Cisco Nexus 7000 Series FCoE Configuration Guide 8.x

First Published: 2016-12-21

Last Modified: 2024-10-08

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

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2016 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

Preface vii

Preface vii

Audience vii

Document Conventions vii

Related Documentation viii

Documentation Feedback ix

Communications, Services, and Additional Information ix

CHAPTER 1

New and Changed Information 1

Change Summary 1

CHAPTER 2

Overview 3

Licensing Requirements 3

Information About Fibre Channel Over Ethernet 3

Fibre Channel Forwarder 4

Fibre Channel Bridge 4

CHAPTER 3

Preparing the Switch for Configuring FCoE 5

Information About FCoE 5

FCoE and FIP 5

FCoE Initiation Protocol 5

FIP Virtual Link Instantiation 6

FCoE Frame Format 6

VLAN Tagging for FCoE Frames 6

FIP Ethernet Frame Format 6

DCBX 7

Data Center Bridging Exchange Protocol 7
Lossless Ethernet 7
Logical Link Up/Down 8
Converged Network Adapters 8
STP Lite 8
Default Settings for FCoE 9
Guidelines and Limitations 9
FCoE 9
Multi-Hop FCoE Distance Configuration 11
Enabling and Disabling FCoE 12
Enabling FCoE 12
Disabling FCoE 13
Configuring the FC-Map 15
Configuring the Fabric Priority 16
Setting the Advertisment Interval 16
Disabling LAN Traffic on an FCoE Link 17
Additional References for FCoE 18

CHAPTER 4 **Storage VDC 21**
Information About Storage VDC 21
Configuring FCoE VDCs 22
Licensing an FCoE Module 22
Creating a Dedicated Storage VDC 23
Allocating the FCoE VLAN Range 25
Allocating Shared Interfaces 26
Example: Storage VDC Configuration 28

CHAPTER 5 **Configuring FCoE VLANs and Virtual Fibre Channel Interfaces 31**
Information About Virtual Fibre Channel (VFC) Interfaces 31
VF Port 31
VE Port 31
VNP Ports 32
Default Settings for FCoE 32
Guidelines and Limitations 32

FCoE 32
Configuring Virtual Interfaces 34
Mapping a VSAN to a VLAN 34
Creating a Virtual Fibre Channel Interface 36
Creating a Virtual Fibre Channel Port Channel Interface 37
Associating a Virtual Fibre Channel Interface to a VSAN 39
Enabling VE Loopback Configuration 40
Verifying the Virtual Fibre Channel Interface 41
Example:Mapping VSANs to VLANs 43
Verifying the FCoE Configuration 44
Additional References for FCoE 45

CHAPTER 6 **Configuring Dynamic FCoE Using FabricPath** 47
Information About Dynamic FCoE Using FabricPath 47
SAN A/B Separation 49
Load-Balancing FCoE Traffic on a Dynamic VFC 51
Supported Dynamic FCoE Using FabricPath Topologies 51
Prerequisites for Dynamic FCoE Using FabricPath 51
Guidelines and Limitations for Dynamic FCoE Using FabricPath 51
Configuration Topology Example 52
Configuring Dynamic FCoE Using FabricPath 53
Instantiation and Initialization of Dynamic VFC 53
Verifying the Dynamic FCoE Using FabricPath Configuration 54
Configuration Output Examples for Dynamic FCoE Using FabricPath 57

CHAPTER 7 **FCoE Over FEX** 59
Information About FCoE Over FEX 59
FCoE Over FEX With Physical Port vPC 59
LAN Shutdown 59
Default Settings for FCoE 60
Guidelines and Limitations 60
FCoE Over FEX 60
FCoE over FEX with Physical Port vPC 61
LAN Shutdown 61

FCoE Over FEX 61
Bringing up FEX in Ethernet VDC 61
Configuring FCoE over FEX 61
Binding VFC to Shared FEX Interface Explicitly 67
Binding VFC to Shared FEX Interface Implicitly 68
Verifying the FCoE over FEX Configuration 69
Example: FCoE Over FEX Configuration 71
Additional References for FCoE over FEX 72

APPENDIX A
Feature History by Platform 75
Supported Features 75

APPENDIX B
Guidelines and Limitations 77
Verified Limits for FCoE 77
FCoE 77
Multi-Hop FCoE Distance Configuration 79

Preface

The preface contains the following sections:

- Preface, on page vii

Preface

This preface describes the audience, organization, and conventions of the Book Title. 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 7000 Series Platform switches.

Document Conventions

Note

- As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.
- The Guidelines and Limitations section contains general guidelines and limitations that are applicable to all the features, and the feature-specific guidelines and limitations that are applicable only to the corresponding feature.

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.

Convention	Description
[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.

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.

This document uses the following conventions:

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

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

Related Documentation

Documentation for Cisco Nexus 7000 Series Switches is available at:

- Configuration Guides

http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-installation-and-configuration-guides-list.html

- Command Reference Guides

  http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-command-reference-list.html

- Release Notes

  http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-release-notes-list.html

- Install and Upgrade Guides

  http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-installation-guides-list.html

- Licensing Guide

  http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-licensing-information-listing.html

Documentation for Cisco Nexus 7000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders is available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-2000-series-fabric-extenders/products-installation-and-configuration-guides-list.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus7k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco DevNet.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

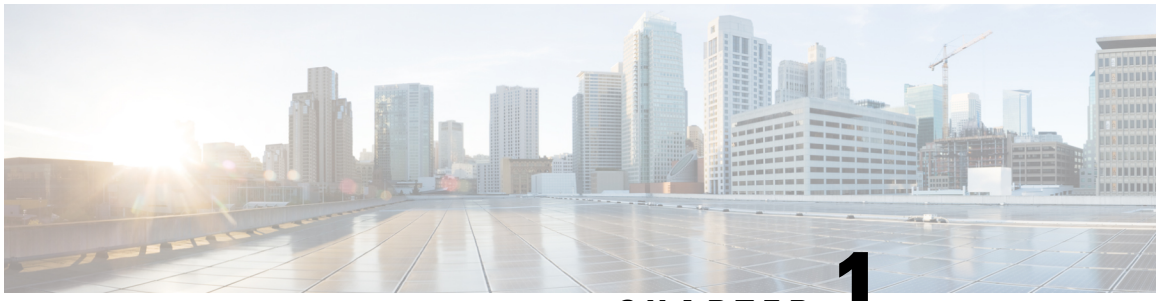

CHAPTER 1

New and Changed Information

• Change Summary, on page 1

Change Summary

This book does not contain any new feature for Cisco MDS NX-OS Release 8.1(1).

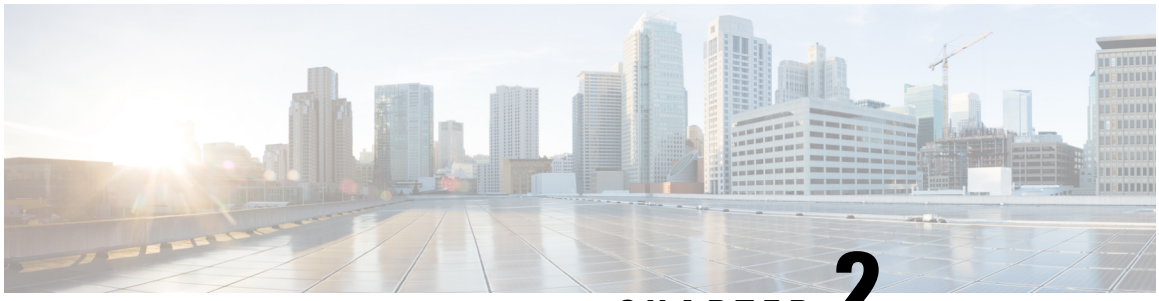

CHAPTER 2

Overview

This chapter contains the following sections:

- Licensing Requirements, on page 3
- Information About Fibre Channel Over Ethernet, on page 3
- Fibre Channel Forwarder, on page 4
- Fibre Channel Bridge, on page 4

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Information About Fibre Channel Over Ethernet

Fibre Channel over Ethernet (FCoE) allows Fibre Channel traffic to be encapsulated over a physical Ethernet link. FCoE and FIP frames use a unique EtherType so that FCoE traffic and standard Ethernet traffic can be carried on the same link.

Classic Ethernet is a best-effort protocol, which means that in a congested network, Ethernet discards packets and relies on higher level protocols to provide retransmission and other reliability mechanisms.

Fibre Channel traffic requires a lossless transport layer; as a data storage protocol, it is unacceptable to lose a single data packet. Native Fibre Channel implements a lossless service at the transport layer using a buffer-to-buffer credit system.

Ethernet links on Cisco NX-OS switches provide two mechanisms to ensure lossless transport for FCoE traffic: link-level flow control and priority flow control.

IEEE 802.3x link-level flow control allows a congested receiver to signal the far end to pause the data transmission for a short period of time. The pause functionality is applied to all traffic on the link.

The priority flow control (PFC) feature on Cisco NX-OS platforms applies pause functionality to specific classes of traffic on the Ethernet link. For example, PFC can provide lossless service for the FCoE traffic and best-effort service for the standard Ethernet traffic using IEEE 802.1p traffic classes.

Fibre Channel Forwarder

Fibre Channel traffic requires a lossless transport layer, which means that it requires a data storage protocol that does not lose a single data packet. A FCoE Forwarder (FCF) communicates with FCoE end devices, such as converged network adapters (CNAs). The FCF accepts the Fibre Channel frame that is encapsulated in an Ethernet packet and forwards that packet over a VLAN across an Ethernet network to a remote FCoE end device. The FCF has a unique FC-MAC address in the FCoE network. An FCF can also assign FC address IDs to the CNAs. An FCF can assign Fabric Provide MAC Addresses (FPMA) to the CNAs consisting of the FC-Map Value for the Fabric and the Fibre Channel IDs (FC IDs) assigned during Fabric Login.

Fibre Channel Bridge

A Fibre Channel Bridge connects an FCoE network to a Fibre Channel network. A Fibre Channel Bridge decapsulates an FCoE frame and sends the Fibre Channel frame to the Fibre Channel network. A Fibre Channel Bridge also encapsulates FC frames from a Fibre Channel network and forwards them to the FCoE network.

Note The Cisco Nexus 7000 series switches do not offer native Fibre Channel ports and do not support FC Bridge functions.

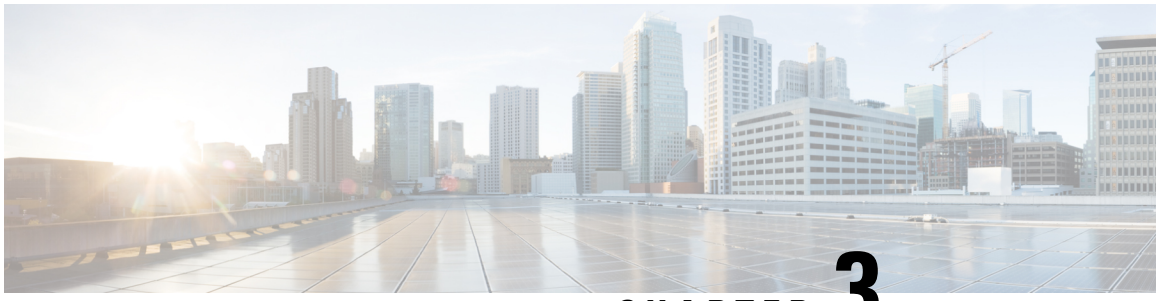

CHAPTER 3

Preparing the Switch for Configuring FCoE

- Information About FCoE, on page 5
- Default Settings for FCoE, on page 9
- Guidelines and Limitations, on page 9
- Enabling and Disabling FCoE, on page 12
- Configuring the FC-Map , on page 15
- Configuring the Fabric Priority, on page 16
- Setting the Advertisment Interval, on page 16
- Disabling LAN Traffic on an FCoE Link, on page 17
- Additional References for FCoE, on page 18

Information About FCoE

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

Note Lossless behavior on Ethernet is provided by using a priority flow control (PFC) mechanism that prevents packet loss during congestion conditions.

The Cisco NX-OS software supports T11-compliant FCoE on all 10-Gigabit and 40-Gigabit Ethernet interfaces.

FCoE and FIP

FCoE Initiation Protocol

The FCoE Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN. Cisco NX-OS switches support the Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) protocol for T11-compliant Generation 2 CNAs.

