwood rate	Displays overall statistics for non-zero traffic

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-21

Use these commands to access various registers and counters:

- Cisco Nexus 5000 Series 10 Gigabit Ethernet PFC:
 - **show hard in gatos asic *gatos_num* registers match mm_CFG_pause\$**
- Cisco Nexus 5000 1G storm control:
 - **show plat fwm info lif ethernet 1/1**
 - **show plat fwm info pif ethernet 1/1**
 - **debug hardware internal gatos asic 0 dump-mem 0x3b9000 20**
- Cisco Nexus 5000 Series 10 Gigabit Ethernet storm control:
 - **show plat fwm info lif ethernet 1/5**
 - **show plat fwm info pif ethernet 1/5**
 - **debug hardware internal gatos asic 1 dump-mem 0x3b9000 20**
- Cisco Nexus 5000 Series storm control counter:
 - **show hardware internal gatos asic 1 counters rx_db 2 | grep storm**
- afm-related CLI commands and tools:
 - **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 ternary content addressable memory (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:
 - **show platform software qosctrl port 0 0 nif {0-48} [sat|switch]:** Displays the platform independent (PI) information for every port (useful if port level configuration exists)
 - **show platform software qosctrl port 0 0 hif {0-48} [sat|switch]:** Displays the PI information for every port (useful if port level configuration exists)
 - **show platform software qosctrl policy hif:** Displays the global Network-QoS and Queuing configurations
 - **show platform software qosctrl global:** Displays global PI level configurations
 - **show platform software qosctrl pss:** Stores persistent storage service (PSS) information
 - **show platform software qosctrl asic mod asic:** Displays per-ASIC level port details
 - **show platform software qosctrl default port mod asic:** Displays default port settings on fabric extender ports
 - **show platform software qosctrl port mod asic port-type port:** Displays per-port level PI and platform dependent (PD) data structures
- N2K-C2148T FEX counters:
 - Use these commands (in the fabric extender shell) in preparation to display the statistics of MAC-level traffic and pause statistics:
 - **show platform software fex info satport fex-interface-id** (for mapping except in the case of network interface [NIF] in RW6)
 - **show platform software redwood sts**
 - **show platform software redwood ss**
 - **show platform software qosctrl port 0 6 hif 1 counters:** Displays counters
 - **show platform software redwood rmon 6 nif0:** Displays statistics of MAC-level traffic and pause statistics of NIF for eth103/1/37
 - **show platform software redwood rmon 6 hif5:** Displays statistics of MAC-level traffic and pause statistics of HIF for eth103/1/37
 - **show platform software redwood rmon 4 nif1:** Displays statistics of MAC-level traffic and pause statistics of NIF for eth103/1/37
 - **show platform software redwood rmon 4 hif5:** Displays statistics of MAC-level traffic and pause statistics of HIF for eth103/1/37
 - **show platform software redwood ss:** Displays mapping of HIF or NIF to SS
 - **show platform software redwood ss 4 3:** Displays statistics of RW4 SS3 - Host Receive from HIF4-7 to NIF0-3
 - **show platform software redwood ss 4 2:** Displays statistics of RW4 SS2 - Host Receive from HIF0-3 to NIF0-3

- **show platform software redwood rate:** Displays overall statistics for non-zero traffic
- **show platform software redwood rmon 6 cif0:** Helps debug traffic going from Common Intermediate Format (CIF) to CPU
- **show platform software qosctrl port 0 6 cif 0 counters:** Helps debug traffic going from CIF to CPU
- Cisco Nexus 5000 Series multicast-optimization
 - **show platform fwm in mco-info**
 - **show platform fwm in vlan 1 all_macgs**
- Cisco Nexus 5000 Series FCoE classification
 - For the FCoE interface, use these commands:
 - **show platform fwm info pif ethernet 1/1 | grep gatos**
 - **debug platform hardware peek lu 7 index 5 pifTable**
 - For the Fibre Channel interface, use these commands, the first of which is used to retrieve the GATOS number and the Fibre Channel number:
 - **show platform fwm info pif fc id**
 - **debug peek lu *gatos* index *fc num*>pifTable**
- Cisco Nexus 5000 Series MTU programming
 - **show hardware internal gatos asic 0 registers match bm_port_CFG.*_max**
- Cisco Nexus 5000 Series interrupt
 - **debug hardware internal gatos asic 0 clear-interrupt**
 - **show hardware internal gatos asic 0 interrupt**
 - **show hardware internal gatos event-history errors**
- Untagged CoS
 - **show platform afm info attachment interface ethernet 3/1**
 - **show system internal ipqos port-node ethernet 3/1**
- Buffer usage and packet drop debugging on N2K-C2232P FEX
 - **show platform software qosctrl asic 0 0**

Troubleshooting CRC Errors

This topic explains how to troubleshoot issues that relate to cyclic redundancy check (CRC) errors on a Cisco Nexus 5000 Series Switch.

Troubleshooting CRC Errors

- Cut-through switching changes how the troubleshooting of problems is performed in the switch.
 - Ethernet CRC is at the end of the frame, so even a CRC error cannot cause a drop on a cut-through port.
 - The frame is already forwarded by the time the CRC value is read.
- The corrupted frame must be forwarded, but is accounted for as an output error.

```
N5k-1# show interface e1/1
...
TX
10157 unicast packets 105 multicast packets 52 broadcast packets
11314 output packets 5317822 bytes
0 jumbo packets
1000 output errors 0 collision 0 deferred 0 late collision
0 lost carrier 0 no carrier 0 babble 0 Txpause
0 interface resets
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-23

There are logical and physical causes for the Cisco Nexus 5000 Series Switch to drop a frame. There are also situations in which 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 action causes subsequent CRC checks to fail later in the path for this frame. A downstream store-and-forward device, or a host, can drop this frame.

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

Troubleshooting CRC Errors (Cont.)

- In addition to receiving errored frames, the Cisco Nexus 5000 Series can generate a bad CRC for several reasons:
 - MTU violation
 - IP length error
 - Ethernet length error when EtherType < 1500 / 0x5dc is interpreted as length
 - Invalid Ethernet preamble
- Received and originated errors count as Tx output errors.
- Only received errors count as Rx CRC errors.
- You are more likely to see CRC errors in a network with a cut-through switch.
- The errors will pass through all cut-through switches and finally drop at the first store-and-forward buffer.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-24

When a CRC error is seen in the FCS on a cut-through port, the receive (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 transmit (Tx) CRC error and propagates it through to the next device in the path.

Finding the Source of CRC Errors

- CRC errors are introduced in three ways:
 - Bad physical connection
 - Copper, fiber, transceiver, physical
 - Stomping caused by intentionally originated errors
 - Received bad CRC stomped from neighboring cut-through switch
- Start by finding any Rx CRC counters.
 - If none, then this switch is responsible for originating errors.
 - Use interrupt counters to find the reason and port.
- Log in to next switch upstream of CRC counters, check for Rx CRC.
 - Use the previous logic to determine whether this switch is originating any errors.
 - Finally, inspect optics and pluggables, fiber and cables and troubleshoot as a Layer 1 issue. Change cable and port to find where the problem follows, and **make sure** there is no Fibre Channel traffic.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-25

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

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

Troubleshooting CPU

This topic explains how to troubleshoot issues that relate to high CPU on a Cisco Nexus 5000 Series Switch.

Troubleshooting CPU

- Hopefully, you have a baseline to compare the current CPU trends with a known nominal state.
- Gather the information by using three commands:
 - **show process cpu sort | exclude 0.0**
 - **show system resources**
 - **show process cpu history**

```
N5K-1# show processes cpu sort | exclude 0.0
```

PID	Runtime(ms)	Invoked	uSecs	1Sec	Process
4230	398	5011881	0	22.0%	snmpd
4204	1467	84869127	0	20.2%	gatosud
4226	433	5601856	0	5.5%	statsclient
4264	1380	391510	3	3.7%	ethpm
4302	254	103	2468	1.8%	netstack

