Send documentation comments to n5kdocfeedback@cisco.com

Cisco Nexus 5000 Series NX-OS Operations Guide

For Cisco Nexus 5000 Platform Switches
and Cisco Nexus 5500 Platform Switches

Cisco NX-OS Release 5.1(3)
December 5, 2011

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

Cisco Nexus 5000 Series NX-OS Operations Guide

CONTENTS

Preface

This preface describes the audience, organization, and conventions of the *Cisco Nexus 5000 Series NX-OS Operations Guide*. It also provides information on how to obtain related documentation.

This chapter includes the following topics:

Audience

This publication is for experienced network administrators who configure and maintain Cisco NX-OS on Cisco Nexus 5000 Platform switches and Cisco Nexus 5500 Platform switches.

Document Organization

This document is organized into the following chapters:

Chapter	Description
Overview	
Virtual Port Channel Operations	Describes best practices and operations for vPCs.
Layer 3 and vPC Operations	Describes virtual port channel (vPC) operations when Layer 3 routing features are enabled on the Cisco Nexus 5500 Platform.
Configuration Synchronization Operations	Describes best practices and operations for the configuration synchronization feature in Virtual Port Channels (vPCs) topologies.
Fibre Channel over Ethernet Operations	Describes considerations and operational guidelines on how to deploy and implement an FCoE solution.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element(keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
variable	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use the following conventions::

Convention	Description
screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

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

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

Related Documentation

Documentation for Cisco Nexus 5000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders is available at the following URL:

http://www.cisco.com/en/US/products/ps9670/tsd_products_support_series_home.html

The following are related Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Fabric Extender documents:

Release Notes

Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes

Cisco Nexus 5000 Series Switch Release Notes

Configuration Guides

Cisco Nexus 5000 Series Configuration Limits for Cisco NX-OS Release 5.0(2)N1(1)

Cisco Nexus 5000 Series Configuration Limits for Cisco NX-OS Release 4.2(1)N1(1) and Release 4.2(1)N2(1)

Cisco Nexus 5000 Series NX-OS Fibre Channel over Ethernet Configuration Guide

Cisco Nexus 5000 Series NX-OS Layer 2 Switching Configuration Guide

Cisco Nexus 5000 Series NX-OS Quality of Service Configuration Guide

Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide

Cisco Nexus 5000 Series NX-OS Security Configuration Guide

Cisco Nexus 5000 Series NX-OS System Management Configuration Guide

Cisco Nexus 5000 Series Switch NX-OS Software Configuration Guide

Cisco Nexus 5000 Series Fabric Manager Configuration Guide, Release 3.4(1a)

Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 4.2

Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide

Maintain and Operate Guides

Cisco Nexus 5000 Series NX-OS Operations Guide

Send documentation comments to n5kdocfeedback@cisco.com

Installation and Upgrade Guides

Cisco Nexus 5000 Series and Cisco Nexus 5500 Platform Hardware Installation Guide

Cisco Nexus 2000 Series Hardware Installation Guide

Cisco Nexus 5000 Series NX-OS Software Upgrade and Downgrade Guide, Release 4.2(1)N1(1)

Regulatory Compliance and Safety Information for the Cisco Nexus 5000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders

Licensing Guide

Cisco NX-OS Licensing Guide

Command References

Cisco Nexus 5000 Series Command Reference

Technical References

Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Fabric Extender MIBs Reference

Error and System Messages

Cisco NX-OS System Messages Reference

Troubleshooting Guide

Cisco Nexus 5000 Troubleshooting Guide

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS Version 2.0.

CHAPTER 1

Overview

This document includes information about maintaining and operating specific features on the Cisco Nexus 5000 Series switch.

This document is part of a documentation set that describes how to use Cisco NX-OS software on the Cisco Nexus 5000 Series switches. It contains information about maintaining and operating specific features. Cisco NX-OS Release 5.0(2)N2(1) supports all the software features previously supported on the Cisco Nexus 5000 Series up through Cisco NX-OS Release 5.0(2)N1(1). Cisco NX-OS Release 5.0(2)N2(1) is compatible with In-Service Software Upgrade (ISSU) supported in Cisco NX-OS Release 4.2(1).

In addition, Cisco NX-OS Release 5.0(2)N2(1) includes enhancements for an improved configuration synchronization procedure with support for the Port Channel force option. Virtual port channel (vPCs) changes include enhancements to improve the resiliency and operation of vPCs, such as the graceful consistency check, auto recovery, and per-VLAN consistency checks.

The Cisco software feature documentation in this guide might include information about features that are shared across software releases and platforms and that are not specific to your particular platform or supported in your software release.

For the latest feature information and caveats, see the release notes for your platform and software release.

Note Any IP addresses that are used in this guide are not intended to be actual addresses. Any examples, configuration sample output, and figures included in this guide are shown for illustrative purposes only.

CHAPTER 2

Virtual Port Channel Operations

This chapter describes the best practices and operational procedures for the virtual port channel (vPC) feature on Cisco Nexus 5000 Series switches that run Cisco NX-OS Release 5.0(2)N2(1) and earlier releases.

This chapter includes the following sections:

Information About vPC Operations

A vPC allows links that are physically connected to two different Cisco Nexus 5000 Series switches to appear as a single port channel to a third switch. The third switch can be a Cisco Nexus 2000 Series Fabric Extender or a switch, server, or any other networking device. A vPC can provide Layer 2 multipath capability which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

For a quick overview of vPC configurations, see the *Virtual PortChannel Quick Configuration Guide* at the following URL:
http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html

vPC Consistency Checks

This section includes the following topics:

- Bypassing a vPC Consistency Check When a Peer Link is Lost, page 2-8

Type 1 and Type 2 Consistency Check Parameters

Before a Cisco Nexus 5000 Series switch brings up a vPC, the two Cisco Nexus 5000 Series switches in the same vPC domain exchange configuration information to verify if both switches have compatible configurations for a vPC topology. Depending on the severity of the impact of possible mismatched configurations, some configuration parameters are considered as Type 1 consistency check parameters while others are considered as Type 2.

When a mismatch in Type 1 parameters occur, the following applies:

- If a graceful consistency check is enabled (default), the primary switch keeps the vPC up while the secondary switch brings it down
- If a graceful consistency check is disabled, both peer switches suspend VLANs on the vPC ports.

Note The graceful consistency check is a new feature introduced in Cisco NX-OS Release 5.0(2)N2(1) and is enabled by default. For more details, see the "Graceful Consistency Check" section on page 2-3.

When Type 2 parameters exist, a configuration mismatch generates a warning syslog message. The vPC on the Cisco Nexus 5000 Series switch remains up and running. The global configuration, such as Spanning Tree Protocol (STP), and the interface-level configurations are included in the consistency check.

The **show vpc consistency-parameters global** command lists all global consistency check parameters. Beginning with Cisco NX-OS Release 5.0(2)N1(1), QoS parameters have been downgraded from Type 1 to Type 2.

This example shows how to display all global consistency check parameters:

```
switch# show vpc consistency-parameters global
Legend:
        Type 1 : vPC will be suspended in case of mismatch
Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
QoS                         2     ([], [3], [], [], [],  ([], [3], [], [], [],
                                  [])                    [])
Network QoS (MTU)           2     (1538, 2240, 0, 0, 0,  (1538, 2240, 0, 0, 0,
                                  0)                     0)
Network Qos (Pause)         2     (T, F, F, F, F, F)     (T, F, F, F, F, F)
Input Queuing (Bandwidth)   2     (50, 50, 0, 0, 0, 0)   (50, 50, 0, 0, 0, 0)
Input Queuing (Absolute     2     (F, F, F, F, F, F)     (F, F, F, F, F, F)
Priority)
Output Queuing (Bandwidth)  2     (50, 50, 0, 0, 0, 0)   (50, 50, 0, 0, 0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F)     (F, F, F, F, F, F)
Priority)
STP Mode                    1     MST                    MST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Enabled,       Normal, Enabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
Allowed VLANs               -     1,10,100-101,200-201   1,10,100-101,200-201,2
```

```
                                                        000
Local suspended VLANs       -     -                     -
```

Use the **show vpc consistency-parameters interface port-channel** *number* command to display the interface-level consistency parameters.

This example shows how to display the interface-level consistency parameters:

```
n5k-1# show vpc consistency-parameters interface port-channel 200

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
STP Port Type               1     Default                Default
STP Port Guard              1     None                   None
STP MST Simulate PVST       1     Default                Default
lag-id                      1     [(7f9b,                [(7f9b,
                                  0-23-4-ee-be-64, 80c8, 0-23-4-ee-be-64, 80c8,
                                   0, 0), (8000,          0, 0), (8000,
                                  0-1e-13-15-7-40, 1, 0, 0-1e-13-15-7-40, 1, 0,
                                   0)]                    0)]
mode                        1     active                 active
Speed                       1     10 Gb/s                10 Gb/s
Duplex                      1     full                   full
Port Mode                   1     trunk                  trunk
Native Vlan                 1     1                      1
Shut Lan                    1     No                     No
Allowed VLANs               -     1-999,1001-3967,4048-4 1-3967,4048-4093
                                  093
```

The Cisco Nexus 5000 Series switch conducts vPC consistency checks when it attempts to bring up a vPC or when you make a configuration change.

In the interface consistency parameters shown in the above output, all configurations except the Allowed VLANs are considered as Type 1 consistency check parameters. The Allowed VLAN (under the trunk interface) is considered as a Type 2 consistency check parameter. If the Allowed VLAN ranges are different on both VLANs that means that only common VLANs are active and trunked for the vPC while the remaining VLANs are suspended for this port channel.

Graceful Consistency Check

Beginning with Cisco NX-OS Release 5.0(2)N2(1) and later releases, when a Type 1 mismatch occurs, by default, the primary vPC links are not suspended. Instead, the vPC remains up on the primary switch and the Cisco Nexus 5000 Series switch performs Type 1 configurations without completely disrupting the traffic flow. The secondary switch brings down its vPC until the inconsistency is cleared.

However, in Cisco NX-OS Release 5.0(2)N2(1) and earlier releases, this feature is not enabled for dual-homed FEX ports. When Type-1 mismatches occur in this topology, the VLANs are suspended on both switches. The traffic is disrupted on these ports for the duration of the inconsistency.

To minimize disruption, we recommend that you use the configuration synchronization feature for making configuration changes on these ports.

To enable a graceful consistency check, use the **graceful consistency-check** command. Use the **no** form of this command to disable the feature. The graceful consistency check feature is enabled by default.

This example shows how to enable a graceful consistency check:

```
switch(config)# vpc domain 10
```

```
switch(config-vpc-domain)# [no] graceful consistency-check
```

This example shows that the vPC ports are down on a secondary switch when an STP mode mismatch occurs:

```
switch(config)# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id                     : 10
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: failed
Per-vlan consistency status       : success
Configuration consistency reason: vPC type-1 configuration incompatible - STP
                                    Mode inconsistent
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1-10
vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        down*  failed      Global compat check failed -
30     Po30        down*  failed      Global compat check failed -
```

Global Mismatch

VLANs suspended on Secondary

237955

This example shows that the vPC ports and the VLANs remain up on the primary switch when an STP mode mismatch occurs:

```
switch(config)# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id                     : 10
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: failed
Per-vlan consistency status       : success
Configuration consistency reason: vPC type-1 configuration incompatible - STP
                                  Mode inconsistent
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1-10
vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        up     failed      Global compat check failed 1-10
30     Po30        up     failed      Global compat check failed 1-10
```

Global Mismatch

VLANs Up on Primary

237956

This example shows that the vPC ports are down on a secondary switch when an interface-level Type 1 inconsistency occurs:

```
switch(config-if)# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 10
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        up     success     success                    1

30     Po30        down*  failed      Compatibility check failed -
                                       for port mode
```

VLANs suspended on secondary Interface

237957

This example shows that the vPC ports and the VLANs remain up on the primary switch when an interface-level Type 1 inconsistency occurs:

```
switch(config-if)# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

<snip>..

vPC status

id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        up     success     success                    1
30     Po30        up     failed      Compatibility check failed 1
                                       for port mode
```

VLANs Up on Primary Interface

237958

Configuring Per-VLAN Consistency Checks

Beginning with Cisco NX-OS Release 5.0(2)N2(1), the Cisco Nexus 5000 Series switch performs Type-1 consistency checks on a per-VLAN basis when you enable or disable STP on a VLAN. VLANs that do not pass this consistency check are brought down on the primary and secondary switches while other VLANs are not affected.

When you enter the **no spanning-tree vlan** *number* command on one peer switch, only the specified VLAN is suspended on both peer switches; the other VLANs remain up.

Note Per-VLAN consistency checks are not dependent on whether graceful consistency checks are enabled.

This example shows the active VLANs before suspending a specified VLAN:

```
switch(config-if)# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
<snip>..
--------------------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ -----------------------------------------------------------
1    Po1    up     1-10
vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        up     success     success                    1-10
30     Po30        up     success     success                    1-10
```

All VLANs are up

237959

This example shows that VLAN 5 is suspended but the remaining VLANs are up:

```
switch(config)# no spanning-tree vlan 5

switch(config)# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
<snip>..
--------------------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ -----------------------------------------------------------
1    Po1    up     1-4,6-10
vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
20     Po20        up     success     success                    1-4,6-10
30     Po30        up     success     success                    1-4,6-10
```

VLAN 5 is suspended

237960

Identifying Inconsistent vPC Configurations

The **show vpc** command displays the vPC status and the vPC consistency check result for the global consistency check and the interface-specific consistency check.

This example shows the global vPC consistency check failed because of the mismatched Network QoS configuration:

```
n5k-1# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 100
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: failed
Configuration consistency reason: QoSMgr Network QoS configuration incompatible
vPC role                          : secondary
<snip>…
```

237970

You can use the **show vpc consistency-parameters global** command to identify the configuration difference between two vPC peer switches.

This example shows the global consistency check failed because the STP mode was configured differently on the two vPC switches:

```
switch# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
QoS                         2     ([], [3], [], [], [],  ([], [3], [], [], [],
                                  [])                    [])
Network QoS (MTU)           2     (1538, 2240, 0, 0, 0,  (1538, 2240, 0, 0, 0,
                                  0)                     0)
Network Qos (Pause)         2     (F, T, F, F, F, F)     (1538, 2240, 0, 0, 0,
                                                         0)
Input Queuing (Bandwidth)   2     (50, 50, 0, 0, 0, 0)   (50, 50, 0, 0, 0, 0)
Input Queuing (Absolute     2     (F, F, F, F, F, F)     (50, 50, 0, 0, 0, 0)
Priority)
Output Queuing (Bandwidth)  2     (50, 50, 0, 0, 0, 0)   (50, 50, 0, 0, 0, 0)
Output Queuing (Absolute    2     (F, F, F, F, F, F)     (50, 50, 0, 0, 0, 0)
Priority)
STP Mode                    1     MST                    Rapid-PVST
STP Disabled                1     None                   None
STP MST Region Name         1     ""                     ""
STP MST Region Revision     1     0                      0
STP MST Region Instance to  1
 VLAN Mapping
STP Loopguard               1     Disabled               Disabled
STP Bridge Assurance        1     Enabled                Enabled
STP Port Type, Edge         1     Normal, Disabled,      Normal, Disabled,
BPDUFilter, Edge BPDUGuard        Disabled               Disabled
STP MST Simulate PVST       1     Enabled                Enabled
Allowed VLANs               -     1-10                   1-2
Local suspended VLANs       -     3-10                   -
```

STP mode Mismatch

237961

You can use the **show vpc** command also shows the vPC consistency check result for each vPC and the reason for the consistency check failure.

This example shows how to display the vPC consistency check status:

```
n5k-1# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

<snip>..
vPC status
----------------------------------------------------------------------------
id     Port          Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
104    Po104         up     success     success                    3000
200    Po200         up     success     success                    1,101-110,1
                                                                   000,3000
201    Po201         down*  success     success                    -
1002   Po1002        up     success     success                    102-103
1003   Po1003        up     success     success                    1,101,3000
1004   Po1004        up     success     success                    102-103
103424 Eth102/1/1    up     failed      Compatibility check failed 1000
                                        for port mode
103425 Eth102/1/2    down*  failed      Consistency Check Not      -
                                        Performed
103426 Eth102/1/3    down*  failed      Consistency Check Not      -
                                        Performed
```

Consistency check passed but interface is down

Consistency check failed

Consistency check never conducted since port was down

237962

If the consistency check fails, the consistency check is not performed on vPC member ports that are down.

If the consistency check has succeeded and the port is brought down, the consistency check shows that it was successful.

You can use the **show vpc consistency-parameters interface ethernet** *slot/port* command to identify the configuration difference that leads to a consistency check failure for a specified interface or port channel.

This example shows how to display configuration differences that lead to consistency check failures.

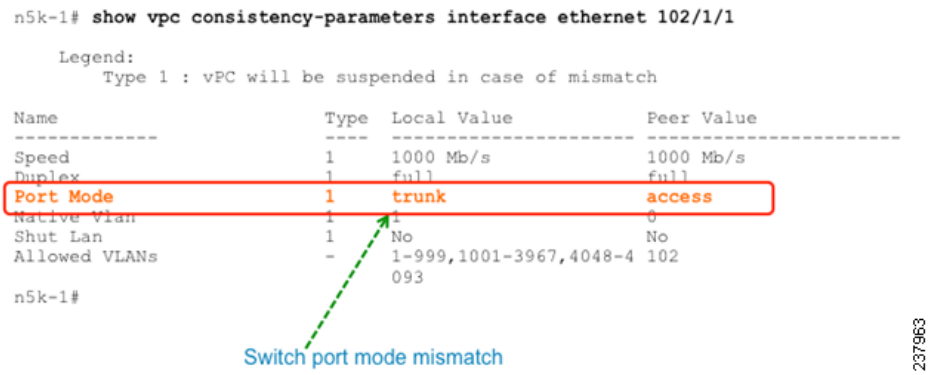

Bypassing a vPC Consistency Check When a Peer Link is Lost

The vPC consistency check message is sent by the vPC peer link. The vPC consistency check cannot be performed when the peer link is lost. When the vPC peer link is lost, the operational secondary switch suspends all of its vPC member ports while the vPC member ports remain on the operational primary switch. If the vPC member ports on the primary switch flaps afterwards (for example, when the switch or server that connects to the vPC primary switch is reloaded), the ports remain down due to the vPC consistency check and you cannot add or bring up more vPCs.

Beginning with Cisco NX-OS Release 5.0(2)N2(1), the auto-recovery feature brings up the vPC links when one peer is down. This feature performs two operations:

- If both switches reload, and only one switch boots up, auto-recovery allows that switch to assume the role of the primary switch. The vPC links come up after a configurable period of time if the vPC peer-link and the peer-keepalive fail to become operational within that time. If the peer-link comes up but the peer-keepalive does not come up, both peer switches keep the vPC links down. This feature is similar to the reload restore feature in Cisco NX-OS Release 5.0(2)N1(1) and earlier releases. The reload delay period can range from 240 to 3600 seconds.
- When you disable vPCs on a secondary vPC switch because of a peer-link failure and then the primary vPC switch fails, the secondary switch reenables the vPCs. In this scenario, the vPC waits for three consecutive keepalive failures before recovering the vPC links.

Note The auto-recovery feature in Cisco NX-OS Release 5.0(2)N2(1) and later releases replaces the reload restore feature in Cisco NX-OS Release 5.0(2)N1(1) and earlier releases.

The auto-recovery feature is disabled by default. To enable auto-recovery, enter the **auto-recovery** command in the vPC domain mode.

This example shows how to enable the auto-recovery feature and to set the reload delay period:

```
switch(config)# vpc domain 10
switch(config-vpc-domain)# auto-recovery ?
  <CR>
  reload-delay  Duration to wait after reload to recover vPCs

switch(config-vpc-domain)# auto-recovery reload-delay ?
  <240-3600>  Time-out for restoring vPC links (in seconds)
switch(config-vpc-domain)# auto-recovery reload-delay 240
Warning:
```

```
 Enables restoring of vPCs in a peer-detached state after reload, will wait for 240
seconds (by default) to determine if peer is un-reachable
```

This example shows how to display the status of the auto-recovery feature:

```
switch(config-vpc-domain)# show running-config vpc
!Command: show running-config vpc
!Time: Tue Dec  7 02:38:44 2010

version 5.0(2)N2(1)
feature vpc
vpc domain 10
  peer-keepalive destination 10.193.51.170
  auto-recovery
```

Configuring Changes in vPC Topologies

One of the challenges with vPC topologies is how to make configuration changes with minimum traffic disruption. Due to the consistency check, the configuration made on one vPC switch could potentially lead to consistency check failure and traffic disruption.

Beginning with Cisco NX-OS Release 5.0(2)N2(1), you can use the following procedure to make configuration changes for Type 1 consistency check parameters on a Cisco Nexus 5000 Series switch. We recommend that you perform the following procedure during a maintenance window because it might reduce the vPC bandwidth by half for a short duration.

Note A graceful consistency-check does not apply to dual-homed FEX ports. As a result, both switches keep the port down for the duration of an inconsistency. Using the configuration synchronization feature reduces the duration of the inconsistency.

To make configuration changes for Type 1 consistency-check parameters, follow these steps:

Step 1 Enable graceful consistency-check in a vPC domain.

```
switch# config term
switch(config)# vpc domain 10
switch(config-vpc-domain)# graceful consistency-check
```

Step 2 Enable the configuration synchronization feature on both vPC peer switches.

For details on using the configuration synchronization feature, see the “Configuration Synchronization Operations” chapter.

Step 3 Perform all configuration changes in the switch profile.

```
switch# config sync
switch(config-sync)# switch-profile abc
switch(config-sync-sp)# interface Port-channel 100
switch(config-sync-sp-if)# switchport mode trunk
switch(config-sync-sp-if)# commit
```

When you commit switch profile configurations on the local switch, the configuration is also sent to the vPC peer switch to reduce misconfigurations when changes are made on only one vPC switch and to reduce the downtime because the configuration is applied rapidly. When there is a short mismatch duration, a graceful consistency-check keeps the primary side forwarding traffic.

Note When you are making a configuration change for a Type 2 consistency check parameter, such as Allowed VLAN for trunk ports, you do not need to follow this procedure.

Replacing a Cisco Nexus 5000 Series Switch or Cisco Nexus 2000 Fabric Extender

This section describes how to replace a Cisco Nexus 5000 Series switch or Cisco Nexus 2000 Series Fabric Extender in a vPC topology with minimal disruption.

This section include the following topics:

Replacing a Cisco Nexus 5000 Series Switch

When you replace a Cisco Nexus 5000 Series switch, you must perform the following procedure on the replacement switch to synchronize the configuration with the existing Cisco Nexus 5000 Series switch. The procedure can be done in a hybrid single/dual-homed Fabric Extender vPC topology.

Note Do not connect a peer-link, vPC, or single/dual homed Fabric Extender topology fabric port to the replacement switch.

Before You Begin

Ensure that you enable pre-provisioning and the configuration synchronization feature on the switch in the vPC topology.

To replace a Cisco Nexus 5000 Series switch in a vPC topology, follow these steps:

Step 1 Boot the replacement switch.

The new switch comes up without a configuration. Ensure the software version is upgraded to match the existing switch.

Step 2 Enable pre-provisioning for all single or dual homed Fabric Extender modules on the replacement switch.

Step 3 Configure the replacement switch as follows:

- If the running configuration was saved offline, go to Step 4 to Step 10 to apply the configuration.
- If the running configuration was not saved offline, you can obtain it from the peer switch if the configuration synchronization feature is enabled. (Create a switch profile and then go to Step 11).
- If neither condition is met, manually add the configuration and then go to Step 11.

Step 4 Edit the configuration file to remove the sync-peer command if using the configuration synchronization feature.

Step 5 Configure the mgmt0 port IP address and download the configuration file.

Step 6 Copy the saved configuration file to the running configuration.

Step 7 Edit the saved configuration file and delete all commands between the **configure sync** command and the **commit** command, including these two commands.

Step 8 Copy the new, edited configuration file to the running configuration again.

Step 9 Verify that the configuration is correct by entering the **show running-config** command and the **show provision failed-config** *slot* command.

Step 10 If switch profile configuration changes were made on the peer switch while the replacement switch was out of service, apply those configurations in the switch profile and then enter the **commit** command.

Step 11 Shut down all single-homed Fabric Extender vPC host ports.

Step 12 Connect the single-homed Fabric Extender topology fabric ports.

Step 13 Wait for single-homed Fabric Extenders to come online.

Step 14 Ensure the vPC role priority of the existing switch is better than the replacement switch.

Step 15 Connect the peer-link ports to the peer switch.

Step 16 Connect the dual-homed Fabric Extender topology fabric ports.

Step 17 Connect the switch vPC ports.

Step 18 Enter the **no shutdown** command on all single-homed Fabric Extender vPC ports.

Step 19 Verify that all vPC switches and the Fabric Extenders on the replacement switch come online and that there is no disruption in traffic.

Step 20 If you are using the configuration synchronization feature, add the sync-peer configuration to the switch profile if this wasn't enabled in Step 3.

Step 21 If you are using the configuration synchronization feature, enter the **show switch-profile** *name* **status** command to ensure both switches are synchronized.

Replacing a Cisco Nexus 2000 Series Fabric Extender

This section describes how to replace a Cisco Nexus 2000 Series Fabric Extender with minimal disruption. This section includes the following topics:

Replacing a Fabric Extender in a Dual-Homed Fabric Extender vPC Topology

Because the hosts behind a Fabric Extender in a dual-homed Fabric Extender vPC topology are by definition singly-connected, traffic disruption will occur for those hosts.

If the replacement Fabric Extender is a different model, the Cisco Nexus 5000 Series switch does not allow you to pre-provision a new type until you disconnect the old Fabric Extender.

To retain the configuration on both Cisco Nexus 5000 Series peer switches in the vPC topology, follow these steps.

Step 1 Save the configuration for the Fabric Extender interfaces to a file.

Step 2 Disconnect the Fabric Extender fabric ports and wait until the Fabric Extender is offline.

Step 3 Pre-provision the slot with the new Fabric Extender model.

Step 4 Modify the configuration file if necessary for the new Fabric Extender if the configurations are incompatible.

Note For vPC ports, this step might affect consistency.

Step 5 Copy the file to the running configuration.

Step 6 Connect the Fabric Extender fabric and host ports and then wait for the Fabric Extender to come online.

Step 7 Verify that all ports are up with the correct configuration.

Replacing a Fabric Extender in a Single-Homed Fabric Extender vPC Topology

If the replacement Fabric Extender is the same model as the original Fabric Extender, then there is no disruption; the configuration on the Fabric Extender interfaces remain unchanged.

If the replacement Fabric Extender is a different model, the Cisco Nexus 5000 Series switch does not allow you to pre-provision a new type until you disconnect the old Fabric Extender.

To replace a Fabric Extender in a single homed Fabric Extender vPC topology, follow the procedure described in "Replacing a Fabric Extender in a Dual-Homed Fabric Extender vPC Topology" section on page 2-12.

Installing a New Cisco Nexus 2000 Series Fabric Extender

With pre-provisioning, you can fully configure the new Fabric Extender before the Fabric Extender is connected to a Cisco Nexus 5000 Series switch.

To install a new Cisco Nexus 2000 Series Fabric Extender, follow these steps:

Step 1 Pre-provision the slot with the Fabric Extender model.

Step 2 Configure the interfaces as though the Fabric Extender is connected.

Step 3 Connect the Fabric Extender and wait for it to come online.

Step 4 Verify that all configurations are applied correctly

Note The switch applies all configurations serially in a best-effort fashion when the Fabric Extender comes online.

vPC Failure Recovery

This section describes different vPC failure scenarios and how to recover from them. This section includes the following topics:

vPC Member Port Failure

Figure 2-1 shows the traffic flow when one vPC member port fails. Once the host MAC_A detects a link failure on one of the port-channel members, it redistributes the affected flows to the remaining port channel members. The return flow from MAC_C to MAC_A could take the path of the left- or the right-side Cisco Nexus 5000 Series switch, depending on the port-channel hash algorithm of the top

switch. For those flows that traverse the right-side Cisco Nexus 5000 Series switch (the red line), the Cisco Nexus 5000 Series switch passes the traffic to the left-side Cisco Nexus 5000 Series switch, because it no longer has the local connection to host MAC_A. This is one of the scenarios where a vPC peer link is used to carry data traffic.

We recommend that you provision enough bandwidth for peer links to accommodate the bandwidth needed for link failure scenarios.

Figure 2-1 vPC Response to a Member Port Failure

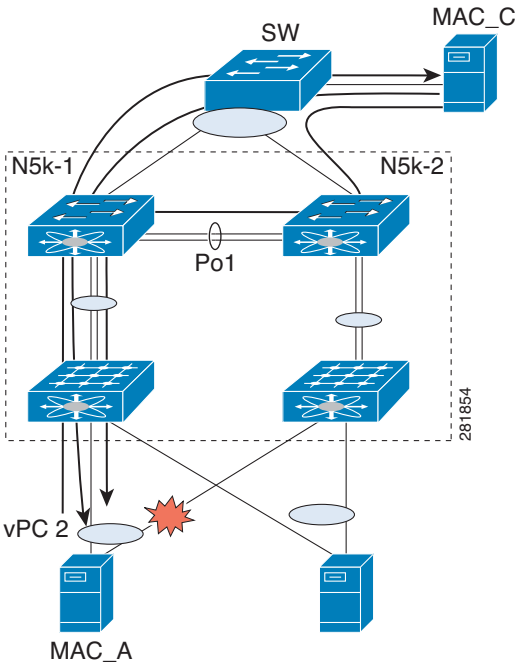

vPC Peer Link Failure

Figure 2-2 shows the vPC response to a peer link failure. In a vPC topology, one vPC peer switch is elected as the vPC primary switch and the other switch is elected as the vPC secondary switch, based on the configured role priority for the switch. In the unlikely scenario where the vPC peer link goes down, the vPC secondary switch shuts down all of its vPC member ports if it can still receive keepalive messages from the vPC primary switch (which indicates that the vPC primary switch is still alive). The vPC primary switch keeps all of its interfaces up. As a result, the hosts or switches that are connected to the Cisco Nexus 5000 Series switch or Cisco Nexus 2000 Series Fabric Extender vPC pair redistributes all the flows to the vPC member ports that are connected to the vPC primary switch.

As a best practice, we recommend that you configure a physical port channel that has at least two 10 Gigabit-Ethernet ports as the vPC peer link.

Figure 2-2 vPC Response to a Peer Link Failure

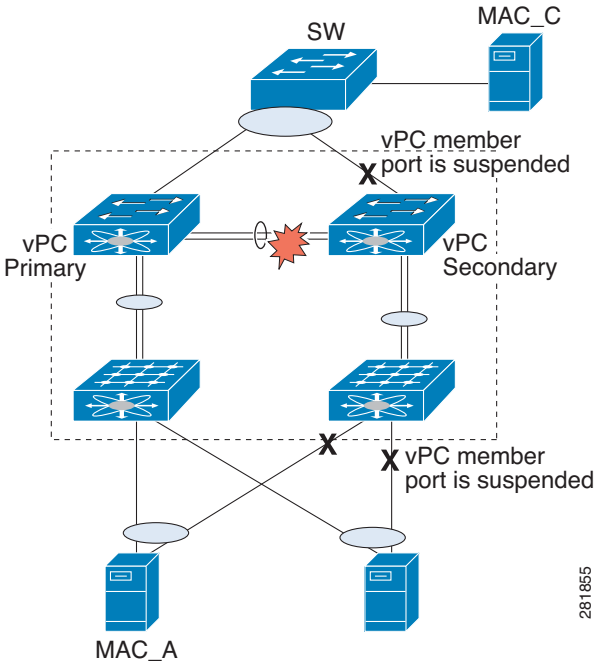

A vPC consistency check cannot be done when a vPC peer-link is down either due to a link failure or when the peer switch is completely down. In either case, any newly configured vPC does not come up because the vPC consistency check cannot proceed, or the existing vPC remains disabled after the link flaps.

Use the reload restore feature that was introduced in Cisco NX-OS Release 5.0(2)N1(1) to fix this problem. The reload restore feature allows a switch to bypass the vPC consistency check and bring up vPC ports when the peer-link or peer switch fails. The reload restore feature has been replaced with the auto-recovery feature in Cisco NX-OS Release 5.0(2)N2(1).

vPC Peer Keepalive Link Failure

The vPC keepalive link carries the heartbeat message between two vPC peer switches. The failure of the vPC keepalive link alone does not impact the vPC operation or data forwarding. Although it has no impact on data forwarding, we recommend that you fix the keepalive as soon as possible to avoid a double failure scenario that could impact the data traffic.

When both switches come up together (such as after power gets restored following a power outage) and only the mgmt/keepalive link fails, the peers are unreachable. However, all other links, including vPC peer links, are up. In this scenario, reaching the vpc-peers through keepalives are achieved through keepalive links while the primary and secondary role election is established through the vpc-peer link. You must establish the first keepalive for the role election to occur in the case when a switch comes up and the vPC-peer link is up.

When keepalives fail to reach the peer switches, role election does not proceed and the primary or secondary role is not established on either vPC peer switch and all vPC interfaces are kept down on both switches.

Note If this scenario occurs again or if the keepalive link goes down after vPC peers are established, the roles do not change and all vPCs remain up.

vPC Peer Switch Failure

When one peer switch fails, half of the network bandwidth is lost and the remaining vPC switch maintains the network connectivity. If the failure occurs on a primary switch, the secondary switch becomes the primary switch.

When one peer switch fails, the remaining peer switch maintains network connectivity for the vPC until it is reloaded. This situation could happen if both vPC peer switches are reloaded and only one switch comes up or both switches loose power and then the power is restored only on one switch. In either case, since the vPC primary election cannot proceed, the Cisco Nexus 5000 Series switch keeps the vPC ports in suspend mode.

To fix these problems, use the reload restore feature and the auto recovery feature as follows:

In NX-OS Release 5.0(2)N1(1), enter the **reload restore** command:

```
switch(config-vpc-domain)# reload restore <timeout in second>
```

In NX-OS Release 5.0(2)N2(1), enter the **auto-recovery reload-delay** command:

```
switch(config-vpc-domain)# auto-recovery reload-delay ?
  <240-3600>  Time-out for restoring vPC links (in seconds)
```

These commands allow the vPC peer switch to bypass the vPC consistency check and bring up vPC ports after the delay timer expires.

vPC Peer Link Failure Followed by a Peer Keepalive Link Failure

If a peer link failure occurs, the vPC secondary switch checks if the primary switch is alive. The secondary switch suspends its vPC member ports after it confirms that the primary switch is up.

In Cisco NX-OS Release 5.0(2)N1(1), if the primary switch fails followed by a peer keepalive failure, the vPC secondary switch remains as the secondary switch and keeps the vPC member ports in the suspend mode.

In Cisco NX-OS Release 5.0(2)N2(1), if you enable the auto-recovery feature and if the vPC primary switch goes down, the vPC secondary switch does not receive messages on the vPC peer keepalive link. Then, after three consecutive keepalive timeouts, the vPC secondary switch changes its role to primary and brings up the vPC member ports.

vPC Keepalive Link Failure Followed by a Peer Link Failure

If the vPC keepalive link fails first and then a peer link fails, the vPC secondary switch assumes the primary switch role and keeps its vPC member ports up.