The following switches do not support Pre-FIP on Generation 1 CNAs:

FIP Virtual Link Instantiation

FIP is used to perform device discovery, initialization, and link maintenance. FIP performs the following tasks:

- FIP VLAN discovery— Discovers the FCoE VLAN that will be used by all other FIP protocols as well as by the FCoE encapsulation for Fibre Channel payloads on the established virtual link. FIP VLAN discovery occurs in the native VLAN used by the initiator or target to exchange Ethernet traffic. The FIP VLAN discovery protocol is the only FIP protocol running on the native VLAN; all other FIP protocols run on the discovered FCoE VLANs.
- FIP FCF discovery—When a FCoE device is connected to the fabric, it sends a Discovery Solicitation message. A Fibre Channel Forwarder (FCF) or a switch responds to the message with a Solicited Advertisement that provides an FCF MAC address to use for subsequent logins.
- FCoE virtual link instantiation— FIP defines the encapsulation of fabric login (FLOGI), fabric discovery (FDISC), logout (LOGO), and exchange link parameter (ELP) frames along with the corresponding reply frames. The FCoE devices use these messages to perform a fabric login.
- FCoE virtual link maintenance— FIP periodically sends maintenance messages between the switch and the CNA to ensure that the connection is still valid.

FCoE Frame Format

FCoE is implemented when the switch encapsulates a Fibre Channel frame in an Ethernet packet with a dedicated Ethernet type, 0x8906. The packet has a 4-bit version field. The other header fields in the frame (the source and destination MAC addresses, VLAN tags, and frame markers) are all standard Ethernet fields. Reserved bits pad the FCoE frame to the IEEE 802.3 minimum packet length of 64 bytes.

A Fibre Channel frame consists of 36 bytes of headers and up to 2112 bytes of data for a total maximum size of 2148 bytes. The encapsulated Fibre Channel frame has all the standard headers, which allow it to be passed to the storage network without further modification. To accommodate the maximum Fibre Channel frame in an FCoE frame, the class-fcoe is defined with a default Maximum Transmission Unit (MTU) of 2240 bytes.

VLAN Tagging for FCoE Frames

The Ethernet frames that are sent by the switch to the adapter include the IEEE 802.1Q tag. This tag includes a field for the class of service (CoS) value used by the priority flow control (PFC). The IEEE 802.1Q tag also includes a VLAN field.

The switch expects frames from a FIP T11-compliant CNA to be tagged with the VLAN tag for the FCoE VLAN. Frames that are not correctly tagged are discarded.

Note You cannot map VLAN 1 or the native VLAN to an FCoE VSAN.

FIP Ethernet Frame Format

FIP is encapsulated in an Ethernet packet with a dedicated EtherType, 0x8914. The packet has a 4-bit version field. Along with the source and destination MAC addresses, the FIP packet also contains a FIP operation code and a FIP operation subcode. The following table describes the FIP operation codes and subcodes.

Table 1: FIP Operation Codes

FIP Operation Code	FIP Subcode	FIP Operation
0x0001	0x01	Discovery Solicitation
	0x02	Discovery Advertisement
0x0002	0x01	Virtual Link Instantiation Request
	0x02	Virtual Link Instantiation Reply
0x0003	0x01	FIP Keepalive
	0x02	FIP Clear Virtual Links
0x0004	0x01	FIP VLAN Request
	0x02	FIP VLAN Notification

DCBX

Data Center Bridging Exchange Protocol

The Data Center Bridging Exchange (DCBX) protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX end points exchange request and acknowledgment messages. For flexibility, parameters are coded in a type-length-value (TLV) format. Cisco NX-OS switches support the Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) is supported on all T11-compliant Generation 2 CNAs.

DCBX runs on the physical Ethernet link between the switch and the CNA. By default, DCBX is enabled on Ethernet interfaces. When an Ethernet interface is brought up, the switch automatically starts to communicate with the CNA.

During the normal operation of FCoE between the switch and the CNA, DCBX provides link-error detection.

DCBX is also used to negotiate capabilities between the switch and the CNA and to send configuration values to the CNA.

The CNAs that are connected to a switch are programmed to accept the configuration values that are sent by the switch, allowing the switch to distribute configuration values to all attached CNAs, which reduces the possibility of configuration errors and simplifies CNA administration.

Lossless Ethernet

Standard Ethernet is a best-effort medium which means that it lacks any form of flow control. In the event of congestion or collisions, Ethernet drops packets. The higher level protocols detect the missing data and retransmit the dropped packets.

To properly support Fibre Channel, Ethernet has been enhanced with a priority flow control (PFC) mechanism.

Logical Link Up/Down

On a native Fibre Channel link, some configuration actions (such as changing the VSAN) require that you reset the interface status. When you reset the interface status, the switch disables the interface and then immediately reenables the interface.

Caution:

If an Ethernet link provides FCoE service, do not reset the physical link because this action is disruptive to all traffic on the link.

The logical link up/down feature allows the switch to reset an individual virtual link. The logical link down is signaled with a FIP Clear Virtual Link message.

Caution:

If the CNA does not support the logical link level up/down feature, the CNA resets the physical link, which means that all traffic on the Ethernet interface is disrupted.

Converged Network Adapters

Cisco NX-OS switches support the following CNA types:

- Hardware adapter
  - Works with the existing Fibre Channel host bus adapter (HBA) driver and Ethernet Network Interface Card (NIC) driver in the server.
  - Server operating system view of the network is unchanged; the CNA presents a SAN interface and a LAN interface to the operating system.
- FCoE software stack
  - Runs on existing 10-Gigabit Ethernet adapters.

The following Cisco NX-OS series and platforms support Generation 2 CNAs that use the FIP to exchange information about its available capabilities and to negotiate the configurable values with the switch:

- Cisco Nexus 7000
- Cisco Nexus 7700

To reduce configuration errors and simplify administration, the switch distributes the configuration data to all the connected adapters.

STP Lite

FCoE does not require full Spanning Tree Protocol (STP) because FCoE has no bridging functionality, which means that no STP loops are created in the network. STP Lite on FCoE interfaces ensures rapid convergence across the network by sending an agreement Bridge Protocol Data Unit (BPDU) whenever it receives a proposal BPDU. The FCoE link sends the identical agreement BPDU in response to either an Multiple Spanning Tree (MST) or a Per VLAN Rapid Spanning Tree Plus (PVRST+) proposal BPDU. Additionally, STP Lite suppresses the MAC address flushing function for FCoE VLANs.

STP Lite is enabled automatically by default across the entire device for FCoE VLANs as soon as the first FCoE VLAN comes up. At the same time, the system automatically converts all FCoE links as the STP-type normal ports. This feature runs only in FCoE VLANs.

Default Settings for FCoE

This table lists the default settings for FCoE parameters.

Table 2: Default FCoE Parameter Settings

Parameters	Default
FCoE feature	Not installed, disabled
FC-Map	0E.FC.00
Fabric priority	128
Advertisement interval	8 seconds

Guidelines and Limitations

FCoE

- VDCs apply only to Cisco Nexus 7000 Series Switches.
- You cannot enable FCoE on default VLAN.
- The QoS policy must be the same on all Cisco FCoE switches in the network.
- Beginning with Cisco NX-OS Release 6.1, FCoE is supported on F2 and F2e Series modules. F3 Series modules are supported from Cisco NX-OS Release 6.2(6) onwards.
  - FCoE supports only F2e (SFP+) modules.
  - FCoE does not support F2e (Copper) modules.

FCoE VDC

FCoE in a dedicated storage VDC has the following guidelines:

- Enable the FCoE feature set in only one VDC.
- Create VLANs in the FCoE allocated VLAN range.
- Do not enable any other features other than storage-related features in the dedicated FCoE VDC.
- Allocate resources for the dedicated FCoE VDC from an F Series module, such as the 32-port 10-Gigabit Ethernet I/O module (PID N7K-F132XP-15) .
- Rollback is not supported in a storage VDC.

- For Cisco NX-OS Release 7.2(0)D1(1), ports from only 24 FEXes can be shared to storage VDC. System will not restrict the user to go beyond 24 but, more than 24 is not tested and not supported.
- FCoE on F2, F2e, and F3 Series modules is supported with the Supervisor 2 module (N7K-SUP2 for Cisco Nexus 7000 Series devices) and the Supervisor 2E module (N77-SUP2E for Cisco Nexus 7700 Series devices and N7K-SUP2E for Cisco Nexus 7000 Series devices).
- In order to enable FCoE over FEX on the storage VDC, you must execute the **allow feature-set FEX** command from the Admin or default VDC beforehand for storage VDC. FCoE over FEX is available from 7.2(0)D1(1) and onwards.
- IVR (Inter VSAN route) zone configuration is not supported for FCoE over FEX.
- F3 Fiber Channel over Ethernet (FCoE) feature licensing is supported from 7.2.0 release onwards. To downgrade to the older version of the image 6.2.x, first uninstall the F3 FCoE license and then proceed. For more information about licensing, refer Cisco NX-OS Licensing Guide.
- F2, F2e, and F3 Series modules can co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F3 Series modules cannot co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F2 series modules cannot exist in the same VDC. This applies to both LAN and storage VDCs.
- Use the **limit-resource module-type** command in the admin or default VDC to assign module resources such as F1, F2, F2e and F3 to a storage VDC. The supported line card modules are F1, F2, F2e and F3.
- When you configure a multi-hop FCoE, ensure that you use the same no-drop classes on both sides. Priority flow control does not work when you use different no-drop classes. Use the **show interface priority-flow-control** command to verify the priority flow control operation.

Shared Interfaces

Note

- Any change in protocol state that flaps the parent port of a shared interface because of any port feature also affects the FCoE traffic on the storage vdc.
- 1500 MTU do not carry FCoE traffic in all FCoE supported platforms.

The following interface config modes are not allowed while sharing an interface from Ethernet vdc to a storage vdc:

- SPAN destination
- Private VLAN mode
- Port-channel interfaces
- Access mode
- mac-packet-classify
- Interfaces that are part of a VLAN that has an associated QoS policy

Shared Ethernet interfaces must be in trunk mode and only shared with one other VDC.

Storage VDC

Configuring a VDC for the Out-Of-Band (OOB) management interface mgmt0 is accomplished with the **vrf context management** command. However, a storage VDC does not support VRF, so configuring mgmt0 requires a different approach.

The following table shows how to configure mgmt 0 for a VDC and for a storage VDC:

Configuring mgmt 0 for VDC	Configuring mgmt 0 for storage VDC

`vrf context management`
`ip route 0.0.0.0/0 default_gateway`
	`interface mgmt 0`
	`ip address mgmt0_ip_address mgmt0_subnet_mask`
	`no shut`
	`ip route 0.0.0.0/0 default_gateway`
	Note The **ip route** command specifies the default route that points to the default gateway.

where

- *mgmt0_ip_address* is the mgmt0 IPv4 address.
- *mgmt0_subnet_mask* is the mgmt0 IPv4 netmask.
- *default_gateway* is the IPv4 address of the default-gateway.

For more information about VDC, see the Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide.

Multi-Hop FCoE Distance Configuration

In Multi-Hop FCoE, when a device sends a pause, the interface that generates the pause frame must have an ingress queue with a buffer space large enough to buffer twice the link distance. This is because, when the pause is generated the wire might get congested. By the time the adjacent device receives or processes the generated pause frame, the wire might get congested again. Therefore, the device that generates the pause must have the ability to buffer twice the link distance.

As per calculations, there can be more than 100 packets traveling on the 10 kilometer link. Due to an ASIC limitation, the F1 series line card does not support lossless FCoE on a link greater than or equal to 10 kilometers.

For more information about Multi-Hop FCoE distance limitations, see http://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/117785-probsol-nexus7000-00.html

The F3 line cards support long haul lossless distance of up to 40 kilometers. In Cisco NX-OS Release 7.2(0) and later, you can change the ingress queuing buffer configuration.

Table 3: Buffer Tuning Table for FCoE Long Distance on F2, F2E, and F3 Line Cards

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
< 5 km	F2/F2e	LR	60% no-drop and 40% drop queue
> 5 km - 10 km	F2/F2e	LR	70% no-drop and 30% drop queue

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
> 10 km - 40 km	F2/F2e	ER	80% no-drop and 20% drop queue
< 10 km	F3	LR	90% no-drop and 10% drop queue
< 40 km	F3	ER	90% no-drop and 10% drop queue

Enabling and Disabling FCoE

Enabling FCoE

You must install the FCoE feature set in the default VDC and enable dependent features in order to enable FCoE in a storage VDC.

Before you begin

- Ensure you are in the default VDC.
- Do not enable FCoE on VLAN 1 or the native VLAN.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **install feature-set fcoe**

Example:

```
switch(config)# install feature-set fcoe
```

Installs the FCoE feature set in the default VDC.

Step 3 **feature lldp**

Example:

```
switch(config)# feature lldp
```

Enables the Link Layer Discovery Protocol (LLDP) feature in the default VDC. This feature is required for FCoE operation.