```
N5K-1# show system resources
Load average: 1 minute: 0.63 5 minutes: 1.54 15 minutes: 1.46
Processes : 281 total, 4 running
CPU states : 26.7% user, 26.7% kernel, 46.5% idle
Memory usage: 2073408K total, 1412172K used, 661236K free
```

© 2012 Cisco and/or its affiliates. All rights reserved.


DCUFT v6.0-5.27

Hardware-accelerated switches do not rely on the CPU for frame forwarding and processing. CPU is critical for control plane activities:

- **Link Aggregation Control Protocol (LACP):** Without keeping up with LACP data units (LACPDUs), 802.3ad port channels would go down.
- **Spanning Tree Protocol (STP) and STP Bridge Assurance:** A downstream switch that is missing bridge protocol data units (BPDUs) will go forwarding on a blocked port. If the CPU cannot keep up with sending BPDUs, loops can form. Bridge assurance helps in some ways. Instead of going forward, a bridge assurance-enabled switch will disable the interface.
- **vPC programming:** MAC addresses that are learned on vPC interfaces must be installed on both switches to prevent flooding as well as to deliver frames to their destination.
- **Redundancy:** During a switch outage, the CPU needs to reprogram state information for all processes and configure MAC addresses on interfaces in their respective VLANs.
- **Configuration and management:** An unresponsive switch is not useful as a troubleshooting tool, and you need a reliable interface with the network.

Troubleshooting CPU (Cont.)

- High CPU utilization is not automatically a problem indication.
- Focus on extended high-average CPU periods.



© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-28

You should have a baseline to compare the current CPU trends with a known nominal state. You can gather information about high CPU by using these three commands:

- **show process cpu sort | exclude 0.0**
- **show system resources**
- **show process cpu history**

High CPU utilization is not automatically a problem indication. Try to focus on extended high-average CPU periods.

Troubleshooting Unified Ports

This topic explains how to resolve issues that relate to unified ports on a Cisco Nexus 5500 Platform Switch.

Unified Ports

- Cisco Nexus unified ports allow you to configure a physical port on a Cisco Nexus 5500 Platform switch as a:
 - 1/10-Gigabit Ethernet
 - FCoE
 - 1-, 2-, 4-, 8-Gigabit native Fibre Channel port
- A unified fabric includes the following:
 - **Unified platform:** Uses the same hardware platform and the same software code level and certifies it once for your LAN and SAN environments
 - **Unified device:** Runs LAN and SAN services on the same platform switch; allows you to connect your Ethernet and Fibre Channel cables to the same device
 - **Unified wire:** Converges LAN and SAN networks on a single CNA and connects them to your server

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.30

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), Cisco introduces unified port technology. Cisco Nexus unified ports allow you to configure a physical port on a Cisco Nexus 5500 Platform switch as a 1/10-Gigabit Ethernet, FCoE, or 1-, 2-, 4-, or 8-Gigabit native Fibre Channel port.

Most networks have two types of switches for different types of networks. For example, LAN switches carry Ethernet traffic up to Cisco Catalyst switches, and SAN switches carry Fibre Channel traffic from servers to Cisco MDS Series switches. With unified port technology, you can deploy a unified platform, unified device, and unified wire approach. Unified ports allow you to move from an existing segregated platform approach, in which you choose LAN and SAN port options, to a single, unified fabric that is transparent and consistent with existing practices and management software. A unified fabric includes these components:

- **Unified platform:** Uses the same hardware platform and software code level and certifies it once for your LAN and SAN environments
- **Unified device:** Runs LAN and SAN services on the same platform switch and allows you to connect your Ethernet and Fibre Channel cables to the same device.
- **Unified wire:** Converges LAN and SAN networks on one converged network adapter (CNA) and connects them to your server.

A unified fabric allows you to manage Ethernet and FCoE features independently with existing Cisco tools.

The Cisco Nexus 5548UP switch and the Cisco Nexus 5596UP switch provide built-in unified port technology. In addition, a unified port expansion module and two Layer 3 modules increase the benefits of a deployed unified fabric.

Port Order

- You must configure Ethernet ports and Fibre Channel ports in a specified order:
 - Fibre Channel ports must be configured from the last port of the module.
 - Ethernet ports must be configured from the first port of the module.
- If the order is not followed, these errors are displayed:
 - **ERROR: Ethernet range starts from first port of the module**
 - **ERROR: FC range should end on last port of the module**

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-31

You must configure Ethernet ports and Fibre Channel ports in a specified order:

- Fibre Channel ports must be configured from the last port of the module.
- Ethernet ports must be configured from the first port of the module.

If the order is not followed, these errors are displayed:

ERROR: Ethernet range starts from first port of the module

ERROR: FC range should end on last port of the module

On a Cisco Nexus 5548UP Switch, the 32 ports of the main slot (slot1) are unified ports. The Ethernet ports start from port 1/1 to port 1/32. The Fibre Channel ports start from port 1/32 backwards to port 1/1.

This example shows how to configure 20 ports as Ethernet ports and 12 ports as Fibre Channel ports:

```
switch# config t
switch(config)# slot 1
switch(config-slot)# port 21-32 type fc
switch(config-slot)# copy running-config startup-config
switch(config-slot)# reload
```

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Use the **show license** commands to display all license information configured on the system.
- Use the **show** commands to identify the impact or potential problems that might occur when performing a Cisco IOS ISSU and to monitor installation upgrades.
- Use the **show switch-profile status** command to view messages about configuration synchronization failures.
- QoS uses policies and flow control to classify the network traffic, police and prioritize the traffic flow, and provide congestion avoidance.
- 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.
- You should have a baseline to compare the current CPU trends with a known nominal state.
- You must configure Ethernet ports and Fibre Channel ports in a specified order.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5-32

Troubleshooting Cisco Nexus 2000 Series Fabric Extenders

Overview

This lesson is designed to provide some examples of common issues that relate to Cisco Nexus 2000 Series Fabric Extenders and the methods to resolve those issues.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that are specific to Cisco Nexus 2000 Series of Fabric Extenders. This ability includes being able to meet these objectives:

- Explain how to troubleshoot issues that relate to fabric-extender integration on a Cisco Nexus switch
- Explain how to troubleshoot issues that relate to packet drops on a fabric extender

Troubleshooting Fabric-Extender Configuration and Management

This topic explains how to troubleshoot issues that relate to fabric-extender integration on a Cisco Nexus switch.

Verifying Fabric-Extender Configuration

- Display all the attached fabric-extender units.

```
N5k-1# show fex
```

FEX Number	FEX Description	FEX State	FEX Model	FEX Serial
100	FEX0100	Online	N2K-C2148T-1GE	JAF1326BBRC
101	FEX0101	Online	N2K-C2232P-10GE	JAF1333ADDD
102	FEX0102	Online	N2K-C2232P-10GE	JAS12334ABC

- Display information about a specific fabric extender.

```
N5k-1# show fex 100
FEX: 100 Description: FEX0100 state: Online
FEX version: 5.0(3)N1(1b) [Switch version: 5.0(3)N1(1b)]
Extender Model: N2K-C2148T-1GE, Extender Serial: JAF1326BBRC
Part No: 73-12009-05
pinning-mode: static Max-links: 1
Fabric port for control traffic: Eth1/3
Fabric interface state:
Eth1/3 -Interface Up. State: Active
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.4

The **show fex** [*FEX-number* [**detail**]] command displays information about a specific fabric extender or all attached units.

Verifying Fabric-Extender Configuration (Cont.)

```
N5k-1# show fex 100 detail
FEX: 100 Description: FEX0100 state: Online
  FEX version: 5.0(3)N1(1b) [Switch version: 5.0(3)N1(1b)]
  FEX Interim version: 5.0(3)N1(1b)
  Switch Interim version: 5.0(3)N1(1b)
  Extender Model: N2K-C2148T-1GE, Extender Serial: JAF1326BBRC
  Part No: 73-12009-05
  Card Id: 70, Mac Addr: 00:0d:ec:d3:b5:c2, Num Macs: 64
  Module Sw Gen: 21 [Switch Sw Gen: 21]
  post level: complete
...

Logs:
05/02/2012 13:09:06.946120: Module register received
05/02/2012 13:09:06.947614: Image Version Mismatch
05/02/2012 13:09:06.947960: Registration response sent
05/02/2012 13:09:06.948392: Requesting satellite to download image
05/02/2012 13:14:54.149480: Image preload successful.
05/02/2012 13:14:55.375447: Deleting route to FEX
05/02/2012 13:14:55.384270: Module disconnected
05/02/2012 13:14:55.386372: Module Offline
05/02/2012 13:16:52.847574: Module register received
05/02/2012 13:16:52.849146: Registration response sent
05/02/2012 13:16:53.419079: Module Online Sequence
05/02/2012 13:17:09.507541: Module Online
```

The example in the figure shows how to display the detailed status of a specific fabric extender. You can also display the fabric-extender interfaces that are pinned to a specific switch interface:

```
switch# show interface port-channel 100 fex-intf
```

Fabric	FEX
Interface	Interfaces
Po100	Eth100/1/48 Eth100/1/47 Eth100/1/46 Eth100/1/45 Eth100/1/44 Eth100/1/43 Eth100/1/42 Eth100/1/41 Eth100/1/40 Eth100/1/39 Eth100/1/38 Eth100/1/37 Eth100/1/36 Eth100/1/35 Eth100/1/34 Eth100/1/33 Eth100/1/32 Eth100/1/31 Eth100/1/30 Eth100/1/29 Eth100/1/28 Eth100/1/27 Eth100/1/26 Eth100/1/25 Eth100/1/24 Eth100/1/22 Eth100/1/20 Eth100/1/19 Eth100/1/18 Eth100/1/17 Eth100/1/16 Eth100/1/15 Eth100/1/14 Eth100/1/13 Eth100/1/12 Eth100/1/11 Eth100/1/10 Eth100/1/9 Eth100/1/8 Eth100/1/7 Eth100/1/6 Eth100/1/5 Eth100/1/4 Eth100/1/3 Eth100/1/2 Eth100/1/1

Verifying Fabric-Extender Configuration (Cont.)