If the peer link and keepalive link fails, there could be a chance that both vPC switches are healthy and the failure occurs because of a connectivity issue between the switches. In this situation, both vPC switches claim the primary switch role and keep the vPC member ports up. This situation is known as a

split-brain scenario. Because the peer link is no longer available, the two vPC switches cannot synchronize the unicast MAC address and the IGMP group and therefore they cannot maintain the complete unicast and multicast forwarding table. This situation is rare.

We recommend that you have a well-planned network design that includes spreading peer links and keepalive links to multiple ASICs or multiple modules and different cabling routes for keepalive and peer links to avoid a double failure.

Tracing Traffic Flow in a vPC Topology

This section describes how to trace a traffic flow in a vPC topology that is similar to a port-channel environment.

Figure 2-3 shows that each hop in the network chooses one vPC member port to carry the traffic flow independently.

Figure 2-3 Traffic Flow in a vPC Topology

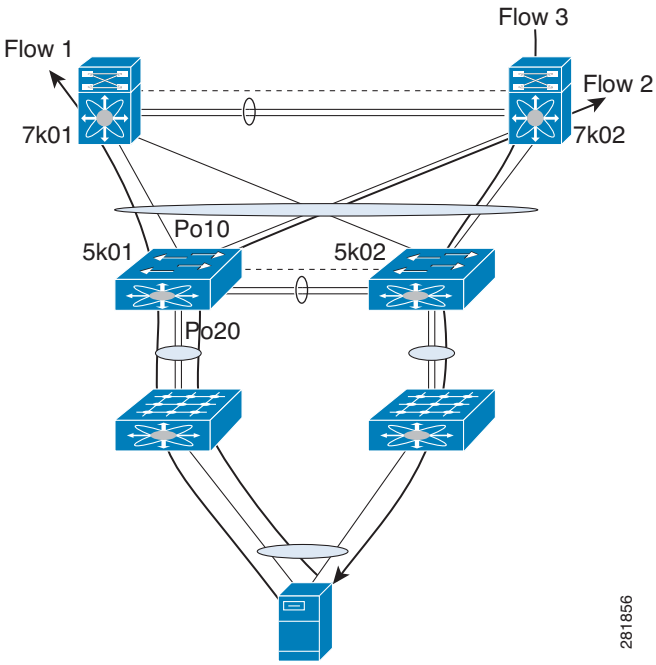

In this example, for flow 1, the host makes a decision whether the traffic flow is sent to the FEX on left or the right side. The FEX runs its hash algorithm to choose one uplink to carry the flow. The N5k determines if the flow should be sent to N7k1 or N7k2. When the egress port for a traffic flow is a vPC, the vPC switch always prefers to use its own vPC member port to carry the traffic in order to minimize the utilization of peer links.

The Cisco NX-OS and Cisco IOS software includes commands to identify the port channel member that carries a particular flow.

This example assumes that the default hash algorithm is used which is src-mac, dst-mac, src-ip and dst-ip. If the hash algorithm also includes the Layer 4 UDP/TCP port, the port information also needs to be provided in the command. The port channel in the command should be the egress port channel.

```
switch# show port-channel load-balance forwarding-path interface Po3 src-interface
ethernet 1/1 vlan 1 src-mac 0000.0000.1111 src-ip 1.1.1.1 dst-mac 001e.1324.4dc0 dst-ip
2.2.2.2
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
crc8_hash: 14   Outgoing port id: Ethernet1/31
Param(s) used to calculate load-balance:
        dst-ip:   2.2.2.2
        src-ip:   1.1.1.1
        dst-mac:  001e.1324.4dc0
        src-mac:  0000.0000.1111
switch#
```

The commands do not show how flows are distributed on the FEX uplink from the FEX to the N5k.

While using the SPAN feature to monitor the traffic flow, the communications between two hosts can be split between two vPC switches. Therefore, you may need to enable SPAN on both vPC switches to obtain a complete trace.

CHAPTER 3

Cisco Nexus 5500 Platform Layer 3 and vPC Operations

This chapter describes virtual port channel (vPC) operations when Layer 3 routing features are enabled on the Cisco Nexus 5500 Platform.

This chapter includes the following sections:

vPC and First Hop Redundancy Protocol

When you use a Cisco Nexus 5548 switch or Cisco Nexus 5596UP switch as a default gateway for hosts, you can deploy the First Hop Redundancy Protocol (FHRP) to provide default gateway redundancy. Beginning with Cisco NX-OS Release 5.0(3)N1(1b), an active FHRP peer and a standby peer can perform Layer 3 forwarding when you enable vPC. This optimization improves bandwidth, avoids sending the Layer 3 traffic over the vPC peer link, and requires no configuration or protocol change. Only the FHRP active peer answers ARP requests. Because both active and standby FHRP peers can forward Layer 3 traffic, you do not need to configure an aggressive timer for FHRP to provide faster failover and convergence time if an active FHRP peer fails.

Figure 3-1 shows that the Layer 3 traffic that originated from the host and is destined to a host several hops away can be routed by both the Host Standby Router Protocol (HSRP) active and the HSRP standby switch.

Figure 3-1 vPC and FHRP

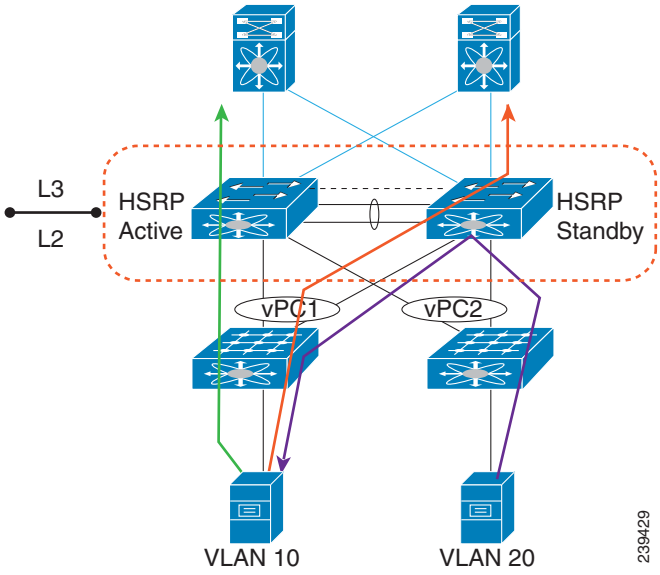

ARP Processing with vPC

When the host connects to a Cisco Nexus 5500 Platform switch and Cisco Nexus 2000 Fabric Extenders in a vPC topology, the host can send an ARP request to the FHRP standby peer due to a hashing algorithm. The ARP request that is received by the standby peer is forwarded to the active peer and the active peer can answer it with an ARP reply.

Similarly, when traffic is moving from north to south, such as when one Cisco Nexus 5500 Platform switch sends an ARP request to a host, the ARP reply might be sent to another switch. In such a case, the ARP reply is forwarded as a Layer 2 frame to the Cisco Nexus 5500 Platform switch that originated the ARP request.

As of Cisco NX-OS Release 5.0(3)N1(1b), ARP synchronization does not occur between two Cisco Nexus 5500 Platform switches. The two switches resolve and maintain their ARP table independently. When one vPC peer switch is reloaded, the switch needs to resolve the ARP by sending ARP requests to the hosts.

Layer 3 Forwarding for Packets to a Peer Switch MAC Address

Typically, a router performs a Layer 3 route table lookup and Layer 3 forwarding when the destination MAC in the Ethernet frame matches its own MAC address. Otherwise, the packets are switched (if Layer 2 functionality is enabled) or dropped. In a topology with Layer 3 and vPC enabled, a vPC peer switch could receive IP packets with the peer's MAC address as the destination MAC rather than the virtual MAC address (when FHRP is enabled) or its own MAC address. In this scenario, a Cisco Nexus 5500 Platform switch can forward the traffic to the peer using a peer link and the peer switch performs the Layer 3 forwarding.

The above scenario often happens with some filers or with Layer 3 peering over vPC. In the case of filers, they may achieves improved load balance and better performance by forwarding traffic to the Burnt-in-Address (BIA) of the routers instead of the HSRP MAC.

Figure 3-2 shows that when the NAS filer sends out packets with N5k-1's MAC RMAC-A as the destination MAC, the packets can be sent over to the N5k-2 switch due to the port channel hashing.

Figure 3-2 vPC and Peer-Gateway

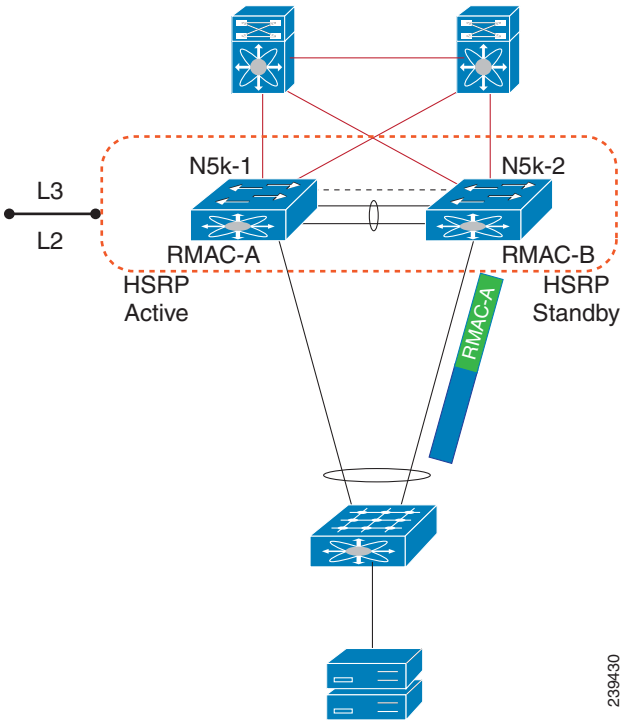

Another scenario that could lead to this situation is when a router is connected to a Cisco Nexus 5500 Platform in a vPC topology.

Figure 3-3 Connecting to a Router in a vPC Topology

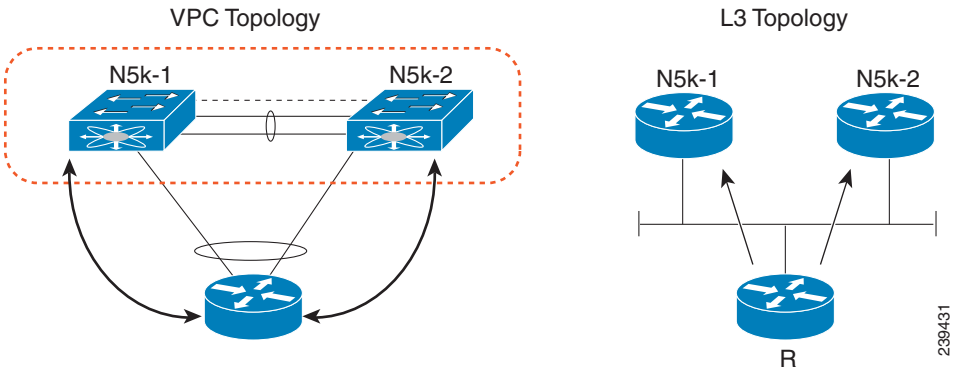

In Figure 3-3, router R considers N5k-1 and N5k-2 as two Layer 3 ECMP next-hop routers and runs ECMP hashing to choose which router to use as the actual next hop for a given flow. Router R connects to N5k-1 and N5k-2 via a vPC. This port channel has an IP address on router R, and Router R performs Layer 3 peering with N5k-1 and N5k-2 over this port channel. It runs the port channel hash algorithm to

choose one physical link to reach the Layer 3 next hop. Because the Layer 3 ECMP and port channel run independent hash calculations there is a possibility that when the Layer 3 ECMP chooses N5k-1 as the Layer 3 next hop for a destination address while the port channel hashing chooses the physical link toward N5k-2. In this scenario,N5k-2 receives packets from R with the N5k-1 MAC as the destination MAC.

Sending traffic over the peer-link to the correct gateway is acceptable for data forwarding, but it is suboptimal because it makes traffic cross the peer link when the traffic could be routed directly.

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), you can use the **peer-gateway** command to allow Cisco Nexus 5500 Platform switches to perform Layer 3 forwarding if the destination MAC of the incoming packet is the MAC of its vPC peer switch. The **peer-gateway** command avoids forwarding such packets to the vPC peer link.

Note You must configure the **peer-gateway** command on both vPC peer switches.

Improved Convergence with a vPC Topology and Layer 3 Routing

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), a delay timer was introduced to avoid the situation where a vPC member port is brought up before the Layer 3 is converged. For example, when one Cisco Nexus 5500 Platform switch is reloaded, the switch starts to receive traffic from hosts once the vPC member ports are up. A delay might occur before the switch establishes a routing protocol adjacency and learns all routes. During this period of the time, received traffic is dropped due to the lack of a route-to-destination address. Figure 3-4 shows an example of where the delay can be used to avoid black hole traffic when a Cisco Nexus 5000 Platform switch is configured for Layer 3 with vPC.

Figure 3-4 vPC Delay Restore

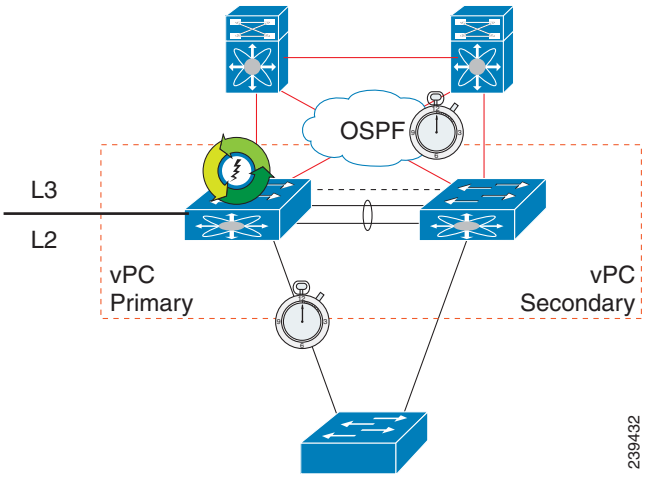

The delay restore feature allows you to configure a timed delay before vPC member ports are brought online. The delay allows the switch to learn all routes, to bring up the vPC member ports, and to forward traffic from hosts. The following example shows how to configure a timed delay of 120 seconds:

```
layer3-switch(config-vpc-domain)# delay restore ?
  <1-3600>  Delay in bringing up the vPC links (in seconds)
layer3-switch(config-vpc-domain)# delay restore 120
layer3-switch(config-vpc-domain)#
```

vPC Peer Link Failure

In addition to suspending vPC member ports, the vPC secondary switch also suspends its switched virtual interface (SVIs) when a vPC peer link is lost. When this occurs, the vPC secondary switch stops advertising the local subnets, which prevents traffic blackholing.

Layer 3 Module Failure

When a Layer 3 module fails on a Cisco Nexus 5500 Platform switch all Layer 3 interfaces are suspended, including Layer 3 port channel and SVI interfaces. As a result, the Layer 3 routing table on the neighboring routers is updated which results in the north to south traffic to be directed towards the peer Nexus 5500 Platform switch. The Layer 2 interfaces, including the Layer 2 port channel and out-of-band management interfaces, remain up.

In a non-vPC topology, when the Layer 3 and SVI interfaces are down, the redundant Cisco Nexus 5500 Platform switch becomes the active peer for all FHRP groups and it continues to forward traffic.

In a vPC topology, although the SVI interfaces are suspended, the vPC member ports are still up on the Cisco Nexus 5500 Platform switch. Even if the switch has a faulty Layer 3 module, Layer 2 traffic forwarding continues.

Figure 3-5 shows a topology where the Layer 3 module on N5k-2 fails. In this scenario, the Layer 3 connection toward the Layer 3 network and all SVI interfaces are suspended. However, the traffic from the hosts can still be sent to N5k-2 depending on the hash results. With the failure of the Layer 3 module, N5k-2 functions as a Layer 2 switch. It forwards the traffic to N5k-1, which forwards the traffic to the Layer 3 network. The return traffic is sent to N5k-1, which sends the traffic directly to the hosts.

Figure 3-5 Layer 3 Module Failure

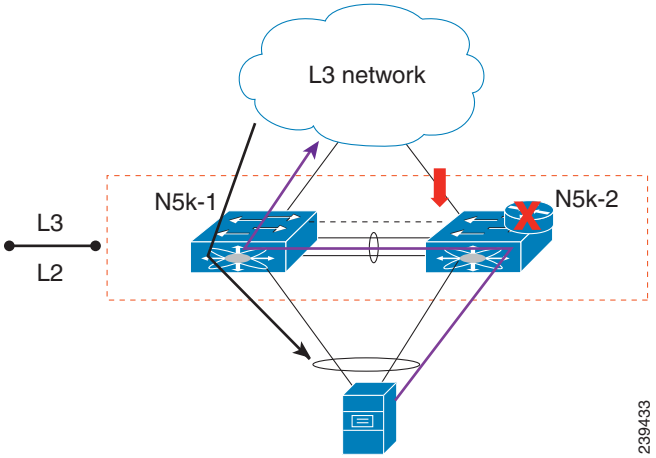

Note Only the Layer 3 traffic needs to cross the peer link. The VLAN traffic is switched by N5k-2 locally.

The peer gateway is disabled on both vPC switches if the Layer 3 module fails on one switch.

For topologies with in-band management, the failure of a Layer 3 module means that the connectivity to the management network and the management system is also lost.

Connecting to a Router in a vPC Topology

When you connect a router to a pair of Cisco Nexus 5500 Platform switches in a vPC topology and enable routing, traffic forwarding may result in suboptimal traffic paths crossing the peer link similar to the situation described in the "Layer 3 Forwarding for Packets to a Peer Switch MAC Address" section on page 3-2. We recommend that you use Layer 3 links for connections between the router and the Nexus 5500 switch, instead of a port channel with an IP address.

Figure 3-6 illustrates the topology that is not recommended. In this topology, control protocol packets may be hashed by the port channel to the wrong Cisco Nexus 5500 Platform switch, which would then forward the control packets to the correct routing peer (1.1.1.1) in the picture.

Figure 3-6 Control Traffic Forwarding in a vPC Topology

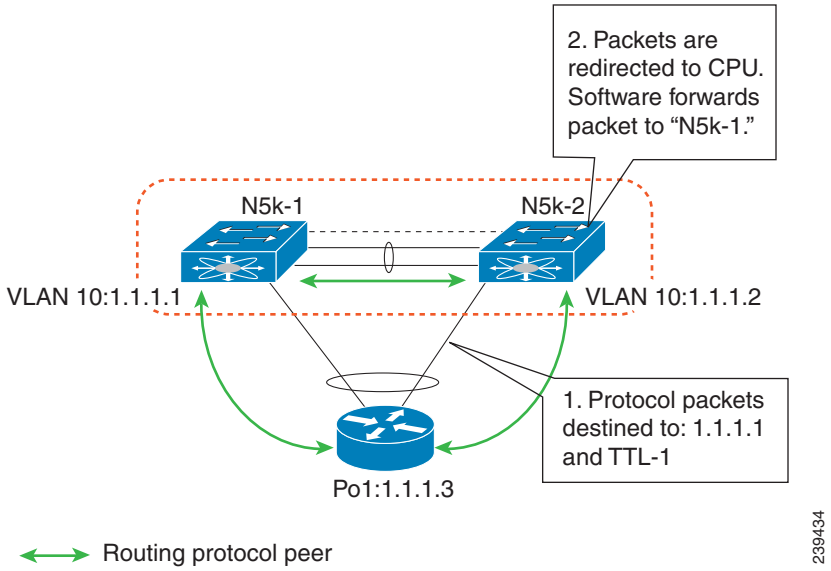

This topology is supported for unicast traffic but not for multicast traffic. In this topology, we recommend that you use Layer 3 interfaces instead of vPC interfaces to connect routers to Cisco Nexus 5500 Platform switches whenever possible.

Figure 3-7, shows the recommended topology for connectivity of routers to a vPC domain. The router connects with Layer 3 interfaces 1.1.1.2 and 2.2.2.2 to the two vPC peers and these interfaces are not part of a vPC port channel.

Figure 3-7 Connecting a Router to a vPC Domain Using Layer 3 Interfaces

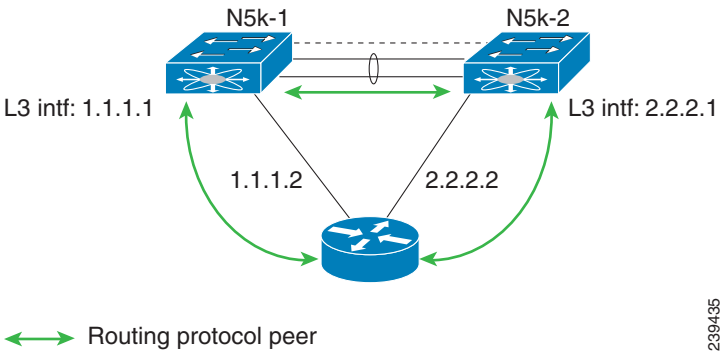

Dedicated VRF For a Keepalive Interface

Beginning in Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform switch supports VRF lite with a Layer 3 module and Enterprise license and you can create a VRF and assign the interface to a VRF. Prior to this release, two VRFs were created by default: the VRF management and VRF default. The management interface(mgmt0) and all SVI interfaces resided in the VRF management and VRF default respectively.

We recommend that you use an out-of-band management interface (mgmt0) as a vPC keepalive interface although you have the option to use the front-panel data port as a vPC keepalive interface. When you choose to use the front panel 10-Gigabit Ethernet port as the vPC keepalive interface, you should create a separate VRF for vPC keepalive packets when Layer 3 is enabled with vPC. This process eliminates the possibility of disrupting the vPC keepalive link by the wrong routes learned by a dynamic routing protocol.

This example shows how to configure a new VRF named vpc_keepalive for the vPC keepalive link and how to display the vPC peer keepalive configuration:

```
vrf context vpc_keepalive
interface Ethernet1/31
  switchport access vlan 123
interface Vlan123
  vrf member vpc_keepalive
  ip address 123.1.1.2/30
  no shutdown
vpc domain 1
  peer-keepalive destination 123.1.1.1 source 123.1.1.2 vrf vpc_keepalive

layer3-switch# show vpc peer-keepalive

vPC keep-alive status           : peer is alive
--Peer is alive for             : (154477) seconds, (908) msec
--Send status                   : Success
--Last send at                  : 2011.01.14 19:02:50 100 ms
--Sent on interface             : Vlan123
--Receive status                : Success
--Last receive at               : 2011.01.14 19:02:50 103 ms
```

```
--Received on interface         : Vlan123
--Last update from peer         : (0) seconds, (524) msec

vPC Keep-alive parameters
--Destination                   : 123.1.1.1
--Keepalive interval            : 1000 msec
--Keepalive timeout             : 5 seconds
--Keepalive hold timeout        : 3 seconds
--Keepalive vrf                 : vpc_keepalive
--Keepalive udp port            : 3200
--Keepalive tos                 : 192
```

The services provided by the Cisco Nexus 5500 Platform switch, such as Ping, SSH, Telnet, and RADIUS, are VRF-aware. You must specify the VRF name in the CLI in order to use the correct routing table.

```
layer3-switch# ping 123.1.1.1 vrf vpc_keepalive
PING 123.1.1.1 (123.1.1.1): 56 data bytes
64 bytes from 123.1.1.1: icmp_seq=0 ttl=254 time=3.234 ms
64 bytes from 123.1.1.1: icmp_seq=1 ttl=254 time=4.931 ms
64 bytes from 123.1.1.1: icmp_seq=2 ttl=254 time=4.965 ms
64 bytes from 123.1.1.1: icmp_seq=3 ttl=254 time=4.971 ms
64 bytes from 123.1.1.1: icmp_seq=4 ttl=254 time=4.915 ms

--- 123.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 3.234/4.603/4.971 ms
```

vPC Consistency Check for Layer 3 Parameters

In a vPC topology, vPC peer switches run routing protocols independently and they maintain the routing table independently. Consistency checks are not performed to verify that Layer 3 configurations in the vPC domain are configured symmetrically.

For example, if you configure a router ACL (RACL) on one SVI and you do not configure the router on the corresponding SVI on the vPC peer, a syslog message is not displayed. You must configure the RACL on both devices. This is consistent with the operation of independent routing devices.

Similarly, if you configure peer gateway on one vPC peer and you want the same peer gateway configuration on the other vPC peer, you must configure the peer gateway on the vPC peer.

To confirm that a vPC domain is correctly configured for Layer 3 operations, the following configurations must be consistent:

- SVI configurations
- RACLs
- Routing protocol configurations

Multicast Interaction in a vPC Topology

This section includes the following topics:

Unsupported Multicast Topology

Figure 3-8 shows an unsupported multicast topology in a vPC configuration.

Figure 3-8 Unsupported Multicast Topology with a vPC

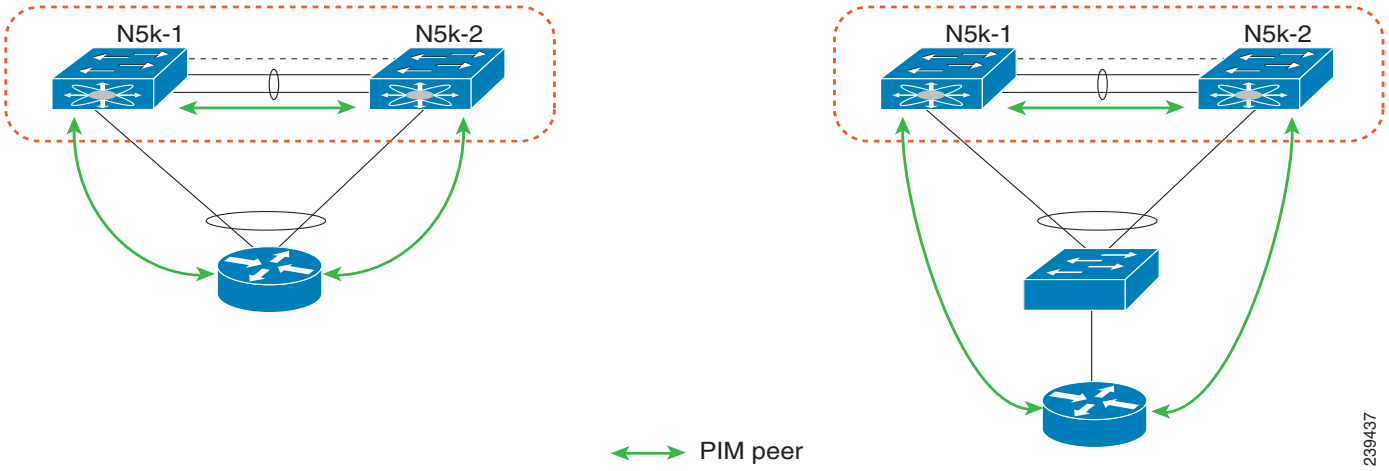

When a PIM router is connected to Cisco Nexus 5500 Platform switches in a vPC topology, the PIM join messages are received only by one switch. The multicast data might be received by the other switch.

Note Multicast forwarding in this topology does not work.

Multicast Routing Table Size

When you enable a vPC on a Nexus 5500 Platform switch, one multicast route (*,G) or (S,G) requires two entries in the routing table; therefore, the multicast routing table size is half the size of what is supported in topologies where vPC is not enabled.

Beginning with Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform multicast routing table size is 2000 entries in non-vPC topologies and 1000 entries in vPC topologies.

Faster Convergence with the Prebuilt Source Tree

In a non-vPC topology, only the designated router (DR) can join the source tree. In a vPC topology, when a receiver is connected to a Cisco Nexus 5500 Platform switch or Fabric Extender (FEX) via vPC, both peer switches initiate a PIM (S,G) join toward the source DR. In a topology where both vPC peer switches have equal costs to the source, the vPC primary switch wins the assert and forwards multicast traffic for receivers connected to the Nexus 5500 Platform switch or FEX using the vPC. The vPC secondary switch also joins the source tree and pulls the multicast data. To prevent data duplication, the vPC secondary switch drops the data due to an empty outgoing interface (OIF) list. Once the vPC secondary switch detects the failure of the vPC primary switch, it adds the receiver VLAN to the OIF list and starts to forward the multicast traffic immediately. Because the vPC secondary switch joins the

source tree before the failure, it does not need to initiate the (S,G) join and waits for the tree to be built. As a result, it improves the convergence time in the case of a failure with the active multicast traffic forwarder.

Figure 3-9 shows one receiver that is connected to a dual-homed FEX. The source and Rendezvous Point (RP) are in the Layer 3 network. N5k-2, which is the VPC primary switch, is the multicast traffic forwarder for receivers in VLAN 10.

Figure 3-9 vPC Switch as the Receiver Designated Router

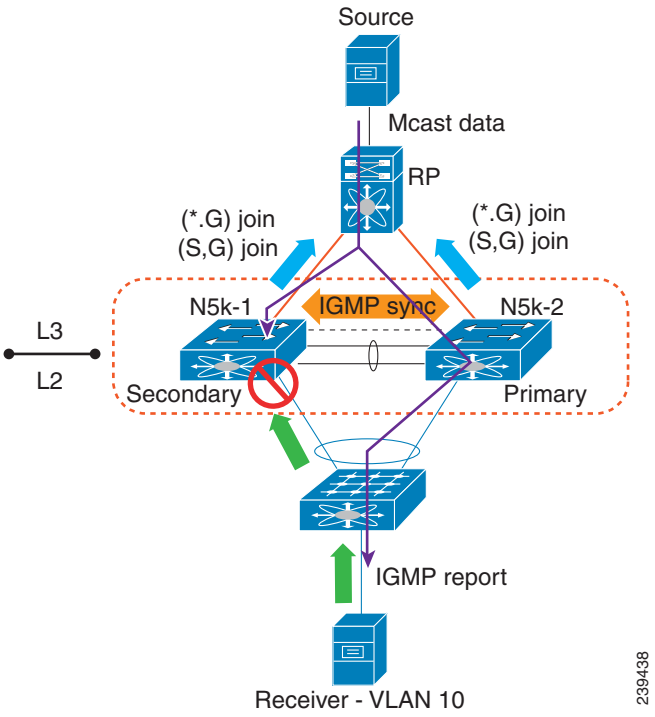

This example shows the output of the multicast routing table and VLAN 10 appears in the OIF list of (S,G) entry on N5k-2. N5k-1 joins the source tree but its OIF list remains empty.

```
N5k-1# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), uptime: 03:03:31, pim ip igmp
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.2
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 03:01:16, igmp

(155.1.3.100/32, 224.1.1.1/32), uptime: 02:13:32, ip pim mrib
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.2
  Outgoing interface list: (count: 0)


N5k-2# show ip mroute 224.1.1.1
IP Multicast Routing Table for VRF "default"

(*, 224.1.1.1/32), uptime: 01:48:07, igmp pim ip
  Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.6
  Outgoing interface list: (count: 1)
    Vlan10, uptime: 01:48:07, igmp

(155.1.3.100/32, 224.1.1.1/32), uptime: 01:00:24, ip pim mrib
```

```
Incoming interface: Ethernet1/6, RPF nbr: 155.1.2.6
Outgoing interface list: (count: 1)
  Vlan10, uptime: 00:55:14, mrib
```

The multicast forwarding algorithm applies to all hosts that are connected to the Cisco Nexus 5500 Platform switch or the FEX in a VPC topology, including hosts directly connected to the switch or hosts connected to straight-through FEX topology.

Using a vPC Switch as a Designated Router (PIM DR)

This section includes the following topics:

DR Election and Source Registration

In vPC topologies, a DR election occurs based on the DR priority and the IP address. The elected DR is responsible for sending the source registration toward the RP. When multicast traffic from a directly connected source is received by the non-DR peer switch, the peer switch notifies the DR switch using a Cisco Fabric Services (CFS) message about the source and group address. The DR generates source registration packets to the rendezvous point (RP).

Multicast Data Forwarding

The Cisco Nexus 5500 Platform switch implements a dual-DR mechanism where both vPC peer switches can forward multicast traffic from directly connected sources. The data forwarding rules are as follows:

- The peer switch receives multicast packets from a directly connected source, performs an mroute lookup, and replicates packets for each interface in the OIF list.
- If the OIF is a VLAN trunked over a vPC peer link, one copy is sent over to the peer link for each VLAN that is present in the OIF list. By default, the vPC peer link is considered an mrouter port. Therefore, the multicast packets are sent over to the peer link for each receiving VLAN. You can use the **no ip igmp snooping mrouter vpc-peer link** command to avoid sending multicast traffic over a peer link for each receiver VLAN when there are no orphan ports.

This example shows how to avoid sending the multicast traffic in this scenario:

```
switch-Layer 3-1(config)# no ip igmp snooping mrouter vpc-peer link
Warning: IGMP Snooping mrouter vpc-peer link should be globally disabled on peer VPC
switch as well.
switch-Layer 3-1(config)#
```

With the above CLI configured, the multicast packet is only sent to peer link for VLANs that have orphan ports.

This example shows how to display the list of all orphan ports:

```
switch-Layer 3-1# show vpc orphan-ports
Note:
--------::Going through port database. Please be patient.::--------

VLAN            Orphan Ports
```

```
-------         -------------------------
1               Eth1/15
switch-Layer 3-1#
```

Note As of Cisco NX-OS Release 5.0(3)N1(1b), the **no ip igmp snooping mrouter vpc-peer link** command cannot be applied with FEX dual-homed topologies due to a software limitation. The command is used only for interfaces on a Cisco Nexus 5500 Platform switch. This software limitation will be removed in a future software release.

One post-routed multicast packet is sent to a vPC peer link using a reserved VLAN. To configure the reserved VLAN, use the follow commands:

```
switch-Layer 3-1(config)# vpc bind-vrf vrf name vlan VLAN ID
switch-Layer 3-1(config)# vpc bind-vrf default vlan 3000
```

One reserved VLAN is required for each VRF. Without these commands, the receivers in non-vPC VLAN and the receivers connected to Layer 3 interfaces may not receive multicast traffic. The non-vPC VLANs are the VLANs that are not trunked over a peer link.

Multicast traffic that is received over a peer link (with a VLAN ID other than the reserved VLAN ID) is not routed. The multicast traffic is treated as Layer 2 frames that are sent to orphan ports only and not to vPC member ports. The multicast traffic that is received over a peer link with a reserved VLAN ID is routed to a non-vPC VLAN (shown as VLAN 13 in Figure 3-10) and receivers behind the Layer 3 interface. The receivers behind the Layer 3 interface can be hosts directly connected to the Cisco Nexus 5500 Platform switch using Layer 3 interfaces or a router joins the source tree.

Figure 3-10 shows the multicast forwarding rules in a vPC dual-DR topology. In this topology, the source in VLAN 10 and receivers in VLAN 11 and VLAN 12 are the vPC hosts (although in this example they are hosts behind a dual-homed FEX topology where the same rule applies to hosts directly to a Cisco Nexus 5500 Platform switch in a vPC topology). VLAN 13 is a non-vPC VLAN and resides only on N5k-2.