Step 4 (Optional) **feature lacp**

Example:

```
switch(config)# feature lacp
```

Enables the Link Aggregation Control Protocol (LACP) feature in the default VDC. This feature is considered a best practice for FCoE operation.

Step 5 **system qos**

Example:

```
switch(config)# system qos
```

Enters Quality of service (QoS) configuration mode.

Step 6 **service-policy type network-qos** *policy-name*

Example:

```
switch(config-sys-qos)# service-policy type
 network-qos default-nq-7e-policy
```

Enables the QoS policy that supports FCoE traffic. The *policy-name* default is default-nq-8e-policy.

Step 7 (Optional) **show feature**

Example:

```
switch(config-sys-qos)# show feature
```

Displays information about the enabled features.

Step 8 (Optional) **copy running-config startup-config**

Example:

```
switch(config-sys-qos)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

What to do next

You must create a storage VDC and allocate resources to finish enabling FCoE.

Disabling FCoE

You can disable or uninstall the FCoE feature set. You can also disallow the FCoE feature set in a VDC.

Before you begin

Ensure you are in the correct VDC.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 Required: **vdc** *vdc_id* **type storage**

Example:

```
switch(config)# vdc fcoe type storage
switch(config-vdc)#
```

Enters VDC configuration mode. The *vdc_id* can be any case-sensitive, alphanumeric string up to 32 characters.

Step 3 (Optional) **no allow feature-set fcoe**

Example:

```
switch(config-vdc)# no allow feature-set fcoe
switch(config-vdc)#
```

Disallows any FCoE in this VDC. The default is allow.

Step 4 **no feature-set fcoe**

Example:

```
switch(config-vdc)# no feature-set fcoe
```

Disables the FCoE feature set. The **no feature-set fcoe** command may take some time to complete if the size of the configuration is very large. The command must clean up all of the configuration associated with the FCoE feature set.

Step 5 Required: **exit**

Example:

```
switch(config-vdc)# exit
switch(config)#
```

Exits VDC configuration mode.

Step 6 Required: **no install feature-set fcoe**

Example:

```
switch(config)# no install feature-set fcoe
```

Uninstalls the FCoE feature set. Use this command in the default VDC after you disable the FCoE feature set.

Step 7 (Optional) **show feature-set**

Example:

```
switch(config)# show feature-set
```

Displays information about the feature sets.

Step 8 (Optional) **copy running-config startup-config**

Example:

```
switch(config)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Configuring the FC-Map

You can prevent data corruption due to cross-fabric talk by configuring an FC-Map that identifies the Fibre Channel fabric for this switch. When the FC-Map is configured, the switch discards the MAC addresses that are not part of the current fabric. An FCF can assign Fabric Provide MAC Addresses (FPMA) to the CNAs consisting of the FC-Map Value for the Fabric and the Fibre Channel ID (FCID) assigned during Fabric Login

Note In most deployments, changing the fc-map is not needed. We recommend that you use VSAN to VLAN mapping. For more information, see the Mapping a VSAN to a VLAN chapter.

Before you begin

For Cisco Nexus 7000 Series, you must be in the storage VDC to configure this feature.

Procedure

Step 1 Required: **switchto vdc** *vdc-id* **type storage**

Example:

```
switch# switchto vdc fcoe type storage
fcoe#
```

Step 2 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 3 **fcoe fcmap** *fabric-map*

Example:

```
fcoe(config)# fcoe fcmap 0x0efc2a
```

Configures the global FC-Map. The default value is 0x0EFC00. The range is from 0x0EFC00 to 0x0EFCFF. Use the **no fcoe map** command to reset to the default value.

Example

This example shows how to configure the global FC-Map on a Cisco Nexus 7000 Series switch:

```
switch# switchto vdc fcoe type storage
fcoe# configure terminal
fcoe(config)# fcoe fcmap 0x0efc2a
```

Configuring the Fabric Priority

The FCoE switch advertises its priority. The priority is used by the CNAs in the fabric to determine the best switch to connect to.

Before you begin

For Cisco Nexus 7000 Series, you must be in the storage VDC to configure this feature.

Procedure

Step 1 Required: **switchto vdc** *vdc-id* **type storage**

Example:

```
switch# switchto vdc fcoe type storage
fcoe#
```

Step 2 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 3 **fcoe fcf-priority** *fabric-priority*

Example:

```
fcoe(config)# fcoe fcf-priority 42
```

Configures the global fabric priority. The default value is 128. The range is from 0 (higher) to 255 (lower). Use the **no fcoe fcf-priority** command to reset the global fabric priority to the default value.

Example

This example shows how to configure the global fabric priority for a Cisco Nexus 7000 Series switch:

```
switch# switchto vdc fcoe type storage
fcoe# configure terminal
fcoe(config)# fcoe fcf-priority 42
```

Setting the Advertisment Interval

You can configure the interval for Fibre Channel fabric advertisement on the switch.

Before you begin

For Cisco Nexus 7000 Series, you must be in the storage VDC to configure this feature.

Procedure

Step 1 Required: **switchto vdc** *vdc-id* **type storage**

Example:

```
switch# switchto vdc fcoe type storage
fcoe#
```

Step 2 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 3 **fcoe fka-adv-period** *interval*

Example:

```
fcoe(config)# fcoe fka-adv-period 8
fcoe#
```

Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds.

Example

This example shows how to configure the advertisement interval for the fabric on a Cisco Nexus 7000 Series switch:

```
switch# switchto vdc fcoe type storage
fcoe# configure terminal
fcoe(config)# fcoe fka-adv-period 42
```

Disabling LAN Traffic on an FCoE Link

You can disable LAN traffic on an FCoE link.

DCBX allows the switch to send a LAN Logical Link Status (LLS) message to a directly-connected CNA. Enter the **shutdown lan** command to send an LLS-Down message to the CNA. This command causes all VLANs on the interface that are not enabled for FCoE to be brought down. If a VLAN on the interface is enabled for FCoE, it continues to carry SAN traffic without any interruption.

Note The **shutdown lan** command is supported in Cisco Nexus 7000 Series switches running Cisco NX-OS Release 6.2(6) and later. See the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference* for more information about this command.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **interface ethernet** *slot*/*port*

Example:

```
switch(config)# interface e 2/1
switch(config-if)#
```

Specifies an interface to configure, and enters interface configuration mode. Use **?** to view a list of supported interfaces.

Step 3 **shutdown lan**

Example:

```
switch(config-if)# shutdown lan
```

Shuts down Ethernet traffic on the interface. If the interface is part of an FCoE VLAN, the shutdown has no impact on the FCoE traffic. Use **no shutdown lan** to reenable Ethernet traffic on this interface.

Step 4 (Optional) **show interface**

Example:

```
switch(config-if)# show interface
```

Displays information about the interface.

Step 5 (Optional) **copy running-config startup-config**

Example:

```
switch(config-if)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Additional References for FCoE

Related Documents

Related Topic	Document Title
Command reference	Cisco NX-OS FCoE Command Reference Guide, Nexus 7000 and MDS 9500

Related Topic	Document Title
Configuration guide	Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guides Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide
Cisco NX-OS licensing	Cisco NX-OS Licensing Guide

Standards and RFCs

Standard/RFC	Title
T11 FC BB-5	Fibre Channel Backbone 5

MIBs

MIB	MIBs Link
	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/support

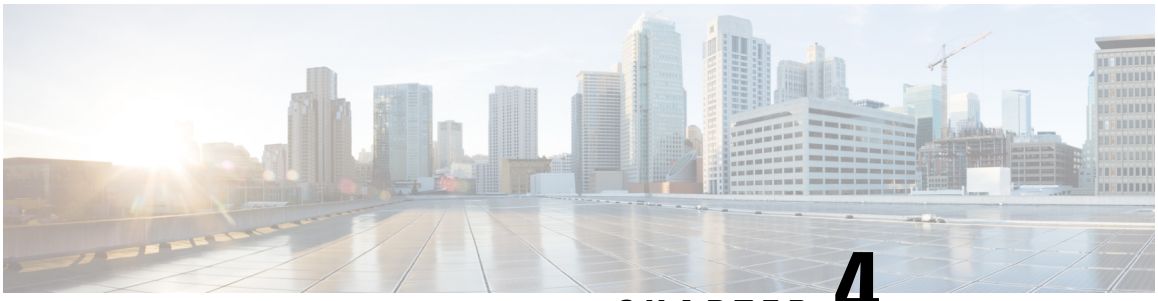

CHAPTER 4

Storage VDC

This chapter contains the following sections:

- Information About Storage VDC, on page 21
- Configuring FCoE VDCs, on page 22
- Example: Storage VDC Configuration, on page 28

Information About Storage VDC

You use a storage virtual device context (VDC) to separate LAN and SAN traffic on the same switch. A VDC allows you to maintain one physical infrastructure but separate logical data paths.

To achieve this configuration, you must perform the following tasks:

- Create a dedicated storage VDC.
- Allocate physical ports to the storage VDC. These can be either ports dedicated to only the storage VDC or ports that are shared between the storage VDC and one other VDC. Dedicated ports can be used to create either VFC E ports (VE ports) or F ports (VF ports). Shared ports can only be used for VFC F ports (VF ports).

Once you share the port to the storage VDC you can create a VFC F-port on top of the shared interface.You cannot modify some details of that port because it must match the underlying shared physical port. If you move the source port to another VDC or delete the VDC, the shared ports are deleted and you must reconfigure them.

Note If the storage VDC restarts or is suspended, any shared Ethernet ports are shut down in the corresponding VDC. These ports come up automatically once the storage VDC is operational.

Configuring FCoE VDCs

Licensing an FCoE Module

You must associate an FCoE license with an FCoE module to configure FCoE. You need one license for each module configured for FCoE.

Before you begin

Ensure you have installed the correct license for FCoE.

SUMMARY STEPS

1. **configure terminal**
2. **license fcoe module** *module-number*
3. (Optional) **show license usage** *module-name*
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`swtich(config)#`
		Enters configuration mode.

Step 2	Required: **license fcoe module** *module-number* **Example:** `swtich(config)#  license fcoe module 2`	Associates an FCoE license to a module.
Step 3	(Optional) **show license usage** *module-name* **Example:** `swtich(config)# show license usage FCOE-N7K-F132XP`	Displays the line card usage used by a storage VDC. For details on license packages, see Cisco NX-OS Licensing Guide. .
Step 4	(Optional) **copy running-config startup-config** **Example:** `swtich(config)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Creating a Dedicated Storage VDC

Note Commands from step 1 to step 7 are performed in the admin VDC.

Commands from step 8 to step 11 are performed in the storage VDC.

Before you begin

- Ensure you have installed the correct license for FCoE.

Note VE ports must exist on dedicated interfaces. VF ports can exist on either dedicated or shared interfaces.

SUMMARY STEPS

1. **configure terminal**
2. **install feature-set fcoe**
3. **system qos**
4. **service-policy type network-qos** *policy-map name*
5. **vdc** *vdc-name* **type storage**
6. **allocate interface ethernet** *int-numb*
7. **switchto vdc** *vdc-name*
8. **feature lldp**
9. (Optional) **feature lacp**
10. (Optional) **show feature-set**
11. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**install feature-set fcoe** **Example:** `switch(config)# install feature-set fcoe`	Installs the FCoE feature-set.
Step 3	**system qos** **Example:**	Enters system qos mode.

	Command or Action	Purpose
	`switch(config)# system qos` `switch(config-sys-qos)#`	
Step 4	**service-policy type network-qos** *policy-map name* **Example:** `switch(config-sys-qos)#service-policy type network-qos default-nq-7e-policy`	Enables no drop queue for FCoE class. **Note** The Cisco MDS 9250i Multiservice Fabric Switch follows a different naming convention which reflects how many ingress queues and egress queues are there in a policy. The name is different, but it is same as the MDS policy. 7e policy is default in MDS and Cisco MDS 9250i Multiservice Fabric Switch, hence it is not shown in **show running-config** command. The 7e policy is not default in Cisco Nexus 7000 Series Switches, hence it is shown in the **show running-config** command.

Step 5	**vdc** *vdc-name* **type storage**
	Example:
	`switch(config)# vdc fcoe-vdc type storage`
	`switch(config-vdc)#`
		Creates a dedicated storage VDC and enters VDC configuration mode. You can only enable storage features in a storage VDC. You do not need to allow the feature-set or enable it in the storage VDC because this process is handled automatically for a storage VDC.
		Note It is expected that, if feature-set fex is not enabled, "**allocate share fex**" will fail, on VDC reload. User has to un-configure the share and configure it back .

Step 6	**allocate interface ethernet** *int-numb*
	Example:
	`switch(config-vdc)# allocate interface ethernet 2/1-2`
	`switch(config-if)#`
		Allocates interfaces to the storage VDC as a dedicated FCoE port. You must allocate all interfaces in the port group. You must configure these interfaces in switchport trunk mode as Spanning Tree Protocol (STP) edge ports.

Step 7	**switchto vdc** *vdc-name*
	Example:
	`switch(config-vdc)# switchto vdc fcoe-vdc`
	`switchport`
	`switch-fcoe-vdc#`
		Switches to the storage VDC.

Step 8	**feature lldp** **Example:** `switch(config)# feature lldp`	Enables the LLDP feature in the storage VDC.
Step 9	(Optional) **feature lacp** **Example:** `switch(config)# feature lacp`	Enables the LACP feature in the storage VDC.
Step 10	(Optional) **show feature-set** **Example:**	Displays the status information about the feature-sets in this VDC.

	Command or Action	Purpose

	`switch# show feature-set`
	`Feature Set Name ID State`
	`-------------------- -------- --------`
	`fcoe 2 enabled`
	`fex 3 disabled`
	`switch#`	

Step 11	(Optional) **copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Allocating the FCoE VLAN Range

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you are in the correct VDC.

SUMMARY STEPS

1. **configure terminal**
2. **vdc** *vdc-name* **type storage**
3. **allocate fcoe-vlan-range** *vlan-range* [**from vdcs** *vdc-name*]
4. (Optional) **show vdc fcoe-vlan-range**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**vdc** *vdc-name* **type storage**
	Example:
	`switch(config)# vdc fcoe-vdc type storage`
	`switch(config-vdc)#`
		Enters VDC configuration mode. You can only enable storage feature in a storage VDC. You do not need to allow the feature-set or enable it in the storage VDC because this process is handled automatically for a storage VDC.

Step 3	**allocate fcoe-vlan-range** *vlan-range* [**from vdcs** *vdc-name*] **Example:**	Allows the VLAN to be used in the storage VDC-I; and allocates the VLANs that can be used for FCoE and mapped to a VSAN. You can optionally allocate the VLANs from another VDC.

	Command or Action	Purpose
	`switch(config-vdc)# allocate fcoe-vlan-range 10-30`	
Step 4	(Optional) **show vdc fcoe-vlan-range** **Example:** `switch(config-vdc)#  show vdc fcoe-vlan-range`	Displays information about the VLAN range allocated for FCoE.
Step 5	(Optional) **copy running-config startup-config** **Example:** `switch(config-vdc)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Related Topics

FCoE, on page 9

Allocating Shared Interfaces

You can share interfaces between a storage VDC and another VDC.

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you are in the correct VDC.
- Ensure any shared interfaces are from an F-series module.
- Ensure you have allocated the FCoE VLAN range.
- Interfaces can only be shared between the storage VDC and one other VDC.

Note

- For shared interface, ensure that LLDP feature is enabled in parent VDC also.
- Only VF ports can exist on shared interfaces. VE ports must be on dedicated interfaces.

SUMMARY STEPS

1. **configure terminal**
2. **interface** *if-range*
3. **switchport mode trunk**
4. **spanning-tree port type edge trunk**
5. **no shutdown**
6. **vdc** *vdc-name* **type storage**
7. **allocate shared interface** *if-range*

8. (Optional) **show vdc shared membership**
9. (Optional) **switchto vdc** *vdc-name*
10. **configure terminal**
11. **feature lldp**
12. **interface** *if-range*
13. **no shutdown**
14. (Optional) **show interface** *if-range*
15. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**interface** *if-range* **Example:** `switch(config)# interface ethernet 2/1`	Enters interface configuration mode for the interface in the Ethernet VDC.
Step 3	**switchport mode trunk** **Example:** `switch(config-if)# switchport mode trunk`	Puts the Ethernet interface into trunk mode.
Step 4	**spanning-tree port type edge trunk** **Example:** `switch(config-if)# spanning-tree port type edge trunk`	Sets the interface to STP-type edge port to support STP Lite for loop prevention.
Step 5	**no shutdown** **Example:** `switch(config-if)# no shutdown`	Administratively enables the Ethernet shared interface.

Step 6	**vdc** *vdc-name* **type storage**
	Example:
	`switch(config-if)# vdc fcoe-vdc type storage`
	`switch(config-vdc)#`
		Enters VDC configuration mode.

Step 7	**allocate shared interface** *if-range* **Example:** `switch(config-vdc)# allocate shared interface ethernet 2/1`	Allocates interfaces that are shared with another VDC for FCoE traffic. You must allocate the shared interfaces to one of the VDC included in the FCoE VLAN allocation. You can only use the **shutdown** or the **switchport trunk allowed vlan** commands on shared interfaces in the storage VDC.

	Command or Action	Purpose
Step 8	(Optional) **show vdc shared membership** **Example:** `switch(config-vdc)#  show vdc shared membership`	Displays the interfaces that are shared for FCoE.

Step 9	(Optional) **switchto vdc** *vdc-name*
	Example:
	`switch(config-vdc)# switchto vdc fcoe-vdc`
	`switch-fcoe-vdc#`
		Switches to the storage VDC.

Step 10	**configure terminal**
	Example:
	`switch-fcoe-vdc# configure terminal`
	`switch-fcoe-vdc(config)#`
		Enters configuration mode.

Step 11	**feature lldp** **Example:** `switch-fcoe-vdc(config)# feature lldp`	Enables the LLDP feature in the storage VDC.
Step 12	**interface** *if-range* **Example:** `switch-fcoe-vdc(config)# interface ethernet 2/1`	Enters interface configuration mode for the shared interface in the storage VDC.
Step 13	**no shutdown** **Example:** `switch-fcoe-vdc(config-if)# no shutdown`	Administratively enables the FCoE shared interface.
Step 14	(Optional) **show interface** *if-range* **Example:** `switch-fcoe-vdc(config-if)#  show interface ethernet 2/1`	Displays information about the shared interface.
Step 15	(Optional) **copy running-config startup-config** **Example:** `switch-fcoe-vdc(config-if)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Example: Storage VDC Configuration

Ethernet VDC Configuration