- Use the **show interface fex-fabric** command to display all fabric-extender fabric interfaces.

```
N5548-2# show interface fex-fabric
```

Fex	Fabric Port	Fabric Port State	Fex Uplink	Model	FEX Serial
101	Eth1/14	Active	3	N2K-C2248TP-1GE	SSI160301N0
102	Eth1/15	Active	1	N2K-C2248TP-1GE	SSI16030660
102	Eth1/16	Active	2	N2K-C2248TP-1GE	SSI16030660

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5-6

You can use the **show interface fex-fabric** command to display all fabric-extender fabric interfaces. The example in the figure shows that three interfaces on a Cisco Nexus 5548UP Switch are connected to two fabric extenders (101 and 102).

Verifying Fabric-Extender Configuration (Cont.)

```
N5k-1# show system internal fex log fport e1/3
Satmgr debug messages for If 0x1a002000:
[19952]05/02/2012 13:08:32.191646: if [0x1a002000]:Phy cleanup rcvd
[19956]05/02/2012 13:08:32.192257: fport [0x1a002000]:Log -Interface Down
[19957]05/02/2012 13:08:32.192266: fport [0x1a002000]:satmgr_fport_fsm: even:t Port
Down. curr state: Discovered
[19958]05/02/2012 13:08:32.192654: fport [0x1a002000]:Log -State changed to: Created
[19962]05/02/2012 13:08:32.192853: fport [0x1a002000]:satmgr_fport_fsm: new state:
Created
[19967]05/02/2012 13:08:32.193991: fport [0x1a002000]:Log -fport phy cleanup retry end:
sending out resp
[19970]05/02/2012 13:08:32.206315: if [0x1a002000]:Pre Cfg rcvd
[19971]05/02/2012 13:08:32.206606: fport [0x1a002000]:Log -pre config: is not a port-
channel member
[19977]05/02/2012 13:08:33.727893: fport [0x1a002000]:Log -Interface Up
[19978]05/02/2012 13:08:33.727904: fport [0x1a002000]:satmgr_fport_fsm: even:t Port
Down. curr state: Created
[19982]05/02/2012 13:08:33.729944: fport [0x1a002000]:Log -Port Bringup rcvd
[19986]05/02/2012 13:08:33.731201: fport [0x1a002000]:Log -Suspending Fabric port.
reason: Fex not configured
[19987]05/02/2012 13:08:33.731216: fport [0x1a002000]:Log -fport bringup retry end:
sending out resp
[19997]05/02/2012 13:08:34.120031: fport [0x1a002000]:Log -Fcot message sent to Ethpm
[19998]05/02/2012 13:08:34.120092: fport [0x1a002000]:Log -Satellite discovered msg
sent
[19999]05/02/2012 13:08:34.120459: fport [0x1a002000]:Log -State changed to: Discovered
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5-7

The example in the figure shows the output of the **show system internal fex log fport interface** command.

Verifying Chassis-Management Information

```
Switch# show inventory fex 100
NAME: "FEX 100 CHASSIS", DESCR: "N2K-C2248TP-1GE CHASSIS"
PID: N2K-C2248TP-1GE , VID: V00 , SN: SSI13380FSM
NAME: "FEX 100 Module 1", DESCR: "Fabric Extender Module: 48x1GE, 4x10GE Supervisor"
PID: N2K-C2248TP-1GE , VID: V00 , SN: JAF1339BDSK
NAME: "FEX 100 Fan 1", DESCR: "Fabric Extender Fan module"
PID: N2K-C2248-FAN , VID: N/A , SN: N/A
NAME: "FEX 100 Power Supply 2", DESCR: "Fabric Extender AC power supply"
PID: NXX-PAC-400W , VID: 000, SN: LIT13370QD6
```

```
Switch# show diagnostic result fex 100
FEX-100: 48x1GE/Supervisor SerialNo      : JAF1339BDSK
Overall Diagnostic Result for FEX-100      : OK
Test results: (. = Pass, F = Fail, U = Untested)
TestPlatform:
0) SPROM: -----> .
1) Inband interface: -----> .
2) Fan: -----> .
3) Power Supply: -----> .
4) Temperature Sensor: -----> .
TestForwardingPorts:
Eth  1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
Port -----
Eth  25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
Port -----
TestFabricPorts:
Fabric 1 2 3 4
Port -----
. . . .
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.8

The **show inventory fex** *FEX-number* command displays inventory information for a fabric extender.

The **show diagnostic result fex** *FEX-number* displays results from the diagnostic test for a fabric extender. In the example in the figure, you can see that all tests were passed successfully.

Verifying Chassis-Management Information (Cont.)

```
Switch# show environment fex 100
Temperature Fex 100:
-----
Module      Sensor      MajorThresh      MinorThres      CurTemp      Status
              (Celsius)      (Celsius)
-----
1      Outlet-1      60      50      33      ok
1      Outlet-2      60      50      38      ok
1      Inlet-1      50      40      35      ok
1      Die-1      100      90      44      ok
Fan Fex: 100:
-----
Fan      Model      Hw      Status
-----
Chassis  N2K-C2148-FAN      --      failure
PS-1      --      --      absent
PS-2      NXX-PAC-400W      --      ok
Power Supply Fex 100:
-----
Voltage: 12 Volts
-----
PS Model      Power      Power      Status
              (Watts)      (Amp)
-----
1  --      --      --
2  NXX-PAC-400W      4.32      0.36      ok
Mod      Model      Power      Power      Power      Power      Status
              Requested      Requested      Allocated      Allocated
              (Watts)      (Amp)      (Watts)      (Amp)
-----
1  N2K-C2248TP-1GE      0.00      0.00      0.00      0.00      powered-up
<->
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.8

To see the environmental sensor status, use the **show environment fex {all | FEX-number} [temperature | power | fan]** command.

Troubleshooting Fabric-Extender Queuing and Packet Drops

This topic explains how to troubleshoot issues that relate to packet drops on a fabric extender.

Troubleshooting Fabric Extender

- Use the **attach fex** command to attach to a fabric extender.

```
N5k-1# attach fex 100
Attaching to FEX 100 ...
To exit type 'exit', to abort type '$.'
fex-100#
```

- Use the **show** commands related to fabric extender.

```
fex-100# show ?
cdp          Show Cisco Discovery Protocol information
cli           Show CLI information
clock         Display current Date
debug         Show debug flags
hardware      Show hardware information
lACP          Show LACP information
lldp          Show information about lldp
logging       Show OBFL log information
platform      Shows list of events received by Platform Manager
process       Show running process information
processes     Show processes
running-config Show running system information
satsyslog     Show information about satsyslog
startup-config Show system startup configuration information
system        System-related show commands
tech-support  Gather information for troubleshooting
terminal      Display terminal configuration parameters
version       Show the software version
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-6-11

Network interface drops can be seen from the **show queuing interface** command output on Cisco Nexus 5000 Series as of version 5.0(3)N1(1).

To get detailed logs, use the **attach fex** command to attach to the fabric extender. Commands are similar to Cisco Nexus 6500 Series or 7000 Series switch linecards:

```
N5k-1# attach fex 100
Attaching to FEX 100 ...
To exit type 'exit', to abort type '$.'
fex-100#
```

You can use several **show** commands that are related to the fabric extender. A fabric extender also has crash logs, its own CPU, and is responsible for communicating link state and offloading some protocols, such as Cisco Discovery Protocol.

Troubleshooting Fabric-Extender Queuing and Packet Drops

- For Cisco Nexus 2248 Fabric Extenders, use the **dbgexec prt** command.

```
fex130# dbgexec prt
prt> drops
PRT_SS_CNT_TAIL_DROP8 : 2 SS0

prt> show rmon 0 ni<0-3>
+-----+-----+-----+-----+-----+-----+
| TX          | Current | Diff | RX          | Current | Diff |
+-----+-----+-----+-----+-----+-----+
| TX_PKT_LT64 | 0        | 0     | RX_PKT_LT64 | 0        | 0     |
| TX_PKT_64   | 5        | 1     | RX_PKT_64   | 8        | 0     |
| TX_PKT_65   | 2062219 | 264039 | RX_PKT_65   | 4073560 | 521532 |
| TX_PKT_128  | 2149866 | 274780 | RX_PKT_128  | 2060397 | 263419 |
| ...         |         |       |             |         |       |
```

- The rmon counters are similar to the **counters detailed** command on the N5k ports.
- Counters are helpful for error tracking and finding packets of a certain size.