Figure 3-10 Multicast Data Forwarding

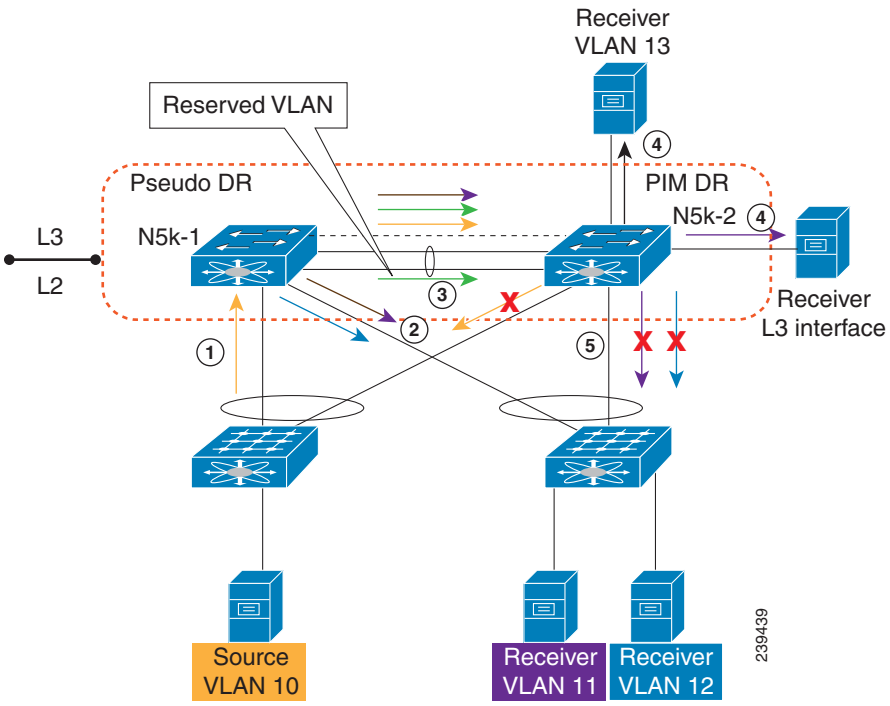

The forwarding process is as follows:

1. IGMP joins from the hosts are synchronized between the two vPC peer switches. N5k-2 is elected as the PIM DR for VLAN 10. Multicast traffic is sent over to N5k-1.
2. The routing engine of N5k-1 performs an mroute lookup and replicates packets to VLAN 11 and VLAN 12. The data packets for VLAN 11 and VLAN 12 are sent to the FEX which in turn sends packets to the two receivers;
3. By default, the replicated packets are sent to the vPC peer link for the source VLAN as well as each receiver VLAN (VLAN 10, VLAN 11, and VLAN 12) in this example. When you use the **no ip igmp snooping mrouter vpc-peer-link** command, the multicast packets are not sent to the peer link for VLAN 10, VLAN 11, and VLAN 12 because there are no orphan ports. One copy of the packets is sent to the peer link with the reserved VLAN 3000 which was configured using the **vpc bind-vrf default vlan 3000** command.

Note In Cisco NX-OS Release 5.0(3)N1(1b), the **no ip igmp snooping mrouter vpc-peer-link** command cannot be applied with a FEX dual-homed topology.

4. For the multicast traffic received from the peer link, if the VLAN ID is the reserved VLAN ID 3000, the N5k-2 route engine performs a Layer 3 lookup and replicates packets to VLAN 13 (a non-vPC VLAN) and receivers behind Layer 3 interfaces.
5. For the multicast packets received over the peer link, VLAN 10, VLAN 11, and VLAN 12 are dropped by N5k-2 to prevent duplicated packets being sent to the vPC hosts. If any orphan ports are in VLAN 10, VLAN 11, and VLAN 12, the packets are bridged to the orphan ports.

Software Upgrade and Downgrade Impact

In Cisco NX-OS Release 5.0(3)N1(1b), the Cisco Nexus 5500 Platform switch does not support ISSUs when Layer 3 modules are installed and Layer 3 features are enabled. Use the **install all** command and the **show install all impact** command to determine the impact of the software upgrade and to indicate whether the software upgrade with Layer 3 features enabled will be disruptive and would require a switch and FEX reload.

show install all impact kickstart

This example shows the output of the **show install all** command:

```
Layer 3-N5548-2# show install all impact kickstart
n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg system n5000-uk9.5.0.3.N1.0.271.bin.upg

Verifying image bootflash:/n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg for boot variable
"kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg for boot variable "system".
[####################] 100% -- SUCCESS

Verifying image type.
[############         ]  50%
[####################] 100% -- SUCCESS

Extracting "system" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:/n5000-uk9-kickstart.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "bios" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Extracting "fexth" version from image bootflash:/n5000-uk9.5.0.3.N1.0.271.bin.upg.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS



Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
------  --------  --------------  ------------  ------
     1       yes      disruptive         reset  Non-disruptive install not supported if
Layer 3 was enabled
   100       yes      disruptive         reset  Non-disruptive install not supported if
Layer 3 was enabled



Images will be upgraded according to following table:
Module       Image         Running-Version             New-Version  Upg-Required
------  ----------  ----------------------  ----------------------  ------------
```

```
     1      system             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1   kickstart             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1        bios      v3.4.0(01/13/2011)     v3.4.0(01/13/2011)           no
   100      fexth             5.0(3)N1(1b)           5.0(3u)N1(1u)           yes
     1   power-seq                    v3.0                   v3.0           no
     2   power-seq                    v1.0                   v1.0           no
     1          uC               v1.0.0.14              v1.0.0.14           no

Layer 3-N5548-2#
```

You can perform a nondisruptive ISSU from an earlier release to NX-OS Release 5.0(3)N1(1b) when upgrading without Layer 3 features enabled.

show spanning-tree issu-impact

To verify that the current STP topology is consistent with ISSU requirements, use the **show spanning-tree issu-impact** command to display the STP configuration and whether or not there are potential STP issues.

This example shows how to display information about the STP impact when performing an ISSU:

```
nexus5010# show spanning-tree issu-impact
For ISSU to Proceed, Check the Following Criteria :
1. No Topology change must be active in any STP instance
2. Bridge assurance(BA) should not be active on any port (except MCT)
3. There should not be any Non Edge Designated Forwarding port (except MCT)
4. ISSU criteria must be met on the VPC Peer Switch as well

Following are the statistics on this switch

No Active Topology change Found!
Criteria 1 PASSED !!

No Ports with BA Enabled Found!
Criteria 2 PASSED!!

No Non-Edge Designated Forwarding Ports Found!
Criteria 3 PASSED !!

ISSU Can Proceed! Check Peer Switch.
```

For information on upgrade procedures, see the *Cisco Nexus 5000 Series NX-OS Upgrade and Downgrade Guide*.

CHAPTER 4

Configuration Synchronization Operations

This chapter provides information about configuration synchronization operations in Virtual Port Channel (vPC) topologies.

This chapter includes the following sections:

Overview

Some Cisco NX-OS software features require consistent configurations across Cisco Nexus 5000 Series switches in the network. For example, vPC topologies require identical configurations on peer switches. As a result, you, as the network administrator, must repeat configurations on both peer switches. This process, which can cause errors due to misconfigurations or omissions, can result in additional service disruptions because of mismatched configurations. Configuration synchronization eliminates these problems by allowing you to configure one switch and automatically synchronize the configuration on the peer switch.

In a vPC topology, an EtherChannel can be formed across two physical switches and vPCs can be connected to any networking device or end host. Because each Cisco Nexus 5000 Series switch forms an EtherChannel bundle to a downstream device, each Cisco Nexus 5000 Series switch must have some matching parameters. You can use a vPC consistency check to verify that both Cisco Nexus 5000 Series switches have the same configuration (Type 1 or Type 2). If they do not match, depending on whether it is a global (for example, spanning-tree port mode), a port-level (for example, speed, duplex, or channel-group type), or even a port-channel interface, the vPC can go into a suspended state or a VLAN can go into a blocking state on both peer switches. As a result, you must ensure that the configuration from one switch is copied identically to the peer switch.

Configuration synchronization allows you to synchronize the configuration between a pair of switches in a network. You use a switch profile to create a configuration file that you can apply locally and you use it to synchronized the configuration to its peer. Configuration synchronization and vPCs are two independent features and configuration synchronization does not eliminate vPC consistency checks. The checks will continue. If there is a configuration mismatch, the vPC can still go into a suspended state. One important benefit of configuration synchronization is that it eliminates the need to manually repeat the same configuration on both switches.

This section includes the following topics:

Benefits of Configuration Synchronization

Configuration synchronization benefits are as follows:

- Provides a mechanism to synchronize configuration from one switch to another switch.
- Merges configurations when connectivity is established between peers.
- Allows you to choose which configuration is synchronized.
- Provides mutual exclusion for commands.
- Provides verify and commit Cisco NX-OS commands.
- Supports existing session and port profile functionality.
- Provides an **import** command to migrate existing vPC configurations to a switch profile.
- Supports Gigabit Expansion Module (GEM) and Fabric Extender (FEX) pre-provisioning.

Requirements

The requirements for configuration synchronization are as follows:

- Cisco NX-OS Release 5.0(2)N1(1) or a later release.
- Cisco Fabric Services over IP (CFSoIP) enabled on each peer.
- Identical switch profiles on each switch.
- Configured peer IP addresses.

Guidelines and Limitations

The guidelines for configuration synchronization are as follows:

- You must configure the following interfaces in a switch profile:
  - Port-channel interfaces
  - Ports that are not channel-group members

- You must configure all port-channel members outside the switch profile in configuration terminal mode.
- You must follow configurations in a specified order.
- Depending on the type of vPC topology (active/active or straight-through) and the type of configuration that is needed (port channel, nonport channel, FEX, QoS, and so on), you must use the switch profile mode or the configuration terminal mode. See the "At-A-Glance Configuration Modes" section on page 4-33 to identify what mode is used for different types of configurations.

Configuration synchronization has the following configuration limitations:

- FCoE in vPC Topologies—FCoE configurations are not supported in switch profiles because configurations are typically different on peer switches. If you enable FCoE on a vPC peer switch, you must not configure the port channel in the switch profile.
- Feature Commands—The **feature** *feature name* commands that enable a conditional feature are not supported in switch profiles. You should independently configure these commands on each peer switch in configuration terminal mode.
- Configuration Rollback and Conditional Features—With configuration synchronization, when a conditional feature is present in a checkpoint and not in the running configuration, a configuration rollback to that checkpoint fails. The workaround is to reconfigure the conditional feature ("feature xyz") before the configuration rollback is executed. This workaround also applies to **the vpc domain** command and **the peer-keepalive** command in vpc-domain mode.

vPC Configurations

Configuration synchronization requires two Cisco Nexus 5000 Series peer switches that are configured in a vPC topology. Figure 4-1 shows a vPC topology configured on two Cisco Nexus 5000 Series switches (N5k-1 and N5k-2).

Figure 4-1 vPC Topology

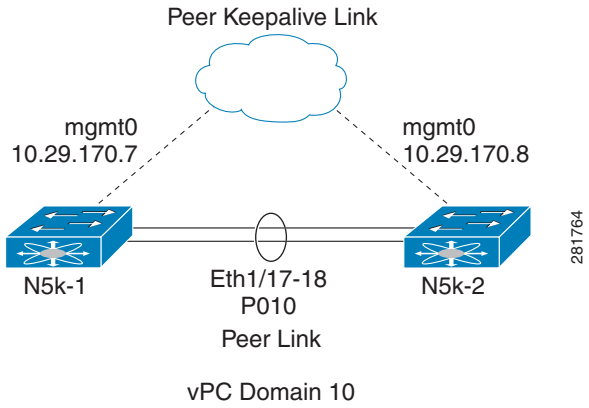

To configure vPC on two Cisco Nexus 5000 Series switches, follow these steps:

Step 1 Create a vPC domain and configure a vPC keepalive link.

You must create identical vPC domain IDs on both vPC peer switches.

The domain ID is used to automatically form the vPC system MAC address.

```
N5k-1(config)# vpc domain 10
```

```
N5k-2(config)# vpc domain 10
```

Step 2 Configure a vPC peer-keepalive link.

You can configure the destination IP for the peer-keepalive link that carries the keepalive messages. Optionally, you can configure other parameters for the keepalive messages.

```
N5k-1(config)# vpc domain 10
N5k-1(config-vpc-domain)# peer-keepalive destination 10.29.170.8 vrf management

N5k-2(config)# vpc domain 10
N5k-2(config-vpc-domain)# peer-keepalive destination 10.29.170.7 vrf management
```

Step 3 Create and configure a vPC peer link.

You can create a peer link by designating an EtherChannel on each switch as the peer link for the specified vPC domain. We recommend that you configure the EtherChannels that you are designating as the vPC peer link in trunk mode and that you use two ports on separate modules on each vPC peer switch for redundancy.

```
N5k-1(config)# interface port-channel 10
N5k-1(config-if)# vpc peer-link

N5k-1(config)# interface ethernet 1/17-18
N5k-1(config-if-range)# switchport mode trunk
N5k-1(config-if-range)# channel-group 10

N5k-2(config)# interface port-channel 10
N5k-2(config-if)# vpc peer-link

N5k-2(config)# interface ethernet 1/17-18
N5k-2(config-if-range)# switchport mode trunk
N5k-2(config-if-range)# channel-group 10
```

Note You can configure a vPC peer link between the two Cisco Nexus 5000 Series switches either manually on both switches or with configuration synchronization from any one of the two peer switches. For information about the configuration synchronization method to configure the vPC peer link, see the "Configuring a vPC Topology Using Configuration Synchronization" section on page 4-10. For additional information on the vPC feature, see Chapter 2, "Virtual Port Channel Operations" and the *Cisco Nexus 5000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.0(2)N2(1) at the following URL:*
http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/layer2/502_n2_1/Cisco_n5k_layer2_config_gd_rel_502_N2_1_chapter8.html

Cisco Fabric Services Over IP

The Cisco Fabric Services over IP (CFSoIP) protocol transports configuration synchronization over the mgmt0 interface (mgmt virtual routing and forwarding [VRF]). You must ensure the connectivity to the mgmt0 interface. CFSoIP and Cisco Fabric Services (CFS) are different protocols. CFS runs across the peer link for a vPC. Although both protocols are based on the CFS protocol, they exchange different control packets.

To use the CFSoIP protocol for configuration synchronization, follow these steps:

Step 1 Enable CFSoIP manually on each peer switch:

```
N5k-1# config terminal
N5k-1(config)# cfs ipv4 distribute

N5k-2# config t
N5k-2(config)# cfs ipv4 distribute
```

Note CFSoIP is not supported on the switch virtual interface (SVI)/default VRF.

Step 2 Establish the peer connection over the mgmt0 transport interface:

```
N5K-1# configure terminal
N5K-1(config)# interface mgmt 0
N5K-1(config-if)# ip address 10.29.170.7/24
N5K-1(config-if)# vrf context management
N5K-1(config-vrf)# ip route 0.0.0.0/0 10.29.170.1

N5K-2# configure terminal
N5K-2(config)# interface mgmt 0
N5K-2(config-if)# ip address 10.29.170.8/24
N5K-2(config-if)# vrf context management
N5K-2(config-vrf)# ip route 0.0.0.0/0 10.29.170.1
```

Switch Profiles

Beginning with Cisco NX-OS Release 5.0(2)N1(1), config-sync mode allows you to create a switch profile. A switch profile contains a predefined configuration that you can use to configure a peer switch so that both peers have the same configuration. In config-sync mode, you define the peer and the configuration in the switch profile. Peers are identified by their IP address and they are local to each switch profile. Commands entered in config-sync mode are buffered until they are committed. Configuration changes made in configuration terminal mode apply only to the local switch.

You must create an identical switch profile on each peer switch in config-sync mode. This configuration is not automatically synchronized and you must configure it on each peer switch.

To create the switch profiles, enter the following commands:

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test
N5k-1(config-sync)# sync-peers destination 10.29.170.8

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
N5k-2(config-sync)# sync-peers destination 10.29.170.7
```

Note The switch profile name must be identical on both peers. You can create only one switch profile on each peer switch.

In Cisco NX-OS Release 5.0(2)N1(1), switch profiles do not support all commands. The config-sync mode commands are limited to vPC configurations.

Switch Profile Commands

The following configuration commands are supported in a switch profile:

- VLAN
- ACL
- Spanning-Tree Protocol (STP)
- Quality of Service (QoS)
- Interface configurations (Ethernet, port channel, or vPC interfaces)

The following commands are not supported in a switch profile:

- Enable feature sets (you must manually enable a specific feature before relevant configurations can be added)
- vPC domain
- vPC peer-keepalive
- FCoE

User-Based Access Controls

In Cisco NX-OS Release 5.0(2)N1(1), only the user who creates a switch profile can edit the switch profile even if another user has the same admin privileges. Beginning with Cisco NX-OS Release 5.0(2)N2(1), you can add, delete, or modify a switch profile configuration based on Role Based Access Control (RBAC) configurations. Users that have the appropriate privilege level to access the switch profile can successfully modify the switch profile and commit the configuration.

For more information on RBAC, see the *Cisco Nexus 5000 Series System Management Configuration Guide*.

As a network administrator, you can restrict a user from accessing a switch profile. When a restricted user has permission to access a switch profile, that user can successfully commit the switch profile on the initiating switch. However, issuing a particular command (for example, the **switchport mode access** command), fails or succeeds in a switch profile according to the RBAC policies and rules assigned to that user.

In addition, the same username and privilege level must exist for a successful commit on the peer switch. If the same username and privilege level does not exist on the peer switch, the commit fails. You must ensure that configuration synchronization peers have the same configured users and roles. Occasionally, the same username can exist but roles might be mismatched. Also, the same user on one peer switch could have a more restricted role on the other peer switch and in that case the commit might fail. You must configure usernames with matching roles on peer switches to avoid these problems. As a best practice, the user with the network administrator role should create the switch profile to reduce the risk of configurations failing to commit due to permission issues.

In vPC topologies, when one peer switch is running Cisco NX-OS Release 5.0(2)N1(1) and a second peer switch is running Cisco NX-OS Release 5.0(2)N2(1), a successful commit depends on which switch was used to issue the commit. On a switch running Cisco NX-OS Release Cisco NX-OS Release 5.0(2)N1(1), only the user who created the switch profile can issue the commit. On a switch running Cisco NX-OS Release 5.0(2)N2(1), users with appropriate privileges can issue the commit.

Verification Checks

To reduce the possibility of overriding switch profile configurations or configurations that are not part of a switch profile, two types of validation checks are performed:

- Mutual Exclusion Check, page 4-7
- Merge Check, page 4-7

Mutual Exclusion Check

The Mutual exclusion check identifies potential conflicts between a switch profile configuration and the global configurations (configurations that are not part of a switch profile). A command that is included in a switch profile cannot be configured outside of the switch profile. The same rules apply on the peer switch.

The mutual exclusion check is done locally and on the peer switch. When entering the **verify** or **commit** command, if the peer switch is reachable using the mgmt0 interface, the check is done on both the local switch and the peer switch. If the peer switch is not reachable, the check is only done on the local switch.

If the mutual exclusion check fails, you must manually correct the configuration and enter the **commit** command again.

The following commands are exceptions and they can exist inside and outside the switch profile without receiving a mutual exchange error:

- Interface configuration commands except port-channel interfaces
- Shutdown/no shut commands
- System Quality of Service (QoS) (**system qos** command) command

For implementations including port channels, consider the following guidelines to minimize mutual exchange errors:

- Port channels created in switch profile mode should not be configured using global configuration (config terminal) mode.
- If a port-channel is created in global configuration mode, channel groups including member interfaces must also be created using global configuration mode.
- Port-channels that are configured within switch profile mode may have members both inside and outside of a switch profile.
- If you want to import a member interface to a switch profile, the port-channel that corresponds with the member interface must also be present within the switch profile.

For more information on configuring port channels, see the *Cisco NX-OS Layer 2 Switching Configuration Guide*. For more information on configuring switch profiles, see the *Cisco NX-OS 5000 System Management Configuration Guide*.

Merge Check

A Merge check is done on the peer switch to ensure that the received configuration does not conflict with the switch profile configuration that already exists on the receiving switch. If a merge check failure occurs, you must manually correct the configuration and enter the **commit** command again.

Commit

Use the **commit** command to synchronize the configuration with the peer switch and to apply the configuration locally. Configurations are stored in the buffer until the **commit** command is issued. A commit can be executed by a vPC primary switch or a secondary switch. The initiator is the switch on which the **commit** command was issued. You can enter the **commit** command only on one switch at any given time. If a commit is attempted while another commit is in progress, it fails and the following syslog message appears:

```
Failed: Session Database already locked, Verify/Commit in Progress
```

As a best practice, you should assign one switch as the initiator, make all configurations on that switch, and then synchronize the configuration on the peer switch to simplify the process and reduce any possible confusion.

All configuration changes (including configuration terminal mode changes for all supported commands) are prevented when a switch profile session is in progress.

Note If the peer switch is reachable and you enter the **commit** command, the configuration is applied locally and to its peer switch. If the commit is unsuccessful, the configuration is not applied on the local or remote switch (atomic behavior).

Commands are executed in the same order in which they are buffered. If there is an order dependency for certain commands (for example, QoS policy commands), the commands must be defined before they are applied. The order of commands can be edited in the buffer. If you are including commands that are part of a feature that requires the feature to be enabled, you must ensure that the feature is enabled and defined manually on each switch.

Note The **feature** command is not synchronized between peers (for example, **feature** *vpc* or **feature** *lacp*.

When you enter the **commit** command, the CLI prompt may not return right away. The length of time it takes to apply the configuration may be longer if the size of the configuration is very large. This operation is normal and we recommend that you do not abort the commit (by pressing **Ctrl-c** or **Ctrl-z**) because it might leave the configuration in an inconsistent state.

Buffering

The switch profile configuration is stored in a buffer until the **commit** command is entered. You can add, delete, or move configurations in the buffer. Once the configuration has been pushed using the **commit** command, it is applied to the system configuration. Use the **show running** command to verify that the configuration has been applied. You can also use the **show running switch-profile** command to specifically check what configuration was synchronized using the switch profile.

Consider the following guidelines for the configuration that is stored in the buffer:

- Configurations are buffered until a successful commit.
- You can add, delete, or move configurations in the buffer.
- Commands are executed in the same order in which they are buffered. You can change the order inside the buffer by using the **buffer-delete** command or the **buffer-move** command.

Import

When you upgrade to Cisco NX-OS Release 5.0(2)N1(1), you have the option to enter the **import** command to copy supported running configuration commands to a switch profile. The **switch -profile import** command allows you to import the entire running configuration or you can choose specific interfaces to merge. Changes are not supported during the import process. If you add new commands in addition to the import configurations, the commit might fail. The commands remain in the buffer. You have the option to correct the buffer and enter the **commit** command again or abort the import mode. If you abort the import, the commands in the buffer are lost.

Configuration Synchronization Best Practices

Configuration synchronization is primarily used in vPC topologies. You should follow the best practice guidelines in this section to ensure a successful configuration synchronization.

Caution You must follow the best practices for configuration synchronization described in this chapter for Cisco NX-OS Release 5.0(2)N1(1). Failure to do so might leave configurations in an inconsistent state.

In addition to configuration synchronization, Cisco NX-OS Release 5.0(2)N1(1) introduced three additional features:

- Pre provisioning—Allows you to configure offline GEM and FEX interfaces.
- Port profiles—Allows you to define consistent interface configurations that are applied to multiple ports.
  - Port profiles apply to ports and switch profiles apply to switch configurations; they are not the same.
  - Port profiles are not required for a configuration synchronization but they can be included in a configuration synchronization.
- Configuration Rollback—Allows you to create checkpoints of the running configuration and then perform a rollback to those checkpoints.

Use pre-provisioning, port profiles, and configuration rollbacks to enhance configuration synchronization and provide maximum benefits in a vPC topology. These features are included in the examples found in this chapter. See the *Cisco Nexus 5000 Series Configuration Guides* for additional information on these features.

Note These features are independent of configuration synchronization; you do not need to enable them to use configuration synchronization.

Configuration Examples

This section describes the following configuration examples:

- Configuring a vPC Topology Using Configuration Synchronization, page 4-10
- Active/Active FEX Topology Examples, page 4-12
- Straight-Through Topology Examples, page 4-19

Note The following examples are current as of Cisco NX-OS Release 5.0(2)N1(1).

Configuring a vPC Topology Using Configuration Synchronization

In Figure 4-2, N5k-1 and N5k-2 are part of vPC Domain 10. The peer keepalive is configured over the mgmt0 interface and Ethernet 1/17-18 are bundled into P010 to form the peer link. Configuration synchronization maintains a consistent configuration on the peer switches and simplifies the switch administration in a vPC topology.

Figure 4-2 vPC Topology

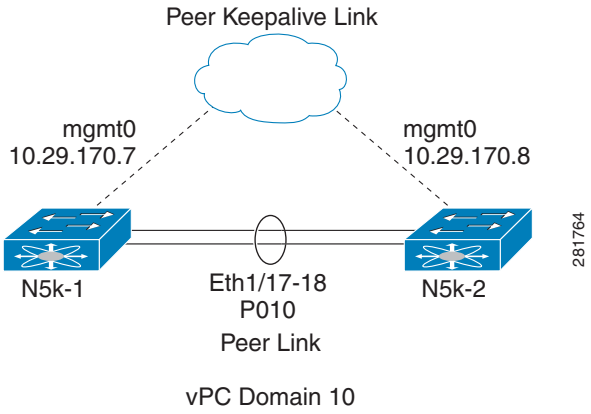

Example 4-1 shows the sample running configuration required for the vPC to become operational.

Example 4-1 Running Configuration of Peer Switches in a vPC Topology

vPC Configuration for N5k-1	vPC Configuration for N5k-2

`vlan 1-10`
 	
`feature vpc`
 	
`vpc domain 10`
`  peer-keepalive destination 10.29.170.8`
`  peer-config-check-bypass`
 	
`interface port-channel10`
`  switchport mode trunk`
`  vpc peer-link`
`  spanning-tree port type network`
 	
`interface Ethernet1/17`
`  switchport mode trunk`
`  channel-group 10`
 	
`interface Ethernet1/18`
`  switchport mode trunk`
`  channel-group 10`
	`vlan 1-10`
	
	`feature vpc`
	
	`vpc domain 10`
	`  peer-keepalive destination 10.29.170.7`
	`  peer-config-check-bypass`
	
	`interface port-channel10`
	`  switchport mode trunk`
	`  vpc peer-link`
	`  spanning-tree port type network`
	
	`interface Ethernet1/17`
	`  switchport mode trunk`
	`  channel-group 10`
	
	`interface Ethernet1/18`
	`  switchport mode trunk`
	`  channel-group 10`

Note Peer-config-check-bypass is a best practice configuration for vPCs. For more information, see the *Cisco Nexus 5000 Series Design Guide* at the following URL: http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/C07-572829-01_Design_N5K_N2K_vPC_DG.pdf

To configure the vPC topology shown in Figure 4-2, follow these steps:

Step 1 Enable the vPC feature on the peer switches.

```
N5K-1(config-vpc-domain)# feature vpc

N5K-2(config-vpc-domain)# feature vpc
```

Step 2 Configure the peer keepalive on both switches using the mgmt0 interface.

```
N5K-1(config)# vpc domain 10
N5K-1(config-vpc-domain)# peer-keepalive destination 10.29.170.8

N5K-2(config)# vpc domain 10
N5K-2(config-vpc-domain)# peer-keepalive destination 10.29.170.7
```

Step 3 Enable CFSoIP on both switches.

```
N5k-1# configuration terminal
N5k-1(config)# cfs ipv4 distribute

N5k-2# configuration terminal
N5k-2(config)# cfs ipv4 distribute
```

Step 4 Configure the switch profile with the same name on both switches.

```
N5K-1(config)# config sync
N5K-1(config-sync)# switch-profile Test

N5K-2(config)# config sync
```

```
N5K-2(config-sync)# switch-profile Test
```

Step 5 Enter the **sync peer destination** command to configure both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test
N5k-1(config-sync-sp)# sync-peers destination 10.29.170.8

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
N5k-2(config-sync-sp)# sync-peers destination 10.29.170.7
```

Step 6 In switch-profile mode, create the port-channel interface for the peer link.

```
N5K-1(config-if)# config sync
N5K-1(config-sync)# switch-profile Test
N5K-1(config-sync-sp)# int po10
N5K-1(config-sync-sp-if)# exit
N5K-1(config-sync-sp)# commit
```

Step 7 In interface mode, associate the port-channel member to PO 10.

```
N5K-1# config t
N5K-1(config)# int ether 1/17-18
N5K-1(config-if)# channel-group 10

N5K-2# config t
N5K-2(config)# int ether 1/17-18
N5K-2(config-if)# channel-group 10
```

Step 8 In switch profile mode, add the appropriate configurations under the port-channel interface to form the peer link.

```
N5K-1(config-if)# config sync
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# interface po10
N5K-1(config-sync-sp-if)# switchport mode trunk
N5K-1(config-sync-sp-if)# vpc peer-link
N5K-1(config-sync-sp-if)# commit

Verification successful...
Proceeding to apply configuration. This might take a while depending on
amount of configuration in buffer.
Please avoid other configuration changes during this time.
Commit Successful
```

Active/Active FEX Topology Examples

This section includes the following examples:

Dual-Homed FEX Topology (Active/Active FEX Topology)

Figure 4-3 shows that each FEX is dual-homed with two Cisco Nexus 5000 Series switches. The FEX-fabric interfaces for each FEX are configured as a vPC on both peer switches. The host interfaces on the FEX appear on both peer switches. If these host interfaces are bundled in a port channel, you must configure the port channel identically on both peer switches. Configuration synchronization helps keep the FEX configuration synchronized between the pair of vPC peer switches.

Figure 4-3 Dual-Homed FEX Active/Active Topology

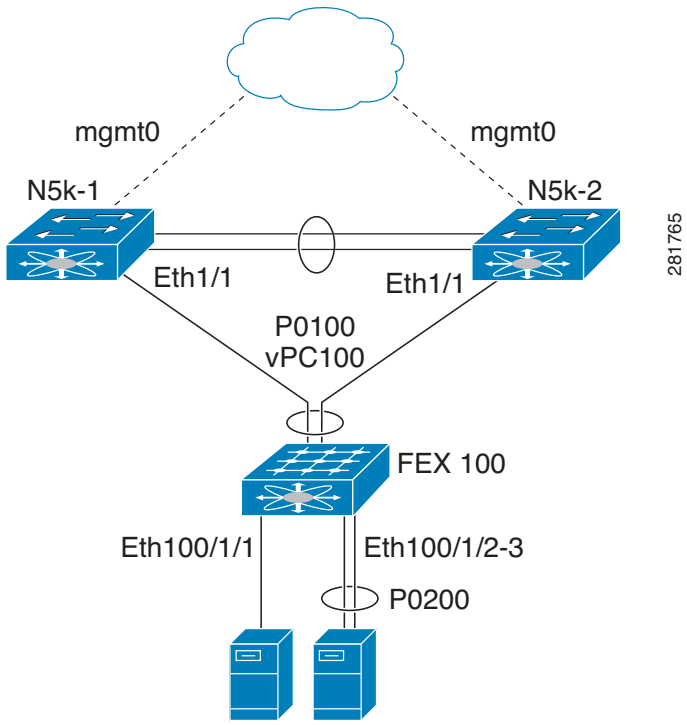

In Figure 4-3, the vPC is already operational**.** FEX 100 is dual-homed to both parent switches: N5k-1 and N5k-2 on FEX-fabric interfaces Ethernet 1/1. Because the FEX is pre-provisioned, there is no existing running configuration on Ethernet 1/1.

Note A port channel within the same FEX is supported on Cisco Nexus 2200 Series Fabric Extenders.

FEX 100 is configured to have two types of host interfaces. One interface is Ethernet100/1/1, which is singly attached to a server (nonport-channel member), and the other interface is Ethernet 100/1/2-3, which is configured in a port channel to the server (port-channel member).

Example 4-2 shows the sample running configuration for the peer switches. Two types of configurations are shown:

- Basic Configuration
- Port profile configuration.

You can use either option or you can use both configurations together.

Note You can use port profiles to reduce operational overhead although they are not required.

Example 4-2 Running Configuration of a FEX in an Active/Active Topology for the Peer Switches

Basic Configuration—No Port Profile	Port Profile Configuration

vlan 1-10
 	
interface port-channel100
  switchport mode fex fabric
  vpc 100
  fex associate 100
 	
interface port-channel 200
  switchport mode trunk
  switchport trunk allowed vlan 1-5
 	
interface Ethernet1/1
  fex associate 100
  switchport mode fex fabric
  channel-group 100
 	
interface Ethernet100/1/1
  switchport mode trunk
  switchport trunk allowed vlan 1-10
 	
interface Ethernet100/1/2
  switchport mode trunk
  switchport trunk allowed vlan 1-5
  channel-group 200
 	
interface Ethernet100/1/3
  switchport mode trunk
  switchport trunk allowed vlan 1-5
  channel-group 200
	vlan 1-10
	
	port-profile type ethernet eth-profile
	  switchport mode trunk
	  state enabled
	
	port-profile type port-channel pc-profile
	  switchport mode trunk
	  state enabled
	
	interface port-channel100
	  switchport mode fex fabric
	  vpc 100
	  fex associate 100
	
	interface port-channel 200
	  inherit port-profile pc-profile
	  switchport trunk allowed vlan 1-5
	
	interface Ethernet1/1
	  fex associate 100
	  switchport mode fex fabric
	  channel-group 100
	
	interface Ethernet100/1/1
	  inherit port-profile eth-profile
	  switchport trunk allowed vlan 1-10
	
	interface Ethernet100/1/2
	  switchport mode trunk
	  switchport trunk allowed vlan 1-5
	  channel-group 200
	
	interface Ethernet100/1/3
	  switchport mode trunk
	  switchport trunk allowed vlan 1-5
	  channel-group 200

New Deployments in an Active/Active FEX Topology

In a new deployment, configuration synchronization is introduced from the beginning to synchronize the configuration across peer switches. As a result, there is no existing running configuration on the FEX ports.

To configure the dual-homed FEX active/active topology shown in Figure 4-3, follow these steps:

Step 1 Enable CFSoIP on both switches.

```
N5k-1# config t
N5k-1(config)# cfs ipv4 distribute

N5k-2# config t
N5k-2(config)# cfs ipv4 distribute
```

Step 2 Create a switch profile on both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test
```

```
N5k-1(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-2>

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
N5k-2(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-1>
```

Step 3 Pre-provision the FEX.

Note In a FEX active/active topology, always pre-provision the FEXs that are dual-homed inside the switch profile. This process helps configuration synchronization when the FEX is not connected to a Cisco Nexus 5000 Series switch.

```
N5k-1(config-sync-sp)# slot 100
N5k-1(config-sync-sp-slot)# provision model N2k-C2232P
N5k-1(config-sync-sp-slot)# exit

N5K-1(config-sync-sp-if)# sh switch-profile buffer

switch-profile  : Test
----------------------------------------------------------
Seq-no  Command
----------------------------------------------------------
1       slot 100
1.1       provision model N2K-C2232P
```

Step 4 Add referred global configuration to the switch profile.

Note Because interface configurations will be synchronized, all policies that are applied on the interface must be synchronized (for example, port profiles, QoS, and ACL policies).