```
!Enable the interface to share:
switch(config-sys-qos)# interface ethernet 2/1

!Initially ethernet 2/1 is allocated to ethernet VDC
switch(config-if)# switchport
```

```
switch(config-if)# switchport mode trunk

!Allocate resources in admin VDC:
Switch(conf)# vdc storage fcoe_vdc
switch(config-if)# allocate fcoe-vlan-range 10-20 from vdc switch
switch(config-vdc)# allocate shared interface ethernet 2/1
Switch(conf)# show vdc shared membership

!Switch to storage VDC and bring up the shared interface:
switch(config-vdc)# switchto vdc fcoe_vdc
switch-fcoe_vdc# configure terminal
switch-fcoe_vdc(config)# interface ethernet 2/1
switch-fcoe_vdc(config-if)# no shutdown

!A VFC interface is created on top of Ethernet interface. The VFC interface can be created
 in 2 ways: implicit and explicit.

!Implicit:
switch-fcoe_vdc(config)# interface vfc 2/1
switch-fcoe_vdc(config-if)# switchport mode f

!Explicit:
switch-fcoe_vdc(config)# interface vfc2
switch-fcoe_vdc(config-if)# bind interface eth2/1
switch-fcoe_vdc(config-if)# switchport mode f
```

Note Ethernet 2/1 must be from an F-series module.

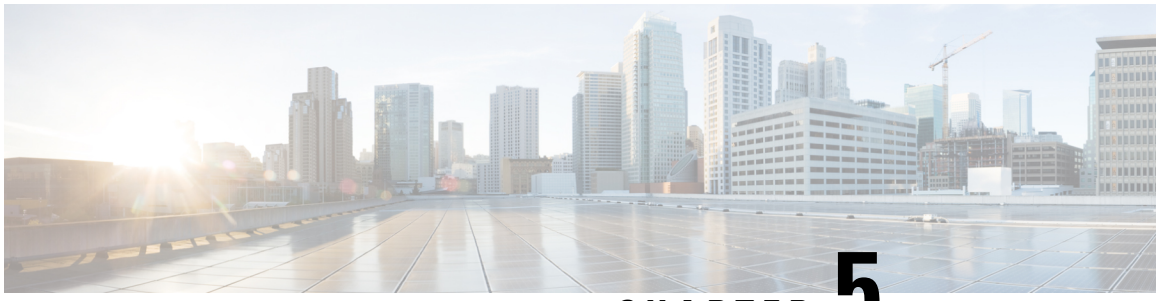

CHAPTER 5

Configuring FCoE VLANs and Virtual Fibre Channel Interfaces

- Information About Virtual Fibre Channel (VFC) Interfaces, on page 31
- Default Settings for FCoE, on page 32
- Guidelines and Limitations, on page 32
- Configuring Virtual Interfaces, on page 34
- Verifying the Virtual Fibre Channel Interface , on page 41
- Example:Mapping VSANs to VLANs , on page 43
- Verifying the FCoE Configuration, on page 44
- Additional References for FCoE, on page 45

Information About Virtual Fibre Channel (VFC) Interfaces

Fibre Channel over Ethernet (FCoE) allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers.

The Fibre Channel portion of FCoE is configured as a virtual Fibre Channel interface. Logical Fibre Channel features (such as interface mode) can be configured on virtual Fibre Channel interfaces.

A virtual Fibre Channel (vfc) interface must be bound to an interface before it can be used. The binding is to a physical Ethernet interface when the converged network adapter (CNA) is directly connected to the switch or port channel when the CNA connects to the Fibre Channel Forwarder (FCF) over a virtual port channel (vPC).

VF Port

A virtual fabric (VF) port in an FCoE network acts as a fabric port that connects to a peripheral device (host or disk) operating as an N port. A VF port can be attached to only one N port.

VE Port

A virtual expansion (VE) port acts as an expansion port in an FCoE network. VE ports can connect multiple FCoE switches together in the network. You can bind a VE port to a physical ethernet port or a port channel.

Traffic is load balanced across equal cost E_Ports and VE_Ports based on SID, DID, and OXID. Traffic across members of a port channel that a VE_Port is bound to is load balanced based on SID, DID, and OXID.

VNP Ports

Connectivity from an FCoE NPV bridge to the FCF is supported only over point-to-point links. These links can be individual Ethernet interfaces or port channel interfaces. For each FCF connected to an Ethernet/port-channel interface, a vFC interface must be created and bound to it. These vFC interfaces must be configured as VNP ports.

On the VNP port, the FCoE NPV bridge emulates an FCoE-capable host with multiple FCoE nodes (ENodes), each with a unique ENodes MAC address. By default, the VNP port is enabled in trunk mode.

Multiple VSANs can be configured on the VNP port. The FCoE VLANs that correspond to the VNP port VSANs must be configured on the bound Ethernet interface.

Default Settings for FCoE

This table lists the default settings for FCoE parameters.

Table 4: Default FCoE Parameter Settings

Parameters	Default
FCoE feature	Not installed, disabled
FC-Map	0E.FC.00
Fabric priority	128
Advertisement interval	8 seconds

Guidelines and Limitations

FCoE

- VDCs apply only to Cisco Nexus 7000 Series Switches.
- You cannot enable FCoE on default VLAN.
- The QoS policy must be the same on all Cisco FCoE switches in the network.
- Beginning with Cisco NX-OS Release 6.1, FCoE is supported on F2 and F2e Series modules. F3 Series modules are supported from Cisco NX-OS Release 6.2(6) onwards.
  - FCoE supports only F2e (SFP+) modules.
  - FCoE does not support F2e (Copper) modules.

FCoE VDC

FCoE in a dedicated storage VDC has the following guidelines:

- Enable the FCoE feature set in only one VDC.
- Create VLANs in the FCoE allocated VLAN range.
- Do not enable any other features other than storage-related features in the dedicated FCoE VDC.
- Allocate resources for the dedicated FCoE VDC from an F Series module, such as the 32-port 10-Gigabit Ethernet I/O module (PID N7K-F132XP-15) .
- Rollback is not supported in a storage VDC.
- For Cisco NX-OS Release 7.2(0)D1(1), ports from only 24 FEXes can be shared to storage VDC. System will not restrict the user to go beyond 24 but, more than 24 is not tested and not supported.
- FCoE on F2, F2e, and F3 Series modules is supported with the Supervisor 2 module (N7K-SUP2 for Cisco Nexus 7000 Series devices) and the Supervisor 2E module (N77-SUP2E for Cisco Nexus 7700 Series devices and N7K-SUP2E for Cisco Nexus 7000 Series devices).
- In order to enable FCoE over FEX on the storage VDC, you must execute the **allow feature-set FEX** command from the Admin or default VDC beforehand for storage VDC. FCoE over FEX is available from 7.2(0)D1(1) and onwards.
- IVR (Inter VSAN route) zone configuration is not supported for FCoE over FEX.
- F3 Fiber Channel over Ethernet (FCoE) feature licensing is supported from 7.2.0 release onwards. To downgrade to the older version of the image 6.2.x, first uninstall the F3 FCoE license and then proceed. For more information about licensing, refer Cisco NX-OS Licensing Guide.
- F2, F2e, and F3 Series modules can co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F3 Series modules cannot co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F2 series modules cannot exist in the same VDC. This applies to both LAN and storage VDCs.
- Use the **limit-resource module-type** command in the admin or default VDC to assign module resources such as F1, F2, F2e and F3 to a storage VDC. The supported line card modules are F1, F2, F2e and F3.
- When you configure a multi-hop FCoE, ensure that you use the same no-drop classes on both sides. Priority flow control does not work when you use different no-drop classes. Use the **show interface priority-flow-control** command to verify the priority flow control operation.

Shared Interfaces

Note

- Any change in protocol state that flaps the parent port of a shared interface because of any port feature also affects the FCoE traffic on the storage vdc.
- 1500 MTU do not carry FCoE traffic in all FCoE supported platforms.

The following interface config modes are not allowed while sharing an interface from Ethernet vdc to a storage vdc:

- SPAN destination
- Private VLAN mode
- Port-channel interfaces
- Access mode
- mac-packet-classify
- Interfaces that are part of a VLAN that has an associated QoS policy

Shared Ethernet interfaces must be in trunk mode and only shared with one other VDC.

Storage VDC

Configuring a VDC for the Out-Of-Band (OOB) management interface mgmt0 is accomplished with the **vrf context management** command. However, a storage VDC does not support VRF, so configuring mgmt0 requires a different approach.

The following table shows how to configure mgmt 0 for a VDC and for a storage VDC:

Configuring mgmt 0 for VDC	Configuring mgmt 0 for storage VDC

`vrf context management`
`ip route 0.0.0.0/0 default_gateway`
	`interface mgmt 0`
	`ip address mgmt0_ip_address mgmt0_subnet_mask`
	`no shut`
	`ip route 0.0.0.0/0 default_gateway`
	Note The **ip route** command specifies the default route that points to the default gateway.

where

- *mgmt0_ip_address* is the mgmt0 IPv4 address.
- *mgmt0_subnet_mask* is the mgmt0 IPv4 netmask.
- *default_gateway* is the IPv4 address of the default-gateway.

For more information about VDC, see the Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide.

Configuring Virtual Interfaces

Mapping a VSAN to a VLAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each virtual fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 10, VLAN 1003 for VSAN 2, and so on). If you enable MST, you must use a separate Multiple Spanning Tree (MST) instance for FCoE VLANs.

Note You must exit VLAN mode to execute the configured commands on the Cisco Nexus 7000 Series Switches.

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you have enabled FCoE.
- Ensure that you are in the storage VDC.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **vsan database**

Example:

```
switch(config)# vsan database
switch(config-vsan-db)#
```

Enters VSAN database configuration mode.

Step 3 **vsan** *vsan-id*

Example:

```
switch(config-vsan-db)# vsan 200
```

Defines the VSAN. The VSAN number range is from 1 to 4094.

Step 4 **vlan** *vlan-id*

Example:

```
switch(config-vsan-db)# vlan 200
switch(config-vlan)#
```

Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.

Step 5 **fcoe** [**vsan** *vsan-id*]

Example:

```
switch(config-vlan)# fcoe vsan 200
```

Enables FCoE for the specified VLAN and configures the mapping from this VLAN to the specified VSAN. If you do not specify a VSAN number, a mapping is created from this VLAN to the VSAN with the same number.

Step 6 **exit**

Example:

```
switch(config-vlan)# exit
switch(config)#
```

Exits VLAN configuration mode. You must exit this mode to execute the configured commands on the Cisco Nexus 7000 Series Switches.

Step 7 (Optional) **show vlan fcoe**

Example:

```
switch(config-vlan)# show vlan fcoe
```

Displays information about the FCoE configuration for a VLAN.

Step 8 (Optional) **copy running-config startup-config**

Example:

```
switch(config-vlan)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Example

This example shows how to map VLAN 300 to VSAN 300 on a Cisco Nexus 7000 Series Switches:

```
switch(config)# switchto vdc fcoe_vdc
switch-fcoe_vdc# configure terminal
switch-fcoe_vdc(config)# vlan 300
switch-fcoe_vd(config-vlan)# fcoe vsan 300
```

Creating a Virtual Fibre Channel Interface

To use FCoE, you must first create Virtual Fibre Channel (VFC) interfaces. Then, you must bind the VFC interfaces to physical interfaces before FCoE can be used.

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you have enabled FCoE.
- Ensure that you have created VDC on Cisco Nexus 7000 Series switches. For information on creating VDC, see Cisco Nexus 7000 Series Virtual Device Context Configuration Guide.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **interface vfc** *vfc-id*

Example:

```
switch(config)# interface vfc 4
switch(config-if)#
```

Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode. The *vfc-id* range is from 1 to 8192.

Step 3 **switchport mode** *mode*

Example:

```
switch(configif)# switchport mode e
```

Configures the switchport mode for a virtual Fibre Channel interface. The *mode* is E or F. The default is F mode.

Step 4 **bind** {**interface** {**ethernet** *slot*/*port* | **ethernet-port-channel** *number*}}

Example:

```
switch(config-if)# bind interface ethernet 1/4
```

Binds the virtual Fibre Channel interface to the specified interface. Use **?** to see the supported interfaces and port channels. Use the **no** form of this command to unbind the virtual Fibre Channel interface from the specified interface.

Step 5 (Optional) **show interface vfc**

Example:

```
switch(config-if)# show interface vfc
```

Displays information about the virtual Fibre Channel interfaces.

Step 6 (Optional) **copy running-config startup-config**

Example:

```
switch(config)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Example

This example shows how to bind a virtual Fibre Channel interface to an Ethernet interface:

```
switch# configure terminal
switch(config)# interface vfc 4
switch(config-if)# bind interface ethernet 1/4
```

This example shows how to delete a virtual Fibre Channel interface:

```
switch# configure terminal
switch(config)# no interface vfc 4
```

Creating a Virtual Fibre Channel Port Channel Interface

You can create a virtual Fibre Channel port channel interface that automatically binds to the port channel with the same interface number.

Before you begin

For the Cisco Nexus 7000 Series, ensure that you create the port channel interface before you create the virtual Fibre Channel port channel interface.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **interface vfc-port-channel** *int-number*

Example:

```
switch(config)# interface vfc-port-channel 2
switch(config-if)#
```

Creates a virtual Fibre Channel interface (if it does not already exist) that is bound to the port channel with the same interface number and enters interface configuration mode.

The *int-number* range is from 1 to 4096.

The default switchport mode for this interface is E.

Note Cisco Nexus 7000 Series switches supports only the Ethernet port channel or the channel group ID numbers ranging from 513 to 4096.

Step 3 **switchport mode** *mode*

Example:

```
switch(config-if)# switchport mode e
```

Configures the switchport mode for a virtual Fibre Channel interface. The *mode* is E or F. The default is F mode.

Step 4 (Optional) **show interface vfc-port-channel** *int-number*

Example:

```
switch(config-if)# show interface vfc-port-channel 2
```

Displays information about the virtual Fibre Channel interfaces bound to port channel interfaces.

Step 5 (Optional) **copy running-config startup-config**

Example:

```
switch(config)# copy running-config startup-config
[########################################] 100%
Copy complete.
```

Copies the running configuration to the startup configuration.

Step 6 (Optional) **show running-config interface vfc***id/slot*

Example:

```
awitch# show running-config interface vfc-po540
!Command: show running-config interface vfc-po540
!Time: Fri Dec 2 15:36:07 2016

version 7.3(0)D1(1)

interface vfc-po540
bind interface ethernet-port-channel540
switchport mode E
no shutdown
```

Displays the virtual Fibre Channel interface configuration information.

Associating a Virtual Fibre Channel Interface to a VSAN

You must configure unique, dedicated VLAN at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If you enable MST, you must use a separate MST instance for FCoE VLANs.

Before you begin

For Cisco Nexus 7000 Series Switches, ensure that you are in the storage VDC.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **vsan database**

Example:

```
switch(config)# vsan database
switch(config-vsan-db)#
```

Enters VSAN configuration mode.

Step 3 (Optional) **vsan** *vsan-id*

Example:

```
switch(config-vsan-db)# vsan 2
```

Creates the VSAN. The *vsan-id* range is from 1 to 4094 and must map to a VLAN on the physical Ethernet interface that is bound to the virtual Fibre Channel interface.

Step 4 Enter one of the following commands:

- **vsan** *vsan-id* **interface vfc** *vfc-id*
- **vsan** *vsan-id* **interface vfc-port-channel** *vfc-id*

Example:

```
switch(config-vsan-db)# vsan 2 interface vfc 4
```

Configures the association between the VSAN and virtual Fibre Channel interface or virtual Fibre Channel port channel. The *vsan-id* range is from 1 to 4094 and must map to a VLAN on the physical Ethernet interface or port channel that is bound to the virtual Fibre Channel interface or virtual Fibre Channel port channel. The *vfc-id* range is from 1 to 8192. Us the **no** form of this command to dissassociate the connection between the VSAN and virtual Fibre Channel interface or virtual Fibre Channel port channel.

Step 5 (Optional) **show vsan**

Example:

```
switch(config-vsan-db)# show vsan
```

Displays information about the VSAN.

Step 6 (Optional) **copy running-config startup-config**

Example:

```
switch(config-vsan-db)# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Example

This example shows how to associate a virtual Fibre Channel interface to a VSAN:

```
switch# configure terminal
switch(config)# vsan database
switch(config-vsan-db)# vsan 2 interface vfc 4
```

Enabling VE Loopback Configuration

The VFID check verifies that the VSAN configuration is correct on both ends of a VE link. You can turn off the VFID check for VE ports to allow VE loopback configuration between to VE ports on the same switch.

Before you begin

For Cisco Nexus 7000 Series, you must be in the storage VDC to configure this feature.

Procedure

Step 1 Required: **switchto vdc** *vdc-id* **type storage**

Example:

```
# switchto vdc fcoe type storage
fcoe#
```

Step 2 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 3 **fcoe veloopback**

Example:

```
fcoe(config)# fcoe veloopback
```

Enables the VFID check for all VE ports.

Example

This example shows how to enable VE loopback for a Cisco Nexus 7000 Series switch:

```
switch# switchto vdc fcoe type storage
fcoe# configure terminal
fcoe(config)# fcoe veloopback
```

Verifying the Virtual Fibre Channel Interface

To display configuration information about virtual Fibre Channel interfaces, perform one of the following tasks:

Command	Purpose
show interface vfc *vfc-id*	Displays the detailed configuration of the specified Fibre Channel interface.
show interface brief	Displays the status of all interfaces.
show vlan fcoe	Displays the mapping of FCoE VLANs to VSANs.

This example shows how to display a virtual Fibre Channel interface bound to an Ethernet interface:

```
switch# show interface vfc 3
vfc3 is up
    Bound interface is Ethernet1/37
    Hardware is Virtual Fibre Channel
    Port WWN is 20:02:00:0d:ec:6d:95:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is F, FCID is 0x490100
    Port vsan is 931
    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
      0 frames input, 0 bytes
```

```
    0 discards, 0 errors
  0 frames output, 0 bytes
    0 discards, 0 errors
Interface last changed at Thu May 21 04:44:42 2009
```

This example shows how to display the status of all the interfaces on the switch (some output has been removed for brevity):

```
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                 Mode   Trunk                          Mode  Speed  Channel
                        Mode                                 (Gbps)
--------------------------------------------------------------------------------
fc3/1      1      auto   on      trunking         swl    TE      2    --
fc3/2      1      auto   on      sfpAbsent        --     --           --
...
fc3/8      1      auto   on      sfpAbsent        --     --           --
--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed    MTU   Port
                                                                   Channel
--------------------------------------------------------------------------------
Ethernet1/1             hwFailure  --                --       1500  --
Ethernet1/2             hwFailure  --                --       1500  --
Ethernet1/3             up         --                10000    1500  --
...
Ethernet1/39            sfpIsAbsen --                --       1500  --
Ethernet1/40            sfpIsAbsen --                --       1500  --
--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed    MTU
--------------------------------------------------------------------------------
mgmt0                   up         172.16.24.41      100      1500
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                 Mode   Trunk                          Mode  Speed  Channel
                        Mode                                 (Gbps)
--------------------------------------------------------------------------------
vfc 1        1      F      --      down             --     --           --
...
```

This example shows how to display the mapping between the VLANs and VSANs on the switch:

```
switch# show vlan fcoe
VLAN      VSAN      Status
--------  --------  --------
15        15        Operational
20        20        Operational
25        25        Operational
30        30        Non-operational
```

Example:Mapping VSANs to VLANs

The following example shows how to configure the FCoE VLAN and a virtual Fibre Channel interface:

Procedure

Step 1 Enable the associated VLAN and map the VLAN to a VSAN. For Cisco Nexus 7000, ensure you are in the storage VDC.

```
switch(config)# vlan 200
switch(config-vlan)# fcoe vsan 200
switch(config-vlan)# exit
```

Step 2 Configure the VLAN on a physical Ethernet interface.