If you know the pattern of the flow of traffic, finding where it is likely to stress the network will be easier.

Packet flow from 10 Gigabit Ethernet links to 1 Gigabit Ethernet links is especially difficult to buffer. You might find that the fabric extender is forced to drop traffic.

The **fex queue-limit** and **buffer-threshold** commands can be adjusted globally, per fabric-extender type, or per fabric extender. You can also use these commands:

- **show ctx:** Driver information
- **show oper:** Link states for Layer 1 status
- **show elog:** Event log chronicling hardware and software interaction (helpful for Layer 1 issues)
- **show ints:** Interrupt counters
- **show bootlog:** Bootup messages
- **show log:** Any other logs

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Use the **show fex** [*FEX-number* [**detail**]] to display information about a specific fabric extender or all attached units.
- Use the **attach fex** command to attach to a fabric extender.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-13

Troubleshooting Cisco MDS Series Switches

Overview

This lesson is designed to provide some examples of common issues that relate to Cisco MDS Series switches and methods to resolve those issues.

Objectives

Upon completing this lesson, you will be able to identify and resolve issues that are specific to Cisco MDS switches. This ability includes being able to meet these objectives:

- Explain how to troubleshoot issues that relate to licensing on a Cisco MDS Series switch
- Explain how to troubleshoot issues that relate to software installation and upgrade on a Cisco MDS Series switch
- Explain how to troubleshoot issues that relate to ports on a Cisco MDS Series switch
- Explain how to troubleshoot issues that relate to Cisco Fabric Services on a Cisco MDS Series switch
- Explain how to troubleshoot issues that relate to VSANs on a Cisco MDS Series switch
- Explain how to troubleshoot issues that relate to zones and zone sets on a Cisco MDS Series switch

Troubleshooting Licensing

This topic explains how to troubleshoot issues that relate to licensing on a Cisco MDS Series switch.

About Troubleshooting Licensing

- Process is similar to troubleshooting Cisco Nexus 7000 Series Switches.
- Use the same licensing guidelines and initial troubleshooting checklist.
- Use the **show license usage** command to verify the license installation.

switch# show license usage	Feature	Installed	Lic Count	Status	Expiry Date	Comments
FM_SERVER_PKG	No	-	Unused	never	Grace 79D 16H	
MAINFRAME_PKG	No	-	Unused		Grace expired	
ENTERPRISE_PKG	Yes	-	InUse	never	-	
DMM_FOR_SSM_PKG	No	0	Unused			
SAN_EXTN_OVER_IP	No	0	Unused			
PORT_ACTIVATION_PKG	No	0	Unused			
SME_FOR_IPS_184_PKG	No	0	Unused		Grace 86D 5H	
SAN_EXTN_OVER_IP_18_4	No	0	Unused			
SAN_EXTN_OVER_IP_IPS2	Yes	1	Unused	never	1 license(s) missing	
SAN_EXTN_OVER_IP_IPS4	No	0	Unused			
10G_PORT_ACTIVATION_PKG	No	0	Unused			
SAN_EXTN_OVER_MPS_184_FIPS	No	0	Unused			
STORAGE_SERVICES_ENABLER_PKG	Yes	1	Unused	never	1 license(s) missing	

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-4


Cisco SAN-OS requires licenses for advanced features. These licenses have two options:

- **Feature-based licensing:** Features are applicable to the entire switch. You need to purchase and install a license for each switch that uses the features in which you are interested. The Enterprise license is an example of a feature-based license.
- **Module-based licensing:** Features require additional hardware modules. You need to purchase and install a license for each module that uses the features in which you are interested. The SAN extension over IP license is an example of a module-based license.

The troubleshooting licensing process is similar to the troubleshooting process for Cisco Nexus 7000 Series Switches. You can use the same licensing guidelines and initial troubleshooting checklist as you use for Cisco Nexus 7000 Series Switches. Use the **show license** commands to display all license information that is configured on this switch.

About Troubleshooting Licensing (Cont.)

- You can also use DCNM to see licensing information.
- Choose **Admin > License** and then click a specific device and choose **Licences** tab.



You can also use Cisco Data Center Network Manager (DCNM) to display information about licensing. Choose **Admin > License**, click a specific switch, and choose the **Licenses** tab to display all license information that is configured on this specific switch.

Licensing Installation Issues

- Symptoms, possible causes and solutions are the same as for Cisco Nexus 7000 Series for these issues:
 - Serial-number issues
 - Grace-period alerts
 - Grace-period warnings after license installation
 - License listed as missing
- Checking in the Fabric Manager Server license from Cisco Device Manager
 - Choose **Admin > Licenses** and choose the **Features** tab.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.6

Symptoms, possible causes, and solutions are the same as those that are used for these issues on Cisco Nexus 7000 Series Switches:

- **Serial number issues:** A common problem with licenses stems from not using the correct chassis serial number when ordering your license. Use the **show license host-id** CLI command to obtain the correct chassis serial number for your switch. When entering the chassis serial number during the license-ordering process, do not use the letter “O” in place of any zeros in the serial number.
- **Grace-period alerts:** The grace period of 120 days stops if you disable a feature that you are evaluating, but if you enable that feature again without a valid license, the grace-period countdown continues where it left off.
- **Grace-period warnings after license installation:** If the license installation does not proceed correctly, or if you are using a feature that exists in a license package that you have not installed, you continue to get grace-period warnings.
- **License that is listed as missing:** After a license is installed and operating properly, it might show up as missing if you modify your system hardware or encounter a bootflash issue.

Incorrect Number of Licenses in Use for Multiple Modules

- Module-based licenses require one installed license per module that uses a licensed feature.
- Installing a SAN extension over IP license while two FCIP instances from different modules are present might cause the system to return this error message:
 - **Installing license failed: Number of License in use is more than the number being installed.**
- The workaround for this scenario includes doing one of the following:
 - Concatenate both licenses into one license file.
 - Manually reduce the usage count by one.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5-7

Module-based licenses require one license that is installed per module that uses a licensed feature. SAN extension over IP is an example of a module-based license. Installing a SAN extension over IP license while two Fibre Channel over IP (FCIP) instances from different modules are present might cause the system to return this error message:

Installing license failed: Number of License in use is more than the number being installed.

This error message is generated because the license grace period is applicable only when no licenses are installed. The installation of one license terminates the grace period and arbitrarily causes the second module to shut down, because this is not allowed by licensing.

The workaround for this scenario includes following one of these steps:

- Concatenate both licenses into one license file.
- Manually reduce the usage count by one.

To concatenate both licenses into one license file, follow these steps:

Step 1 Open both license files by using WordPad.

Step 2 Copy both license files to one file; for example:

```
SERVER this_host ANY
VENDOR cisco
INCREMENT SAN_EXTN_OVER_IP_IPS2 cisco 1.0 permanent 1 \
VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M9500EXT12EK9=</SKU> \
HOSTID=VDH=FOXXXXXXXX \
NOTICE="<LicFileID>2005082204514XXXX</LicFileID><LicLineID>1</LicLineID> \
<PAK>MDS-1X-JAB-0F1A81</PAK>" SIGN=F0652E02XXXX
INCREMENT SAN_EXTN_OVER_IP_IPS2 cisco 1.0 permanent 1 \
VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>M9500EXT12EK9=</SKU> \
HOSTID=VDH=FOXXXXXXXX \
NOTICE="<LicFileID>2005082204572XXXX</LicFileID><LicLineID>1</LicLineID> \
<PAK>MDS-1X-JAB-0F1AD1</PAK>" SIGN=D222AE4AXXXX
```

- Step 3** Save the new concatenated license file.
- Step 4** Upload and install the concatenated license file on the Cisco MDS Series switch.

To reduce the usage count to one, follow these steps:

- Step 1** Manually bring down one of the modules to reduce the usage count by one.
- Step 2** Reinsert the module after installing both licenses.

Troubleshooting Installs, Upgrades, and Reboots

This topic explains how to troubleshoot issues that relate to software installation and upgrade on a Cisco MDS Series switch.