```
N5k-1(config-sync-sp)# port-profile type ethernet eth-profile
N5k-1(config-sync-port-prof)# switchport mode trunk
N5k-1(config-sync-port-prof)# state enabled

N5k-1(config-sync-sp)# port-profile type port-channel pc-profile
N5k-1(config-sync-port-prof)# switchport mode trunk
N5k-1(config-sync-port-prof)# state enabled
```

Step 5 Configure the Ethernet interfaces (the non-port-channel members) inside the switch profile.

```
N5k-1(config-sync-sp)# interface Ethernet100/1/1
N5k-1(config-sync-sp-if)# inherit port-profile eth-profile
N5k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-10
```

Step 6 Create the port-channel interface inside the switch profile.

Note You must configure port-channel interfaces in the switch profile, not in configuration terminal mode.

This example shows that port channel 100 (vPC 100) is the EtherChannel from N5k to N2k:

```
N5k-1(config-sync-sp)# interface Port-channel100
```

This example shows that port channel 200 is the EtherChannel from N2k to the end device:

```
N5k-1(config-sync-sp)# interface Port-channel200
```

Step 7 Commit the configuration inside the switch profile.

```
N5k-1(config-sync-sp)# commit
```

Step 8 Add members to the port channel in configuration terminal mode on both switches.

Note The configuration must be done on both switches in configuration terminal mode.

This example shows that N5k-1- Ethernet1/1 is a FEX-fabric member of port channel 100:

```
N5k-1(config)# int ether1/1
N5k-1(config-if)# channel-group 100 force
```

This example shows that N5k-1- Ethernet1/100/2-3 are members of port channel 200:

```
N5k-1(config)# interface Ethernet100/1/2-3
N5k-1(config-if-range)# channel-group 200 force
```

This example shows that N5k-2- Ethernet1/1 is a FEX-fabric interface that is in port channel 100:

```
N5k-2(config)# int ether1/1
N5k-2(config-if)# channel-group 100 force
```

This example shows that N5k-2- Ethernet1/100/2-3 are members of port channel 200:

```
N5k-2(config)# interface Ethernet100/1/2-3
N5k-2(config-if-range)# channel-group 200 force
```

Note In Cisco NX-OS Release 5.0(2)N2(1), if you do not use the **channel-group 200 force** command on the Ethernet interfaces, a problem will occur on pre-provisioned interfaces that are offline. In this example, if module 100 is offline, the configuration on PO 200 in Step 8 must be specifically configured on each member interface, in addition to the **channel-group** command. The **channel-group 200 force** command is not supported in Cisco NX-OS Release 5.0(2)N1(1) and earlier releases.

```
N5k-1(config)# interface Ethernet100/1/2-3
N5k-1(config-if-range)# switchport mode trunk
N5k-1(config-if-range)# switchport trunk allowed vlan 1-5

N5k-2(config)# interface Ethernet100/1/2-3
N5k-2(config-if-range)# switchport mode trunk
N5k-2(config-if-range)# switchport trunk allowed vlan 1-5
```

Step 9 Modify the port-channel configuration in the switch profile.

```
N5k-1(config-sync-sp-if)# interface Port-Channel100
N5k-1(config-sync-sp-if)# switchport mode fex-fabric
N5k-1(config-sync-sp-if)# fex associate 100
N5k-1(config-sync-sp-if)# vpc 100

N5k-1(config-sync-sp)# interface Port-channel200
N5k-1(config-sync-sp-if)# inherit port-profile pc-profile
N5k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
```

Step 10 Commit the configuration in the switch profile.

```
N5k-1(config-sync-sp)# commit
```

Existing Deployment with an Active/Active FEX Topology

In an existing deployment, the configurations are already present and configuration synchronization is used to simplify future configuration modifications.

To configure peer switches in the vPC topology shown in Figure 4-3, follow these steps:

Step 1 Enable CFSoIP on both switches.

```
N5k-1# config t
N5k-1(config)# cfs ipv4 distribute

N5k-2# config t
N5k-2(config)# cfs ipv4 distribute
```

Step 2 Create a switch profile on both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
```

Step 3 Pre-provision the FEX on both switches.

Note In a FEX active/active topology, always pre-provision the FEXs that are dual-homed inside the switch profile.

```
N5k-1(config-sync-sp)# slot 100
N5k-1(config-sync-sp-slot)# provision model N2k-C2232P
N5k-1(config-sync-sp-slot)# exit

N5k-2(config-sync-sp)# slot 100
N5k-2(config-sync-sp-slot)# provision model N2k-C2232P
N5k-2(config-sync-sp-slot)# exit
```

Step 4 Commit the configuration in the switch profile on both switches.

```
N5k-1(config-sync-sp)# commit

N5k-2(config-sync-sp)# commit
```

Step 5 Import the running configuration.

```
N5k-1(config-sync-sp)# import running-config
N5k-1(config-sync-sp-import)# show switch-profile Test buffer
```

Import the configuration to the switch profile on both switches. You can import the configuration using one of the following three methods:

- Running configuration—All configurations that are allowed inside a switch profile are imported. You must remove unwanted configurations. For example, you must remove port-channel member configurations if the member interfaces do not match on the peer switches.
- Interface configuration—Only specified interface configurations are imported.
- Manual mode—Selected configurations are imported. If the configuration that needs to be imported is small, use the manual mode to paste the desired configuration.

Table 4-1 shows the command sequence to import the running configuration:

Table 4-1 Command Sequence to Import the Running Configuration

Sequence Number	Command
1	vlan 1-10
2 2.1 2.2 2.3	interface port-channel100 switchport mode fex-fabric vpc 100 fex associate 100
3 3.1 3.2	interface port-channel200 switchport mode trunk switchport trunk allowed vlan 1-5
4 4.1 4.2 4.3	interface Ethernet1/1 fex associate 100 switchport mode fex-fabric channel-group 100
5 5.1 5.2	interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-10
6 6.1 6.2 6.3	interface Ethernet100/1/2 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200
7 7.1 7.2 7.3	interface Ethernet100/1/3 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

Step 6 Remove member interfaces of PO 100 and PO 200 from the buffer.

```
N5k-1(config-sync-sp-import)# buffer-delete 4, 6, 7
```

Use the **buffer-delete** command to delete the unwanted configuration from the buffer.

Step 7 Commit the configuration in the switch profile on both switches.

```
N5k-1(config-sync-sp-import)# commit

N5k-2(config-sync-sp-import)# commit
```

Step 8 Add the sync peer on both switches.

Note When importing the configuration, you must use the **sync-peers** command after the configurations are imported independently on both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile sp
N5k-1(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-2>

N5k-2# config sync
N5k-2(config-sync)# switch-profile sp
N5k-2(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-1>
```

Caution When you remove a switch profile using the **no switch-profile** *name* **[all-config | local-config]** command, the configuration in the switch profile is immediately removed from the running configuration. This disrupts the configurations that were present in the switch profile, for example port channel and vPC configurations. For current information about this issue, refer to CSCtl87240 and CSCtl87260 in the *Cisco Nexus 5000 Series Switch and Cisco Nexus 2000 Series Fabric Extender Release Notes* and in the Cisco Bug Toolkit located at the following URL: http://tools.cisco.com/Support/BugToolKit/action.do?hdnAction=searchBugs.

Straight-Through Topology Examples

This section includes the following examples:

Switch vPC Topology and Straight-Through FEX Topologies (Host vPC)

In Figure 4-4, the Cisco Nexus 5000 Series switch ports are directly connected to another switch or host and are configured as part of a port channel that becomes part of a vPC.

Figure 4-4 shows that vPC 20 is configured on port channel 20, which has Eth1/10 on N5k-1 and Eth2/1 on N5K-2 as members.

Figure 4-4 Switch vPC Topology

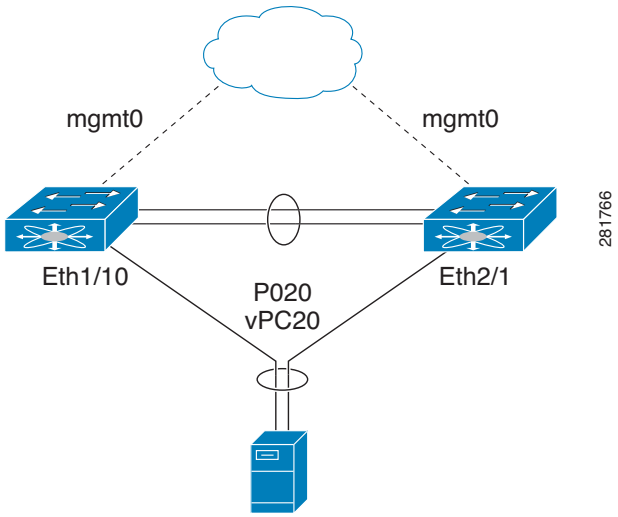

In Figure 4-5, each FEX is single-homed (straight-through FEX topology) with a Cisco Nexus 5000 Series switch. The host interfaces on this FEX are configured as port channels and those port channels are configured as vPCs.

Eth100/1/1 on N5K-1 and Eth102/1/5 on N5K-2 are configured as members of PO200 and PO200 is configured for vPC 200.

Figure 4-5 FEX Straight-Through Topology (Host vPC)

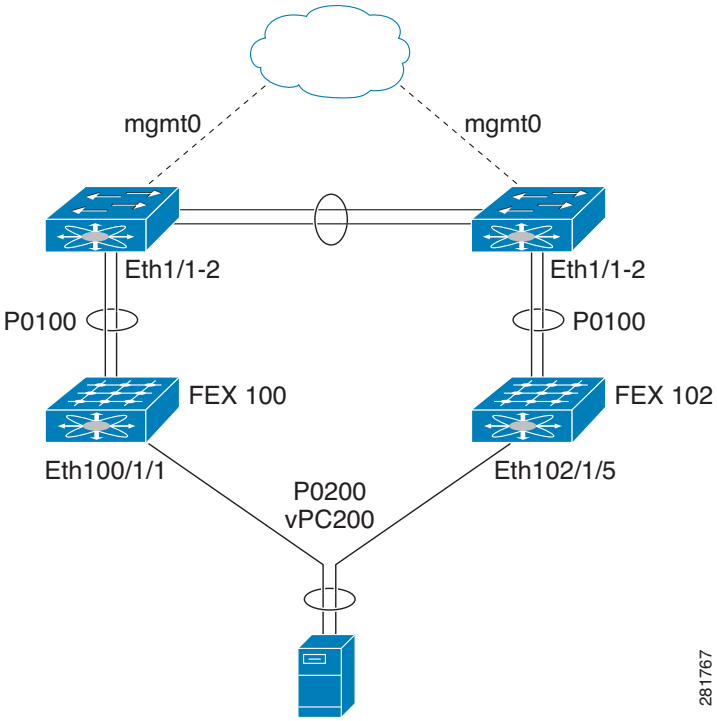

In both topologies, port channels P020 and P0200 must be configured identically on the peer switches and configuration synchronization is used to synchronize the configurations of the vPC switches.

Example 4-3 shows the sample running configuration that must be configured for the peer switches shown in the vPC topologies in Figure 4-4 and Figure 4-5.

Example 4-3 Running Configuration Example for the Nexus 5000 Series Switches in a vPC Straight-Through Topology.

Basic Configuration—No Port Profile	Port Profile Configuration

`vlan 1-10`
 	
`interface port-channel 20`
`  switchport mode trunk`
`  vpc 20`
`  switchport trunk allowed vlan 1-5`
 	
`interface port-channel 200`
`  switchport mode trunk`
`  vpc 200`
`  switchport trunk allowed vlan 1-5`
 	
`interface Ethernet1/10`
`  switchport mode trunk`
`  switchport trunk allowed vlan 1-5`
`  channel-group 20`
 	
`interface Ethernet100/1/1`
`  switchport mode trunk`
`  switchport trunk allowed vlan 1-5`
`  channel-group 200`
	`vlan 1-10`
	
	`port-profile type port-channel pc-profile`
	`  switchport mode trunk`
	`  state enabled`
	
	`interface port-channel 20`
	`  inherit port-profile pc-profile`
	`  vpc 20`
	`  switchport trunk allowed vlan 1-5`
	
	`interface port-channel 200`
	`  inherit port-profile pc-profile`
	`  vpc 200`
	`  switchport trunk allowed vlan 1-5`
	
	`interface Ethernet1/10`
	`  switchport mode trunk`
	`  switchport trunk allowed vlan 1-5`
	`  channel-group 20`
	
	`interface Ethernet100/1/1`
	`  switchport mode trunk`
	`  switchport trunk allowed vlan 1-5`
	`  channel-group 200`

New Deployment in a vPC Topology and Straight-Through FEX Topology

In a new deployment, configuration synchronization is introduced initially to synchronize the new configuration. Because it is a new deployment, there is no existing running configuration on the FEX ports.

Note In a straight-through FEX topology, you must use configuration terminal mode to pre-provision FEXs or GEMs.

To configure the peer switches in the topologies shown in Figure 4-4 and Figure 4-5, follow these steps:

Step 1 Pre-provision the FEX configuration in configuration terminal mode for both switches as follows:

Provision the N5k-1- slot 100 for FEX 100.

```
N5K-1(config)# slot 100
N5K-1(config-slot)# provision model N2K-C2232P
N5K-1(config)# int ether 1/1-2
N5K-1(config-if-range)# channel-group 100
N5K-1(config-if-range)# int po100
N5K-1(config-if)# fex associate 100
N5K-1(config-if)# switchport mode fex-fabric
```

Provision the N5k-2- slot 102 for FEX 102.

```
N5K-2(config)# slot 102
N5K-2(config-slot)# provision model N2K-C2232P
```

```
N5K-2(config)# int ether 1/1-2
N5K-2(config-if-range)# channel-group 102
N5K-2(config-if-range)# int po102
N5K-2(config-if)# fex associate 102
N5K-2(config-if)# switchport mode fex-fabric
```

Provision the N5k-2- slot 2 for a GEM.

```
N5K-2(config)# slot 2
N5K-2(config-slot)# provision model N55-M16P
```

Step 2 Enable CFSoIP on both switches.

```
N5k-1# config t
N5k-1(config)# cfs ipv4 distribute

N5k-2# config t
N5k-2(config)# cfs ipv4 distribute
```

Step 3 Create a switch profile and configure the peer on both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test
N5k-1(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-2>

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
N5k-2(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-1>
```

Step 4 Add the referred global configuration to the switch profile. Because the configuration on the interfaces will be synchronized, all policies that are applied on the interface must be synchronized (for example, port profiles, QoS and ACL policies).

```
N5k-1(config-sync-sp)# port-profile type port-channel pc-profile
N5k-1(config-sync-port-prof)# switchport mode trunk
N5k-1(config-sync-port-prof)# state enabled
```

Step 5 Create port-channel interfaces inside the switch profile.

Note Use switch profile mode to create the port-channel interfaces.

```
N5k-1(config-sync-sp)# interface port-channel 20
N5k-1(config-sync-sp)# interface port-channel 200
```

Step 6 Commit the configuration in the switch profile.

```
N5k-1(config-sync-sp)# commit
```

Step 7 Add members to the port channel in configuration terminal mode on both switches. When the configuration is done in configuration terminal mode, both switches must be configured independently.

Note In this topology, port-channel members must not be identical on the peer switches. For Cisco NX-OS Release 5.0(2)N1(1), port-channel members should only be configured in configuration terminal mode, not in the switch profile.

```
N5k-1(config)# interface Ethernet1/10
N5k-1(config-if)# channel-group 20 force
N5k-1(config)# interface Ethernet100/1/1
```

```
N5k-1(config-if)# channel-group 200 force

N5k-2(config)# interface Ethernet2/1
N5k-2(config-if)# channel-group 20 force
N5k-2(config)# interface Ethernet102/1/5
N5k-2(config-if)# channel-group 200 force
```

Note In Cisco NX-OS Release 5.0(2)N2(1), if you do not use the **channel-group 200 force** command on the Ethernet interfaces, a problem will occur on pre-provisioned interfaces that are offline. In this example, if module 100 is offline, the configuration on P0200 in Step 7 must be configured on the member interfaces. The **channel-group 200 force** command is not supported in Cisco NX-OS Release 5.0(2)N1(1) and earlier releases.

```
N5k-1(config)# interface Ethernet100/1/1
N5k-1(config-if)# switchport mode trunk
N5k-1(config-if)# switchport trunk allowed vlan 1-5

N5k-2(config)# interface Ethernet2/1
N5k-2(config-if)# switchport mode trunk
N5k-2(config-if)# switchport trunk allowed vlan 1-5

N5k-2(config)# interface Ethernet102/1/5
N5k-2(config-if)# switchport mode trunk
N5k-2(config-if)# switchport trunk allowed vlan 1-5
```

Note Ethernet 1/10 is not included in the list because it is not pre-provisioned (it is an offline interface).

Step 8 Modify the port-channel configuration in the switch profile.

```
N5k-1(config-sync-sp)# interface port-channel 20
N5k-1(config-sync-sp-if)# inherit port-profile pc-profile
N5k-1(config-sync-sp-if)# vpc 20
N5k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
N5k-1(config-sync-sp)# interface port-channel 200
N5k-1(config-sync-sp-if)# inherit port-profile pc-profile
N5k-1(config-sync-sp-if)# vpc 200
N5k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
```

Step 9 Commit the configuration in the switch profile.

```
N5k-1(config-sync-sp)# commit
```

Existing Deployments in a vPC Topology and Straight-Through FEX Topology

In an existing deployment, the configurations are already present and configuration synchronization is used to simplify future configuration modifications.

Note In a straight-through FEX topology, use configuration terminal mode to pre-provision FEXs and GEMs.

To configure the peer switches in the topologies shown in Figure 4-4 and Figure 4-5, follow these steps:

Step 1 Pre-provision the FEXs in configuration terminal mode on both switches.

```
N5K-1(config)# slot 100

N5K-2(config)# slot 102
```

Step 2 Enable CFSoIP on both switches.

```
N5k-1# config t
N5k-1(config)# cfs ipv4 distribute

N5k-2# config t
N5k-2(config)# cfs ipv4 distribute
```

Step 3 Create a switch profile on both switches.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
```

Step 4 Import the running configuration.

```
N5k-1(config-sync-sp)# import running-config
N5k-1(config-sync-sp-import)# show switch-profile Test buffer
```

Import the configuration to the switch profile on both switches. You can import the configuration using one of the following three methods:

- Running configuration—All configurations that are allowed inside a switch profile are imported. You must remove unwanted configurations. For example, you must remove port-channel member configurations.
- Interface configuration—Only specified interface configurations are imported.
- Manual mode—Selected configurations are imported. If the configuration that needs to be imported is small, use the manual mode to paste the desired configuration.

Table 4-2 shows the command sequence to import the running configuration for Step 4:

Table 4-2 Command Sequence to Import the Running Configuration (Step 4)

Sequence Number	Command
1	vlan 1-10
2 2.1 2.2 2.3	interface port-channel20 switchport mode trunk vpc 20 switchport trunk allowed vlan 1-5
3 3.1 3.2	interface port-channel100 switchport mode fex-fabric fex associate 100
3 3.1 3.2	interface port-channel200 switchport mode trunk vpc 200 switchport trunk allowed vlan 1-5

Table 4-2 Command Sequence to Import the Running Configuration (Step 4) (continued)

Sequence Number	Command
4 4.1 4.2 4.3	interface port-channel20 switchport mode trunk vpc 20 switchport trunk allowed vlan 1-5
5 5.1 5.2 5.3	interface Ethernet1/1 fex associate 100 switchport mode fex-fabric channel-group 100
6 6.1 6.2 6.3	interface Ethernet1/2 fex associate 100 switchport mode fex-fabric channel-group 100
7 7.1 7.2 7.3	interface Ethernet1/10 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 20
8 8.1 8.2 8.3	interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

```
N5k-2(config-sync-sp)# import running-config
```

Step 5 (Optional) If you do not want to synchronize the fabric configuration, remove the fabric configuration and the member interfaces of PO 20 and PO 200 from the buffer.

```
N5k-1(config-sync-sp-import)# buffer-delete 3,5,6-8
```

The **buffer-delete** command deletes the unwanted configuration from the buffer.

Step 6 Commit the configuration in the switch profile on both switches.

```
N5k-1(config-sync-sp-import)# commit

N5k-2(config-sync-sp-import)# commit
```

Step 7 Add the sync peer on both switches.

Note When importing a configuration, use the **sync-peers** command after you import configurations on both switches independently.

```
N5k-1# config sync
N5k-1(config-sync)# switch-profile Test
N5k-1(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-2>

N5k-2# config sync
N5k-2(config-sync)# switch-profile Test
N5k-2(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N5k-1>
```

Caution When you remove a switch profile using the **no switch-profile** *name* **[all-config | local-config]** command, the configuration in the switch profile is immediately removed from the running configuration. This disrupts the configurations that were present in the switch profile, for example port channel and vPC configurations. For current information about this issue, refer to CSCtl87240 and CSCtl87260 in the *Cisco Nexus 5000 Series Switch and Cisco Nexus 2000 Series Fabric Extender Release Notes* and in the Cisco Bug Toolkit located at the following URL: http://tools.cisco.com/Support/BugToolKit/action.do?hdnAction=searchBugs.

Reloading a Cisco Nexus 5000 Series Switch

In this deployment, the N5k-2 switch reboots and a new configuration was committed on N5k-1 using a switch profile.

Example 4-4 shows the configuration that was synchronized between the peers prior to the N5k-2 reload.

Example 4-4 Synchronized Configuration for Peer Switches Prior to the N5k-2 Reload

N5k-1 Running Configuration	N5k-2 Running Configuration
`interface Ethernet100/1/11` `  switchport mode trunk`	`interface Ethernet100/1/11` `  switchport mode trunk`

This example shows the configuration change that was made on the N5k-1 during the N5k-2 reload:

Note If the peer is unreachable once the commit is issued (for example, on the N5k-1 switch), the configuration is applied locally.

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# int ether 100/1/11
N5K-1(config-sync-sp-if)# switchport trunk allowed vlan 5,6
N5K-1(config-sync-sp)# commit

Verification successful...
Proceeding to apply configuration. This might take a while depending on
amount of configuration in buffer.
Please avoid other configuration changes during this time.
Commit Successful

N5K-1(config-sync)# show run int ether 100/1/11

interface Ethernet100/1/11
  switchport mode trunk
  switchport trunk allowed vlan 5-6
```

This example shows how to display the vPC consistency parameters:

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# int ether 100/1/11
N5K-1# show vpc consistency-parameters int ether 100/1/11

    Legend:
        Type 1 : vPC will be suspended in case of mismatch
```

```
Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Allowed VLANs               -     5-6                    1-3967,4048-4093
Local suspended VLANs       -     -                      -
```

Synchronizing the Peer Switches After a Switch Reload

To synchronize the configurations on the peer switches after one of the peer switches reloads, follow these steps:

Step 1 Reapply the configurations that were changed on N5k-1.

```
N5K-2(config-sync)# switch-profile Test
N5K-2(config-sync-sp)# interface ethernet100/1/11
N5K-2(config-sync-sp-if)# switchport trunk all vlan 5-6
```

Step 2 Enter the **commit** command on N5k-2.

```
N5K-2(config-sync)# switch-profile Test
N5K-2(config-sync-sp)# commit
Verification successful...
Proceeding to apply configuration. This might take a while depending on
amount of configuration in buffer.
Please avoid other configuration changes during this time.
Commit Successful
```

Step 3 Verify that the configuration is applied correctly and is synchronized on the peers.

```
N5K-2# show vpc consistency-parameters int ether 100/1/11
 Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Allowed VLANs               -     5-6                    5-6
Local suspended VLANs       -     -                      -


N5K-2# show run int ether 100/1/11

!Command: show running-config interface Ethernet100/1/11

interface Ethernet100/1/11
  switchport mode trunk
  switchport trunk allowed vlan 5-6
```

Note All configurations are applied serially in a best-effort fashion when the FEX comes online.

vPC Peer-Link Failures

When there is a peer-link failure and both switches are operational, the secondary switch shuts down its vPC ports. In a FEX active/active topology, this situation disconnects the active/active FEX on the secondary switch. If the switch profile configuration is changed on the primary switch, the configuration will not be accepted on the secondary switch unless the active/active FEX is pre-provisioned. We recommend that you pre-provision all active/active FEXs when using configuration synchronization.

Note Even if the FEXs have been originally configured in configuration terminal mode and they are operational, you should provision the FEXs in the switch profile to qualify as a provisioned FEX.

In this topology, FEX 100 is provisioned and FEX 101 is not provisioned, and both FEX 100 and 101 are already operational.

Example 4-5 shows the sample running configuration that is present for FEX 100 (which is in an operational state).

Example 4-5 Running Configuration for FEX 100

Running Configuration for FEX 100

```
fex 100
    pinning max-links 1
    description "FEX0101"

interface port-channel200
    switchport mode fex-fabric
    vpc 200
    fex associate 100

interface Ethernet1/10
    description connects to Fex 100
    fex associate 100
    switchport mode fex-fabric
    channel-group 200
```

The next step is to provision FEX 100 inside the switch profile.

Example 4-6 shows the running configuration when FEX 100 is provisioned.

Example 4-6 Running Configuration for FEX 100 Provisioned

Running Configuration for FEX 100 with Provisioning

```
fex 100
    pinning max-links 1
    description "FEX0101"

slot 100
    provision model N2K-C2148T

interface port-channel200
    switchport mode fex-fabric
    vpc 200
    fex associate 100

interface Ethernet1/10
    description connects to Fex 100
    fex associate 100
    switchport mode fex-fabric
channel-group 200
```

Working Example

This example shows how to provision the FEX:

```
N5K-1(config-sync-sp)# slot 100
```

```
N5K-1(config-sync-sp-slot)# provision model N2K-C2148T
N5K-1(config-sync-sp-slot)# exit
N5K-1(config-sync-sp)#
N5K-1(config-sync-sp)# commit
```

This example shows that the vPC peer link fails:

```
N5K-1(config-sync-sp)# sh vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 10
Peer status                       : peer link is down
vPC keepalive status             : peer is alive
```

This examples shows that in the switch profile, a configuration is added under Ethernet 100/1/1:

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# int ether 100/1/1
N5K-1(config-sync-sp-if)# switchport mode trunk
N5K-1(config-sync-sp-if)# commit
Verification successful...
Proceeding to apply configuration. This might take a while depending on
amount of configuration in buffer.
Please avoid other configuration changes during this time.
Commit Successful
```

This example shows how to verify that both switches are synchronized:

```
N5K-1(config-if)# sh run int ether 100/1/1
interface Ethernet100/1/1
  switchport mode trunk

N5K-2(config-if)# sh run int ether 100/1/1
interface Ethernet100/1/1
  switchport mode trunk
```

Nonworking Example

Example 4-7 shows the running configuration for FEX 101 that is not provisioned inside the switch profile.

Example 4-7 Running Configuration for FEX 101

Running Configuration for FEX 101

```
fex 101
    pinning max-links 1
    description "FEX0101"

interface port-channel201
    switchport mode fex-fabric
    vpc 201
    fex associate 101

interface Ethernet1/11
    description connects to Fex 101
    fex associate 101
    switchport mode fex-fabric
    channel-group 201
```

This example shows that the vPC peer link fails:

```
N5K-1(config-sync-sp)# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 10
Peer status                       : peer link is down
vPC keepalive status              : peer is alive
```

This example shows that configuration changes made on N5k-1 for Ethernet 101/1/1 fail because FEX 101 is not provisioned inside the switch profile:

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# int ethernet 101/1/1
N5K-1(config-sync-sp-if)# switchport mode trunk
N5K-1(config-sync-sp-if)# commit
Verification successful...
Failed to Commit: Commit Failed
```

This example shows how to correct the issue by provisioning FEX 101 inside the switch profile. If FEX 101 is not provisioned inside the switch profile, the configuration changes must be done manually on both switches:

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# slot 101
N5K-1(config-sync-sp-slot)# provision model N2K-C2148T
N5K-1(config-sync-sp-slot)# commit
```

This example shows how to make the same configuration change again:

```
N5K-1(config-sync)# switch-profile Test
Switch-Profile started, Profile ID is 1
N5K-1(config-sync-sp)# int ether 101/1/1
N5K-1(config-sync-sp-if)# switchport mode trunk
N5K-1(config-sync-sp-if)# commit

Verification successful...
Proceeding to apply configuration. This might take a while depending on
amount of configuration in buffer.
Please avoid other configuration changes during this time.
Commit Successful
```

```
!Command: show running-config interface Ethernet101/1/1

interface Ethernet101/1/1
  switchport mode trunk

!Command: show running-config interface Ethernet101/1/1
!Time: Tue Oct 19 01:18:14 2010

version 5.0(2)N1(1)

interface Ethernet101/1/1
  switchport mode trunk
  speed 1000
```

mgmt0 Interface Connectivity is Lost

Configuration synchronization sends switch profile configurations over the mgmt0 interface to the peer switch. When the mgmt0 interface connectivity is lost and the configuration needs to be changed, configure the switch profile on both switches. When the mgmt0 interface is restored, both switches become synchronized.

Note If you make configuration changes when the mgmt0 interface is down, the configurations that are applied on each switch must be identical. If the configurations are not identical, when the mgmt0 interface comes up and you enter a **commit** command on either switch, the commit fails because of a configuration mismatch.

If you enter the **commit** command when the mgmt0 interface is up and then the mgmt0 interface goes down, the commit eventually fails when both switches detect that the peer switch is no longer reachable from the mgmt0 interface.

Rollback Failures with Conditional Features

With configuration synchronization, when a conditional feature is present in a checkpoint and not in the running configuration, a rollback to that checkpoint fails. As a workaround, you can reconfigure the conditional feature before a rollback is executed. The workaround applies to the **vpc** *domain* and **peer-keepalive** commands in vpc-domain mode.

This example shows the running configuration of the system when a checkpoint called chkpt is created:

```
feature vpc

vpc domain 100
    vpc peer-keepalive destination 10.0.0.1

interface Ethernet 1/1
    switchport mode trunk
    channel-group 100

switch-profile Test
    interface port-channel 100
        switchport mode trunk
        vpc peer-link
```

If you perform a write-erase at this point and you reload the switch and attempt to perform a rollback to the checkpoint chkpt, the rollback fails. This example shows a rollback failure when this situation occurs:

Note To avoid the rollback failure, preconfigure the **feature** *vpc*, **vpc domain**, and **peer-keepalive** command before performing the rollback.

```
N5k-1# rollback running-config checkpoint chkpt verbose
```

Note Applying a configuration in parallel might cause a rollback verification to fail.

```
Collecting Running-Config
Generating Rollback patch for switch profile
Executing Rollback patch for switch profiles. WARNING - This will change the configuration
of switch profiles and will also affect any peers if configured
==========================================================
`config sync `
`switch-profile Test`
Switch-Profile started, Profile ID is 1
`interface port-channel100`
`switchport mode trunk`
Syntax error while parsing 'vpc peer-link'

==========================================================
Generating Running-config for verification

Verification failed, rolling back to previous configuration
Collecting Running-Config
2010 Oct 14 07:43:12 switch %$ VDC-1 %$ %ASCII-CFG-2-ACFG_OPER_FAIL: Operation failed
because of Rollback Patch is not Empty
…
```

Channel Group Failures

The **channel-group** command fails for port profiles or pre-provisioned interfaces if the port channel does not exist (auto-creation is not supported). The workaround is to explicitly create the port channel first using the **interface port-channel** *xxx* command.

Note Port-channel members must be configured in configuration terminal mode.

Nonworking Example

This example shows the error message that appears when the port-channel interface is not created first:

```
N5K-1(config-if-range)# int ether 100/1/2-3
N5K-1(config-if-range)# channel-group 200
Pre-provisioned interface: port channel must exist first
```

The **channel-group** command fails when a module comes online if you make a configuration change on a port channel but not on the pre-provisioned interfaces. The failure does not occur in Cisco NX-OS Release 5.0(2)N2(1) which supports the **channel-group** xxx **force** command.

At-A-Glance Configuration Modes

Table 4-3 shows which mode is used to configure features and interfaces in an active/active or straight-through topology and for the switches and hosts in a vPC topology.

For example, to configure a port-channel interface in an active/active topology, use the switch profile configuration mode.

Table 4-3 At-A-Glance Topology Configuration Modes

Configuration	FEX Active/Active Topology	Straight-Through Topology	Switch / Host vPC Topology
Nonport-channel member	Switch profile mode	Switch profile mode	Switch profile mode
Port-channel interface (interface port channel command)	Switch profile mode	Switch profile mode	Switch profile mode
Port-channel member interfaces	Configuration terminal mode	Configuration terminal mode	Configuration terminal mode
FEX/GEM pre-provisioning	Switch profile mode	Configuration terminal mode	Configuration terminal mode
Global configurations (for example, port profiles, QoS, ACLs, and so on)	Switch profile mode	Switch profile mode	Switch profile mode

Terminology

- Port-channel interface—Interface that is part of a port channel.
- Nonport-channel member—Stand-alone interface that is not part of any port channel.
- Port-channel member—Interface that is a member of a port channel.
- Switch profile—Predefined configuration profile that is used to synchronize a consistent configuration across peer switches. Switch profiles are used in the configuration synchronization feature.
- Port profile—Interface-command profile that can be applied to a range of interfaces (for example, Ethernet, VLAN network interface, or port channel).
- Configuration synchronization—Feature that uses a switch profile to synchronize consistent configurations between two peer switches.
- config-sync mode—Configuration mode that is used to define and access a switch profile.
- Configuration terminal mode (config-t)—Configuration mode used to commit configurations locally on a switch.
- Pre-provisioning—Ability to configure offline interfaces before they are connected (or brought online). Pre-provisioning can be done on Cisco Nexus 2000 Fabric Extenders (FEXs) and/or Generic Expansion Module (GEMs).

CHAPTER 5

Fibre Channel over Ethernet Operations

This chapter includes the following sections:

Introduction

The Cisco Nexus 5000 Series switch has supported FCoE since 2009. As the adoption of FCoE increases within the data center, there are design and operational considerations to take into effect. This document discusses these considerations and provides operational guidelines on how to deploy and implement an FCoE solution with Cisco Nexus 5000 Series switches.

FCoE Considerations

This section includes the following topics:

Preserving SAN Fabric Isolation

High availability (HA) is a requirement in any data center design—whether it is accomplished through HA at the port level, supervisor level, or even at the physical network level. Fibre Channel Storage Area Networks (FC SANs) achieve high availability by building out two identical but physically separate networks commonly referred to as SAN A and SAN B (also called Fabric A and Fabric B). These networks, unlike Data Center LAN networks, are completely physically isolated from one another and have no knowledge of each other. Depending on host operating systems and drivers, traffic is able to be load balanced or "multi-pathed" between the two isolated networks, from the application side, in order to provide better service to the storage traffic. This required isolation is an important element in building FCoE networks along side the data center Ethernet LANs.

This section includes the following topics:

Maintaining Different FC-MAPs Per Fabric

FC-MAP is a characteristic of a FCoE switch that identifies which fabric the switch belongs to. For instance, there can be an FC-MAP for Fabric A and a different FC-MAP for Fabric B. By configuring a specific FC-MAP value on a FCoE switch, it is possible to designate certain switches to belong to one fabric or another.

In order to maintain fabric isolation in an FCoE environment, it is recommended to use different FC-MAP values per SAN Fabric. Because the FC-MAP value of the Cisco Nexus 5000 Series switch is used in the addressing for FCoE-enabled devices, changing the FC-MAP value is a disruptive process to all hosts that are logged into the switch. Due to this disruption, it is recommended that the FC-MAP is configured as part of the initial switch set up.