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# spanning-tree port type edge trunk
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1,200
switch(config-if)# exit
```

Step 3 Create a virtual Fibre Channel interface and bind it to a physical Ethernet interface. For Cisco Nexus 7000, ensure you are in the storage VDC.

```
switch(config)# interface vfc 4
switch(config-if)# bind interface ethernet 1/4
```

Note By default, all virtual Fibre Channel interfaces reside on VSAN 1. If the VLAN to VSAN mapping is to a VSAN other than VSAN 1, then proceed to Step 4.

Step 4 Associate the virtual Fibre Channel interface to the VSAN. For Cisco Nexus 7000, ensure you are in the storage VDC.

```
switch(config)# vsan database
switch(config-vsan)# vsan 200 interface vfc 4
switch(config-vsan)# exit
```

Step 5 (Optional) Display membership information for the VSAN.

```
switch# show vsan 200 membership
vsan 200 interfaces
           vfc 4
```

Step 6 (Optional) Display the interface information for the virtual Fibre Channel interface.

```
switch# show interface vfc 4

vfc4 is up
Bound interface is Ethernet1/4
Hardware is Virtual Fibre Channel
Port WWN is 20:02:00:0d:ec:6d:95:3f
Port WWN is 20:02:00:0d:ec:6d:95:3f
snmp link state traps are enabled
Port WWN is 20:02:00:0d:ec:6d:95:3f
APort WWN is 20:02:00:0d:ec:6d:95:3f
snmp link state traps are enabled
Port mode is F, FCID is 0x490100
Port vsan is 200
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
0 frames input, 0 bytes 0 discards, 0 errors
0 frames output, 0 bytes 0 discards, 0 errors
Interface last changed at Thu Mar 11 04:44:42 2010
```

Verifying the FCoE Configuration

To display FCoE configuration information, perform one of these tasks in the context of storage VDC:

Command	Purpose
show fcoe	Displays whether FCoE is enabled on the switch.
show fcoe database	Displays the contents of the FCoE database. **Note** This command is effective when interfaces are configured.
show interface vfc [*number*]	Displays information about the vFC interfaces.
show interface [*interface number*] **fcoe**	Displays the FCoE settings for an interface or all interfaces.

This example shows how to verify that the FCoE capability is enabled:

```
switch# show fcoe
Global FCF details
        FCF-MAC is 00:0d:ec:6d:95:00
        FC-MAP is 0e:fc:00
        FCF Priority is 128
        FKA Advertisement period for FCF is 8 seconds
```

This example shows how to display the FCoE database:

```
switch# show fcoe database
------------------------------------------------------------------------------
INTERFACE       FCID            PORT NAME               MAC ADDRESS
------------------------------------------------------------------------------
vfc3            0x490100        21:00:00:1b:32:0a:e7:b8 00:c0:dd:0e:5f:76
```

This example shows how to display the FCoE settings for an interface.