Troubleshooting a Nondisruptive Upgrade and Software Installation

- Similar to the Cisco Nexus 7000 Series Switch troubleshooting process
- If a service cannot allow the upgrade to proceed, then the service aborts the upgrade.
- You are prompted to enter the **show install all failure-reason** command to determine the reason why the upgrade cannot proceed.

```
Do you want to continue with the installation (y/n)? [n] y
Install is in progress, please wait.
Notifying services about the upgrade.
>[#          ] 0% -- FAIL. Return code 0x401E0066 (request timed out).
Please issue "show install all failure-reason" to find the cause of the failure.
Install has failed. Return code 0x401E0066 (request timed out).
Please identify the cause of the failure, and try 'install all' again.
```

```
switch# show install all status
There is an on-going installation...
Enter Ctrl-C to go back to the prompt.
Verifying image bootflash:/b-4.0.0.104
-- SUCCESS
Verifying image bootflash:/i-4.0.0.104
-- SUCCESS
Extracting system version from image bootflash:/i-4.0.0.104.
-- SUCCESS
<...>
```

Prompt failure-reason

When a nondisruptive upgrade begins, the system notifies all services that an upgrade is about to start, and finds out whether the upgrade can proceed. If a service cannot allow the upgrade to proceed at this time (for example, Fabric Shortest Path First [FSPF] timers are not configured to the default value, or a Cisco Fabric Services operation is in progress), then the service aborts the upgrade. If this occurs, you are prompted to enter the **show install all failure-reason** command to determine why the upgrade cannot proceed.

If there are any failures for whatever reason (a save runtime state failure or line-card upgrade failure) when the upgrade is already in progress, then the switch is rebooted disruptively because the changes cannot be rolled back. In such cases, the upgrade has failed; you are not prompted to enter the **show install all failure-reason** command, and entering it will not yield any useful information.

If you need additional assistance to determine why an upgrade is unsuccessful, collect the details from the **show tech-support** command output and the console output from the installation, if available.

Troubleshooting Software Upgrades and Downgrades

- **Symptom:** The software installation reports an incompatibility.
- **Possible cause:** The running image might have an enabled feature that is incompatible with the proposed new image.
- **Solutions:**
 - Review the incompatibility issues displayed by the **install all** command. Correct any problems and retry the installation.
 - Verify which features are enabled on your switch and disable any features that might be incompatible with your new image.

```
Warning: The startup config contains commands not supported by the system image;  
as a result, some resources might become unavailable after an install.  
Do you wish to continue? (y/ n) [y]: n  
  
switch# show incompatibility system bootflash:new-image  
The following configurations on active are incompatible with the system image  
1) Feature Index : 67 , Capability : CAP_FEATURE_SPAN_FC_TUNNEL_CFG  
Description : SPAN - Remote SPAN feature using fc-tunnels  
Capability requirement : STRICT  
2) Feature Index : 119 , Capability : CAP_FEATURE_FC_TUNNEL_CFG  
Description : fc-tunnel is enabled  
Capability requirement : STRICT
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-10

To view the results of a dynamic compatibility check, use the **show incompatibility system bootflash:filename** CLI command. Use the **show incompatibility** CLI command for diagnosis when the **install all** CLI command warns of compatibility issues. During an attempted upgrade, the **install all** CLI command might return the warning that is shown in the figure.

Message 1 indicates that the remote SPAN (RSPAN) feature is in use, but it is not supported by the image that was installed. The incompatibility is strict because continuing the upgrade might cause the switch to move into an inconsistent state; that is, configured features might stop working.

Message 2 indicates that the Fibre Channel tunnel feature is not supported in the new image. The RSPAN feature uses Fibre Channel tunnels.

Troubleshooting Ports

This topic explains how to troubleshoot issues that relate to ports on a Cisco MDS Series switch.

Initial Troubleshooting Checklist

- Check the physical media to ensure that there are no damaged parts.
- Verify that the small form-factor pluggable (SFP) devices in use are authorized by Cisco and are not faulty.
- Verify that you have enabled the port by using the **no shutdown** CLI command.
- Use the **show interface** CLI command to verify the state of the interface.
- Verify that you if you have one host-optimized port configured as an ISL, you have not connected to the other three ports in the port group.
- Verify that no ports on a Generation 2 module are out of service.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5-12

Troubleshooting a port problem involves gathering information about the configuration and connectivity of individual devices and the entire SAN fabric. For port interfaces, begin your troubleshooting activity as shown in the figure.

You must administratively enable a port by using the **no shutdown** command. When the interface is enabled, the administrative state of the port is up. If you administratively disable an interface by using the **shutdown** command, the administrative state of the port is down, and the physical link layer state change is ignored.

For a port to be in an up operational state so that it can transmit or receive traffic, the interface must be administratively up, the interface link layer state must be up, and the interface initialization must be complete.

The interface cannot transmit or receive data when a port operational state is down. The interface is operating in trunking expansion port (TE Port) mode when a port operational state is trunking.

Troubleshooting Port States

```
switch# show interface fc1/3
fc1/3 is trunking
Hardware is Fibre Channel, SFP is short wave laser
Port WWN is 20:03:00:0b:fd:8c:f8:80
Peer port WWN is 20:10:00:0b:fd:2c:8c:00
Admin port mode is auto, trunk mode is on
Port mode is TE
Port vsan is 161
Speed is 2 Gbps
Transmit B2B Credit is 255
Receive B2B Credit is 255
Receive data field Size is 2112
```

```
NPI2# show flogi
INTERFACE VSAN FCID PORT NAME NODE NAME
-----
fc1/5 1 0x9f0100 50:06:04:82:c3:a0:98:5c 50:06:04:82:c3:a0:98:5c
fc1/9 1 0x9f0020 21:00:00:e0:8b:08:dd:22 20:00:00:e0:8b:08:dd:22
fc1/12 1 0x9f0040 50:06:04:82:c3:a0:98:52 50:06:04:82:c3:a0:98:52
fc1/13 1 0x9f0300 21:00:00:e0:8b:08:a2:21 20:00:00:e0:8b:08:a2:21
fc8/6 1 0x9f0101 20:00:00:e0:69:40:8d:63 10:00:00:e0:69:41:a0:12
fc8/14 1 0x9f0003 50:06:04:82:c3:a0:98:4c 50:06:04:82:c3:a0:98:4c
```

```
NPI2# show fcns database
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
-----
0x9f0100 N 50:06:04:82:c3:a0:98:5c (EMC) scsi-fcp:target 250
0x7e0200 N 21:00:00:e0:8b:08:d3:20 (QLogic) scsi-fcp:init
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-13

To display complete information for an interface, use the **show interface** command. In addition to the state of the port, the command displays this information:

- Port world wide name (pWWN)
- Speed
- Trunk virtual SAN (VSAN) status
- Transmit (Tx) and receive (Rx) buffer-to-buffer credits that are configured and remaining
- Maximum receive buffer size
- Number of frames that are sent and received
- Transmission errors, including discards, errors, cyclic redundancy checks (CRCs), and invalid frames

If ports are online, use the **show flogi** command to verify that the Fibre Channel ports for the host and storage have performed a fabric login (FLOGI) and are communicating with their respective switches. If you do not see the ports in the **show flogi** output, use the **debug flogi event interface** command to isolate the FLOGI issue.

If the ports are in the **show flogi** output, use the **show fcns database** command to verify that the assigned Fibre Channel ID (FCID) during FLOGI exists in the name server database.

At this point, the host bus adapter (HBA) and subsystem ports have successfully established link level connectivity and each one can communicate with its locally attached switch in the fabric.

Port in ErrDisabled State

- **Possible causes:**
 - Port is flapping.
 - Switch detected a high number of bad frames (CRC errors), potentially indicating something is wrong with the media.
- **Solution:**
 - Verify the SFP, cable, and connections.

```
mds# show logging logfile
<>
Jun 4 06:54:04 switch %PORT_CHANNEL-5-CREATED: port-channel 17 created
Jun 4 06:54:24 switch %PORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-
channel 17 is down (No operational members)
Jun 4 06:54:40 switch %PORT_CHANNEL-5-PORT_ADDED: fc1/8 added to port-channel 7
Jun 4 06:54:56 switch %PORT-5-IF_DOWN_ADMIN_DOWN: Interface fc1/7 is down
(Administratively down)
Jun 4 06:54:59 switch %PORT_CHANNEL-3-COMPAT_CHECK_FAILURE: speed is not
compatible
Jun 4 06:55:56 switch%PORT_CHANNEL-5-PORT_ADDED: fc1/7 added to port-channel 7
<>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DOC# 7500-5-14

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

Troubleshooting Cisco Fabric Services

This topic explains how to troubleshoot issues that relate to Cisco Fabric Services on a Cisco MDS Series switch.

Verifying Cisco Fabric Services

- Process is similar to the Cisco Nexus 7000 Series troubleshooting process.
- To verify that an application is listed and enabled, issue the **show cfs application** command to all switches.
- Verify the set of switches in which an application is registered with Cisco Fabric Services, using the **show cfs peers name application-name** command for physical scope applications.