By default, when the **feature fcoe** command is used to enable FCoE on a Cisco Nexus 5000 Series switch, a default FC-MAP is assigned to the switch. The simplest way to ensure SAN A and SAN B isolation between FCoE-enabled switches in the Ethernet fabric is to change the FC-MAP value to something other then the default for all switches belonging to Fabric B. This will prohibit FCoE switches from joining the wrong fabric and aide to providing the SAN isolation that is a requirement for FC and FCoE traffic.

To change the FC-MAP of a switch:

```
switch# configure terminal
switch(config)# fcoe fcmap 0e.fc.2a
```

Note Changing the FC-MAP value of a switch is disruptive to all attached FCoE hosts and it requires the hosts to login to the fabric again. Therefore, it is recommended to change the FC-MAP when the switch is installed and initially configured or during a maintenance window.

Note The default value of the FC-MAP on a Cisco Nexus 5000 Series switch is 0E.FC.00. The configurable values for FC-MAP ranges from OE.FC.00 to 0E.FC.FF.

VLAN to VSAN Numbering

When configuring an FCoE fabric, the first step is to create a VLAN to VSAN mapping which allows the FC traffic in a single VSAN to traverse the Ethernet network. It is a best practice to have dedicated VLANs for FCoE traffic in order to separate the storage traffic from all other Ethernet VLANS. It is also recommended not to assign VLAN 1, VSAN 1 or the configured native VLAN to the FCoE network. Typically those VLAN/VSANs are utilized for management traffic or for devices that have no other VLAN or VSAN assigned to them. Using VLAN 1 as an FCoE VLAN will not be supported on the Cisco Nexus 5000 Series switch running Cisco NX-OS release 5.0(1)N1(2) or a later release.

VLAN to VSAN mapping is a one-to-one relationship. Mapping multiple VSANs to a single VLAN instance is not supported. Note that both the VLAN instance and VSAN instance in an FCoE VLAN/VSAN mapping take up a hardware VLAN resource. Currently, there can be up to 31 VLAN/VSAN mappings supported on the Cisco Nexus 5000 Series switch. VLAN and VSAN numbering can range from 1-4096.

FCoE VLANS are different from typical Ethernet VLANs in that it acts more of a container for the storage traffic than anything else. MAC learning, broadcasts, or flooding do not occur and it does not map to a subnet. FCoE VLANs are simply used to carry the traffic for a specified FC VSAN and keep it separate from any other Ethernet VLANs that may be traversing the network.

In order to avoid confusion and service disruption in the event of a misconfiguration, it is recommended that you configure different FCoE VLAN and VSAN numbers for both SAN A and SAN B. Using the same VLAN or VSAN numbering between the two fabrics could result in the merging of both SAN fabrics in the event of a miss-configuration or miss-cabling. It is also best practice to only define SAN A VLANs on SAN A switches and vice-versa.

Host-facing FCoE ports must be configured as trunk ports carrying the native VLAN, FCoE VLAN and any other Ethernet VLANs necessary for the host application. These host facing ports should also be configured as spanning tree edge ports using the **spanning-tree port type edge** [**trunk**] interface-level command.

Note

- FCoE Initialization Protocol (FIP) uses the native VLAN and therefore all FCoE links should be trunked to carry the FCoE VLAN as well as the native VLAN.
- The FCoE VSAN must be configured and in the VSAN database of the Cisco Nexus 5000 Series switch prior to mapping it to a VLAN
- Enabling FCoE on VLAN 1 is NOT supported

FCoE and Spanning Tree Protocol Considerations

Native FC has no concept of a looped environment and therefore has no need for a protocol similar to the Spanning Tree Protocol (STP) in the Ethernet world. However, when placing FCoE onto an Ethernet fabric, STP is run on the FCoE VLANs connecting to a host (VF port) over a lossless Ethernet cloud. This lossless cloud could be made up of DCB bridges or FIP snooping devices. Because of this, there are certain recommendations for STP configurations that should be followed when deploying FCoE. The goal is to have isolated STP topologies between SAN A, SAN B, and the Ethernet fabric. This eliminates any Ethernet topology changes from affecting storage traffic.

Note STP is not run on FCoE VLANs on VE port connections between two FCFs.

Note Beginning with Cisco NXOS Release 5.0(1)N1(1) for the Cisco Nexus 5000 Series switch, STP is not run on FCoE VLANs on VF ports connecting directly to attached hosts (including host connections to a Cisco Nexus 2232 Fabric Extender). STP will continue to run on VF ports that connect to hosts through a DCB cloud or FIP snooping device.

Note In Cisco NXOS Release 4.2(1)N2(1a) and earlier releases, STP runs on FCoE VLANs for any VF port connection (either direct attached hosts or hosts connected over a DCB cloud). Because of this, it is required to configure the VF port as a spanning-tree port type edge trunk

This section includes the following topics:

MST Instances For Dual Fabric FCoE Deployments

When running multi-instance STP in an Ethernet environment, it is required that all switches in the same MST region have the identical mapping of VLANs to instances. This does not require that all VLANs be defined on all switches. When running FCoE over an environment using MST, it is recommended to have a dedicated MST instances for the FCoE VLANs belonging to SAN A and a dedicated MST instance for the FCoE VLANs belonging to SAN B. These instances should be separate from any instances that include regular Ethernet VLANs. This example shows the FCoE VLANS in Fabric A are VLANs 20-29 and the FCoE VLANS in Fabric B are VLANs 30-39:

Spanning-tree MST configuration:

- name FCoE-Fabric
- revision 5
- instance 5 vlan 1-19,40-3967,4048-4093
- instance 10 vlan 20-29
- instance 15 vlan 30-39

In the above configuration, instance 5 maps to native Ethernet VLANs, instance 10 maps to the VLANs for Fabric A (20-29) and instance 15 maps to the VLANs for Fabric B (30-39).

Due to the MST configuration requirement, it will be necessary to have the same MST configuration, containing both the SAN A and SAN B instance, on all switches within the same MST region. This means that switches participating in SAN A will also contain an MST configuration with a separate instance for SAN B VLANs even though those SAN B VLANs will not be defined on the SAN A switch.

PVST+ for Dual Fabric FCoE Deployments

When running PVST, each VLAN already has its own spanning tree topology. Because FCoE traffic in each SAN fabric is defined by different individual VLANs, PVST+ will automatically isolate the spanning tree domains for the VLANs in SAN A, SAN B, as well as the Ethernet fabric.

FCoE and Virtual Port Channel (vPC) Considerations

Virtual Port Channeling (vPC) is an Ethernet feature that allows a single device to connect to multiple upstream devices and forward out all available links without the implications of spanning tree blocking paths due to Ethernet loops. vPC is useful in three situations:

1. Connecting a server to two upstream switches
2. Connecting a FEX to two upstream Nexus 5X00s
3. Connecting a switch to two upstream switches

The upstream switches in all scenarios must support the virtual port channel feature. The downstream device has no knowledge of the vPC and simply views the connection as a standard Ethernet port channel.

Though it is not possible to run FCoE traffic on top of a vPC because of the SAN A and SAN B physical isolation requirement in native FC, it is possible to run FCoE and vPC side-by-side on the same physical infrastructure from the host to the first-hop FCoE device. To configure this topology, the following must be considered:

- A host must connect to the upstream Cisco Nexus 5000 Series vPC pair switches using only 2 10G links – one attaching to a Cisco Nexus 5000 Series switch in Fabric A and one attaching to a Cisco Nexus 5000 Series switch in fabric B. This is commonly referred to a *single-port vPC* because only one port goes to each switch.
- Generation 2 CNAs are required in the host in order to support vPC topologies.

Note

- FCoE and vPC can run side-by-side only on single-port host-connected vPCs. FCoE and vPC's between a FEX and a Cisco Nexus 5000 Series switch or between two layers of switches is not supported.
- FCoE and vPCs containing more that one link to each access device is not supported. vPCs which coexist with FCoE must contain only a single link to each vPC peer device.
- vPC's across switches (FCFs) within the same SAN fabric is not supported. Each vPC peer must be part of different fabrics—one peer in SAN A and one peer in SAN B.

This section includes the following topics:

Required Teaming Drivers for vPC With CNAs

When connecting a host to an upstream vPC switch pair, the only requirement from the host side is to support link aggregation on the NIC interfaces. This can be accomplished using link aggregation control protocol (LACP) or standard 802.3ad *port channel mode on* behavior. It is important to check that either the host operating system or the native CNA hardware supports one of these options.

Second Generation CNA Requirement

When connecting a host containing a CNA to upstream Cisco Nexus 5000 Series switches configured in a vPC, 2nd generation CNAs are required from both Emulex and QLogic. This is regardless of the presence of FCoE traffic on the host connections. These 2nd generation CNAs are also required when connecting to a Cisco Nexus 2232 Fabric Extender with a vPC (Ethernet only), FCoE, or FCoE+vPC configuration from a host connection.

View Of Ethernet Traffic And FC Traffic Through A CNA

Currently CNAs present two different types of adapters to the host operating system: Ethernet NICs and Fibre Channel HBAs. Though these adapters physically correspond to the same 10GE port on a CNA, to the operating system, it will appear as two completely separate and physically isolated interfaces. Because of this adapter port virtualization, it is possible to build two separate topologies based on traffic type: one for the Ethernet fabric using the NICs and one for the FC fabric using the HBAs.

For FCoE and vPC to run side-by-side from the host, the port channel would be configured on the NICs interfaces presented and SAN multi-pathing or other SAN HA mechanisms would be configured on the FC HBAs presented to the OS by the CNA. Today, it is required that only 2X10GE links be used in a host side vPC port channel when running FCoE on the same wires. Each 10GE link will be used to provide a single connection to each upstream vPC switch.

Figure 5-1 Ethernet And FC Traffic Through A CNA

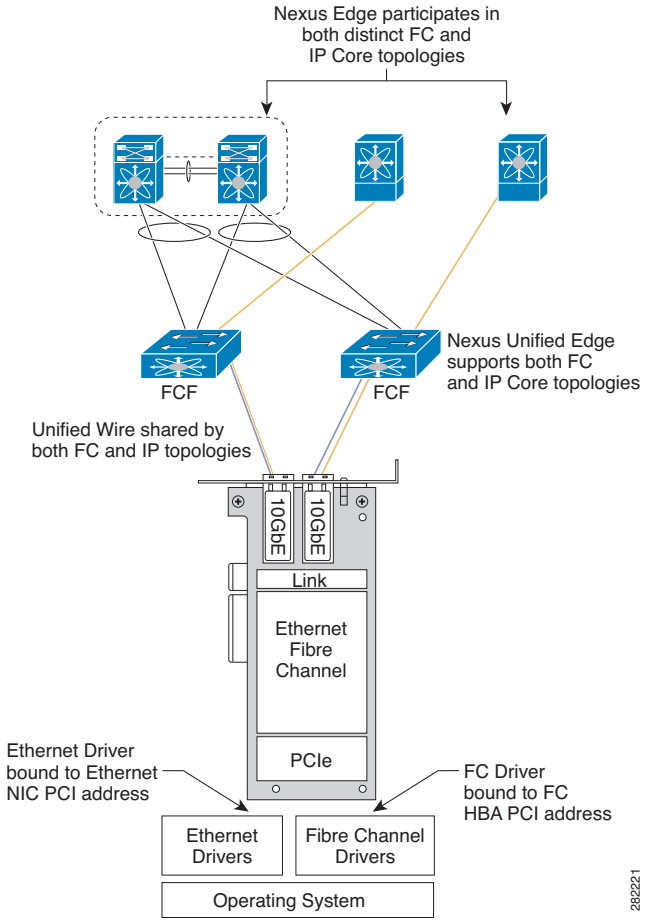

Figure 5-2 Adapter Control Panel Display

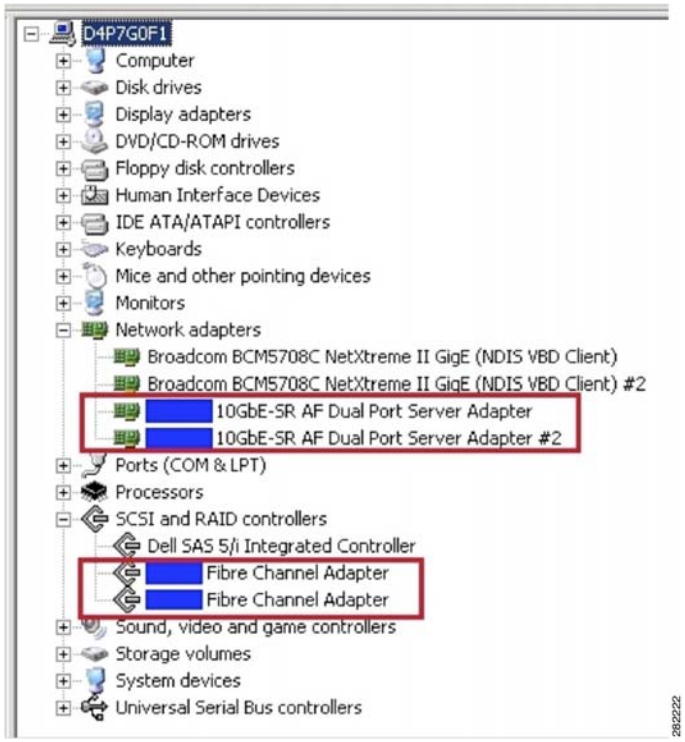

Note

- VPC + FCoE over a consolidated wire from the host requires the host supports port channels capabilities (LACP or “port channel mode ON”). Please check with specific CNA and OS vendors for a support matrix.
- VPC + FCoE over a consolidated wire are only supported between a host and either the first hop Nexus 5000 or Nexus 5000/2232 pair. VPC and FCoE on a consolidated wire is NOT supported beyond the access layer or when connecting a host to the Nexus 7000 platform.
- vPC and FCoE can not coexist on the same wire beyond any first hop access device.

FCoE VLAN Configuration On A vPC

Typically, interfaces belonging to the same port channel are required to have the same port configuration. This includes VLAN configuration. However, in order to maintain fabric separation alongside vPC connections, it is necessary to declare the FCoE VLAN for SAN A on one uplink and the FCoE VLAN for SAN B on the other uplink. This is a recommended best practice configuration. Figure 5-3 shows the hosts connected to a Cisco Nexus 5000 Series switch running vPC and FCoE simultaneously.

Figure 5-3 FCoE VLAN Configuration In A vPC Topology

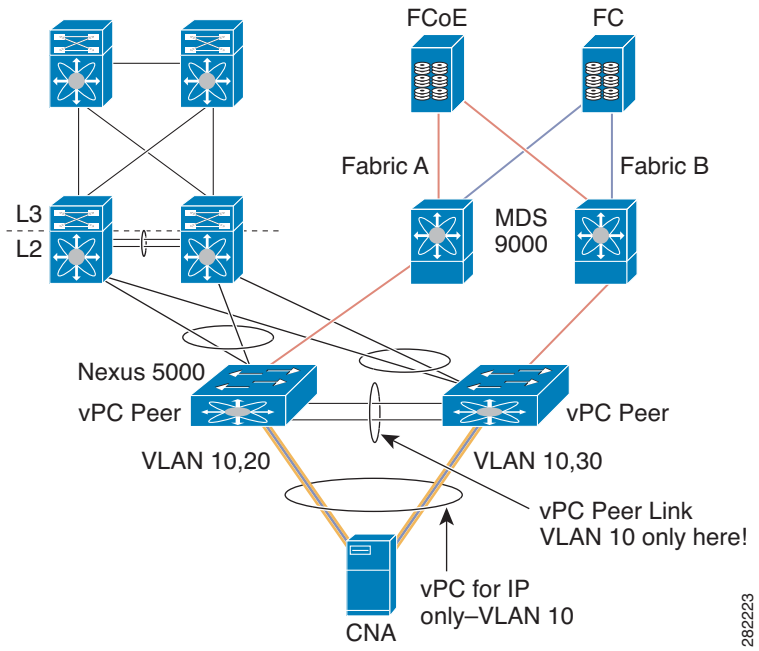

Changing Buffer Allocations for Longer Distance FCoE

Beginning with the Cisco NXOS Release 5.0(1)N1(1) for the Cisco Nexus 5000 Series switch, it is possible to tune the port buffer allocation and xon and xoff thresholds in order to support increased distance between VE ports. The default distance configured for each port when configured to carry FCoE traffic (or any “no-drop” traffic) is 300 meters. This supported distance is based on the amount of available buffer space allocated to catch frames in flight between the time a PAUSE is initiated towards a downstream device and the time that downstream devices processes that PAUSE frame and stops sending frames. This per port buffer allocation and configuration must match between the two ports on either end of the link (including host CNA ports as well). This is similar to the way buffer-to-buffer credits is initialized between two devices in a native FC environment.

The current xon threshold and buffer size allocated for FCoE is such that buffer-size - xon = ~300 meters worth of FCoE frames. The default configuration parameters for the class-fcoe (or any no-drop class) on the Nexus 5000 series switch is shown below:

- qos-group 1
- q-size: 76800, HW MTU: 2400 (2240 configured)
- drop-type: no-drop, xon: 128, xoff: 240

In order to support a distance of 3000m for FCoE traffic between two FCoE capable switches (connecting two FCFs with VE ports), the buffer allocation as well as the xon and xoff values need to be altered for the FCoE class of services: class-fcoe. This can be accomplished by editing the quality of service configuration. An example of this configuration can be found in the “Configuring NO-Drop Buffer Threshold” section of the Nexus 5000 Configuration Guide:

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/502_n1_1/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_Rel_502_N1_1.pdf

The necessary thresholds for support no-drop service up to 3000m is outlined in the table below:

Configuration for 3000m no-drop class	Buffer size	Pause Threshold (XOFF)	Resume Threshold (XON)
Nexus 5000 Series	143680 bytes	58860 bytes	38400 bytes
Nexus 5500 Platform	152000 bytes	103360 bytes	83520 bytes

Consolidated Links And Dedicated Links for FCoE

Because FCoE uses the Ethernet fabric for transport, there is the possibility of consolidating both Ethernet LAN traffic and Storage SAN traffic onto the same infrastructure. There are multiple levels of consolidation; wire consolidation and device consolidation are two of the most common and are discussed below.

Link Consolidation refers to when Ethernet LAN traffic and Storage SAN traffic are sharing the same physical wire between host and switch or between two switches.

Device consolidation refers to when Ethernet LAN traffic and Storage SAN traffic are passing through the same switching device but maintain isolation through the use of dedicated wires or switch ports.

The topologies discussed throughout this guide will mention two terms to describe the scope of FCoE traffic: consolidated link – where FCoE and native Ethernet traffic simultaneously use the same link -- and dedicated link – where FCoE and native Ethernet traffic use two separate DCB Ethernet links. The following sections will discuss the different places in the Data Center Network where consolidated and dedicated links make sense.

Figure 5-4 shows an example of consolidated vs dedicated links. The wires running from the host to the access devices are consolidated links carrying both Ethernet and FCoE traffic. Moving from the access to aggregation, there are dedicated links: blue wires dedicated to the Ethernet traffic and orange wires dedicated to FCoE traffic only.

Figure 5-4 Consolidated And Dedicated Links

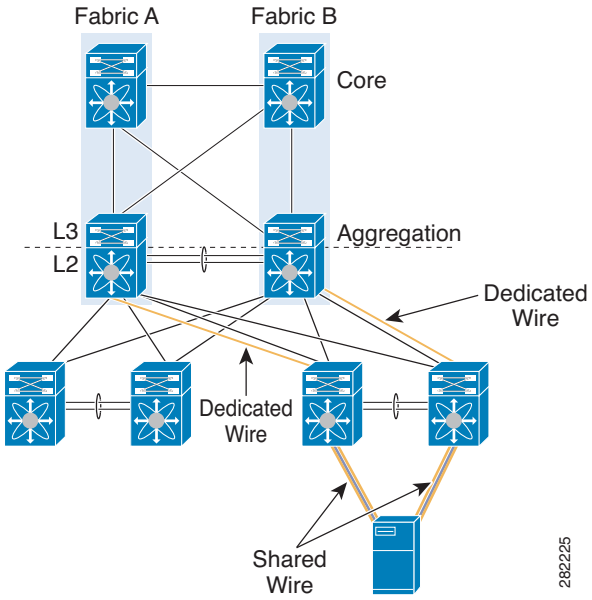

This section includes the following topics:

Where Consolidated Links Makes Sense

One of the benefits of FCoE at the access layer is the ability to consolidate the FC SAN and Ethernet LAN onto the same physical wires and same physical devices. This consolidation lends to a large CapEx savings by reducing the number of access switches, host adapters, cables and optics required to run both LAN and SAN networks within the data center. This consolidation is made possible due to the excess bandwidth that 10GE to the server is able to provide. Because very few servers in the Data Center today are pushing 10-Gigabit Ethernet of Ethernet-only traffic, there is room for the added storage traffic to share these common wires without impacting the performance of the host application.

Also, due to the CNA behavior and ability to present to the host application different physical devices corresponding to both LAN and SAN networks, it is possible to separate Ethernet HA from FC HA at the host level. This is accomplished by being able to use separate Ethernet teaming options on the NICs while using separate FC multi-pathing options on the HBAs. Depending on the operating system and CNA being used, these teaming options will vary.

Moving beyond the access layer, oversubscription ratios and Ethernet bandwidth provisioning will determine the amount of excess bandwidth available and the benefit of running consolidated links vs. dedicated links within the Data Center.

Where Dedicated Wires Makes Sense

High Availability requirements in LAN and SAN networks differ considerably. Where in Ethernet, HA is achieved by multi-homing devices to one another (partial/full Mesh), in Fibre Channel (and FCoE), HA is achieved by building two physically isolated networks. Both of these requirements must be me in a network that combines FCoE and Ethernet.

There have been multiple enhancements to the Ethernet HA model that improves on Ethernet Data Center design by overcoming some of the challenges of the Spanning Tree protocol. One example of this is the virtual Port Channeling feature found in the Nexus product suite. The nature of vPC is to be able to forward out multiple paths to multiple upstream devices without spanning tree blocking any of the uplinks. While this is great for Ethernet traffic, it breaks the SAN A/SAN B isolation required for FC/FCoE.

Therefore, it is often beneficial to use dedicated wires for Ethernet traffic and Storage traffic independently. With dedicated wires, the Ethernet links can be configured to take advantage of advance Ethernet features such as vPC and the storage links can be configured based on the fabric isolation requirement. This is especially common when connected access switches to upstream LAN aggregation/SAN core devices.

Cisco Nexus 5000 Series Switch FCoE Considerations

The Cisco Nexus 5000 Series switches include a Unified Port Controller (UPC) ASIC responsible for the handling the forwarding decisions and buffering for multiple 10-Gigabit Ethernet ports:

- The Cisco Nexus 5000 Platform switches (the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch) include the first generation UPC ASIC.
- The Cisco Nexus 5500 Platform switches (the Cisco Nexus 5548P switch, Nexus 5548UP switch, and Nexus 5596UP switch) include the second generation UPC ASIC.

The following sections discuss the differences between the first and second generation architectures that relate to FCoE configuration and supported topologies.

This section includes the following topics:

VLAN Scalability

One of the differences between the first and second generation ASICs is the number of available VLAN resources available. The first generation ASICs support up to 512 VLANs (507 of which are user configurable). With the second generation ASIC, the available VLAN number has increased from 512 to 4096. Currently, 31 VLANs and 31 VSANs are supported for FCoE VLAN/VSAN mappings on both generations.

Note The VLAN and the VSAN in an FCoE VLAN/VSAN mapping consume a hardware VLAN resources.

FCoE QoS Configuration

The Nexus 5000 Series switches always reserve some buffer space for FCoE traffic. When you enable the FCoE feature on Nexus 5000 Series switch, Nexus automatically configures the necessary QoS policy and buffer allocations using the reserved buffers.

The Nexus 5500 Series switches allow all available port buffers to be configured based on traffic needs. This allows you to create a custom FCoE policy that can use any available buffers.

When you enable FCoE on a Nexus 5500 Series switch, the system looks for a custom QoS policy. If it does not find one, it automatically uses the default QoS configuration shown below:

```
switch(config-sys-qos)# service-policy type qos input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing output fcoe-default-out-policy
switch(config-sys-qos)# service-policy type network-qosfcoe-default-nq-policy
```

For more information, see the Cisco Nexus 5000 Series NX-OS Quality of Service Configuration Guide, which is available from:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Unified Port Options

Unified ports are capable of operating a 1- and 10-Gigabit Ethernet or 1-, 2-, and 4-Gigabit or 2-, 4-, and 8-Gigabit FC (depending on the transceiver used) which provide more configuration flexibility. Unified ports no longer require you to purchase a set number of FC ports through an expansion module. Unified Ports are available in the expansion module on the Cisco Nexus 5548P switch and the Nexus 5548UP platform as well as on all base ports of the Cisco Nexus 5596UP switch. There are configuration requirements that must be carefully followed when utilizing unified ports.

Ports of a similar type, either Ethernet or FC, must be configured in a contiguous sequence. Changes to the port-type require a switch reboot or expansion module reboot depending on where the unified ports are configured. For this reason, careful planning should be done when first configuring the switch. Cisco

recommends as a best practice to start Ethernet port configurations at one end of the platform (from Eth1/1 counting up) and the necessary Fibre Channel ports configured from the opposite end of the platform (Eth 1/48 counting down).

For additional information on configuring unified ports, see the *Unified Port Configurations on Cisco Nexus 5500 Platform Switches* documentation.

Priority Flow Control and Enhanced Transmission Selection Considerations

Both Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS) are part of the IEEE 802.1Q Enhance Ethernet Standards that are currently in the final stages of standardization. Both PFC and ETS are support on all Cisco Nexus 5000 Series switches. PFC is class of service (COS) based PAUSE allowing for FCoE traffic assigned to a specific COS value to retain the lossless qualities which are required for the FC protocol. ETS is a mechanism for dividing a 10-Gigabit Ethernet link into multiple lanes based on the COS value and allocating the necessary bandwidth requirements which are honored in the presence of congestion. ETS prevents situations where default traffic would interfere with higher priority traffic.

PFC and ETS are often used in today's FCoE networks to provide lossless transport and dedicated bandwidth for FCoE traffic. However, they are not specific to FCoE and have many uses outside of an FCoE environment for providing specific levels of service to specific traffic classes.

This section includes the following topics:

Default PFC and ETS Settings

PFC and ETS both use the Class of Service (COS) bits in order to classify between traffic types. There are 8 COS values in the IEEE 802.1Q standard trunking header for Ethernet frames. The Cisco Nexus 5000 Series switch allows you to manually configure 6 classes. Up to 4 of the 6 user configurable classes can be designated as no-drop classes of service, meaning that in the event of port congestions, traffic belonging to the no-drop classes will pause to prohibit packet drop.

By default, the Nexus 5000 Platform as well as other vendor's FCoE products have decided on COS value of 3 for FCoE traffic. When FCoE is enabled on the Cisco Nexus 5000 Series switch, COS 3 is automatically configured for no-drop service (PFC setting) as well as a guarantee of 50% of the bandwidth in the case of congestion (ETS setting). It is best practice to leave the default COS value of 3 for FCoE traffic due to the agreement between vendors to support this as a "no-drop" class.

In the event that other traffic already exists within the network that is using the COS value of 3 or there is another reason to move FCoE traffic from COS 3, this can be changed through a Quality of Service configuration.

Changing PFC and ETS Settings

PFC and ETS settings are configured and changed in the Quality of Service configuration on the Nexus 5000 Series switch. This example shows a QoS configuration that changes the FCoE no-drop class of service to COS 4 as the reserved bandwidth for FCoE to 20% of the 10-Gigabit Ethernet link:

Step 1 Create classification rules first by defining and applying policy-map type qos:

```
N5k(config)# class-map type qos class-lossless
N5k(config-cmap-qos)# match cos 4
N5k(config-cmap-qos)# policy-map type qos policy-lossless
N5k(config-pmap-qos)# class type qos class-lossless
N5k(config-pmap-c-qos)# set qos-group 7
N5k(config-pmap-uf)# system qos
N5k(config-sys-qos)# service-policy type qos input policy-lossless
```

Step 2 Define and apply policy-map type network:

```
N5k(config-pmap-qos)# class type network-qos policy-lossless
N5k(config-cmap-uf)# match qos-group 7
N5k(config-cmap-uf)# policy-map type network-qos policy-lossless
N5k(config-pmap-uf)# class type network-qos class-lossless
N5k(config-pmap-uf-c)# pause no-drop
N5k(config-pmap-uf)# system qos
N5k(config-sys-qos)# service-policy type network-qos policy-lossless
```

Step 3 Create classification rules first by defining and applying policy-map type qos:

```
N5k(config)# class-map type queuing class-voice
N5k(config-cmap-que)# match qos-group 2
N5k(config-cmap-que)# class-map type queuing class-high
N5k(config-cmap-que)# match qos-group 3
N5k(config-cmap-que)# class-map type queuing class-low
N5k(config-cmap-que)# match qos-group 7
N5k(config-cmap-que)# exit
```

Step 4 Create classification rules for the individual classes:

```
N5k(config)# policy-map type queuing policy-BW
N5k(config-pmap-que)# class type queuing class-voice
N5k(config-pmap-c-que)# priority
N5k(config-pmap-c-que)# class type queuing class-voice
N5k(config-pmap-c-que)# bandwidth percent 20
N5k(config-pmap-c-que)# class type queuing class-high
N5k(config-pmap-c-que)# bandwidth percent 40
N5k(config-pmap-c-que)# class type queuing class-low
N5k(config-pmap-c-que)# bandwidth percent 10
N5k(config-pmap-c-que)# class type queuing class-fcoe
N5k(config-pmap-c-que)# bandwidth percent 30
N5k(config-pmap-c-que)# class type queuing class-default
N5k(config-pmap-c-que)# bandwidth percent 0
N5k(config-pmap-c-que)# system qos
N5k(config-sys-qos)# service-policy type queuing output policy-BW
```

Host-Side Considerations For Altering PFC And ETS Settings

Data Center Bridging eXchange (DCBX) protocol is another portion of the IEEE 802.1Q Data Center Bridging (DCB) standard currently in review by the Ethernet standards body. DCBX is a protocol that runs between DCB-capable devices to ensure that PFC and ETS settings are configured consistently between DCB peers. DCB can also be used as a way to configure DCB peer devices from a central switching location. CNAs that support DCB-*willing* are configured to accept the DCB configurations (including PFC and ETS settings) of the upstream DCB switching device. This greatly simplifies management and configuration of DCB and FCoE devices.

If changing the default configuration for FCoE traffic on the Cisco Nexus 5000 Series switch, it is possible for the switch to relay these configuration changes to any connected CNAs using the DCBX protocol. It is necessary that the CNA vendor and platform support DCBX in a *willing* mode in order for this to take place. Please check with the individual CNA vendors on whether they support receiving DCBX configurations for a network device.

If the CNA does not support a method of DCB-willing, in order to change from a default PFC and ETS configuration, it is required to manually alter the configuration of the Nexus 5000 Series as well as the downstream CNA device so that they are the same. Depending on the CNA, different tools or commands will be used to change these settings.

Note If the DCBX negotiation fails between a host and switch or between a switch and switch, the PFC setting will not be set on the Nexus 5000 Series switch and the vFC interfaces will remain down until the DCB configuration matches.

Note Though the DCBX standard states that there are 8 possible no-drop lanes, CNA vendors differ on the number of COS values that are supported for FCoE and no-drop service today. Check with the CNA vendor for the correct number of supported FCoE and no-drop classes.

Cisco Nexus Interoperability

For information on interoperability, see the *Cisco Data Center Interoperability Support Matrix*.

FCoE Supported Topologies

This section includes the following topics:

Single-Hop FCoE Deployment Topologies

There are two possible single-hop solutions when deploying FCoE with a Cisco Nexus 5000 Series switch and Cisco Nexus 2000 Series Fabric Extender. The first solution is referred to as “direct connect” where a host is directly connected to the first hop converged access switch. The second single hop solution deploys a FEX between the server and the first hop switch. Because the FEX acts as a remote line card to the parent switch and has no local switching capabilities, it is not considered a hop in the Ethernet or Storage topologies. The following section outlines in detail the current single hop deployment options and configurations which are supported with the switch and FEX today.

This section includes the following topics:

Switch Mode and NPV Mode

The Cisco Nexus 5000 Series switch has two modes of operation relating to storage traffic forwarding: switch mode and N-Port Virtualizer (NPV) mode. This is the same as the modes of operation available on the Cisco Multiprotocol Director Series (MDS) Fibre Channel switches. The default mode on both platforms is “switch” mode. In the following topologies, the Cisco Nexus 5000 Series switch can either be in switch or NPV mode. The only requirement for a Cisco Nexus 5000 Series switch in NPV mode is that the upstream device supports the standard N-Port ID Virtualization (NPIV) functionality.

When the Cisco Nexus 5000 Series switch is operating in switch mode, all fabric services, for example, FSPF, zoning or DNS, are native on the access device. This means that all forwarding decisions are made by FSPF running on the switch. This mode also means that the switch consumes a Domain ID within the Fibre Channel Fabric. Limitations exist as to the number of Domain IDs that are supported within a single fabric. Specific domain ID limitations are defined by the storage vendors and OSM partners.

NPV defines the ability for a Fibre Channel switch to act as a proxy for both FLOGIs and forwarding decision and pass those duties to an upstream device. This upstream device must be capable of running NPIV which is an FC standard allowing multiple FCiDs to be handed out a single FC port. The benefit of an NPV device in a FC network is the elimination of the domain ID and therefore the ability to add more FC switches to a fabric without exceeding the supported Domain ID limitation.

The Cisco Nexus 5000 Series switch can also operate in NPV mode. When NPV is enabled on the switch, no FC fabric services are run locally on the platform and instead, forwarding and zoning services are handled by the upstream NPIV device. To avoid interoperability challenges when connecting a switch to a non-Cisco SAN core switch, Cisco recommends that the switch be configured in NPV mode.

Enabling NPV on the switch is a disruptive process and should be done at the time of initial set up to avoid any disruption to the fabric. Because enabling NPV requires a switch reboot and erases the current running configuration, be sure to save the current running configuration to an external text file so that it can be reapplied after the reboot occurs if enabling NPV after the initial set up of the switch.

Changing between switch mode and NPV mode can be done using the following commands:

To enable NPV mode:

```
switch# feature npv
```

To disable NPV mode (return to switch mode):

```
switch# no feature npv
```

Note Running NPV on the switch requires that the upstream connected device has NPIV functionality enabled

Note FC or FCoE hosts conversing with an FC or FCoE storage devices connected to the same switch in NPV is NOT supported.

vPC and Active/Standby

Host facing interfaces on the Nexus 5000 Series switch can provide connections to servers in a couple of different ways: single attached NICs for single attached hosts, active-standby NIC teaming for dual-homed servers and vPC for dual-homed servers. This guide focuses on the dual-homed server options as FC requires two independent paths to storage: Fabric A and Fabric B.

Active/Standby connections refer to servers that are dual-homed to an Ethernet LAN but only actively forwarding out one link. The second link is used as back-up in case of a failure but does not actively forward traffic unless a failure occurs. vPC is a technology introduced by Cisco Nexus products that allows a dual homed server to actively forward out both Ethernet links simultaneously. The benefits of vPC is that it gives servers access to twice as much bandwidth as in an active/standby configuration and also has the ability to converge faster than Spanning-tree in the event of a failure.