```
switch# show interface ethernet 1/37 fcoe
Ethernet1/37 is FCoE UP
    vfc3 is Up
        FCID is 0x490100
       PWWN is 21:00:00:1b:32:0a:e7:b8
       MAC addr is 00:c0:dd:0e:5f:76
```

Additional References for FCoE

Related Documents

Related Topic	Document Title
Command reference	Cisco NX-OS FCoE Command Reference Guide, Nexus 7000 and MDS 9500
Configuration guide	Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guides Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide
Cisco NX-OS licensing	Cisco NX-OS Licensing Guide

Standards and RFCs

Standard/RFC	Title
T11 FC BB-5	Fibre Channel Backbone 5

MIBs

MIB	MIBs Link
	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/support

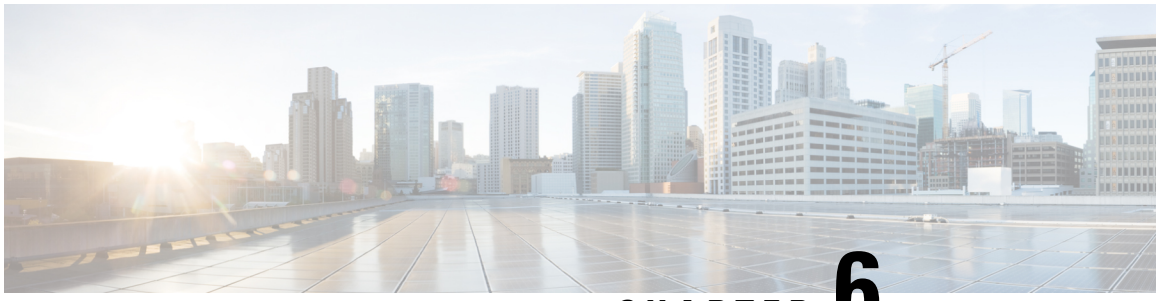

CHAPTER 6

Configuring Dynamic FCoE Using FabricPath

This chapter contains the following sections:

- Information About Dynamic FCoE Using FabricPath, on page 47
- Prerequisites for Dynamic FCoE Using FabricPath, on page 51
- Guidelines and Limitations for Dynamic FCoE Using FabricPath, on page 51
- Configuration Topology Example, on page 52
- Configuring Dynamic FCoE Using FabricPath, on page 53
- Instantiation and Initialization of Dynamic VFC, on page 53
- Verifying the Dynamic FCoE Using FabricPath Configuration, on page 54
- Configuration Output Examples for Dynamic FCoE Using FabricPath, on page 57

Information About Dynamic FCoE Using FabricPath

Fibre Channel over Ethernet (FCoE) enables I/O consolidation. It permits both LAN and SAN traffic to coexist on the same switch and the same wire. This feature enables you to consolidate multiple separate networks into a single converged infrastructure.

Key values of I/O consolidation using traditional FCoE are as follows:

- Elimination of separate network infrastructures for SAN and LAN traffic.
- Reduction in hardware requirements, such as cabling and server interface cards (NICs and HBAs), and lowering capital expense.
- Reduction in power and cooling requirements for fewer physical assets.
- Increasing deployment agility for multiprotocol networks, which preserves long-term investments while preparing for future uncertainty in protocol needs.

By using FabricPath Ethernet technology, you can take FCoE consolidation even further:

- Create a logical, rather than physical, SAN A/B separation.
- Efficiently load balance multiprotocol traffic within the data center.
- Dynamically establish relationships between switches, reducing the possibility for human error during configurations.
- Improved high availability percentages as the scale increases.

The FabricPath architecture provides an inherent multipath capability with redundancy to handle node failures. Fabric level redundancy is provided through a double fabric model (SAN A/SAN B). The separation of the two SANs is logically implemented as two different VSANs that map to two different VLANs (VLAN A and B). Fibre channel traffic in SAN A becomes the FCoE traffic in VLAN A, the Fiber Channel traffic in SAN B becomes the FCoE traffic in VLAN B, and the LAN traffic is carried on one or more additional VLANs over the converged Ethernet infrastructure. In this logical environment, the VSAN A/VSAN B configuration protects against fabric-wide control plane failures.

The traditional method of hosts that connect to two separate SANs is still supported with the FCoE over FabricPath architecture. The host is connected to two different leaf nodes that host a disjointed set of VSANs. Beyond these leaf nodes, the fabric is converged on the same infrastructure, but the host continues to see two SAN fabrics.

The following figure shows a FabricPath topology with n spines (S) and m leafs (L). The m leafs communicate to each other through the n spines using FabricPath encapsulation.

Figure 1: FabricPath Topology

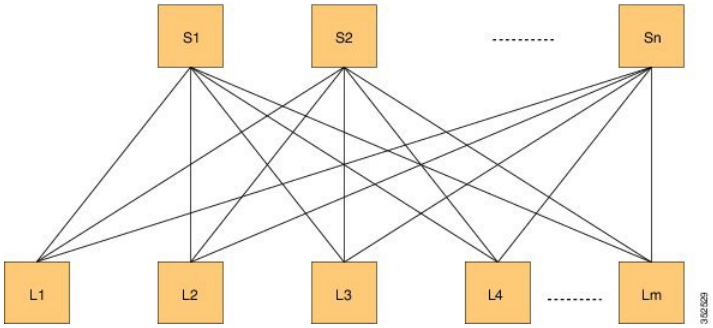

FCoE creates an overlay of FCoE virtual links on top of the underlying Ethernet topology, irrespective of how that Ethernet topology is constructed and which protocol is used to compute the MAC address routes.

In a dynamic FCoE environment, the topology is developed using the leafs as FCoE Forwarder (FCF) switches that are forwarded through transparent spines.

FCoE hosts and FCoE storage devices are connected to a FabricPath topology through the leaf switches. In this configuration, only the leaf switches perform FCoE forwarding (only the leaf switches behave as FCFs); the spine switches just forward MAC-in-MAC encapsulated Ethernet frames that are based on the outer destination MAC address.

The following figure shows the logical FCoE overlay topology of VE_Port to VE_Port virtual links on a FabricPath topology.

Figure 2: FCoE Overlay of VE_Port to VE_Port Virtual Links

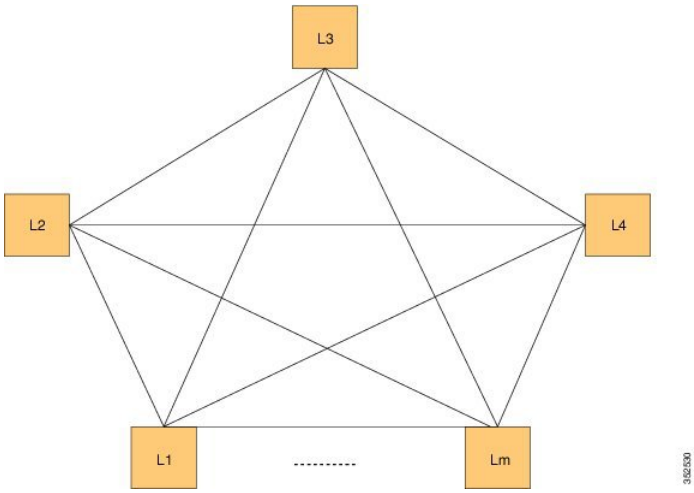

Only the FCFs, that are implemented by the leaf switches are part of this overlay topology. This topology is seen by Fabric Shortest Path First (FSPF), for each FCoE VLAN. FSPF computes over which virtual link to forward an FCoE frame based on its DomainID (D_ID). A virtual link is uniquely identified by the pair of MAC addresses associated with the two VE_Ports logically connected by it. Identifying the virtual link is equivalent to identifying which MAC addresses to use for the FCoE encapsulation on the transport network.

Use L_m as the number of leafs that are feature enabled. The feature might not be enabled on all leafs. The FCoE mesh is basically the leafs where FCoE or FabricPath is enabled.

SAN A/B Separation

For Dynamic FCoE, SAN A/B separation is realized in a logical manner across the backbone. As shown in the following illustration, physical SAN A/B separation is maintained from the FCF leafs to the end devices. Beyond the leafs, FCoE traffic for SANs A and B are carried by FabricPath Equal Cost Multipathing (ECMP) links across all spines, maintaining logical SAN A/B separation.

Figure 3: Physical Topology Diagram

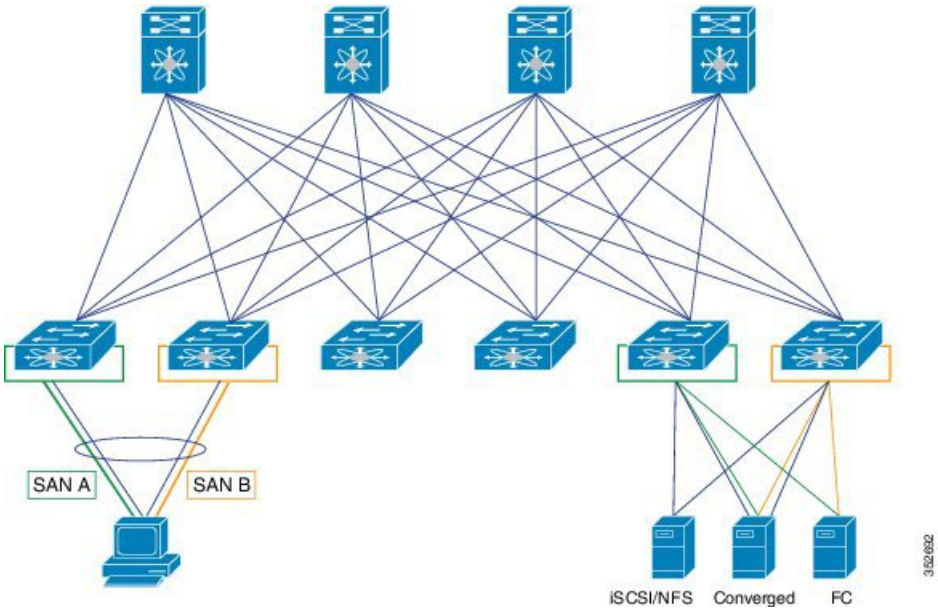

In the previous figure, the physical connectivity for the topology follows typical leaf/spine CLOS architectural best practices. Logically, SAN A and SAN B are isolated at the Top of Rack (ToR) switches physically. Once the traffic enters the FabricPath network, the storage traffic is logically separated (see the following figure) across the network where it is physically separated once more to the storage device edge.

Figure 4: Logical Topology Diagram

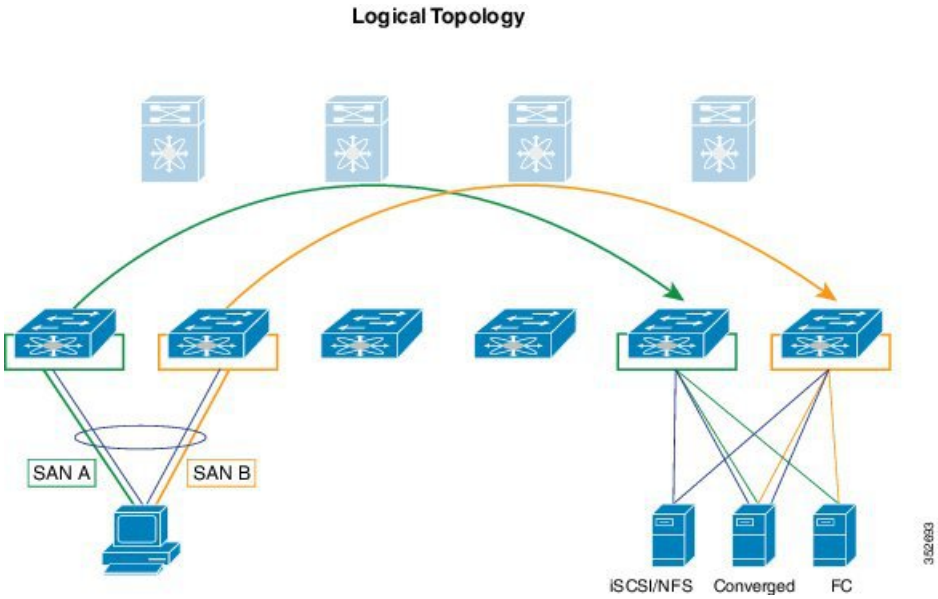

Dynamic FCoE gains the additional redundancy that is inherent in the FabricPath network by using the increased spine connectivity. A larger network with a large number of spines means increased reliability and stability for the storage network. This is achieved while retaining the best practices requirements for storage environments.

Load-Balancing FCoE Traffic on a Dynamic VFC

FabricPath provides redundant paths between a source and destination. Because FCoE traffic traverses the FabricPath network with one or more FCoE and non-FCoE nodes (spines, leafs), you must ensure in-order delivery through proper port-channel hashing across the redundant paths. All FabricPath nodes have port-channel hashing enabled that includes the exchange ID. Traffic from a single flow always traverses through only one set of nodes through the network to maintain in-order delivery.

Supported Dynamic FCoE Using FabricPath Topologies

The supported topologies for Dynamic FCoE Using FabricPath are as follows:

- FCoE devices that are directly connected to an FCF leaf
- Traditional FCoE VE_Port connectivity to an FCF leaf
- Legacy FC fabric connected to an FCF leaf
- NPV and FCoE NPV devices that are connected to an FCF leaf
- Native FC devices that are directly connected to an FCF leaf

Note Although physical separation is possible through a multi-topology configuration of FabricPath, it is not required.

Prerequisites for Dynamic FCoE Using FabricPath

Dynamic FCoE prerequisites are as follows:

- You must enable FabricPath.
- You must enable feature fcoe for the FCF leafs.
- You must assign the highest FabricPath cost to the MCT if there is a vPC+ MCT on the FCF leafs.
-
- You must enable mode fabric path on the VLANs that are mapped to VSANs in all the nodes (leafs and spines).

Guidelines and Limitations for Dynamic FCoE Using FabricPath

Dynamic FCoE Using FabricPath has the following guidelines and limitations:

- You must enable feature FCoE on the FabricPath leaf node.
- You must enable mode FabricPath on FCoE VLANs used for storage traffic.

- The minimum number of switches for a FabricPath deployment is one switch. However, if you are going to have a separation of SAN A/B, you need to have two spine switches. Otherwise, there is no separation at all.
- You must statically define the FabricPath switch ID. Changing a switch ID is required for a dynamic vFC. Some traffic loss might occur during a switch ID change. We recommend that you statically configure switch IDs.
- A multichassis EtherChannel trunk (MCT) must be of the highest Intermediate System-to-Intermediate System (IS-IS) cost which is 16777215. FCoE VLANs do not come up as an MCT. Fabric IS-IS should be high so that FCoE/FTP traffic does not go through.
- You should ensure the following:
  - Define the FCoE VLAN in a separate topology and explicitly prune the MCT links.
  - Configure a higher cost on MCT to avoid using it for regular forwarding.
- Shutting a VFC dynamically is not recommended because a Layer 2 Multipathing (L2MP) loop might occur and result in traffic loss.
- If you want to take a certain data path for a VSAN, use a FabricPath multitopology in the Dynamic FCoE Using FabricPath topology.

Configuration Topology Example

The following figure represents the configuration example that will be described in the following sections.

Figure 5: Configuration Example

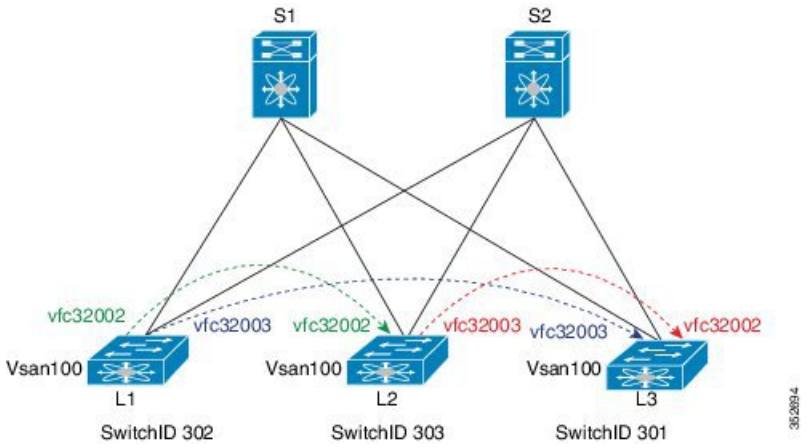

Note The component labels in the previous diagram are for illustrative purposes only.

Configuring Dynamic FCoE Using FabricPath

SUMMARY STEPS

1. Establish the FabricPath infrastructure.
2. Configure non-FCoE leafs for FCoE traffic.
3. Configure leafs for FCoE (FCF) processing.
4. Configure ports on leafs for FC/FCoE.

DETAILED STEPS

Procedure

Step 1 Establish the FabricPath infrastructure.

Step 2 Configure non-FCoE leafs for FCoE traffic.

Step 3 Configure leafs for FCoE (FCF) processing.

A leaf needs this configuration for failover cases.

Step 4 Configure ports on leafs for FC/FCoE.

Instantiation and Initialization of Dynamic VFC

Dynamic FCoE enables the capability of creating both a virtual Fibre Channel port (VFC), as well as instantiating the Inter-Switch Link port type (VE_Port/TE Port). Enabling FCoE and FabricPath on the same VLAN should serve as a trigger to instantiation and initialization of the Dynamic VFCs in TE mode. The process is as follows:

1. Every FCF leaf is uniquely identified by a global FCF-MAC address.
2. Every FCF leaf floods an FIP unsolicited multicast discovery advertisement to ALL-FCF MAC addresses and source MAC addresses that are set to its global FCF-MAC address on the FabricPath-enabled FCoE VLANs. This is triggered by two factors:
   a. Feature FCoE is enabled on the leaf.
   b. FabricPath is enabled on the FCoE VLANs.
3. All FCF leafs on this FabricPath cloud should receive this multicast advertisement on the corresponding FCoE-enabled FP VLAN. Upon receiving this FIP multicast frame, a dynamic VFC in VE mode is created between the two FCF leaf nodes.
4. Only one dynamic VFC in TE mode is between any two FCF leafs.
5. The dynamic VFCs can be differentiated based on their VFC ID range. All dynamic VFCs obtain an ID that is greater than 32001.

6. The VFC might have multiple FabricPath FCoE VLANs up. The VLANs might or might not be in the same topology.
7. Every FCF leaf is one hop away. For all VE paths that use FabricPath, a default fixed FSPF cost value is used.

Verifying the Dynamic FCoE Using FabricPath Configuration

Command	Purpose
show interface brief	Displays a brief summary of the interface configuration information.
show interface vfc	Displays the configuration information of virtual Fibre Channel interfaces.
show vpc	Displays the configuration information of virtual port channels.
show topology	Displays topology information for connected SAN switches.
show fcoe	Displays the status of FCoE parameters on the switch.
show running-config	Displays the configuration that is currently running on the switch.
show fcoe dce	Displays the Dynamic FCoE database using FabricPath.

show interface brief Command

```
switch# show interface brief
--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Por
t
Interface                                                                     Ch
#
--------------------------------------------------------------------------------
Eth1/1        1       eth  access up      none                       10G(D) --
Eth1/2        1       eth  access down    Link not connected         10G(D) --
Eth1/3        1       eth  access up      none                       10G(D) --
Eth1/4        1       eth  access up      none                       10G(D) --
Eth1/5        1       eth  access up      none                       10G(D) --
Eth1/6        1       eth  access up      none                       10G(D) --
Eth1/7        1       eth  access up      none                       10G(D) --
Eth1/8        1       eth  access down    SFP not inserted           10G(D) --
Eth1/9        1       eth  access down    SFP validation failed      10G(D) --
Eth1/10       1       eth  access down    SFP not inserted           10G(D) --
Eth1/11       1       eth  f-path up      none                       10G(D) --
Eth1/12       1       eth  access down    SFP not inserted           10G(D) --
Eth1/13       1       eth  access up      none                       10G(D) --
Eth1/14       1       eth  access up      none                       10G(D) --
Eth1/15       1       eth  access down    SFP validation failed      10G(D) --
Eth1/16       1       eth  access down    Link not connected         10G(D) --
```

```
Eth1/17       1       eth  access up      none                          10G(D) --
Eth1/18       1       eth  access up      none                          10G(D) --
Eth1/19       1       eth  access down    SFP validation failed         10G(D) --
Eth1/20       1       eth  access up      none                          10G(D) --
Eth1/21       1       eth  access down    SFP validation failed         10G(D) --
Eth1/22       1       eth  access up      none                          10G(D) --
Eth1/23       1       eth  access down    SFP validation failed         10G(D) --
Eth1/24       1       eth  access down    SFP not inserted              10G(D) --
Eth1/25       1       eth  access up      none                          10G(D) --
Eth1/26       1       eth  access up      none                          10G(D) --
Eth1/27       1       eth  access up      none                          10G(D) --
Eth1/28       1       eth  access up      none                          10G(D) --
Eth1/29       1       eth  access up      none                          10G(D) --
Eth1/30       1       eth  access down    SFP not inserted              10G(D) --
Eth1/31       1       eth  access down    SFP not inserted              10G(D) --
Eth1/32       1       eth  access down    SFP not inserted              10G(D) --

--------------------------------------------------------------------------------
Port   VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0  --           up     10.193.52.117                           1000     1500

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status      Bind                Oper    Oper
                  Mode   Trunk               Info                Mode    Speed
                         Mode                                            (Gbps)
--------------------------------------------------------------------------------
vfc32002   1      E     on     trunking    54:7f:ee:b1:8a:00       TE   10
vfc32003   1      E     on     trunking    54:7f:ee:73:e8:00       TE   10
```

show interface vfc Command