```
Switch# show cfs peers name dpvm
Scope : Physical
-----
Switch WWN IP Address
-----
20:00:00:0e:d7:0e:bf:c0 10.76.100.51 [Local]
20:00:00:0e:d7:00:3c:9e 10.76.100.52
Total number of entries = 2
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--5-10

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

```
Switch# show cfs application
```

```
-----
Application      Enabled      Scope
-----
ivr              Yes         Physical
ntp              No          Physical
dpvm             Yes         Physical
fscm             Yes         Physical
role             Yes         Physical
radius           Yes         Physical
fctimer          No          Physical
syslogd          No          Physical
callhome         No          Physical
device-alias     Yes         Physical
port-security    Yes         Logical
```

Total number of entries = 11

A physical scope means that Cisco Fabric Services applies the configuration for that application to the entire switch. A logical scope means that Cisco Fabric Services applies the configuration for that application to a specific VSAN.

Verify the set of switches in which an application is registered with Cisco Fabric Services, by using the **show cfs peers name application-name** for physical scope applications and the **show cfs peers name application-name vsan vsan-id** for logical scope applications.

Recovering from a Merge Failure

```
Switch# show cfs merge status name ntp

Physical Merge Status:Failure [ Mon Jun 04 06:49:52 2012 ]
Failure Reason: Conflicting entries in the compared databases

Local Fabric
-----
Switch WWN IP Address
-----
20:00:00:05:30:00:6b:9e 10.76.100.167 [Merge Master]
20:00:00:0e:d7:00:3c:9e 10.76.100.52

Remote Fabric
-----
Switch WWN IP Address
-----
20:00:00:0d:ec:06:55:c0 10.76.100.205 [Merge Master]
```

- For a more detailed description of the merge failure, issue the **show cfs internal session-history name application name detail** command.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-5.17

During a merge, the merge managers in the merging fabrics exchange their configuration databases with each other. The application on one manager merges the information, decides whether 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, then there might be an active session for that application in some switch. Check the lock status for that application on all the switches by using the **show cfs lock** CLI command. If there are any locks, then the merge will not proceed. Commit the changes or clear the session lock so that the merge can proceed.

To recover from a merge failure by using the CLI, follow these steps:

- Step 1** To identify a switch that shows a merge failure, issue the **show cfs merge status name application-name** command; for example:

```
Switch# show cfs merge status name ntp

Physical Merge Status:Failure [ Mon Jun 04 06:49:52 2012 ]
Failure Reason: Conflicting entries in the compared databases

Local Fabric
-----
Switch WWN IP Address
-----
20:00:00:05:30:00:6b:9e 10.76.100.167 [Merge Master]
20:00:00:0e:d7:00:3c:9e 10.76.100.52

Remote Fabric
-----
Switch WWN IP Address
-----
20:00:00:0d:ec:06:55:c0 10.76.100.205 [Merge Master]
```

Step 2 For a more detailed description of the merge failure, issue the **show cfs internal session-history name *application name* detail** command; for example:

```
switch# show cfs internal session-history name ntp detail
```

```
-----  
Time Stamp                Source WWN                Event  
User Name                Session ID  
-----  
Mon Jun 04 04:30:19 2012   20:00:00:0d:ec:04:99:c0  LOCK_REQUEST  
admin 3848  
Mon Jun 04 04:30:19 2012   20:00:00:0d:ec:04:99:c0  LOCK_ACQUIRED  
admin 3848  
Mon Jun 04 04:30:19 2012   20:00:00:0d:ec:04:99:c0  COMMIT  
admin 3849  
Mon Jun 04 04:30:19 2012   20:00:00:0d:ec:04:99:c0  LOCK_RELEASE_REQUEST  
admin 3848  
Mon Jun 04 04:30:19 2012   20:00:00:0d:ec:04:99:c0  LOCK_RELEASED  
admin 3848  
Mon Jun 04 04:33:07 2012   20:00:00:0d:ec:04:99:c0  LOCK_REQUEST  
admin 3868  
Mon Jun 04 04:33:07 2012   20:00:00:0d:ec:04:99:c0  LOCK_ACQUIRED  
admin 3868  
-----
```

Step 3 Enter configuration mode and issue the ***application-name* commit** command to restore all peers in the fabric to the same configuration database; for example:

```
Switch# config terminal  
Switch(config)# ntp commit  
Switch(config)#
```

Troubleshooting VSANs

This topic explains how to troubleshoot issues that relate to VSANs on a Cisco MDS Series switch.

Initial Troubleshooting Checklist

- Verify the FSPF parameters for switches in the VSAN.
- Verify the domain parameters for switches in the VSAN.
- Verify the physical connectivity for any problem ports or VSANs.
- Verify that you have both devices in the name server.
- Verify that you have both end devices in the same VSAN.
- Verify that you have both end devices in the same zone.
- Verify that the zone is part of the active zone set.

© 2012 Cisco and/or its affiliates. All rights reserved.DOC# 750-5-19

Most VSAN problems can be avoided by following the best practices for VSAN implementation.

When suspending or deleting VSANs, make sure that you suspend and unsuspend one VSAN at a time, and that you wait a minimum of 60 seconds after you issue the **vsan suspend** command before you issue any other configuration command. Failure to do so might result in some Fibre Channel interfaces or member ports in a port channel becoming 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. For VSANs, begin your troubleshooting activity as the list in the figure shows.

VSAN Issues

- Host cannot communicate with storage because they are not in the same VSAN.

```
switch# show vsan membership
vsan 1 interfaces:
fc2/7   fc2/8   fc2/9   fc2/10  fc2/11  fc2/12  fc2/13  fc2/14
fc2/15  fc2/16  fc7/1   fc7/2   fc7/3   fc7/4   fc7/5   fc7/6
fc7/7   fc7/8   fc7/9   fc7/10  fc7/11  fc7/12  fc7/13  fc7/14
fc7/15  fc7/16  fc7/17  fc7/18  fc7/19  fc7/20  fc7/21  fc7/22
fc7/25  fc7/26  fc7/27  fc7/28  fc7/29  fc7/30  fc7/31  fc7/32
vsan 2 interfaces:
fc2/6   fc7/23  fc7/24
vsan 3 interfaces:
fc2/1   fc2/2   fc2/5
vsan 4 interfaces:
fc2/3   fc2/4
```

- Use the **vsan database vsan vsan-id interface** command to move the ports into the same VSAN.
- xE Port connecting to the remote switch is isolated.
 - Use the **show port internal info interface fc slot/port** command to determine the root cause of the VSAN isolation.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-20

Use the **show vsan membership** command to view all the ports that are connected to your host and storage and to verify that appropriate devices are in the same VSAN. Use this command on the switches that connect to your host or storage devices.

If the host and storage are in different VSANs, use the **vsan database vsan vsan-id interface** command to move the interface that is connected to the host and storage devices into the same VSAN.

Use the **show interface** command to verify that the trunks that connect the end switches are configured to transport the appropriate VSANs. If the trunk is not configured for the VSAN, then use the **interface** command and then the **switchport trunk allowed vsan** command in interface mode to add the VSAN to the allowed VSAN list for the interface that connects the host and storage devices.

TE Ports are like expansion ports (E Ports) except that they carry traffic for multiple VSANs. E Ports carry traffic for a single VSAN. Because TE Ports carry traffic for multiple VSANs, Inter-Switch Link (ISL) isolation can affect one or more VSANs. For this reason, on a TE Port you must troubleshoot for ISL isolation on each VSAN.

To resolve VSAN isolation on a TE Port, use the **show interface** command on the TE Port to verify that you have an isolated VSAN.

Use the **show interface fc slot/port trunk vsan vsan-id** command to verify the reason for VSAN isolation. Use the **show port internal info interface fc slot/port** command to determine the root cause of the VSAN isolation.

```
switch# show port internal info interface fc2/14
fc2/14 - if_index: 0x0109C000, phy_port_index: 0x3c
Admin Config - state(up), mode(TE), speed(auto), trunk(on)
beacon(off), snmp trap(on), tem(false)
rx bb_credit(default), rx bb_credit multiplier(default)
rxbufsize(2112), encap(default), user_cfg_flag(0x3)
description()
Hw Capabilities: 0xb
```

```

trunk vsans (up) (7)
.
.
.
trunk vsans (isolated) (1,8)
TE port per vsan information
fc2/29, Vsan 1 - state(down), state reason(Isolation due to domain other side
eport isolated), fcid(0x000000)
    port init flag(0x10000), current state [TE_FSM_ST_ISOLATED_DM_ZS]
fc2/29, Vsan 7 - state(up), state reason(None), fcid(0x690202)
    port init flag(0x38000), current state [TE_FSM_ST_E_PORT_UP]
fc2/29, Vsan 8 - state(down), state reason(Isolation due to vsan not
configured on peer), fcid(0x000000)
    port init flag(0x0), current state [TE_FSM_ST_ISOLATED_VSAN_MISMATCH]

```

The last few lines of the command output provide a description of the reason for VSAN isolation for every isolated VSAN. In this example, VSAN 7 is up, and two VSANs are isolated. VSAN 1 is isolated because of domain ID misconfiguration, and VSAN 8 is isolated because of VSAN misconfiguration.