Based on the Ethernet high availability requirement, LAN admins may choose to attached servers using active/standby connections or vPC connections. Regardless of the method use to dual home a server, FCoE can co-exist with both of these topologies.

Direct Attached CNAs With Active/Standby Ethernet Topologies

Figure 5-5 shows a topology where a dual-port CNA is connecting to two switches in an active/standby configuration. Although Ethernet traffic will only traverse one link in this configuration, the FCoE traffic will be forwarded out both paths to the fabric. This is because of the way the CNA is able to differentiate between the NIC adapters for Ethernet and FC adapters for FC/FCoE. For more information on the CNA view of the Ethernet NICs and storage HBAs, see the "View Of Ethernet Traffic And FC Traffic Through A CNA" section on page 5-6.

Figure 5-5 Dual-Port CNA Connecting To Two Cisco Nexus 5000 Series Switches In An Active/Standby Topology

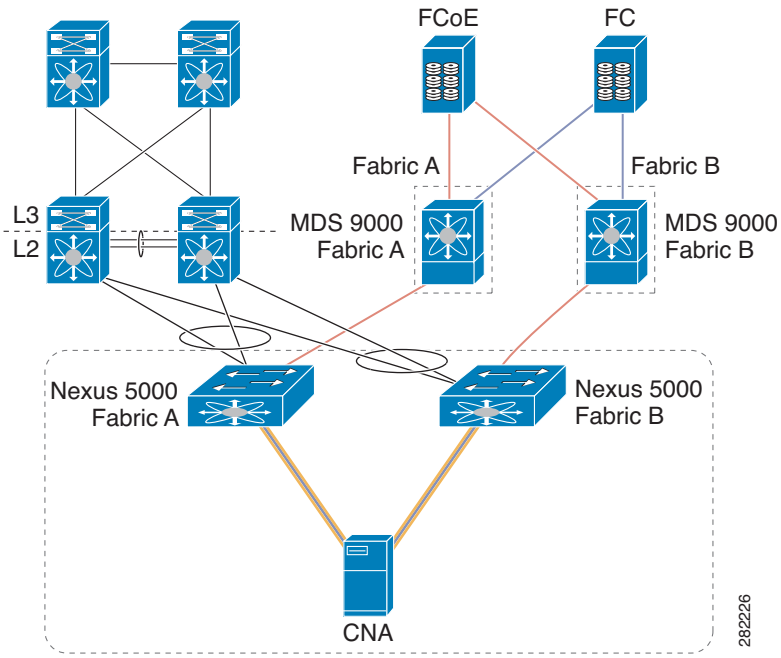

Direct Attached CNAs With vPC Ethernet Topologies

Figure 5-6 shows a topology where a dual-port CNA is connecting to two switches in a vPC configuration where only a single port connects the CNA to each switch. The operating system is able to see the Ethernet aspects of these two physical ports and port channel the Ethernet traffic coming out of the server. The FC traffic is still mapped to each link separately – one 10-Gigabit link transporting Fabric A traffic and the other 10-Gigabit link transporting Fabric B traffic. For more information on the CNA view of the Ethernet NICs and Storage HBAs, see the "View Of Ethernet Traffic And FC Traffic Through A CNA" section on page 5-6.

Figure 5-6 Dual-Port CNA Connecting To Two Cisco Nexus 5000 Series Switches In A vPC Topology

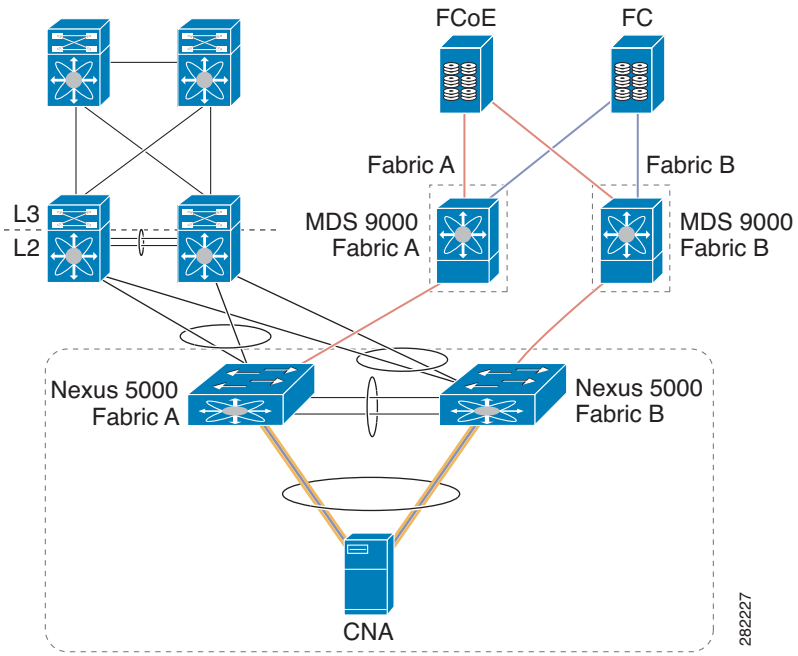

Note Direct-connect FCoE (a CNA that is directly connected to a Cisco Nexus 5000 Series switch switchport) is not supported on a port channel interface configured to have more then one member port. Directly connected FCoE devices are supported over virtual port channels where a single link from each CNA port connects through to each upstream switch or fabric extender.

Cisco Nexus 5000 Series Switch and Cisco Nexus 2000 Fabric Extender Topologies

The Nexus 2232 Fabric Extender acts as a remote line card to the parent Cisco Nexus 5000 Series switch. The Nexus 2232 Fabric Extender has 32 10-Gigabit Ethernet host facing interfaces, all of which support lossless Ethernet and FCoE. Supporting FCoE over a Cisco Nexus 5000 Series switch and FEX topology has the following requirements:

- Each Nexus 2232 Fabric Extender running FCoE must be single-homed to the upstream parent switch.
- Generation 2 (FIP Enabled) CNAs are required for host connections to the Cisco Nexus 2232 Fabric Extender host interfaces.

Adding the Cisco Nexus 2232 Fabric Extender into the FCoE topology does not change the supported configurations. Hosts can be connected to the Cisco Nexus 2232 Fabric Extender using active/standby Ethernet connections or over vPC connections. Figure 5-7 shows the supported topology.

Figure 5-7 Hosts Connected To The Cisco Nexus 2232 Fabric Extender Using Active/Standby Ethernet Connections or vPC Connections

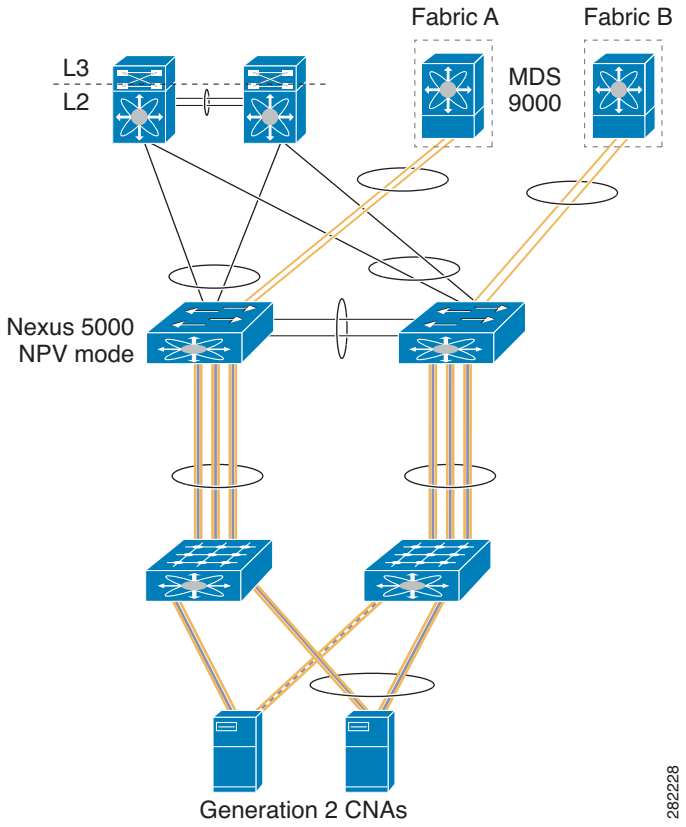

Note FCoE is not supported on a FEX interface or port channel interfaces when the FEX is connected to two switches in a FEX active-active topology.

FIP Snooping Bridges

FIP Snooping Bridges (FSBs) are lossless Ethernet bridges that are capable of watching a FIP conversation between a CNA and FCF. They have no FC/FCoE forwarding logic capabilities but instead "snoop" FIP packets and watch the FIP conversation, including FLOGI/LOGIN, between the CNA and FCF. Once a FIP snooping bridge sees a CNA login to the FC/FCoE fabric through a specific FCF, it dynamically creates an access list to guarantee that the communication between that CNA and FCF will remain point-to-point. FIP snooping is a security precaution used when transversing lossless Ethernet bridges to ensure that rogue devices can not enter the data center network and pretend to be an FCF.

It is important to note that FSBs are Layer 2 Lossless Ethernet bridges that have been enhanced to dynamically create ACLs based on the FIP communication that is seen within the fabric. FSBs have no knowledge of FC/FCoE protocols or services and do not forward FCoE traffic based on FSPF. Instead, all traffic runs over the Layer 2 protocol (STP) and is switched based on MAC address.

The Cisco Nexus 4000 Series switch is a FIP Snooping device for IBM blade chassis and must be connected to a Cisco Nexus 5000 Series FCF switch in order to support passing FCoE frames. The Cisco Nexus 4000 Series switch has 14 down-facing 10-Gigabit ports connecting to each of the 14 blade servers and 6 10-Gigabit Ethernet uplink ports used to connect to a Cisco Nexus 5000 Series switch. Figure 5-8 and Figure 5-9 shows the two supported configurations when connecting a Cisco Nexus 4000 Series switch FIP Snooping bridge to a Cisco Nexus 5000 Series FCF switch:

Cisco Nexus 4000 Series Switch To Cisco Nexus 5000 Series Switch FCoE With Consolidated Links

Figure 5-8 shows a Cisco Nexus 4000 Series switch connected to a Cisco Nexus 5000 Series switch using consolidated links where both FCoE and Ethernet traffic are utilizing the same link simultaneously. Because FCoE requires fabric separation, the Ethernet traffic must also only follow one path and can not take advantage of other Ethernet HA technologies such as vPC.

Figure 5-8 Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Consolidated Links

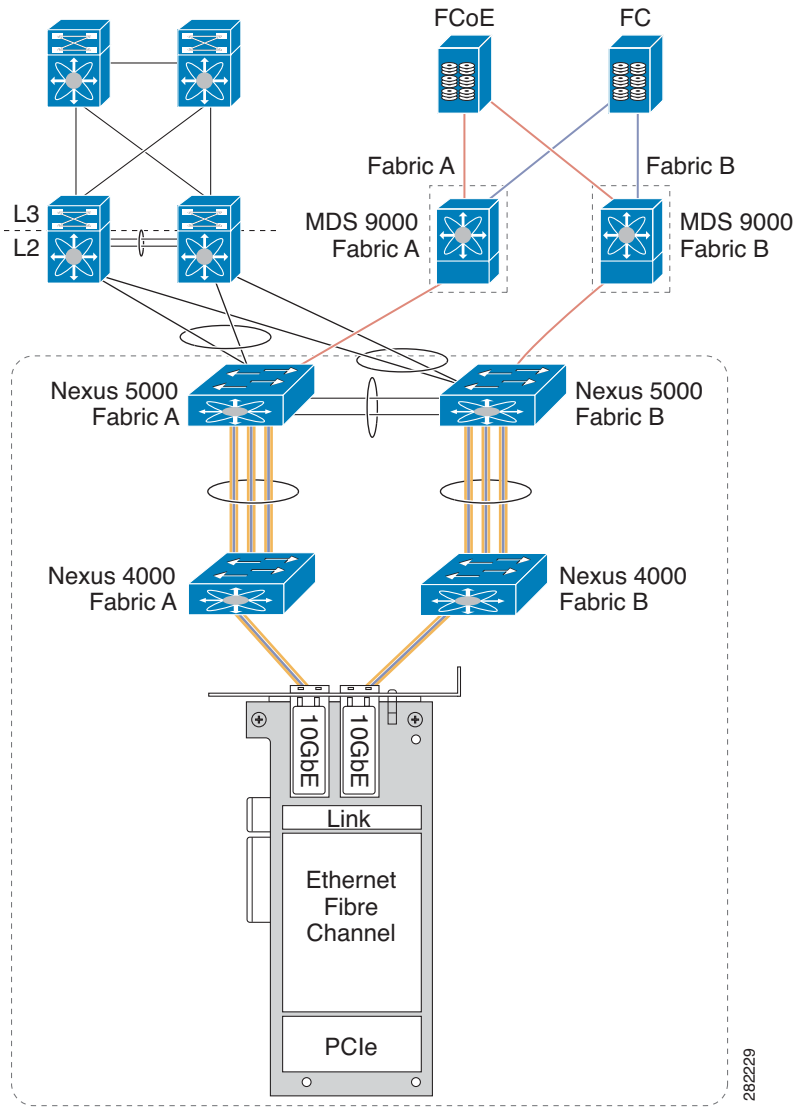

Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Dedicated Wires

Figure 5-9 shows the Cisco Nexus 4000 Series switches connecting to Cisco Nexus 5000 Series switches using dedicated links; blue links are Ethernet ONLY links and pink and blue links are FCoE-only links. There are no consolidated links shown in Figure 5-9. The benefit of running dedicated links between the Cisco Nexus 4000 Series switches and Cisco Nexus 5000 Series switches in this topology is the fact that both storage and Ethernet traffic are able to take advantage of their respective HA models. Ethernet traffic is multi-homed to the upstream switches and using vPC to forward out all available paths while FCoE is maintaining fabric isolation through the Ethernet network.

Figure 5-9 Cisco Nexus 4000 Series Switch Connected To A Cisco Nexus 5000 Series Switch FCoE With Dedicated Wires

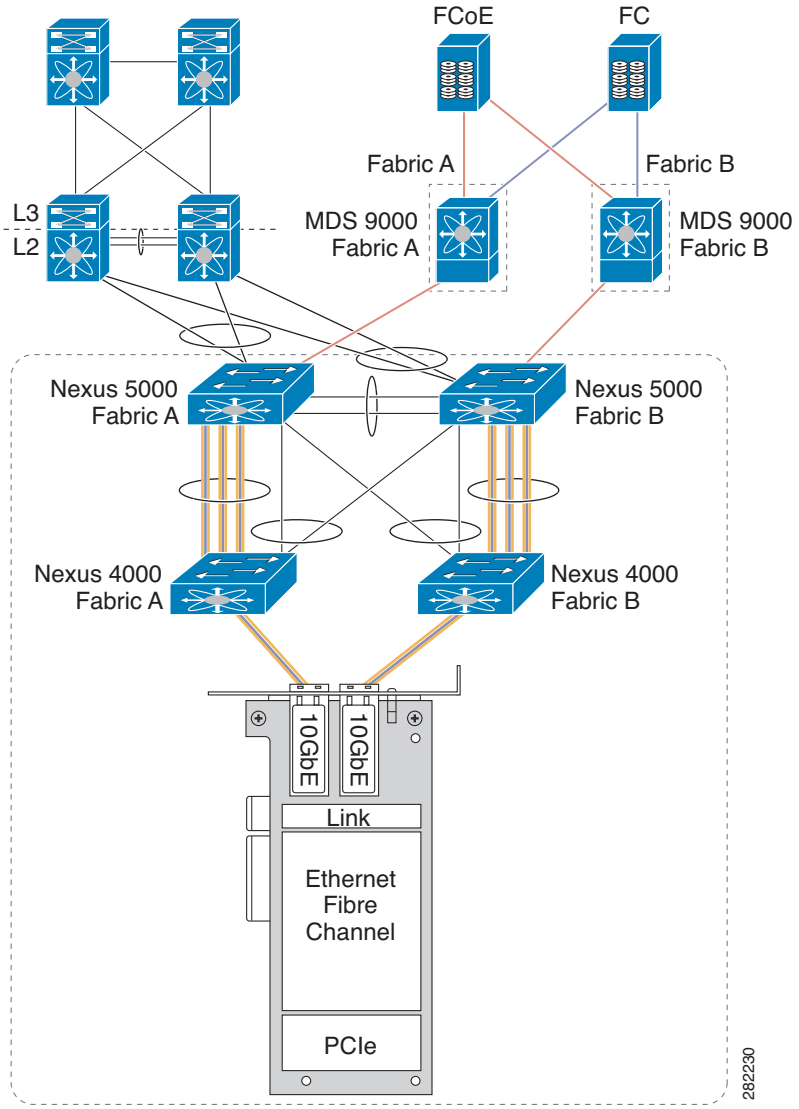

Multi-Hop FCoE Solutions

Multi-Hop FCoE is achieved with the support of Virtual E-ports (VE ports) connection two FCFs. Like E_Ports in native FC, VE ports are use to expand the FCoE fabric. VE ports are supported on the Nexus 5000 Series switch as of the NXOS Release 5.0(1)N2(2). There are two options for connecting Nexus 5000 Series switches with the use of VE ports: using single-links or over a port channel. For configuration examples of VE ports, see Appendix B, "Port Configuration Examples".

In order to maintain fabric isolation, the Cisco Nexus 5000 FCF switches in each fabric should be configured to have the same FC-MAP value. The FC-MAP values should be different between Fabric A and Fabric B. For additional information on FC-MAP configurations, see Appendix B, "Port Configuration Examples". VE ports brought up between two Cisco Nexus 5000 Series switches with differing FC-MAPs are not supported which ensures that fabrics are not merged by connecting FCFs in Fabric A to FCFs in Fabric B. Figure 5-10 shows FCF connections using VE ports.

Figure 5-10 VE Ports And FCF Mapping

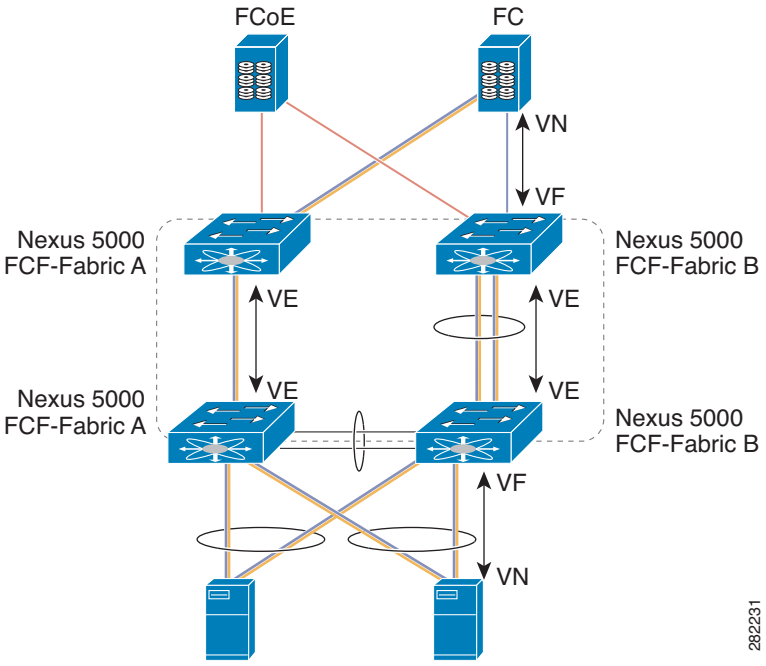

Note VE ports are not supported over vPCs.

FCoE Operations

This section includes the following topics:

- Tracking FCoE Statistics, page 5-22
- SPAN for FC and FCoE Traffic, page 5-23
- Roles Based Access Control, page 5-24

Tracking FCoE Statistics

FCoE statistics for FCoE traffic transversing an interface on a Cisco Nexus 5000 Series switch can be seen by monitoring the statistics on the vFC interface which is bound to the physical Ethernet interface or port channel interface.

This section includes the following topics:

Tracking VE Port Statistics

The following example shows how to monitor VE port statistics:

```
switch(config-if)# show inter vfc 300
vfc300 is trunking
    Bound interface is port-channel300
    Hardware is Virtual Fibre Channel
    Port WWN is 21:2b:00:05:9b:77:f5:7f
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Trunk vsans (admin allowed and active) (3,5)
    Trunk vsans (up) (3,5)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()
    1 minute input rate 15600 bits/sec, 1950 bytes/sec, 21 frames/sec
    1 minute output rate 43664 bits/sec, 5458 bytes/sec, 21 frames/sec
    51295547 frames input, 10484381916 bytes
    0 discards, 0 errors
    39089018 frames output, 10620127132 bytes
    0 discards, 0 errors
    last clearing of "show interface" counters never
    Interface last changed at Mon Jan 17 19:05:27 2011
```

Tracking VF Port Statistics

The following example shows how to monitor VF port statistics:

```
switch(config-if)# show inter vfc 31
vfc31 is trunking (Not all VSANs UP on the trunk)
    Bound interface is Ethernet1/1
    Hardware is Virtual Fibre Channel
    Port WWN is 20:1e:00:05:9b:77:f5:7f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 3
    Trunk vsans (admin allowed and active) (3)
    Trunk vsans (up) (3)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()
    1 minute input rate 6912756368 bits/sec, 864094546 bytes/sec, 8640880 frames/sec
    1 minute output rate 6963590568 bits/sec, 870448821 bytes/sec, 396313 frames/sec
    789408333283 frames input, 78940833327276 bytes
    0 discards, 0 errors
    36207053863 frames output, 79510690165704 bytes
    0 discards, 0 errors
```

```
    last clearing of "show interface" counters never
    Interface last changed at Mon Jan 17 19:05:21 2011
```

SPAN for FC and FCoE Traffic

This section includes the following topics:

Possible SPAN Sources

Following are possible SPAN sources:

- FC interface (only rx-source on 5500 platform)
- VFC interface
- VSAN (not supported on 5500 platform)
- VLAN
- Ethernet interface
- Port channel interface
- SAN port channel interface

Possible SPAN Destinations

Following are possible SPAN destinations:

- FC interface
- Ethernet interface

SPAN Configuration Examples

This example shows how to display configuration information on Ethernet 1/1:

```
switch(config)# show running-config interface eth 1/1
interface Ethernet1/1
    switchport monitor
```

This example shows how to display the health monitoring of all interfaces for failover purposes:

```
switch(config)# show running-config monitor all
monitor session 1 type local
    no description
    source interface vfc33 both
    destination interface Ethernet1/1
    no shut
```

This example shows the health monitoring of session 1:

```
switch(config)# show monitor session 1
    session 1
---------------
type : local
```

```
state : up
source intf :
    rx : vfc33
    tx : vfc33
    both : vfc33
source VLANs :
    rx :
source VSANs :
    rx :
destination ports : Eth1/1
Legend: f = forwarding enabled, l = learning enabled
```

This example shows the health monitoring configuration:

```
switch(config)# show running-config monitor
monitor session 1
    source interface fc3/1 tx
    destination interface Ethernet1/1
    no shut
```

This example shows the health monitoring of all sessions:

```
switch(config)# show monitor session all
session 1
---------------
type : local
state : up
source intf :
    rx : fc3/1
    tx : fc3/1
    both : fc3/1
source VLANs :
    rx :
source VSANs :
    rx þ:
destination ports : Eth1/1
Legend: f = forwarding enabled, l = learning enabled
```

Roles Based Access Control

With the Cisco Nexus Family of switches deploying unified I/O capabilities, the roles of LAN and SAN administrators are converging. To help manage these two different roles on the Cisco Nexus Series Family of switches, the Roles Based Access Control (RBAC) feature facilitates various administrative operations.

When deploying unified I/O within a data center, Cisco recommends defining the following three roles:

- Unified Administrator—This role includes all actions that impact both LAN and SAN operations. This role is sometimes referred to as a global administrator.
- LAN Administrator—This role includes a set of actions that impact LAN operation while denying any actions that could impact SAN operations.
- SAN Administrator—This role includes a set of actions that impact SAN operation while denying any actions that could impact LAN operations.

These are general roles that are used to enforce the operational model where separate LAN and SAN administrative teams retain management control of their perspective networks without interference. More specific roles may be added if operations need to be more tightly defined.

This section includes the following topics:

Unified Administrator Role

The Unified Administrator role may perform all actions. In addition, the Unified Administrator plays a large role in the initial set up of the unified network.

Before implementing a unified network design, the physical interfaces and VLANs used for unified traffic should be identified and defined. Standard implementation of FCoE requires binding a virtual Fibre Channel interface (vFC) to either a physical Ethernet interface or MAC-Address. It is also required to map the VSAN used to carry the FC traffic to a corresponding Ethernet VLAN. While Ethernet interfaces and VLANs normally fall under the scope of a LAN administration, the unified interfaces and FCoE VLANs must be identified so that they can be separated from the LAN administration domain.

Cisco recommends that you identify the interfaces used for Unified I/O, and that you designate a range of VLANs for FCoE use before implementation begins. The Unified Administrator role will configure these unified interfaces and FCoE VLANs.

LAN Administrator Role

This role is assigned all the permissions that impact LAN traffic. This role also denies any actions that would possibly impact SAN traffic (FCoE and FC). One of the main difference between the LAN administrator role and a LAN administrator in a legacy data center without unified I/O is the inability to shut down a physical Ethernet port carrying FCoE traffic. Potentially, both FC and Ethernet traffic could be traveling over the link simultaneously and, therefore, shutting the port could have an impact on SAN operations.

A list of commands which can impact SAN operations, and therefore should be limited from the role of the LAN Administrator, can be found in Appendix A, “RBAC Configuration”. Individual network designs may require additional limited commands.

SAN Administrator Role

This role is assigned all the permissions that impact SAN traffic. The role also denies actions that would impact LAN traffic.

SAN administration in a unified environment and a legacy SAN environment are similar. Today, unified I/O runs only between the servers and the top-of-rack Cisco Nexus 5000 switch, where FC links are run back into the core of the existing SAN infrastructure. The FC module inside the Cisco Nexus 5000 switch can operate in either NPV or switch mode. The switch most commonly operates in NPV mode and, from a management perspective, looks identical to a FC blade or fabric switch operating in NPV mode.

A list of commands which can impact LAN operations and therefore should be limited from the role of the SAN Administrator can be found in Appendix A, “RBAC Configuration”. Individual network designs may require additional limited commands.

FCoE Limitations

This section includes the following topics:

Generation 1 And Generation 2 CNA Limitations

When FCoE was introduced on the Cisco Nexus 5000 Series switch, Cisco worked with QLogic and Emulex to create the first generation of CNA adapters. These CNAs used a pre-standard implementation of the DCBX protocol nicknamed CIN-DCBX. These adapters also did not support the standard FIP implementation as defined in the FCoE Standard (FC-BB-5) and they are often referred to as Pre-FIP adapters.

Starting in 2009, after the ratification of the FCoE standard, second generation CNAs were put out by both QLogic and Emulex that supported standard FIP and FCoE. These CNAs also used a pre-standard version of the DCBX protocol nicknamed CEE-DCBX which has been decided on by multiple venders to be the de-facto standard until IEEE DCBX is ratified.

Topologies and Platforms Which Require Generation 2 CNAs

While the Cisco Nexus 5010 switch and Nexus 5020 switch are backwards compatible with both Generation 1 and Generation 2 CNAs and support, the Nexus 2000 Fabric Extenders and the Nexus 5500 Platform switches only support Generation 2 CNA connections. Also, Generation 2 CNAs are required when connecting a host using vPC into a fabric, whether the host is running FCoE or just native Ethernet.

LACP and FCoE To The Host

Today, when deploying FCoE over a host-facing vPC, the vFC interface is bound to the port channel interfaces associated with the vPC. This requires that the port channel interface be up and forwarding before FCoE traffic can be switched. Cisco recommends when running vPC in an Ethernet environment is to use LACP in order to negotiate the parameters on both sides of the port channel to ensure that configurations between both sides is consistent.

However, if there are inconsistencies in any of the Ethernet configuration parameters LACP uses to bring up the port channel interface, both sides of the virtual port channel will remain down. This means that FCoE traffic from the host is now dependent on the correct configuration on the LAN/Ethernet side. When this dependency occurs, Cisco recommends that you use the static port channel configuration (channel-group # mode on) when deploying vPC and FCoE to the same host.

Deploying a Cisco Nexus 5000 Series Switch as an NPIV Core

The Nexus 5000 Series switch supports both NPV and NPIV functionality. If acting as an NPIV core switch with downstream NPV switches attached to it, it is important to note that hosts and targets which are communicating to one another can not be attached to the same downstream NPV device.

VE Ports on a Cisco Nexus 5010 Switch or Cisco Nexus 5020 Switch

Cisco Nexus 5000 Series and Cisco Nexus 5500 Platform switches support VE port connections. On Cisco Nexus 5010 and Nexus 5020 switches, VE ports can be configured between two switches using a single port channel or multiple individual links. VE ports configured between two switches using multiple port channels is not supported. This has to do with the number of MAC addresses available for the VE port on the Cisco Nexus 5010 switch and the Cisco Nexus 5020 switch. This limitation does not apply to the Cisco Nexus 5500 Platform.

Additional Information

See "Configuring FCoE NPV" in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Cisco Nexus 5000 Series Switch overview information:
http://www.cisco.com/en/US/products/ps9670/index.html

Cisco Nexus 5000 Series Configuration Guides:
http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

Fibre Channel over Ethernet information: www.fcoe.com

APPENDIX A

RBAC Configuration

This chapter includes information about RBAC configurations in relation to FCoE operations and it includes the following sections:

Global Administrator Actions

The Global administrator role is unrestricted and all commands are available.

LAN Administrator Actions

This section lists the commands that the LAN administrator may not perform. Commands that are not listed are implicitly permitted.

Global Level Deny Actions

```
switch(config)# feature lacp
switch(config)# feature tacacs+
switch(config)# feature udld
switch(config)# feature fcoe
switch(config)# aaa *
switch(config)# boot *
switch(config)# cfs *
switch(config)# class-map *
switch(config)# device-alias *
switch(config)# diagnostic *
switch(config)# fex *
switch(config)# hw-module logging onboard *
switch(config)# license *
switch(config)# line *
switch(config)# lldp *
switch(config)# monitor session *
switch(config)# ntp *
switch(config)# policy-map *
switch(config)# privilege *
switch(config)# radius-server *
```

```
switch(config)# role *
switch(config)# snmp-server *
switch(config)# spanning-tree
    bridge assurance *
    loopguard *
    mode *
    mst *
    pathcost *
    port type *
    vlan <fcoe-vlan>
switch(config)# ssh *
switch(config)# system
    core *
    default switchport *
    jumbomtu *
    qos *
switch(config)# tacacs+ *
switch(config)# telnet server enable
switch(config)# trunk protocol enable
switch(config)# username *
switch(config)# vrf  *
switch(config)# xml server *
```

This section includes the following topics:

VLAN-Level Deny Actions

Deny Actions for All VLANs

```
switch(config)# vlan vlan
switch(config-vlan)# fcoe
```

Deny Actions for Pre-determined FCoE VLANs

```
switch(config)# no vlan fcoe-vlan
switch(config)# vlan fcoe-vlan *
switch(config)# spanning-tree vlan fcoe-vlan *
switch(config-mst)# instance n vlan fcoe-vlan
switch(config)# mac-address-table aging-time t vlan fcoe-vlan
switch(config)# mac-address-table static aaaa.bbbb.cccc vlan fcoe-vlan
switch(config-monitor)# source vlan fcoe-vlan
switch(config)# vlan fcoe-vlan
switch(config-vlan)# ip igmp snooping *
```

Interface-Level Deny Actions

Interface-Level Deny Actions

Note Access to the management interface is limited to the unified administrator.

```
switch(config)# interface mgmt *
```

Deny Actions for Pre-determined Ethernet Interfaces Designated to Carry FCoE Traffic

```
switch(config-if)# bandwidth *
switch(config-if)# fcoe *
switch(config-if)# flowcontrol *
switch(config-if)# link debounce *
switch(config-if)# lldp *
switch(config-if)# priority-flow-control *
switch(config-if)# service-policy *
switch(config-if)# shutdown
switch(config-if)# shutdown force
switch(config-if)# spanning-tree bpdufilter
switch(config-if)# spanning-tree bpduguard
switch(config-if)# spanning-tree cost *
switch(config-if)# spanning-tree guard *
switch(config-if)# spanning-tree link-type *
switch(config-if)# spanning-tree mst *
switch(config-if)# spanning-tree port type *
switch(config-if)# spanning-tree port-priority *
switch(config-if)# speed *
switch(config-if)# switchport host
switch(config-if)# switchport mode *
switch(config-if)# switchport monitor
switch(config-if)# switchport trunk native vlan <fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan <range>
switch(config-if)# switchport trunk allowed vlan add <fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan all
switch(config-if)# switchport trunk allowed vlan except *
switch(config-if)# switchport trunk allowed vlan none
switch(config-if)# switchport trunk allowed vlan remove <fcoe-vlan>
```

FC Deny Actions

FC Deny Actions

Note The LAN administrator may not execute SAN-related commands.

```
switch(config)# fabric-binding *
switch(config)# fcalias *
switch(config)# fcdomain *
switch(config)# fcdroplatency *
switch(config)# fcflow *
switch(config)# fcid-allocation *
switch(config)# fcinterop *
switch(config)# fcns *
switch(config)# fcroute *
switch(config)# fcs *
switch(config)# fcsp *
switch(config)# fctimer *
switch(config)# fdmi *
switch(config)# fspf *
switch(config)# in-order-guarantee
switch(config)# interface fc *
switch(config)# interface san-port-channel *
switch(config)# interface vfc *
switch(config)# npiv *
switch(config)# npv *
switch(config)# port-security enable
switch(config)# port-track enable
switch(config)# rib *
```

```
switch(config)# rlir *
switch(config)# rscn *
switch(config)# scsi-target *
switch(config)# system default zone *
switch(config)# vsan database  *
switch(config)# wwn *
switch(config)# zone *
switch(config)# zoneset *
```

SAN Administrator Actions

This section lists the commands that the SAN administrator may not perform. Commands that are not listed are implicitly permitted.