Troubleshooting DPVM

- Use the **show dpvm** command in EXEC mode to verify that Cisco Fabric Services distribution is enabled for DPVM.
 - Optionally, use the **dpvm distribute** command in config mode to enable Cisco Fabric Services distribution if required.
- Use the **show dpvm status** command in EXEC mode to verify that autolearning is disabled.
 - Optionally, use the **no dpvm auto-learn** command in config mode if you need to disable autolearning before activating the database.
- Use the **show dpvm pending-diff** command in EXEC mode to compare the active and pending databases.
 - Optionally, use the **dpvm commit** command in config mode to commit any pending entries to the config database.
- Use the **dpvm activate** command in config mode to activate the database.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-21

You can dynamically assign VSAN membership to ports, by assigning VSANs based on the device WWN. Dynamic Port VSAN Membership (DPVM) offers flexibility and eliminates the need to reconfigure the VSAN to maintain fabric topology when a host or storage device connection is moved between two switches or between ports on the same switch. DPVM retains the configured VSAN regardless of where a device is connected or moved.

Verify these requirements when using DPVM:

- The interface through which the dynamic device connects to the Cisco MDS Series switch must be configured as a fabric port (F Port). Fabric loop ports (FL Ports) do not support DPVM and no entries will be learned through an FL Port.
- The static port VSAN of the F Port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN that is configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).

Troubleshooting FSPF

```
switch1# show fspf database
FSPF Link State Database for VSAN 2 Domain 1
LSR Type = 1
Advertising domain ID = 1
LSR Age = 81
LSR Incarnation number = 0x80000098
LSR Checksum = 0x2cd3
Number of links = 2
NbrDomainId IfIndex NbrIfIndex Link Type Cost
-----
237 0x00010002 0x00010001 1 1000
238 0x00010003 0x00010002 1 1000

FSPF Link State Database for VSAN 2 Domain 237
LSR Type = 1
Advertising domain ID = 237
LSR Age = 185
LSR Incarnation number = 0x8000000c
LSR Checksum = 0xe0a2
Number of links = 2
NbrDomainId IfIndex NbrIfIndex Link Type Cost
-----
239 0x00010000 0x00010003 1 1000
1 0x00010001 0x00010002 1 1000
<...>
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-6.22

The implementation of VSANs dictates that each configured VSAN support a separate set of Cisco Fabric Services. One such service is the FSPF routing protocol, which can be independently configured per VSAN. Therefore, within each VSAN topology, FSPF can be configured to provide a unique routing configuration and resulting traffic flow. Using the traffic engineering capabilities that VSANs offer allows greater control over traffic within the fabric and higher utilization of the deployed fabric resources.

To troubleshoot FSPF by using the CLI, follow these steps:

1. Use the **show fspf database vsan** command to verify that each path is in the FSPF database.
2. Use the **show fspf vsan vsan-id interface** command to verify that the FSPF parameters are correct for each interface and verify that the interface is in the FSPF active state.
3. Use the **show fspf internal route vsan** command to verify that all Fibre Channel routes are available.

Loss of Two-Way Communication

- FSPF hello interval, retransmit time, or dead interval is misconfigured.

```
switch1# show fspf vsan 1 interface fc1/16
FSPF interface fc1/16 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
xStatistics counters :
  Number of packets received : LSU 0 LSA 0 Hello 2 Error packets 1
  Number of packets transmitted : LSU 0 LSA 0 Hello 4 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

Default Settings

FSPF is not in full state,
indicating a problem.

- There is a region mismatch on the switch.
 - Use the **debug fspf all** command and look for nonexistent region messages.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-9-23

If FSPF is misconfigured, the switches will not reach the two-way state. These events occur when two-way communication is lost:

- The port enters initial state (INIT state), removes its neighbor domain ID from the Recipient Domain ID field, and inserts 0xFFFFFFFF.
- FSPF removes the ISL from the topology database.
- New link state records (LSRs) are flooded to adjacent switches to notify them that the FSPF database has changed.

To resolve a wrong hello interval, retransmit time, or dead interval on an ISL, follow these steps:

- Step 1** Use the **debug fspf all** command and look for a wrong hello interval, retransmit time, or dead interval messages.
- Step 2** Use the **undebg all** command to turn off debugging.
- Step 3** Use the **show fspf vsan vsan-id interface** command to show FSPF information.
- Step 4** Use the **interface** command and then the **fspf hello-interval**, **fspf retransmit-interval** or **fspf dead-interval** command in interface mode to change the intervals.

To identify a region mismatch problem on a switch, follow these steps:

- Step 1** Use the **show fspf vsan** command to display the currently configured region in a VSAN.
- Step 2** Use the **debug fspf all** command and look for nonexistent region messages.

```
switch1# debug fspf all
```

```
Jun 5 00:39:31 fspf: FC2 packet received for non existent region 0 in VSAN 1
Jun 5 00:39:45 fspf: Interface fc1/2 in VSAN 1 : Event INACTIVITY , State
change INIT -> INIT
```

Use the **show fspf** command to check the autonomous region. The region must match on all switches in the VSAN.

Troubleshooting Zones and Zone Sets

This topic explains how to troubleshoot issues that relate to zones and zone sets on a Cisco MDS Series switch.

Initial Troubleshooting Checklist

- Verify that you have an active zone set.
- Verify that you have the correct hosts and storage devices in the same zone.
- Verify that the zone is part of the active zone set.
- Verify that the default zone policy is permit if you are not using zoning.
- Verify that you have only pWWN-based zoning if you have a Cisco MDS 9020 Fabric Switch in your fabric.

© 2012 Cisco and/or its affiliates. All rights reserved.

DOC# 500-5-25

Zoning enables access control between storage devices and user groups. Creating zones increases network security and prevents data loss or corruption.

Zone sets consist of one or more zones in a VSAN. A zone set can be activated or deactivated as a single entity across all switches in the fabric, but only one zone set can be activated at any time in a VSAN.

Zones can be members of more than one zone set. A zone consists of multiple zone members. Members in a zone can access one another; members in different zones cannot access one another.

The criteria that is listed in the figure must be met for zoning to function properly.

Zone and Zone Set Issues

- A host cannot see a storage device for these reasons:
 - The default zone policy does not allow the devices to communicate.

```
switch# show zone status vsan 1
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow session: none
hard-zoning: enabled
Default zone:
qos: low broadcast: disabled only: disabled
Full Zoning Database :
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Name: Database Not Available
```

- Storage devices and host interfaces do not belong to the same zone or the zone is not part of the active zone set.

```
switchA# show zone
zone name NewZoneName vsan 2
pwwn 22:35:00:0c:85:e9:d2:c2
pwwn 10:00:00:00:c9:32:8b:a8
zone name Zone2 vsan 4
pwwn 10:00:00:e0:02:21:df:ef
pwwn 20:00:00:e0:69:a1:b9:fc
zone name zone-cc vsan 5
pwwn 50:06:0e:80:03:50:5c:01
pwwn 20:00:00:e0:69:41:a0:12
pwwn 20:00:00:e0:69:41:98:93
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-28

To verify that the host is not communicating with storage, use the CLI to verify that the host and storage device are in the same VSAN.

Then configure zoning, if necessary, by using the **show zone status vsan-id** command to determine whether the default zone policy is set to **deny**. The default zone policy of **permit** means that all nodes can see all other nodes. **Deny** means that all nodes are isolated when not explicitly placed in a zone.

Use the **show zone member** command for the host and storage device to verify that both are in the same zone. Use the **show zoneset active** command to determine whether the zone and the host and disk appear in the active zone set.

If there is no active zone set, use the **zoneset activate** command to activate the zone set. Finally, verify that the host and storage can now communicate.

Troubleshooting Zone Set Activation

- Two main problems can occur with activating a zone set:
 - No zone set is active.
 - Zone set activation fails.
- Error message: **ZONE-2-ZS_CHANGE_ACTIVATION_FAILED: Activation failed.**
 - Use the **zoneset activate** CLI command to activate the zone set.
- Use the **show zoneset active vsan-id** command to display the active zones.

```
switchA# show zoneset active vsan 2
zoneset name ZoneSet1 vsan 2
zone name NewZoneName vsan 2
 * pwwn 22:35:00:0c:85:e9:d2:c2
 * pwwn 10:00:00:00:c9:32:8b:a8
```

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v6.0-6-27

When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning; this copy is called the active zone set. A zone that is part of an active zone set is called an active zone. Two main problems can occur with activating a zone set:

- No zone set is active.
- Zone set activation fails.

Zone activation can fail if a new switch joins the fabric. When a new switch joins the fabric, it acquires the existing zone sets.

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 the 2048-KB limit. If the size exceeds the limit, you must remove some zones or devices within a zone.