Global Level Deny Actions

```
switch(config)# feature * (except feature fcoe)
switch(config)# aaa *
switch(config)# boot *
switch(config)# cfs *
switch(config)# class-map *
switch(config)# device-alias *
switch(config)# diagnostic *
switch(config)# fex *
switch(config)# hw-module logging onboard *
switch(config)# ip *
switch(config)# ipv6 *
switch(config)# license *
switch(config)# line *
switch(config)# lldp *
switch(config)# mac-address-table *
switch(config)# monitor session *
switch(config)# ntp *
switch(config)# policy-map *
switch(config)# privilege *
switch(config)# radius-server *
switch(config)# role *
switch(config)# snmp-server *
switch(config)# spanning-tree
    bridge assurance *
    loopguard *
    mode *
    mst *
    pathcost *
    port type *
    vlan <non-fcoe-vlan>
switch(config)# ssh *
switch(config)# system
    core *
    default switchport *
    jumbomtu *
    qos *
switch(config)# tacacs+ *
switch(config)# telnet server enable
switch(config)# trunk protocol enable
switch(config)# username *
switch(config)# vrf  *
switch(config)# xml server *
```

This section includes the following topics:

VLAN Level Deny Actions

Deny Actions for Pre-determined Non-FCoE VLANs

```
switch(config)# no vlan <non-fcoe-vlan>
switch(config)# vlan <non-fcoe-vlan>*
switch(config)# spanning-tree vlan <non-fcoe-vlan>*
switch(config-mst)# instance n vlan <non-fcoe-vlan>
switch(config)# mac-address-table aging-time t vlan <non-fcoe-vlan>
switch(config)# mac-address-table static aaaa.bbbb.cccc vlan <non-fcoe-vlan>
switch(config-monitor)# source vlan <non-fcoe-vlan>
switch(config)# vlan <non-fcoe-vlan>
switch(config-vlan)# ip igmp snooping *
switch(config-if)# spanning-tree vlan <non-fcoe-vlan>
```

Interface Level Deny Actions

Interface Level Deny Actions

Note Access to the management interface is limited to the unified administrator.

```
switch (config)# interface mgmt *
```

Deny Actions for Pre-determined Ethernet Interfaces Designated to not Carry FCoE Traffic

The SAN administrator may execute **no** commands on these interfaces.

Deny Actions for Pre-determined Ethernet Interfaces Designated to Carry FCoE Traffic

This deny-list applies to Ethernet, port-channel, and vEthernet interfaces that are designated to carry FCoE traffic.

```
switch(config-if)# bandwidth *
switch(config-if)# fcoe *
switch(config-if)# flowcontrol *
switch(config-if)# link debounce *
switch(config-if)# lldp *
switch(config-if)# priority-flow-control *
switch(config-if)# service-policy *
switch(config-if)# shutdown
switch(config-if)# shutdown force
switch(config-if)# shutdown lan // TBD. This is a new command to shut stop LAN VLANs
switch(config-if)# spanning-tree bpdufilter
switch(config-if)# spanning-tree bpduguard
switch(config-if)# spanning-tree cost *
switch(config-if)# spanning-tree guard *
switch(config-if)# spanning-tree link-type *
switch(config-if)# spanning-tree mst *
switch(config-if)# spanning-tree port type *
switch(config-if)# spanning-tree port-priority *
switch(config-if)# speed *
switch(config-if)# switchport host
```

```
switch(config-if)# switchport mode *
switch(config-if)# switchport monitor
switch(config-if)# switchport trunk native *
switch(config-if)# switchport trunk allowed vlan <range>
switch(config-if)# switchport trunk allowed vlan add <non-fcoe-vlan>
switch(config-if)# switchport trunk allowed vlan all
switch(config-if)# switchport trunk allowed vlan except *
switch(config-if)# switchport trunk allowed vlan none
switch(config-if)# switchport trunk allowed vlan remove <non-fcoe-vlan>
```

LAN Deny Actions

LAN Deny Actions

The SAN administrator can not execute LAN-related commands.

```
switch(config)# cdp *
switch(config)# ip igmp snooping *
switch(config)# port-channel load-balance ethernet
switch(config)# rmon
switch(config)# track
```

Sample Configurations

The following configurations are used to create both LAN and SAN administrative roles. These configurations follow the outline listed above concerning the commands that are assigned to or withheld from each role. Configuration is not needed for the Global Administrator who automatically has access to all configuration commands.

Note This configuration assumes that vFC 1 is mapped to Ethernet 1/1 and that VLAN 100 has been designated the FCoE VLAN. This configuration is based on the specific environment and which Ethernet ports and VLANs have been pre-determined to carry FCoE traffic.

LAN-Admin Configuration

role name LAN-admin

description assume vlan 100 is fcoe enabled and eth1/1 is an vfc bound (fcoe) interface

rule 97 deny command config t ; feature lacp

rule 96 deny command config t ; feature tacacs+

rule 95 deny command config t ; feature udld

rule 94 deny command config t ; feature fcoe

rule 93 deny command config t ; aaa *

rule 92 deny command config t ; boot *

rule 91 deny command config t ; cfs *

rule 90 deny command config t ; class-map *

rule 89 deny command config t ; device-alias *

rule 88 deny command config t ; diagnostic *

rule 87 deny command config t ; fex *

rule 86 deny command config t ; hw-module logging onboard *
rule 85 deny command config t ; license *
rule 84 deny command config t ; line *
rule 83 deny command config t ; lldp *
rule 82 deny command config t ; monitor session *
rule 81 deny command config t ; ntp *
rule 80 deny command config t ; policy-map *
rule 79 deny command config t ; privilege *
rule 78 deny command config t ; radius-server *
rule 77 deny command config t ; role *
rule 76 deny command config t ; snmp-server *
rule 75 deny command config t ; ssh *
rule 74 deny command config t ; system *
rule 73 deny command config t ; no system *
rule 72 deny command config t ; tacacs+ *
rule 71 deny command config t ; telnet server enable
rule 70 deny command config t ; trunk protocol enable
rule 69 deny command config t ; username *
rule 68 deny command config t ; vrf *
rule 67 deny command config t ; xml server *
rule 66 deny command config t ; fabric-binding *
rule 65 deny command config t ; fcalias *
rule 64 deny command config t ; fcdomain *
rule 63 deny command config t ; fcdroplatency *
rule 62 deny command config t ; fcflow *
rule 61 deny command config t ; fcid-allocation *
rule 60 deny command config t ; fcinterop *
rule 59 deny command config t ; fcns *
rule 58 deny command config t ; fcroute *
rule 57 deny command config t ; fcs *
rule 56 deny command config t ; fcsp *
rule 55 deny command config t ; fctimer *
rule 54 deny command config t ; fdmi *
rule 53 deny command config t ; fspf *
rule 52 deny command config t ; in-order-guarantee
rule 51 deny command config t ; npiv *
rule 50 deny command config t ; npv *
rule 49 deny command config t ; port-security enable

rule 48 deny command config t ; port-track enable
rule 47 deny command config t ; rib *
rule 46 deny command config t ; rlir *
rule 45 deny command config t ; rscn *
rule 44 deny command config t ; scsi-target *
rule 43 deny command config t ; vsan database *
rule 42 deny command config t ; wwn *
rule 41 deny command config t ; zone *
rule 40 deny command config t ; zoneset *
rule 39 deny command config t ; vlan * ; fcoe *
rule 38 deny command config t ; vlan * ; no fcoe *
rule 37 deny command config t ; spanning-tree vlan 100
rule 36 permit command config t ; spanning-tree vlan *
rule 35 deny command config t ; spanning-tree *
rule 34 deny command config t ; mac-address-table aging-time * vlan 100
rule 33 deny command config t ; mac-address-table static * vlan 100 *
rule 32 deny command config t ; monitor session * ; source vlan 100
rule 31 deny command config t ; vlan 100 *
rule 30 deny command config t ; no vlan 100 *
rule 29 deny command config t ; interface Ethernet1/1 ; bandwidth *
rule 28 deny command config t ; interface Ethernet1/1 ; fcoe *
rule 27 deny command config t ; interface Ethernet1/1 ; flowcontrol *
rule 26 deny command config t ; interface Ethernet1/1 ; link debounce *
rule 25 deny command config t ; interface Ethernet1/1 ; lldp *
rule 24 deny command config t ; interface Ethernet1/1 ; priority-flow-control *
rule 23 deny command config t ; interface Ethernet1/1 ; service-policy *
rule 22 deny command config t ; interface Ethernet1/1 ; shutdown
rule 21 deny command config t ; interface Ethernet1/1 ; shutdown force
rule 20 deny command config t ; interface Ethernet1/1 ; spanning-tree bpdufilter *
rule 19 deny command config t ; interface Ethernet1/1 ; spanning-tree bpduguard *
rule 18 deny command config t ; interface Ethernet1/1 ; spanning-tree cost *
rule 17 deny command config t ; interface Ethernet1/1 ; spanning-tree guard *
rule 16 deny command config t ; interface Ethernet1/1 ; spanning-tree link-type *
rule 15 deny command config t ; interface Ethernet1/1 ; spanning-tree mst *
rule 14 deny command config t ; interface Ethernet1/1 ; spanning-tree port type *
rule 13 deny command config t ; interface Ethernet1/1 ; spanning-tree port-priority *
rule 12 deny command config t ; interface Ethernet1/1 ; speed *
rule 11 deny command config t ; interface Ethernet1/1 ; switchport host

```
  rule 10 deny command config t ; interface Ethernet1/1 ; switchport mode *
  rule 9 deny command config t ; interface Ethernet1/1 ; switchport monitor
  rule 8 deny command config t ; interface Ethernet1/1 ; switchport trunk native vlan 100
  rule 7 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan *
  rule 6 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan add 100
  rule 5 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan all
  rule 4 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan except *
  rule 3 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan none
  rule 2 deny command config t ; interface Ethernet1/1 ; switchport trunk allowed vlan remove 100
  rule 1 permit read-write
  interface policy deny
    permit interface eth1/1-40
  vlan policy deny
    permit vlan 100-200
  vsan policy deny
```

SAN-Admin Configuration

```
role name SAN-admin
  description assuming vlan 100 is fcoe enabled and vfc1 has been bound to eth1/1
  rule 83 permit command config t ; vlan * ; fcoe *
  rule 82 deny command config t ; vlan * ; *
  rule 81 deny command config t ; ip igmp snooping *
  rule 80 deny command config t ; cdp *
  rule 79 deny command config t ; port-channel load-balance ethernet *
  rule 78 deny command config t ; rmon *
  rule 77 deny command config t ; track *
  rule 76 deny command config t ; no ip igmp *
  rule 75 deny command config t ; no cdp *
  rule 74 deny command config t ; no port-channel load-balance *
  rule 73 deny command config t ; no rmon *
  rule 72 deny command config t ; no track *
  rule 71 deny command config t ; interface * ; switchport trunk native *
  rule 70 deny command config t ; interface * ; switchport trunk allowed vlan *
  rule 69 deny command config t ; interface * ; switchport trunk allowed vlan add 100
  rule 68 deny command config t ; interface * ; switchport trunk allowed vlan all
  rule 67 deny command config t ; interface * ; switchport trunk allowed vlan except *
  rule 66 deny command config t ; interface * ; switchport trunk allowed vlan none
  rule 65 deny command config t ; interface * ; switchport trunk allowed vlan remove 100
```

rule 64 deny command config t ; interface * ; bandwidth *
rule 63 deny command config t ; interface * ; fcoe *
rule 62 deny command config t ; interface * ; flowcontrol *
rule 61 deny command config t ; interface * ; link debounce *
rule 60 deny command config t ; interface * ; lldp *
rule 59 deny command config t ; interface * ; priority-flow-control *
rule 58 deny command config t ; interface * ; service-policy *
rule 57 deny command config t ; interface * ; shutdown
rule 56 deny command config t ; interface * ; shutdown force
rule 55 deny command config t ; interface * ; shutdown lan
rule 54 deny command config t ; interface * ; spanning-tree bpdufilter
rule 53 deny command config t ; interface * ; spanning-tree bpduguard
rule 52 deny command config t ; interface * ; spanning-tree cost *
rule 51 deny command config t ; interface * ; spanning-tree guard *
rule 50 deny command config t ; interface * ; spanning-tree link-type *
rule 49 deny command config t ; interface * ; spanning-tree mst *
rule 48 deny command config t ; interface * ; spanning-tree port type *
rule 47 deny command config t ; interface * ; spanning-tree port-priority *
rule 46 deny command config t ; interface * ; speed *
rule 45 deny command config t ; interface * ; switchport host
rule 44 deny command config t ; interface * ; switchport mode *
rule 43 deny command config t ; interface * ; switchport monitor
rule 42 deny command config t ; no vlan 100 *
rule 41 permit command config t ; feature fcoe
rule 40 deny command config t ; feature *
rule 39 deny command config t ; aaa *
rule 38 deny command config t ; boot *
rule 37 deny command config t ; cfs *
rule 36 deny command config t ; class-map *
rule 35 deny command config t ; device-alias *
rule 34 deny command config t ; diagnostic *
rule 33 deny command config t ; fex *
rule 32 deny command config t ; hw-module logging onboard *
rule 31 deny command config t ; ip *
rule 30 deny command config t ; ipv6 *
rule 29 deny command config t ; license *
rule 28 deny command config t ; line *
rule 27 deny command config t ; lldp *

```
rule 26 deny command config t ; mac-address-table *
rule 25 deny command config t ; monitor session *
rule 24 deny command config t ; ntp *
rule 23 deny command config t ; policy-map *
rule 22 deny command config t ; privilege *
rule 21 deny command config t ; radius-server *
rule 20 deny command config t ; role *
rule 19 deny command config t ; snmp-server *
rule 18 deny command config t ; spanning-tree bridge assurance *
rule 17 deny command config t ; spanning-tree loopguard *
rule 16 deny command config t ; spanning-tree mode *
rule 15 deny command config t ; spanning-tree mst *
rule 14 deny command config t ; spanning-tree pathcost *
rule 13 deny command config t ; spanning-tree port type *
rule 12 deny command config t ; ssh *
rule 11 deny command config t ; system core *
rule 10 deny command config t ; system default switchport *
rule 9 deny command config t ; system jumbomtu *
rule 8 deny command config t ; system qos *
rule 7 deny command config t ; tacacs+ *
rule 6 deny command config t ; telnet server enable
rule 5 deny command config t ; trunk protocol enable
rule 4 deny command config t ; username *
rule 3 deny command config t ; vrf  *
rule 2 deny command config t ; xml server *
rule 1 permit read-write
vlan policy deny
  permit vlan 100-100
interface policy deny
  permit interface fc3/1-4
  permit interface Ethernet1/1
  permit interface vfc1
```

APPENDIX B

Port Configuration Examples

This appendix describes port configuration examples relating to FCoE topologies and it includes the following sections:

VE Port Configuration Example

This section provides a sample configuration of the Cisco Nexus 5000 Series switch FCoE VE Port implementation. The configuration covers the switches in switch mode. FCoE initiators are used in this lab. You can attach either FC F Port storage directly to a Nexus 5000 Series switch FC GEMs, or use an FCoE target.

Note This example can be used for configuring VE ports between two Cisco Nexus 5000 Series switches in both fabrics. It does not include server configurations.

FCoE VE Port Topology Example

Figure B-1 shows the topology that was used for the configuration example. The following configuration parameters are used in this topology:

- FCoE VLAN for Fabric A: 10
- FCoE VSAN for Fabric A: 10
- FCoE VLAN for Fabric B: 20
- FCoE VSAN for Fabric B: 20
- Ethernet Only VLAN across both fabrics: 200

You should choose these values before the time of configuration.

Figure B-1 FCoE VE Port Topology

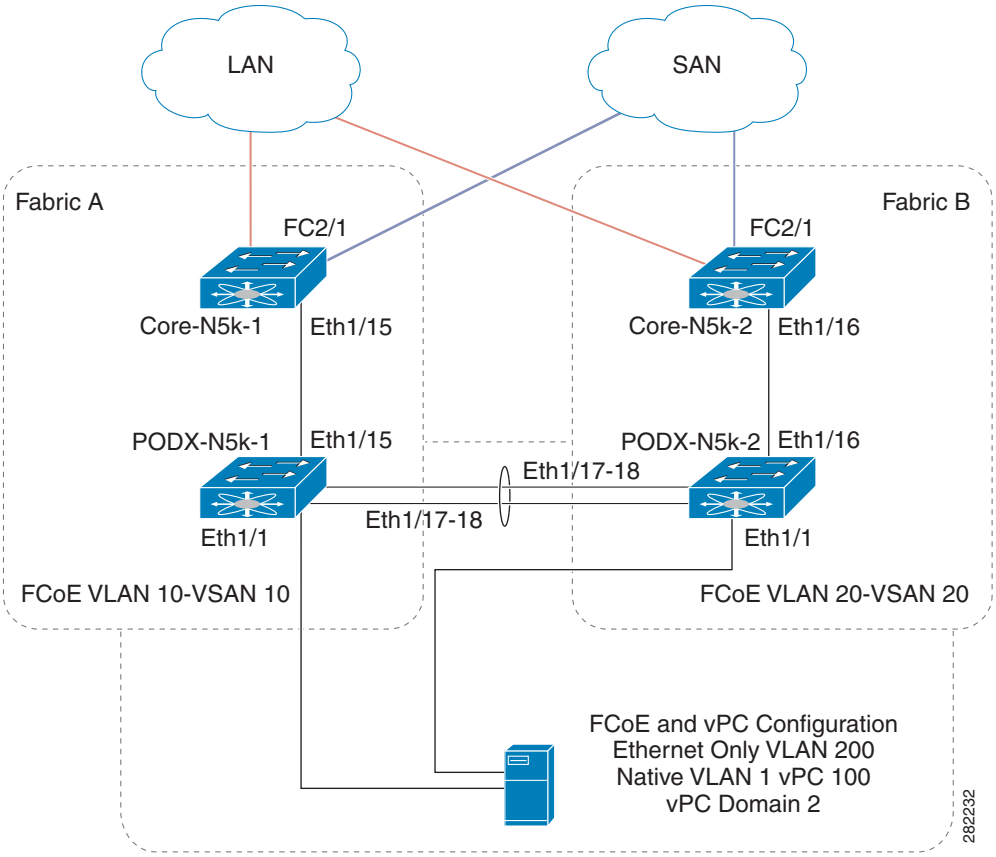

Note The FCoE VLAN/VSAN numbering does not have to be the same within the fabric. As a best practice, use different FCoE VLANs and VSAN numbers between the two fabrics to avoid confusion. Configurations have often been set up to assign ODD VLAN/VSANs for one fabric and EVEN VLANs/VSANs for the other fabric. This is just one example of keeping the numbers separate between the two fabrics

Enabling FCoE and Verifying QoS Configuration

Step 1 Enable FCoE.

```
switch# configure terminal
switch(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
```

Step 2 (Optional) If you do not want to use the default Quality of Service (QoS) settings, specify your own policies:

Note Note: if you use custom policies, **class-fcoe** must be included in your QoS policies.

```
switch(config) system qos
switch(config-sys-qos)# service-policy type qos input fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing input
fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing output
fcoe-customized-out-policy-name
switch(config-sys-qos)# service-policy type network-qos fcoe-customized-nq-policy-name
```

Step 3 Verify that the FCoE policy maps can be found in the running configuration:

Note Note: If you specified customized QoS policy map names in Step 2, make sure you replace the default map names with your customized map names.

```
switch(config-sys-qos)# show policy-map system

  Type network-qos policy-maps
  ============================

  policy-map type network-qos system
    class type network-qos class-fcoe
      match qos-group 1

      pause no-drop
      mtu 2158
    class type network-qos class-default
      match qos-group 0

      mtu 1500

  Service-policy (qos) input:   system
    policy statistics status:   disabled

    Class-map (qos):   class-fcoe (match-any)
      Match: cos 3
      set qos-group 1

    Class-map (qos):   class-default (match-any)
      Match: any
      set qos-group 0

  Service-policy (queuing) input:   default-in-policy
    policy statistics status:   disabled

    Class-map (queuing):   class-fcoe (match-any)
      Match: qos-group 1
      bandwidth percent 50

    Class-map (queuing):   class-default (match-any)
      Match: qos-group 0
      bandwidth percent 50

  Service-policy (queuing) output:   default-out-policy
    policy statistics status:   disabled

    Class-map (queuing):   class-fcoe (match-any)
      Match: qos-group 1
      bandwidth percent 50
```

```
    Class-map (queuing):   class-default (match-any)
      Match: qos-group 0
      bandwidth percent 50
```

Quality of Service configuration on the Nexus 5000 series consists of three main constructs:

- Class-map and policy-map type qos: for classification purposes
- Class-map and policy-map type network: for network properties such as drop and no drop, queue size
- Class-map and policy-map type queueing: for bandwidth allocation

This exercise consists of changing the bandwidth allocation and the COS settings for FCoE.

Without proper configuration of class-fcoe in QoS, the following problems may occur:

- vFC interfaces do not come up (CNAs require advertisement of DCB parameters for FCoE)
- Drops noticed for I/Os

Note QoS has the following guidelines:

- A classification policy-map only applies in input
- A network policy-map applies globally (system)
- A queueing policy-map normally is meaningful in output, but since the exercise uses it to control the bandwidth allocation from CNA to the Cisco Nexus 5000 Series switch, in this case it is applied in input

Beginning in Cisco NX-OS Release 5.0(2)N1(1), you can modify the buffer allocation for no-drop classes:

```
switch(config-pmap-nq)# policy-map type network-qos nqos_policy
switch(config-pmap-nq)# class type network-qos nqos_class
switch(config-pmap-nq-c)# pause no-drop buffer-size <size> pause-threshold <threshold>
resume-threshold <threshold>
```

Step 4 Verify the FCoE system class is active:

```
switch(config-sys-qos)# show queuing interface ethernet 1/1
Ethernet1/1 queuing information:
TX Queuing
        qos-group sched-type oper-bandwidth
        0 WRR 50
        1 WRR 50
RX Queuing
        qos-group 0
        q-size: 370240, HW MTU: 1500 (1500 configured)
        drop-type: drop, xon: 0, xoff: 2314
        Statistics:
        Pkts received over the port : 0
        Ucastpkts sent to the cross-bar : 0
        Mcastpkts sent to the cross-bar : 0
        Ucastpkts received from the cross-bar : 0
        Pkts sent to the port : 0
        Pkts discarded on ingress : 0
        Per-priority-pause status : Rx (Inactive), Tx (Inactive)
        qos-group 1
        q-size: 79360, HW MTU: 2158 (2158 configured)
        drop-type: no-drop, xon: 128, xoff: 252
```

```
        Statistics:
        Pkts received over the port : 0
        Ucastpkts sent to the cross-bar : 0
        Mcastpkts sent to the cross-bar : 0
        Ucastpkts received from the cross-bar : 0
        Pkts sent to the port : 0
        Pkts discarded on ingress : 0
        Per-priority-pause status : Rx (Inactive), Tx (Inactive)
Total Multicast crossbar statistics:
        Mcastpkts received from the cross-bar : 0
```

Step 5 Repeat Step 1 through Step 4 on both upstream Cisco Nexus 5000 Series switches (CORE_N5k-1 and CORE_N5k-2 in this example).

Configuring VE Ports

FCoE VLAN and VSAN numbering in this example is as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20

Note There are two switches in Fabric A and two switches in Fabric B. The FCoE VLAN/VSANs must match between the switches in the same fabric in order to bring up the VE port between them.

Step 1 Configure the VSAN on the Nexus 5000 Series switch for Fabric A:

```
switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 10
```

Step 2 Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric A:

```
switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 10
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID        Translated VSAN ID      Association State
----------------      ------------------      -----------------
10                        10 Operational
switch(config-vlan)#
```

Step 3 Repeat Step 1 and Step 2 on the upstream Nexus 5000 Series switch in Fabric A.

Step 4 Configure the VSAN on the Nexus 5000 for Fabric B:

```
switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 20
```

Step 5 Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric B

```
switch(config)# vlan 20
switch(config-vlan)# fcoe vsan 20
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID        Translated VSAN ID      Association State
----------------      ------------------ -----------------
20                        20 Operational
switch(config-vlan)#
```

Step 6 Repeated Step 1 and Step 2 on the upstream Nexus 5000 Series switch in Fabric B.

Step 7 Configure the underlying 10-Gigabit Ethernet port that the vFC interface will be bound to. The VE port will use this interface as the physical transport for FCoE traffic between the two switches. This interface needs to be configured to trunk the appropriate FCoE VLAN as well as the Ethernet VLAN (in this example, we are using VLAN 200 to carry Ethernet traffic).

The 10-Gigabit Ethernet interfaces connecting the switches in this lab are shown in the topology above:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- PODX-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- PODX-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2

Configuration for both switches in Fabric A:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/15
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 10, 200
switch(config-if)#
```

Configuration for both switches in Fabric A:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/16
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 20, 200
switch(config-if)#
```

Step 8 Configure the vFC interface on the switch that will be bound to the VE port and add this vFC interface to VSAN 44 in the VSAN database:

The vFC numbers for the VE ports are as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- POD1-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- POD1-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2
- POD1-N5K-1 (Fabric A) uses vfc 15 and binds it to Ethernet 1/15
- POD1-N5K-2 (Fabric B) uses vfc 16 and binds it to Ethernet 1/16

Configuration for both switches in Fabric A:

```
switch(config)# int vfc 15
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 10
switch(config-if)# bind interface eth 1/15
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 10 interface vfc 15
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1           fc2/2           fc2/3           fc2/4
fc2/5           fc2/6           fc2/7           fc2/8
vsan 10 interfaces:
vfc15
vsan 4079(evfp_isolated_vsan) interfaces:
```

```
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db)# exit
```

Configuration for both switches in Fabric B:

```
switch(config)# int vfc 16
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 20
switch(config-if)# bind interface eth 1/16
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 20 interface vfc 16
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1             fc2/2             fc2/3             fc2/4
fc2/5             fc2/6             fc2/7             fc2/8
vsan 20 interfaces:
vfc16
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db)# exit
```

Note Don't forget that these interface configurations must be configured on both sides of the ISL connecting the two switches in the same fabric.

Step 9 Verify that the vFC is up and operational. By default, the vFC will show as trunking. Make sure that it is bound to the correct physical interface and that VSAN 44 shows as allowed and active as well as up on the vFC interface.

Verify both switches in Fabric A:

```
switch(config)# show int vfc 15
        vfc15 is trunking
        Bound interface is Ethernet1/15
            Hardware is Virtual Fibre Channel
            Port WWN is 20:0e:00:0d:ec:b4:43:7f
            Peer port WWN is 00:00:00:00:00:00:00:00
            Admin port mode is E, trunk mode is on
        snmp link state traps are enabled
            Port mode is TE
            Port vsan is 10
        Trunk vsans (admin allowed and active) (10)
            Trunk vsans (up)                        (10)
            Trunk vsans (isolated)                  ()
            Trunk vsans (initializing)              ()
            1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
            1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
              13 frames input, 1028 bytes
                0 discards, 0 errors
              13 frames output, 1180 bytes
                0 discards, 0 errors
        last clearing of "show interface" counters never
            Interface last changed at Sat Nov  6 17:58:39 2010
```

Verify both switches in Fabric A:

```
switch(config)# show int vfc 16
        vfc16 is trunking
        Bound interface is Ethernet1/16
            Hardware is Virtual Fibre Channel
            Port WWN is 20:0e:00:0d:ec:b4:43:7d
            Peer port WWN is 00:00:00:00:00:00:00:00
```

```
    Admin port mode is E, trunk mode is on
snmp link state traps are enabled
    Port mode is TE
    Port vsan is 20
Trunk vsans (admin allowed and active) (20)
    Trunk vsans (up)                        (20)
    Trunk vsans (isolated)                  ()
    Trunk vsans (initializing)              ()
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      13 frames input, 1028 bytes
        0 discards, 0 errors
      13 frames output, 1180 bytes
        0 discards, 0 errors
last clearing of "show interface" counters never
    Interface last changed at Sat Nov  6 17:58:39 2010
```

APPENDIX C

FCoE with vPC Configuration Example

Beginning with Cisco NX-OS Release 4.1(3)N1(1), the Cisco Nexus 5000 Series switch supports vPCs which can be configured to increase bandwidth and increased load-balancing to the Ethernet fabric. This appendix includes a sample configuration on how to configure FCoE when using vPCs on the Cisco Nexus 5000 Series switch and includes the following sections:

Figure C-1 shows the topology used in the examples described in this appendix.

Figure C-1 Nexus 5000 FCoE and vPC Lab Topology

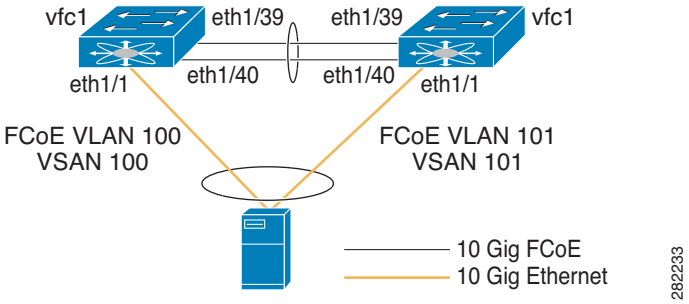

The configuration example includes the following parameters:

switchname: n5k-tme-1

switchname: n5k-tme-2

mgmt ip: 172.25.182.66

mgmt ip: 172.25.182.67

The configuration example includes the following hardware:

- Dell Server PE2950
- QLogic QLE8142 (Schultz) Generation-2 CNA
- 2 Cisco Nexus 5010 switches running Cisco NX-OS Release 4.1(3)N1(1)

The configuration example includes the following considerations and requirements:

1. Generation 2 CNAs that support DCBX are required.
2. Single host CNA port channel connection to a separate switch. FCoE interfaces will not be brought up if the port channel on a single switch contains more than one member port in a port channel or vPC.

3. Cisco NX-OS Release 4.1(3)N1(1) or a later release.
4. FC Features Package (FC_FEATURES_PKG) is necessary for running FCoE. If this is not installed, there will be a temporary license that will last 90 days.

This appendix includes the following sections:

Cisco Nexus 5000 Series Switch vPC Configuration Example

This example presumes that the basic configuration has been completed on the switch (for example, IP Address (mgmt0), switchname, and password for the administrator).

This example shows how to configure the basic vPC configuration. For more information on configuring vPC, refer to the *Cisco Nexus 5000 Series vPC Quick Configuration Guide*.

Note The configuration must be done on both peer switches in the vPC topology.

Step 1 Enable the vPC feature on both peer switches.

```
tme-n5k-1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
tme-n5k-1(config)# feature vpc
tme-n5k-1(config)#

tme-n5k-2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
tme-n5k-2(config)# feature vpc
tme-n5k-2(config)#
```

Step 2 Configure the vPC domain and peer-keep alive destinations:

```
tme-n5k-1(config)# vpc domain 2
tme-n5k-1(config-vpc-domain)# peer-keepalive destination 192.165.200.229

tme-n5k-2(config)# vpc domain 2
tme-n5k-2(config-vpc-domain)# peer-keepalive destination 192.165.200.230
```

Note In this set up, switch tme-n5k-1 has the mgmt IP address of 192.165.200.229 and switch tme-n5k-2 has the mgmt IP address of 192.165.200.230.

Step 3 Configure the port channel interface that will be used as the vPC peer-link:

```
tme-n5k-1(config)# int port-channel 1
tme-n5k-1(config-if)# vpc peer-link
```

Note The spanning tree port type is changed to network port type on vPC peer-link. This will enable STP Bridge Assurance on vPC peer-link provided that the STP Bridge Assurance (which is enabled by default) is not disabled.

```
tme-n5k-2(config)# int port-channel 1
tme-n5k-2(config-if)# vpc peer-link
```

Step 4 Verify that the peer-keepalive can be reached:

```
tme-n5k-1(config)# show vpc peer-keepalive
vPC keep-alive status           : peer is alive
--Destination                   : 172.25.182.167
--Send status                   : Success
--Receive status                : Success
--Last update from peer         : (0   ) seconds, (975 ) msec
tme-n5k-1(config)#

tme-n5k-2(config)# show vpc peer-keepalive
--PC keep-alive status           : peer is alive
--Destination                   : 172.25.182.166
--Send status                   : Success
--Receive status                : Success
--Last update from peer         : (0   ) seconds, (10336 ) msec
tme-n5k-2(config)#
```

Step 5 Add member ports to the vpc-peer link port channel and bring up the port channel interface:

```
tme-n5k-1(config-if-range)# int po 1
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)# exit
tme-n5k-1(config)# int eth 1/39-40
tme-n5k-1(config-if-range)# switchport mode trunk
tme-n5k-1(config-if-range)# channel-group 1
tme-n5k-1(config-if-range)# no shut
tme-n5k-1(config-if-range)#

tme-n5k-2(config-if-range)# int po 1
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)# exit
tme-n5k-2(config)# int eth 1/39-40
tme-n5k-2(config-if-range)# switchport mode trunk
tme-n5k-2(config-if-range)# channel-group 1
tme-n5k-2(config-if-range)# no shut
tme-n5k-2(config-if-range)#

tme-n5k-1(config-if-range)# show int po1
port-channel 1 is up
Hardware: Port-Channel, address: 000d.ecde.a92f (bia 000d.ecde.a92f)
MTU 1500 bytes, BW 20000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/39, Eth1/40
Last clearing of "show interface" counters never
1 minute input rate 1848 bits/sec, 0 packets/sec
1 minute output rate 3488 bits/sec, 3 packets/sec
tme-n5k-1(config-if-range)#

tme-n5k-2(config-if-range)# show int po1
port-channel1 is up
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae)  MTU 1500 bytes,
BW 20000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
```

```
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/39, Eth1/40
Last clearing of "show interface" counters never
minute input rate 1848 bits/sec, 0 packets/sec
minute output rate 3488 bits/sec, 3 packets/sec
tme-n5k-2(config-if-range)#
```

Step 6 Create the vPC and add member interfaces:

```
tme-n5k-1(config)# int po 11
tme-n5k-1(config-if)# vpc 11
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)# int eth 1/1
tme-n5k-1(config-if)# switchport mode trunk
tme-n5k-1(config-if)# channel-group 11
tme-n5k-1(config-if)# spanning-tree port type edge trunk
tme-n5k-1(config-if)#
```

Warning **Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration**

```
tme-n5k-2(config)# int po 11
tme-n5k-2(config-if)# vpc 11
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)# int eth 1/1
tme-n5k-2(config-if)# switchport mode trunk
tme-n5k-2(config-if)# channel-group 11
tme-n5k-2(config-if)# spanning-tree port type edge trunk
```

Warning **Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration.**

Note To run FCoE over a vPC topology, the port channel can only have a single member interface.

Note The vPC number configured under the port channel interface must match on both Nexus 5000 switches. The port channel interface number does not have to match on both switches.

Step 7 Verify that the vPC interfaces are up and operational:

```
tme-n5k-1(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
```

```
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
minute input rate 4968 bits/sec, 8 packets/sec
minute output rate 792 bits/sec, 1 packets/sec
tme-n5k-1(config-if)#

tme-n5k-2(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae)
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s
Beacon is turned off
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
minute input rate 4968 bits/sec, 8 packets/sec
minute output rate 792 bits/sec, 1 packets/sec
tme-n5k-1(config-if)#
```

Cisco Nexus 5000 Series Switch FCoE Configuration Example

Once the vPC is set up between the two Nexus 5000s, we can move on to configuring the FCoE topology. This cheat sheet presumes that basic configuration has been executed on the Nexus 5000 switch that will provide IP Address (mgmt0), switchname, password for admin, etc. and that the vPC configuration has been completed as outlined in the previous section. The following steps will walk through the basic FCoE configuration necessary to set up an FCoE topology in conjunction with the vPC topology.

Step 1 Enable FCoE on the Nexus 5000:

```
tme-n5k-1(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
tme-n5k-1(config)#

tme-n5k-2(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
fme-n5k-2(config)#
```

Note This can take a few moments to complete.

Step 2 Create a VSAN and map it to a VLAN that has been designated to carry FCoE traffic:

```
tme-n5k-1(config)# vsan database
tme-n5k-1(config-vsan-db)# vsan 100
tme-n5k-1(config-vsan-db)# exit
tme-n5k-1(config)# vlan 100
me-n5k-1(config-vlan)# fcoe vsan 100
tme-n5k-1(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
--------  --------  --------
100       100       Operational
tme-n5k-1(config-vlan)#

tme-n5k-2(config)# vsan database
tme-n5k-2(config-vsan-db)# vsan 101
tme-n5k-2(config-vsan-db)# exit
tme-n5k-2(config)# vlan 101
tme-n5k-2(config-vlan)# fcoe vsan 101
tme-n5k-2(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
--------  --------  --------
101       101       Operational
tme-n5k-2(config)#
```

Note VLAN and VSAN numbers are not required to be the same.

Step 3 Configure the VLANs that are allowed to transverse the vPC links:

```
tme-n5k-1(config)# int po 11
tme-n5k-1(config-if)# switchport trunk allowed vlan 1, 100
tme-n5k-1(config-if)# show int trunk
--------------------------------------------------------------------------------
Port          Native    Status        Port
--------------------------------------------------------------------------------
Eth1/1         1       trnk-bndl     Po11
Eth1/39        1       trnk-bndl     Po1
Eth1/40        1       trnk-bndl     Po1
Po1            1       trunking      --
Po11           1       trunking      --

--------------------------------------------------------------------------------
Port          Vlans Allowed on Trunk
--------------------------------------------------------------------------------
Eth1/1        1,100
Eth1/39       1-3967,4048-4093
Eth1/40       1-3967,4048-4093
Po1           1-3967,4048-4093
Po11          1,100

--------------------------------------------------------------------------------
Port          Vlans Err-disabled on Trunk
--------------------------------------------------------------------------------
Eth1/1        none
Eth1/39       100
Eth1/40       100
Po1           100
Po11          none

--------------------------------------------------------------------------------
Port          STP Forwarding
--------------------------------------------------------------------------------
Eth1/1        none
Eth1/39       none
Eth1/40       none
Po1           1
```

```
Po11          1,100
tme-n5k-1(config-if)#

tme-n5k-2(config)# int po 11
tme-n5k-2(config-if)# switchport trunk allowed vlan 1, 101
tme-n5k-2(config-if)# show int trunk
--------------------------------------------------------------------------
Port          Native    Status        Port
--------------------------------------------------------------------------
Eth1/1         1       trnk-bndl     Po11
Eth1/39        1       trnk-bndl     Po1
Eth1/40        1       trnk-bndl     Po1
Po1            1       trunking      --
Po11           1       trunking      --

--------------------------------------------------------------------------
Port          Vlans Allowed on Trunk
--------------------------------------------------------------------------
Eth1/1        1,101
Eth1/39       1-3967,4048-4093
Eth1/40       1-3967,4048-4093
Po1           1-3967,4048-4093
Po11          1,101

--------------------------------------------------------------------------
Port          Vlans Err-disabled on Trunk
--------------------------------------------------------------------------
Eth1/1        none
Eth1/39       101
Eth1/40       101
Po1           101
Po11          none

--------------------------------------------------------------------------
Port          STP Forwarding
--------------------------------------------------------------------------
Eth1/1        none
Eth1/39       none
Eth1/40       none
Po1           1
Po11          1,101
tme-n5k-2(config-if)#
```

Step 4 Create a virtual Fibre Channel interface (vfc) and add it to the VSAN that was created in the previous step:

```
tme-n5k-1(config)# int vfc 1
tme-n5k-1(config-if)# bind interface po11
Warning: VFC will not come up for pre-FIP CNA
tme-n5k-1(config-if)# no shut
tme-n5k-1(config-if)#

tme-n5k-2(config)# int vfc 1
tme-n5k-2(config-if)# bind interface po11
Warning: VFC will not come up for pre-FIP CNA
tme-n5k-2(config-if)# no shut
tme-n5k-2(config-if)#

tme-n5k-1(config)# vsan database
tme-n5k-1(config-vsan-db)# vsan 100 interface vfc 1
tme-n5k-1(config)# show vsan membership
vsan 1 interfaces:
fc2/1             fc2/2             fc2/3             fc2/4
fc2/5             fc2/6             fc2/7             fc2/8
```

```
vsan 100 interfaces:
vfc1

vsan 4079(evfp_isolated_vsan) interfaces:

vsan 4094(isolated_vsan) interfaces:
tme-n5k-1(config)#

tme-n5k-2(config)# vsan database
tme-n5k-2(config-vsan-db)# vsan 101 interface vfc 1
tme-n5k-2(config)# show vsan membership
vsan 1 interfaces:
fc2/1             fc2/2             fc2/3             fc2/4
fc2/5             fc2/6             fc2/7             fc2/8

vsan 101 interfaces:
vfc1

vsan 4079(evfp_isolated_vsan) interfaces:

vsan 4094(isolated_vsan) interfaces:
tme-n5k-2(config)#
```

Step 5 Verify that the vfc is up and operational:

```
tme-n5k-1(config-if)# show int brief
--------------------------------------------------------------------------------
Ethernet      VLAN   Type    Mode   Status   Reason           Speed
--------------------------------------------------------------------------------
Eth1/1        1      eth    trunk   up       none            10G(D)
Eth1/2        1      eth    access  up       none            10G(D)
Eth1/38       1      eth   access  down   SFP not inserted    10G(D)
Eth1/39       1      eth   trunk   up       none            10G(D)
Eth1/40       1      eth   trunk   up       none            10G(D)

--------------------------------------------------------------------------------
Port-channel VLAN  Type Mode   Status  Reason                 Speed
--------------------------------------------------------------------------------
Po1           1      eth   trunk  up      none            a-10G(D)  none
Po11          1      eth   trunk  up      none            a-10G(D)  none

--------------------------------------------------------------------------------
Port   VRF          Status IP Address                       Speed    MTU
--------------------------------------------------------------------------------
mgmt0  --           up     172.25.182.166                   1000     1500

--------------------------------------------------------------------------------
Interface   Vsan       Admin   Admin   Status       SFP   Oper  Oper  Port
--------------------------------------------------------------------------------
vfc1        100    F      on      up    --     F    auto   --
tme-n5k-1(config-if)#

tme-n5k-2(config-if)# show int brief
--------------------------------------------------------------------------------
Ethernet      VLAN   Type    Mode   Status   Reason           Speed    Port
--------------------------------------------------------------------------------
Eth1/1        1      eth    trunk   up       none            10G(D)   11
Eth1/2        1      eth    access  up       none            10G(D)   --
Eth1/38       1      eth   access  down   SFP not inserted    10G(D)   --
Eth1/39       1      eth   trunk   up       none            10G(D)   1
Eth1/40       1      eth   trunk   up       none            10G(D)   1

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

```
Port-channel VLAN  Type Mode   Status  Reason                 Speed  Protocol
------------------------------------------------------------------------
Po1           1      eth   trunk  up      none             a-10G(D)  none
Po11          1      eth   trunk  up      none             a-10G(D)  none

------------------------------------------------------------------------
Port   VRF          Status IP Address                          Speed    MTU
------------------------------------------------------------------------
mgmt0  --           up     172.25.182.167                      1000     1500

------------------------------------------------------------------------
Interface   Vsan       Admin   Admin   Status       SFP   Oper  Oper
------------------------------------------------------------------------
vfc1        101     F      on      up    --     F    auto   --
tme-n5k-2(config-if)#
```

Step 6 Verify that the virtual Fibre Channel interface has logged into the fabric:

```
tme-n5k-1# show flogi database
--------------------------------------------------------------------------
INTERFACE      VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------
vfc1     100    0x540000 21:00:00:c0:dd:11:2a:01 20:00:00:c0:dd:11:2a:01

Total number of flogi = 1.
tme-n5k-2# show flogi database
--------------------------------------------------------------------------
INTERFACE      VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------
vfc1           101  0x540000  21:00:00:c0:dd:11:2a:01 20:00:00:c0:dd:11:2a:01

Total number of flogi = 1.
```

Step 7 Verify that the vPC is up and operational:

```
tme-n5k-1(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908)
    MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA
    Port mode is trunk
    full-duplex, 10 Gb/s
    Beacon is turned off
    Input flow-control is off, output flow-control is off
    Switchport monitor is off
    Members in this channel: Eth1/1
    Last clearing of "show interface" counters never
    1 minute input rate 4968 bits/sec, 8 packets/sec
    1 minute output rate 792 bits/sec, 1 packets/sec

tme-n5k-2(config-if)# show vpc statistics vpc 11
port-channel11 is up
vPC Status: Up, vPC number: 11
Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae)
    MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA
    Port mode is trunk
    full-duplex, 10 Gb/s
    Beacon is turned off
```

```
Input flow-control is off, output flow-control is off
Switchport monitor is off
Members in this channel: Eth1/1
Last clearing of "show interface" counters never
1 minute input rate 4968 bits/sec, 8 packets/sec
1 minute output rate 792 bits/sec, 1 packets/sec
```

APPENDIX D

FCoE with Cisco Nexus 4000 Series Switch Configuration Example

This section includes a configuration example on how to configure an IBM blade server connecting to a Cisco Nexus 4000 Series switch which is then connected to a Cisco Nexus 5000 Series switch which accesses FC storage on a Cisco MDS 9000 Series Family switch using FCoE. Because the Cisco Nexus 4000 Series switch is a FIP snooping bridge, the FLOGI done by the CNAs do not login on the Cisco Nexus 4000 Series switch but onto the Cisco Nexus 5000 Series switch, which is the FCF. Creation of the vFC interface for the Cisco Nexus 4000 Series switch blade servers does not change whether the Cisco Nexus 5000 Series switch is in switching or NPV mode. Where the actual fabric login happens is determined by the mode on the Cisco Nexus 5000 Series switch.

- Cisco Nexus 5000 Series switch in switching mode—Login is on the Cisco Nexus 5000 Series switch.
- Cisco Nexus 5000 Series switch in NPV mode—Login will be on the Cisco MDS 9000 Series Family switch or any FC switch upstream with NPIV configured.

In this example, the Cisco Nexus 5000 Series switch in switching mode. Figure D-1 shows the topology used in the example.

Figure D-1 Nexus 4000 FCoE Lab Topology

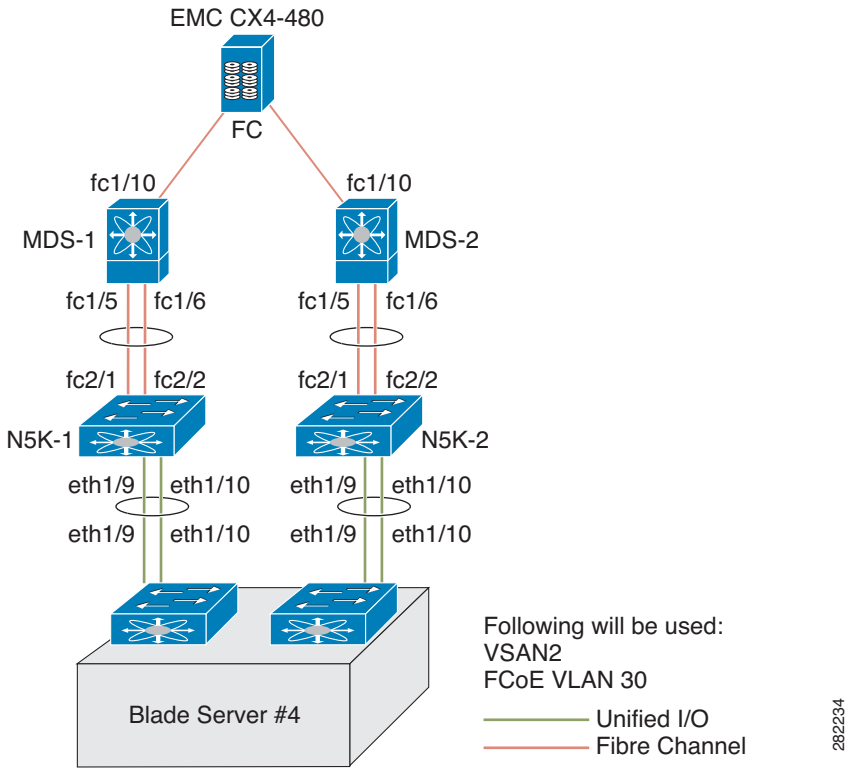

The following hardware was used:

- IBM Blade Chassis model BCH
- IBM HS22 blade server running Windows 2003 using the Qlogic QMI8142
- Cisco Nexus 4000 Series switch running Cisco NX-OS Release 4.1(2)E1(1)
- Cisco Nexus 5010 switch running Cisco NX-OS Release 4.1(3)N1(1)
- Cisco MDS 9124 Director switch running Cisco SAN-OS Release 4.1(3a)
- EMC CX4-480

This appendix includes the following sections:

Cisco Nexus 5000 Series Switch in Switching Mode

Before following the steps in this example, be sure to complete a basic configuration on the Cisco Nexus 5000 Series switch (for example, IP Address (mgmt0), switch name, and password for the administrator) and FCoE has not been enabled.

To use this configuration example in production, you must have the FC Features Package license installed otherwise there will be a temporary license that expires after 90 days. When the license expires, the feature is disabled.

On the Cisco Nexus 5000 Series switch, by default FCoE is not enabled.

This example shows how to enable FCoE:

```
n5k-2# show interface brief
--------------------------------------------------------------------------------
----------------------
Ethernet VLAN Type Mode Status Reason Speed Port
Interface                 Ch #
Eth1/1 1 eth access up none 10G(D) --
Eth1/2 1 eth access up none 10G(D) --
[snip]
Eth2/4 1 eth access down SFP not inserted 10G(D) --
--------------------------------------------------------------------------------
Port VRF Status IP Address Speed MTU
--------------------------------------------------------------------------------
mgmt0 -- up 172.25.182.164 1000 1500
```

Note There are no FC interfaces, even though there is a 4x4 GEM card installed in the Cisco Nexus 5010 switch.

```
n5k-2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
n5k-2(config)# feature fcoe
FC license checked out successfully fc_plugin extracted successfully FC plugin loaded
successfully FCoE manager enabled successfully FC enabled on all modules successfully
```

Note Beginning with Cisco NX-OS Release 4.1(3)N1(1), the switch does not need to be reboot when you enable FCoE. The Cisco Nexus 5000 Series switch is in switching mode by default when FCoE is enabled.

```
n5k-2(config)# show feature
Feature Name Instance State
fcsp 1 disabled
fcoe 1 enabled
fex 1 enabled

n5k-2(config)# show interface brief
--------------------------------------------------------------------------------
----------------------
Interface Vsan Admin Admin Status SFP Oper Oper Port
Mode Trunk Mode Speed Channel
Mode (Gbps)
--------------------------------------------------------------------------------
----------------------
fc2/1 1 auto on down swl -- --
fc2/2 1 auto on down swl -- --
fc2/3 1 auto on down swl -- --
fc2/4 1 auto on sfpAbsent -- -- --
```

```
--------------------------------------------------------------------------------
--------------------
Ethernet VLAN Type Mode Status Reason Speed Port
Interface Ch #
--------------------------------------------------------------------------------
--------------------
Eth1/1 1 eth access up none 10G(D) --
Eth1/2 1 eth access up none 10G(D) --
```

Note Use the **show interface brief** command to show the FC interfaces.

Configuring a SAN Port Channel on the Cisco Nexus 5000 Series Switch to the Cisco MDS Directory Series

This example shows how to configure a SAN port channel on the Cisco Nexus 5000 Series switch that is connected to a Cisco MDS 9000 Director. For redundancy, Cisco recommends that you create a SAN port channel from the FC interfaces.

Step 1 Configure a SAN port channel on the Cisco Nexus 5000 Series switch.

```
Fn5k-2# configure terminal
n5k-2(config)# interface san-port-channel 1
n5k-2(config-if)# interface fc2/1-2
n5k-2(config-if)# channel-group 1
```

Note After you add fc2/1 fc2/2 to san-port-channel 1 you need to disable the port channel. This must also be done on the switch at the other end of the port channel. Then, shut the interfaces at both ends to bring them up.

```
n5k-2(config-if)# no shut
n5k-2(config-if)# interface san-port-channel 1
n5k-2(config-if)# no shut
n5k-2(config-if)# show san-port-channel database
san-port-channel 1
Administrative channel mode is on Operational channel mode is on Last membership
update is successful 2 ports in total, 0 ports up Age of the port-channel is
0d:00h:17m:14s
Ports: fc2/1 [down] fc2/2 [down]
n5k-2(config-if)#
```

Note The SAN port channel is currently down because the Cisco MDS 9000 Series Director has not been configured.

Step 2 Configure the Cisco MDS 9124 switch to create a port channel between the Cisco Nexus 5000 Series switch and the Cisco MDS 9124 switch.

Note With the SAN port channel on the Cisco Nexus 5000 configured to the MDS, you will need to perform the same configuration on the Cisco MDS 9000 Series switch. A SAN port channel configuration on the Cisco MDS 9000 Series switch is called a port channel.

```
mds9124-2# configure terminal
```

```
mds9124-2(config)# interface port-channel 1
mds9124-2(config-if)# interface fc1/5, fc1/6
mds9124-2(config-if)# channel-group 1 force
```

Note After you add fc1/5 fc1/6 to port-channel 1 you need to disable the port channel. This must also be done on the switch at the other end of the port channel. Then, shut the interfaces at both ends to bring them up.

Note

```
mds9124-2(config-if)# no shut
mds9124-2(config-if)# interface port-channel 1
mds9124-2(config-if)# no shut
```

Step 3 Verify that the SAN port channel on the Cisco Nexus 5000 Series switch is up and running. Use the **show san-port-channel database** command to show the SAN port channel configuration.

```
n5k-2(config-if)# show san-port-channel database
san-port-channel 1
Administrative channel mode is on
Operational channel mode is on
Last membership update is successful
2 ports in total, 2 ports up
First operational port is fc2/2
Age of the port-channel is 0d:00h:25m:10s
Ports: fc2/1 [up]
fc2/2 [up] *
```

Configuring a Port Channel on a Cisco Nexus 5000 Series Switch to a Cisco Nexus 4000 Series Switch

This example shows how to configure a port channel on the Cisco Nexus 5000 Series switch that is connected to the Cisco Nexus 4000 Series switch.

Step 1 Configure the port channel on the Cisco Nexus 5000 Series switch.

The port channel is configured to provide redundancy for traffic coming from the Cisco Nexus 4000 Series switch to the Cisco Nexus 5000 Series switch. In this example, all VLANs can traverse the port channel. The FCoE VLAN and the native VLAN must be allowed to traverse the port channel.In production environments, Network Administrators may designate other VLANs to traverse this network.

```
n5k-2# configure terminal
n5k-2(config)# feature lacp
n5k-2(config)# interface port-channel 2 mode active
n5k-2(config-if)# interface eth1/9-10
n5k-2(config-if)# channel-group 2
n5k-2(config)# interface port-channel 2
n5k-2(config-if)# switchport mode trunk
n5k-2(config-if)# no shut
n5k-2#
```

Step 2 Configure the port channel on the Cisco Nexus 4000 Series switch.

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# feature lacp
bch1-n4k-b9(config)# interface port-channel 20
```

```
bch1-n4k-b9(config-if)# interface eth1/15-16
bch1-n4k-b9(config-if)# channel-group 2 mode active
bch1-n4k-b9(config)# interface port-channel 2
bch1-n4k-b9(config-if)# switchport mode trunk
bch1-n4k-b9(config-if)# no shut
bch1-n4k-b9(config-if)#
```

Configuring a Virtual Fibre Channel Interface on a Cisco Nexus 4000 Series Switch

This example shows how to configure a vFC interface on a Cisco Nexus 4000 Series switch.

Step 1 On the Cisco Nexus 5000 Series switch, configure a VSAN to match the production VSAN on the Cisco MDS 9000 Series switch. This is a one-time configuration.

Step 2 On the Cisco Nexus 5000 Series switch, configure an FCoE VLAN to map to the VSAN (VLAN-to-VSAN mapping). This is one-time configuration.

Step 3 On the Cisco Nexus 4000 Series switch, configure a FIP snooping VLAN that matches the FCoE VLAN on the Nexus 5000 Series switch. This is a one-time configuration.

Step 4 On the Cisco Nexus 4000 Series switch, configure the uplinks to allow FCoE traffic (FIP snooping).

Step 5 On the Cisco Nexus 4000 Series switch blade server, configure the Ethernet interfaces for FCoE traffic.

Step 6 On the Cisco Nexus 5000 Series switch, configure the vFCs.

Step 7 On the Cisco Nexus 4000 Series switch blade server, bind the vFC to the MAC address of the blade server.

Step 8 Verify that the vFC is in the correct VSAN.

Note Completing the above tasks ensure that the connection to an FCoE CNA on the blade server from the Nexus 4000 is successful.

Configuring a VSAN on the Cisco Nexus 5000 Series Switch

You can configure a VSAN on the Cisco Nexus 5000 Series switch using Fabric Manager, Device Manager, or the CLI. This example shows CLI configuration tasks and Fabric Manager or Device Manager GUI tasks.

This example shows the storage on the Cisco MDS 9000 Series resides on VSAN 2. Configure the VSAN to ensure that the vFCs configured on the Cisco Nexus 5000 Series switch can communicate with the storage device.

```
n5k-2# configure terminal
n5k-2(config)# vsan database
n5k-2(config-vsan-db)# vsan 2
n5k-2(config-vsan-db)# show vsan vsan 1 information
name:VSAN0001 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:up
vsan 2 information
```

```
name:VSAN0002 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
 operational state:down
        vsan 4079:evfp_isolated_vsan
        vsan 4094:isolated_vsan
```

Configuring An FCoE VLAN on the Cisco Nexus 5000 Series Switch

You can configure a VLAN and then map the VLAN to a particular VSAN using the CLI. Fabric Manager and Device Manager can not be used for this configuration. Cisco recommends that you configure a separate VLAN for FCoE traffic and separate VLANs for standard Ethernet traffic.

This example shows how to create the FCoE VLAN:

```
n5k-2# configure terminal
        n5k-2(config)# vlan 30
        n5k-2(config-vlan)# fcoe vsan 2
        n5k-2(config-vlan)# show vlan fcoe
VLAN
VSAN
Status
--------
--------
--------
30
2
Operational
```

Configuring a FIP Snooping VLAN on the Cisco Nexus 4000 Series Switch

On the Cisco Nexus 4000 Series switch, by default the FIP snooping feature is disabled. Cisco recommends that during the basic configuration, when prompted, you should enable FCoE and FIP snooping and configure, for example, the appropriate Class of Service (CoS) no drop, MTU, and QoS, without having to manually configure these features after the initial configuration.

The example shows how to verify that FIP snooping is enabled:

```
bch1-n4k-b9# show feature
Feature Name Instance State
tacacs 1 disabled lacp 1 enabled [snip] fipsm 1 enabled
```

With the FCoE VLAN configured on the Cisco Nexus 5000 Series switch as VLAN 30, then the same VLAN number must be used to create the VLAN on the Cisco Nexus 4000 Series switch and the VLAN must be configured as a FIP snooping VLAN.

This example shows how to configure the VLAN on the Cisco Nexus 4000 Series switch:

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# vlan 30
bch1-n4k-b9(config-vlan)# fip-snooping enable
```

Configuring the Cisco Nexus 4000 Series Switch Uplinks To Allow FCoE Traffic

In this example, we have already created the port channel that allows all VLANs to traverse the uplink between the Cisco Nexus 4000 Series switch and the Cisco Nexus 5000 Series switch from the previous section. The uplink (in this case a port channel) must be enabled to do FIP snooping with a port type mode of fcf.

This example shows how to configure the uplink:

```
bch1-n4k-b9# configure terminal
bch1-n4k-b9(config)# interface port-channel 20
bch1-n4k-b9(config-if)# fip-snooping port-mode fcf
```

Configuring Blade Server Ethernet Interfaces on the Cisco Nexus 4000 Series Switch For FCoE Traffic

You can configure the blade server using the CLI. Fabric Manager and Device Manager can not be used for this configuration.

Ensure that the FCoE VLAN (VLAN 30) can traverse the Ethernet interface on the blade server (Ethernet 1/4). In most cases, the CNA ports allow for both regular Ethernet traffic and FCoE traffic that resides on different VLANs. By default, all Ethernet interfaces on the Cisco Nexus 4000 Series switch is in access mode and resides on VLAN 1.

This example shows how to configure the Ethernet interface to allow multiple VLANs (trunk):

```
bch1-n4k-b9#configure terminal
bch1-n4k-b9(config)#interface ethernet 1/4
bch1-n4k-b9(config-if)# switchport mode trunk
bch1-n4k-b9(config-if)# switchport trunk allowed vlan 1,30
```

Note The above command is not needed but if you want to specify the allowed VLANs, make sure the FCoE VLAN is on the allowed list as shown in the example.

```
bch1-n4k-b9(config-if)# spanning-tree port type edge trunk
Warning: Edge port type (portfast) should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this  interface when
edge port type (portfast) is enabled, can cause temporary bridging loops.
Use with CAUTION

Creating vFC Interfaces on the Nexus 5000 - CLI
```

When the trunk configuration is complete, create the vFC interface on the Cisco Nexus 5000 Series switch. You can use Device Manager or the CLI to configure the vFC interface.

Because the CNA is connected on Ethernet interface eth1/4 on the Cisco Nexus 4000 Series switch and is not physically connected to the Cisco Nexus 5000 Series switch, you must bind the vFC to the MAC address of the CNA that is doing FCoE. At this time, Qlogic is the only vendor that does FCoE on the blade server that is interoperable with the Cisco Nexus 4000 Series switch. Qlogic provides 2 separate MAC addresses, one for the standard Ethernet traffic and another specifically for FCoE.

This example shows how to identify the MAC address from the specific blade server in the IBM blade chassis.

```
bch1-n4k-b9# show fip-snooping vlan-discovery
 Legend:
Interface VLAN FIP MAC
Eth1/4 1 00:c0:dd:04:0c:df
 Eth1/5 1 00:c0:dd:04:0d:13
```

Use the MAC address that has been identified on the blade server to create the vFC for this blade server on the Cisco Nexus 5000 series switch.

This example shows that the vFC is moved into VSAN 2. As a best practice in creating the vFC number to devices on the Cisco Nexus 4000 Series switch, you should create a numbering scheme that can easily identify where the vFCs are mapped to which blade server on which blade chassis. For this example, we are using the blade server in slot 4 on the first IBM blade chassis, which we have named BCH1. In this example, the vFC for this blade server is interface vfc104.

```
n5k-2# configure terminal
n5k-2(config)# interface vfc 104
n5k-2(config-if)# bind mac-address 00:c0:dd:04:0c:df
n5k-2(config-if)# no shutdown
n5k-2(config-if)# show vsan membership
vsan 1 interfaces:
fc2/1 fc2/2 fc2/3 fc2/4 san-port-channel 1 vfc104
vsan 2 interfaces:
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:
n5k-2(config-if)# vsan database ð
this will get to the VSAN database
 n5k-2(config-vsan-db)# vsan 2 interface vfc104
n5k-2(config-vsan-db)# show vsan membership
vsan 1 interfaces: fc2/1 fc2/2 fc2/3 fc2/4 san-port-channel 1
 vsan 2 interfaces:
vfc104
vsan 4079(evfp_isolated_vsan) interfaces:
n5k-2# show interface vfc104
vfc104 is up
```

Bound MAC is 00:c0:dd:04:0c:df FCF priority is 128 Hardware is Virtual Fibre Channel Port WWN is 20:67:00:0d:ec:b2:b9:bf Admin port mode is F, trunk mode is on snmp link state traps are enabled Port mode is F, FCID is 0xcd0000 Port vsan is 2 [snip]

Configuring The vFC Interface Using Device Manager

This example shows how to use Device Manager to create the vFC interface.

Step 1 Open Device Manager and login to the Cisco Nexus 5000 Series switch.

Figure D-2 Device Manager Login Window

Step 2 From the Device Manager menu, choose Interface > Virtual Interfaces > Fibre Channel to configure one vFC. You can also use the Quick Configuration Tool to configure multiple vFCs and bind them to physical interfaces at one time.

Figure D-3 Device Manager Menu

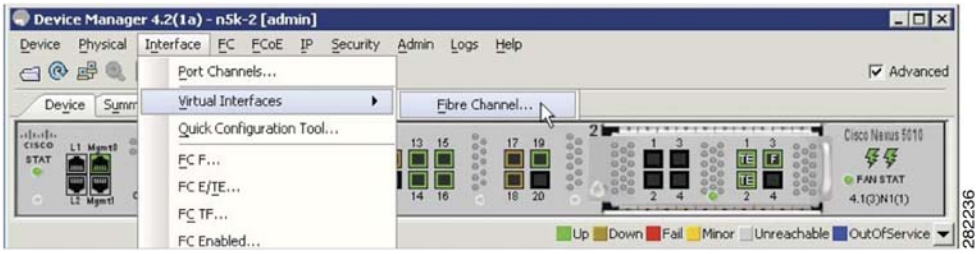

Step 3 From the Virtual FC Interfaces window, click Create to create the vFC.

Figure D-4 Virtual FC Interfaces Window

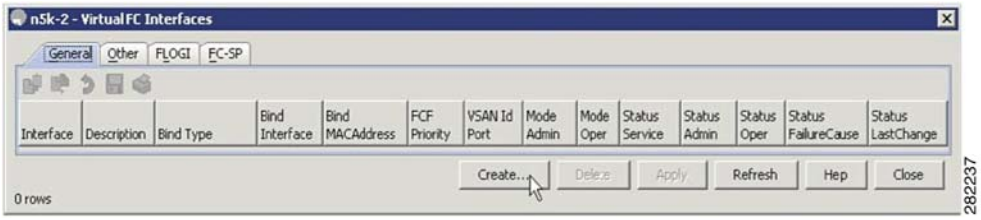

Step 4 In the Create Virtual FC Interfaces General window, enter the VFC Id, Bind Type and the interface (physical or MAC address depending on the bind type) and click Create. The window is redisplayed showing the vFCs with the new vFC ID.

Figure D-5 Create Virtual FC Interfaces General Window

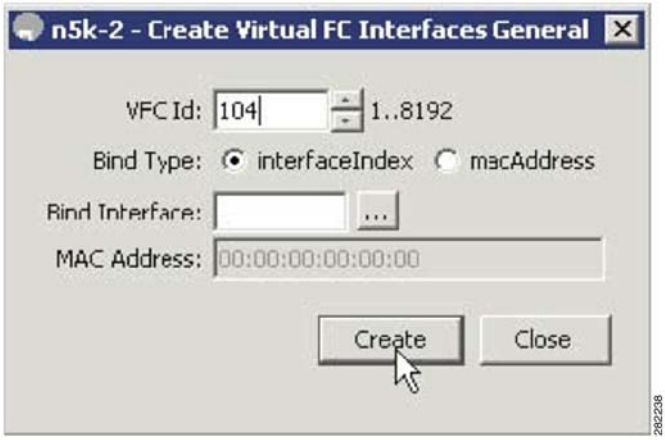

Note As a best practice, create a vFC that is recognizable of the vFC back to the blade server. For example, 104 correlates to BCH1 on blade server 4.

Step 5 From the Virtual FC Interfaces window, choose Bind Type > macAddress.

Figure D-6 Changing The Bind Type From Interface to Mac Address

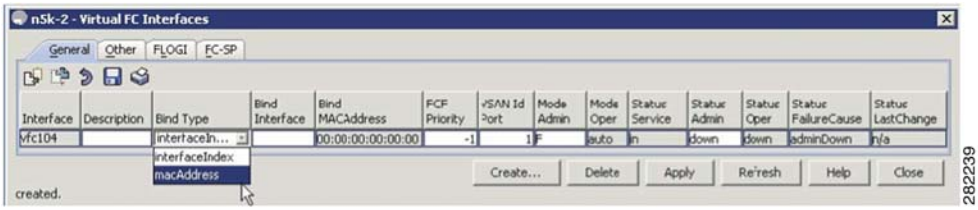

Once the Bind Type is set to macAddress, you can enter the MAC address for the blade server in the Bind MAC Address column. In this example, 00:c0:dd:04:0c:df is the MAC address. By default, the VSAN membership is set down and VSAN 1. You can edit these sections for example, VSAN 2 and up).

Step 6 Click on Apply to commit the changes and then click Refresh to validate the vFC is up.

Figure D-7 The Configured vFC MAC Address in Device Manager

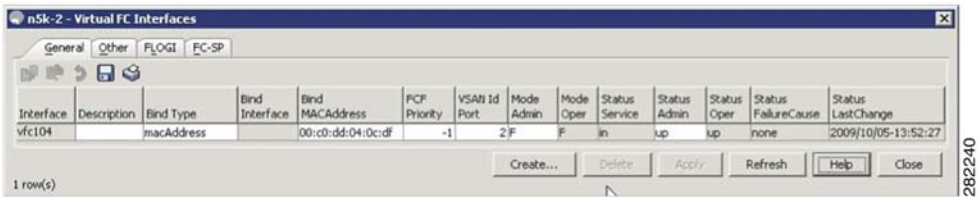

This completes the configuration of FCoE on the Cisco Nexus 4000 Series switch uplinked to the Cisco Nexus 5000 Series switch. The fabric management, for example, zoning and LUN masking, is managed with the existing SAN administrators tools. The vFC appear in Fabric Manager as a normal FC device but instead of a solid line to the host, a dash line is shown from the Cisco Nexus 5000 Series switch to the host.

Figure D-8 Fabric Manager View With FCoE Devices

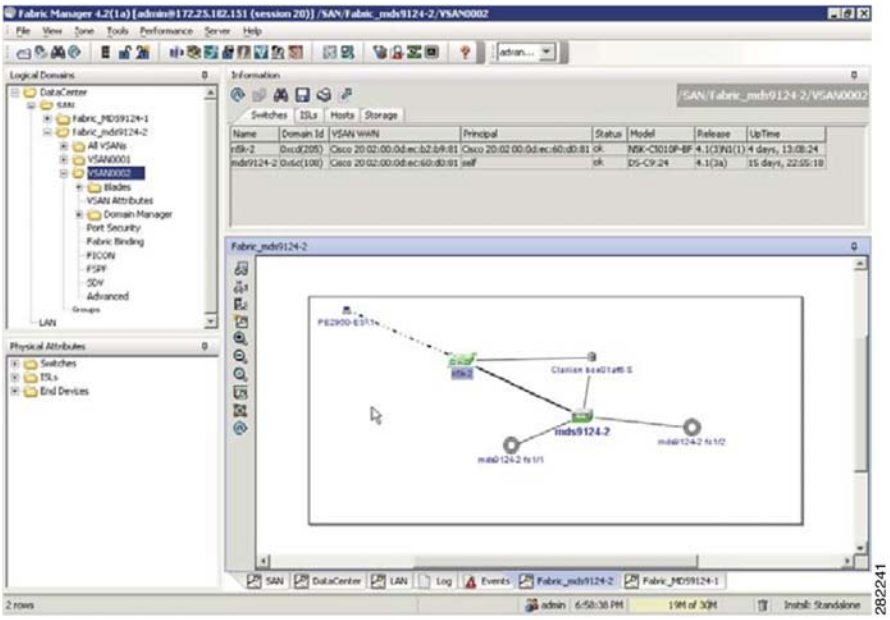

INDEX

Numerics

A

B

C

D

E

F

M

N

P

Q

R

S

T

U

V