```
switch# show zone analysis active vsan 1
Zoning database analysis vsan 1
Active zoneset: zs1 [*]
Activated at: 08:03:35 UTC Nov 17 2005
Activated by: Local [ GS ]
Default zone policy: Deny
Number of devices zoned in vsan: 0/2 (Unzoned: 2)
Number of zone members resolved: 0/2 (Unresolved: 2)
Num zones: 1
Number of IVR zones: 0
Number of IPS zones: 0
Formattted size: 38 bytes / 2048 Kb
```

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

```
switch(config)# zoneset activate ZoneSet1 vsan 2
```

If you still experience zone set activation failure, use the **show zone internal change event-history vsan vsan-id** command to determine the source of the zone set activation problem.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Use the **show license usage** command to verify the license installation.
- Enter the **show install all failure-reason** command to determine why the upgrade cannot proceed.
- Use the **show port internal info** CLI command to verify the port status is in link-failure.
- To verify that an application is listed and enabled with Cisco Fabric Services, issue the **show cfs application** command to all switches.
- Use the **show vsan membership** command to see all the ports connected to your host and storage, and verify that appropriate devices are in the same VSAN.
- Use the **show zoneset active vsan-id** command to display the active zones.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0--9-28

Module Summary

This topic summarizes the key points that were discussed in this module.

Module Summary

- Maintain a consistent Cisco NX-OS release across all your devices. Review the Cisco NX-OS release notes for the latest features, limitations, and caveats. Enable system message logging.
- Troubleshoot any new configuration changes after implementing the change, and gather information that defines the specific symptoms.
- Verify the physical connectivity between your device and end devices, the Layer 2 connectivity, the end-to-end connectivity, and the routing configuration.
- If your troubleshooting attempts do not resolved the problem, contact Cisco TAC or your technical support representative.

© 2012 Cisco and/or its affiliates. All rights reserved.

DCUFT v5.0-5.1

This module uses a symptom-based troubleshooting approach that allows you to diagnose and resolve your Cisco Nexus Operating System (NX-OS) problems by comparing the symptoms that you observed in your network with the symptoms that are listed in each lesson. By comparing the symptoms, you should be able to diagnose and correct software-configuration issues and inoperable hardware components so that the problems are resolved with minimal disruption to the network. Address those problems with corrective actions such as these:

- Identify key Cisco NX-OS troubleshooting tools.
- Obtain and analyze protocol traces by using Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) or Ethalyzer on the CLI.
- Identify or rule out physical port issues.
- Identify or rule out switch module issues.
- Diagnose and correct Layer 2 issues.
- Diagnose and correct Layer 3 issues.
- Recover from switch upgrade failures.
- Obtain core dumps and other diagnostic data for use by Cisco Technical Assistance Center (TAC) or your customer support representative.

To troubleshoot your network, follow these general steps:

- Maintain a consistent Cisco NX-OS release across all your devices.
- See the Cisco NX-OS release notes for your Cisco NX-OS release for the latest features, limitations, and caveats.
- Enable system message logging.
- Troubleshoot any new configuration changes after implementing the change.
- Gather information that defines the specific symptoms.
- Verify the physical connectivity between your device and end devices.
- Verify the Layer 2 connectivity.
- Verify the end-to-end connectivity and the routing configuration.
- After you have determined that your troubleshooting attempts have not resolved the problem, contact Cisco TAC or your technical support representative.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which command should you use to verify that Cisco Fabric Services is enabled for the application on all devices in the network or Cisco Fabric Services region? (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) **show cfs status**
 - B) **show cfs application**
 - C) **show cfs lock**
 - D) **show radius session status**
- Q2) Which three options are potential causes when you cannot create a VDC? (Choose three.) (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) There are not enough resources available to create the VDC.
 - B) Your user account does not have the vdc-admin role.
 - C) You are attempting to create more than two VDCs.
 - D) The Cisco Fabric Services protocol has placed a lock on the VDCs.
 - E) You are not logged in to the default VDC.
 - F) You forgot to switch to the VDC by using the **switchto** command before creating the VDC.
 - G) Your user account does not have the network-admin role.
- Q3) Which issue is a potential cause when you cannot delete a VDC? (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) The VDC is in use.
 - B) You are trying to delete the default VDC.
 - C) The Cisco Fabric Services protocol has placed a lock on the VDC.
 - D) Licensed features have been enabled in the VDC.
 - E) Interfaces are still allocated to the VDC.
- Q4) Which issue is a potential cause when you cannot allocate an interface to a VDC? (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) You are trying to allocate an interface that is part of a port-group on a N7K-M132XP-12 or N7K-F132XP-15 I/O module to a VDC without also allocating the other interfaces in the port group.
 - B) The Cisco Fabric Services protocol has placed a lock on the VDC.
 - C) The interface is assigned to another nondefault VDC.
 - D) The interface has IP configuration on it.
- Q5) Which issue is a potential cause when a VDC remains in the failed state? (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) A higher priority VDC has claimed resources that were assigned to the VDC.
 - B) No interfaces are allocated to the VDC.
 - C) There were not enough available resources when the VDC was created.
 - D) The high availability policy for the VDC was set to bringdown and a VDC failure has occurred.

- Q6) Which Cisco NX-OS command is used to connect to an I/O module in slot 1 of a Cisco Nexus 7000 Series Switch? (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) **session slot 1**
 - B) **switchto module 1**
 - C) **attach module 1**
 - D) **connect slot 1**
- Q7) The network administrator can configure the severity and log file size for the NVRAM logging on Cisco Nexus 7000 Series Switches. (Source: Troubleshooting Cisco Nexus 7000 Series Switches)
- A) True
 - B) False
- Q8) Which license is required on the Cisco Nexus 5000 Series or Nexus 5500 Platform Switch to support FCoE? (Source: Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches)
- A) Enterprise License
 - B) Storage Protocol Services License
 - C) Advanced Services License
 - D) Storage Protocol Advanced Services License
- Q9) High CPU utilization does not automatically indicate a problem. (Source: Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches)
- A) True
 - B) False
- Q10) What does the **show install all status** command do? (Source: Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches)
- A) displays detailed logs of the past five installation-related commands, from oldest to newest
 - B) displays the system and configuration information that you can provide to Cisco TAC when reporting a problem
 - C) displays the fabric extender status during a Cisco IOS ISSU
 - D) displays a high-level log of the installation
- Q11) Which multifunction adapter integrates with the Cisco Nexus 5000 Series and Nexus 5500 Platform Switches to provide Cisco Unified Fabric convergence? (Source: Troubleshooting Cisco Nexus 5000 Series and Nexus 5500 Platform Switches)
- A) network interface card
 - B) host bus adapter
 - C) consolidated network adapter
 - D) converged network adapter
- Q12) Which command is used to display inventory information for a fabric extender? (Source: Troubleshooting Cisco Nexus 2000 Series Fabric Extenders)
- A) **show interface port-channel 100 fex-intf**
 - B) **show diagnostic result fex FEX-number**
 - C) **show inventory fex FEX-number**
 - D) **show environment fex all**

- Q13) For what should you use the switch# **show incompatibility system bootflash:file-name** command? (Source: Troubleshooting Cisco MDS Series Switches)
- A) upgrading switch software
 - B) upgrading Flash
 - C) upgrading BIOS
 - D) downgrading BIOS
- Q14) Which three CLI commands verify that the storage array can receive Fibre Channel frames from the switch? (Choose three.) (Source: Troubleshooting Cisco MDS Series Switches)
- A) **fctrace**
 - B) **fcroute**
 - C) **fcping**
 - D) **ping**
- Q15) Which command is valid for activating a zone set in VSAN 10? (Source: Troubleshooting Cisco MDS Series Switches)
- A) switch(config)# **zone name Zoneset1 vsan 10 activate**
 - B) switch# **zone name Zoneset1 vsan 10 activate**
 - C) switch(config-zoneset)# **zone activate name Zoneset1 vsan 10**
 - D) switch(config)# **zone activate name Zoneset1 vsan 10**
- Q16) Which command can you use to analyze the active zone set database for VSAN 10? (Source: Troubleshooting Cisco MDS Series Switches)
- A) switch(config)# **zoneset import interface fc 1/3 vsan 10**
 - B) switch(config)# **import zoneset interface fc 1/3 vsan 10**
 - C) switch# **show zoneset analysis active vsan 10**
 - D) switch# **zoneset activate name vsan 10**

Module Self-Check Answer Key

Q1)	B
Q2)	A, E, G
Q3)	B
Q4)	A
Q5)	D
Q6)	C
Q7)	B
Q8)	B
Q9)	A
Q10)	D
Q11)	D
Q12)	C
Q13)	A
Q14)	A, C, D
Q15)	D
Q16)	C