```
switch# show interface vfc 32002
vfc32002 is trunking
    Dynamic VFC Peer MAC is 54:7f:ee:b1:8a:00
    Hardware is Ethernet
    Port WWN is 2d:01:54:7f:ee:73:e6:78
    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) (1,100)
    Trunk vsans (up)                       (100)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1)
    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
      0 frames input, 0 bytes
      0 frames output, 0 bytes
    Interface last changed at Mon Feb 14 19:46:53 2011

switch# show interface vfc 32003
vfc32003 is trunking
    Dynamic VFC Peer MAC is 54:7f:ee:73:e8:00
    Hardware is Ethernet
    Port WWN is 2d:02:54:7f:ee:73:e6:78
    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) (1,100)
```

```
    Trunk vsans (up)                        (100)
    Trunk vsans (isolated)                  ()
    Trunk vsans (initializing)              (1)
    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
      0 frames input, 0 bytes
      0 frames output, 0 bytes
    Interface last changed at Mon Feb 14 19:49:23 2011


================================================================================
```

show vpc Command

```
switch# show vpc
   vPC domain id : 300 vPC+ switch id : 1550
   vPC Peer-link status
   ---------------------------------------------------------------------
   id   Port   Status Active vlans
   --   ----   ------ --------------------------------------------------
   1   Po1    up     -
```

show topology Command

```
switch# show topology
FC Topology for VSAN 100 :
-------------------------------------------------------------------------------
       Interface  Peer Domain Peer Interface     Peer IP Address(Switch Name)
-------------------------------------------------------------------------------
         vfc32002  0x0b(11)         vfc32002  10.193.52.108(nc-9)
         vfc32003 0x64(100)         vfc32003  10.193.52.118(o2-356)
```

show fcoe Command

```
switch# show fcoe
Global FCF details
        FCF-MAC is 54:7f:ee:73:e6:20
        FC-MAP is 0e:fc:00
        FCF Priority is 128
        FKA Advertisement period for FCF is 8 seconds

VFC MAC details
```

show fcoe dce Command

```
switch# show fcoe dce

Dynamic VFC MAC details :
------------------------------------------------------------
   Interface    Peer-swid       Peer-mac
------------------------------------------------------------
   vfc32002     303             54:7f:ee:b1:8a:00
   vfc32003     301             54:7f:ee:73:e8:00
```

Configuration Output Examples for Dynamic FCoE Using FabricPath

The following output examples show how to configure Dynamic FCoE using FabricPath. You must enter the **feature fabricpath** command and configure the appropriate links as FabricPath core ports.

This example covers VSAN 100 and VSAN 200.

The following is a description of the topology example:

- S1 and S2 are FabricPath spines.
- L1 through L4 are FCF leafs.
- L5 and L6 are non-FCoE leafs.

Figure 6: Sample Dynamic FCoE Configuration

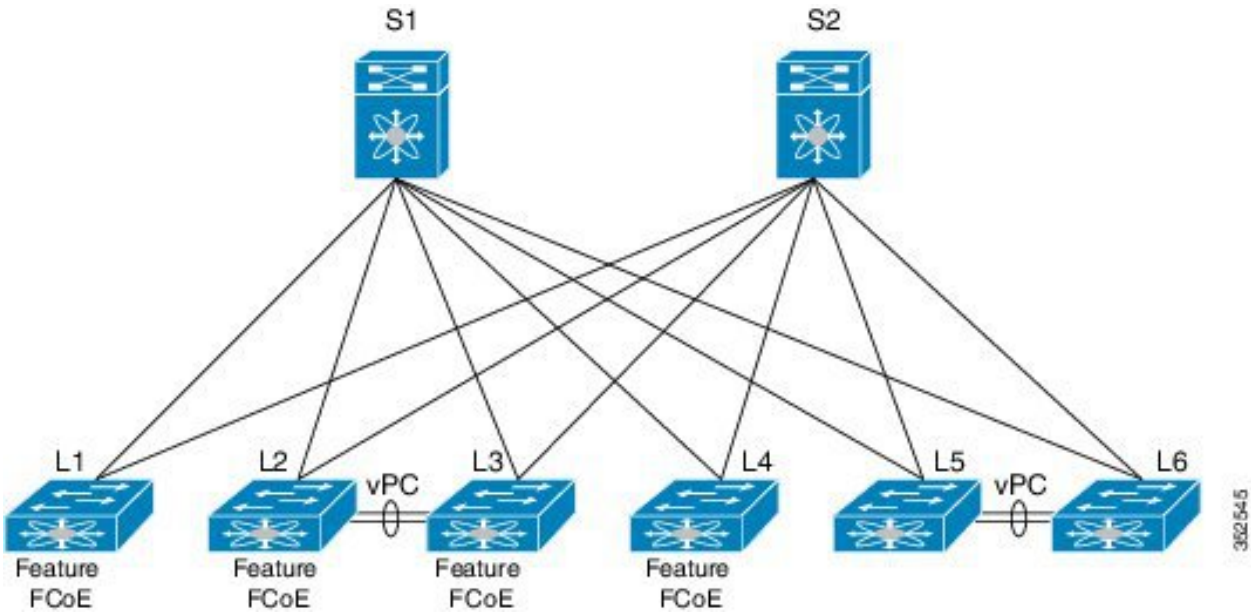

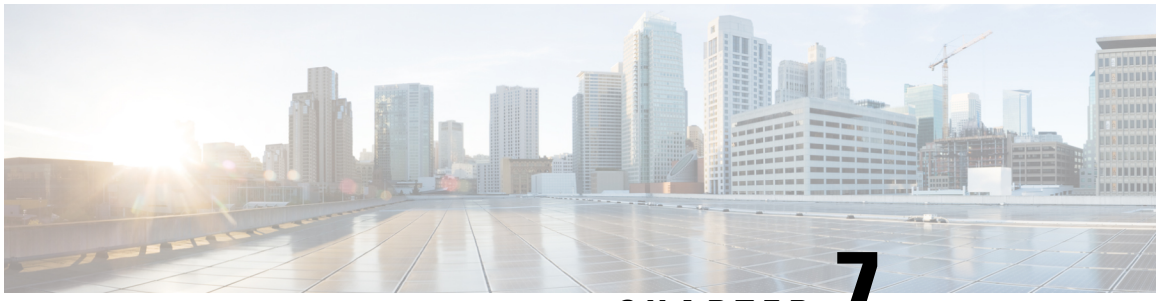

CHAPTER 7

FCoE Over FEX

- Information About FCoE Over FEX, on page 59
- Default Settings for FCoE, on page 60
- Guidelines and Limitations, on page 60
- FCoE Over FEX, on page 61
- Verifying the FCoE over FEX Configuration, on page 69
- Example: FCoE Over FEX Configuration, on page 71
- Additional References for FCoE over FEX, on page 72

Information About FCoE Over FEX

The Fibre Channel over Ethernet (FCoE) over Fabric Extenders (FEX) feature allows Fibre Channel traffic to be carried on a FEX port. To enable this feature, the FEX port is shared with the storage Virtual Device Context (VDC). The FEX is connected to the Cisco Nexus 7000/7700 device through a Fabric Port Channel (FPC). FCoE over FEX enables provision of FCoE on host connections.

FCoE Over FEX With Physical Port vPC

The FCoE over FEX with Virtual PortChannel (vPC) feature allows Fibre Channel traffic to be carried over a FEX using a physical port virtual PortChannel (vPC). To enable this feature, the FEX port is shared with the storage VDC.

LAN Shutdown

The LAN Shutdown feature detects the capability of the FCoE host to support Data Center Bridging Exchange (DCBX). DCBX allows the switch to send LAN Logical Link Status (LLS) messages in a type-length-value (TLV) format. The LAN shutdown feature enables bring up and bring down of LAN links on a unified link carrying both FCoE and LAN traffic. When you enable the **shutdown lan** command, only the LAN traffic stops while the FCoE traffic continues.

Default Settings for FCoE

This table lists the default settings for FCoE parameters.

Table 5: Default FCoE Parameter Settings

Parameters	Default
FCoE feature	Not installed, disabled
FC-Map	0E.FC.00
Fabric priority	128
Advertisement interval	8 seconds

Guidelines and Limitations

FCoE Over FEX

FCoE over FEX has the following guidelines and limitations:

- FCoE over FEX is supported with 2232PP, B22-HP, N2348UPQ FEX.
- To enable FCoE over FEX on the storage VDC, use the **feature-set fex** command in the storage VDC.
- FCoE license should be enabled on the FEX uplink port on the parent switch.
- To enable support of FCoE over FEX on the parent switch explicitly, use the **hardware qos fcoe-fex** command.
- Inter VSAN route (IVR) zone configuration is not supported for FCoE over FEX.
- Custom QoS policy cannot be applied on the uplink interfaces of the FEX Port channel.
- FCoE over host interface port channel (HIF-PC) is not supported.
- If you connect a FEX with FPC ports that is sharing its ports with a storage VDC to an F2e module that is also sharing its ports with the same storage VDC, the port manager times out when you reload the F2e module. The ports on the F2e module that were shared with the storage VDC will be in the removed state in the storage VDC. You will have to reload the switch to recover from this state.
- FCoE over FEX is not supported on Active/Active FEX.

Note For more information about FEX, refer http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus2000/sw/configuration/guide/rel_4_0_1a/NX2000CLIConfig/FEX-features.html

FCoE over FEX with Physical Port vPC

FCoE over FEX with Physical Port vPC has the following guidelines and limitations:

- FCoE over physical port vPC+ is supported.
- FCoE over physical port vPC is supported on shared interfaces.
- FCoE over physical port vPC is supported only on host facing links.
- Port channel vPC does not support FCoE.
- Both vPC legs should be shared between the ethernet VDC and the storage VDC.

LAN Shutdown

The LAN Shutdown feature has the following guidelines and limitations:

- LAN shutdown is supported only on shared interfaces in the parent virtual device context (VDC).
- LAN shutdown requires Link Layer Discovery Protocol (LLDP) to run in the parent VDC.

FCoE Over FEX

Bringing up FEX in Ethernet VDC

Refer Configuring the Fabric Extender section.

Configuring FCoE over FEX

SUMMARY STEPS

1. **configure terminal**
2. **hardware qos fcoe-fex**
3. **license fcoe module** *module-number*
4. **exit**
5. **switchto vdc** *ethernet-vdc-name*
6. **configure terminal**
7. **interface port-channel** *channel-number*
8. **switchport**
9. **switchport mode fex-fabric**
10. **priority-flow-control mode** {**auto** | **off** | **on**}
11. **fex associate** *chassis-id*
12. **no shutdown**
13. **interface ethernet** *slot/port-list*
14. **switchport**
15. **switchport mode fex-fabric**

16. **priority-flow-control mode** {**auto** | **off** | **on**}
17. **fex associate** *chassis-id*
18. **channel-group** *number*
19. **no shutdown**
20. **interface** *if-range*
21. **switchport**
22. **switchport mode trunk**
23. **spanning tree port type edge trunk**
24. **no shutdown**
25. **exit**
26. **feature lldp**
27. **switchback**
28. **switchto vdc** *storage-vdc-name*
29. **configure terminal**
30. **feature-set fex**
31. **feature lldp**
32. **switchback**
33. **configure terminal**
34. **vdc** *storage-vdc-name*
35. **allocate fcoe-vlan-range** *vlan-range* **from vdc** *ethernet-vdc-name*
36. [**no**] **allocate shared interface** *if-range*

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode in the default VDC.

Step 2	**hardware qos fcoe-fex** **Example:** `switch(config)# hardware qos fcoe-fex`	Enables FCoE over FEX support and Priority Flow Control (PFC) behavior on the FEX port. Use the **no** form of the command to disable FCoE over FEX support.
Step 3	**license fcoe module** *module-number* **Example:** `switch(config)# license fcoe module 3`	Associates an FCoE license to a module in which the fabric port channel (FPC) ports are present. .
Step 4	**exit** **Example:** `switch(config)# exit`	Exits the default VDC configuration mode.

	Command or Action	Purpose
Step 5	**switchto vdc** *ethernet-vdc-name* **Example:** `switch# switchto vdc ethernet-vdc`	Switches to the ethernet VDC from the default VDC.

Step 6	**configure terminal**
	Example:
	`switch-ethernet-vdc# configure terminal`
	`switch-ethernet-vdc(config)#`
		Enters configuration mode.

Step 7	**interface port-channel** *channel-number*
	Example:
	`switch-ethernet-vdc(config)# interface port-channel 4`
	`switch-ethernet-vdc(config-if)#`
		Creates a port-channel interface and enter interface configuration mode.

Step 8	**switchport** **Example:** `switch-ethernet-vdc(config-if)# switchport`	Sets the interface as a Layer 2 switching port.
Step 9	**switchport mode fex-fabric** **Example:** `switch-ethernet-vdc(config-if)# switchport mode fex-fabric`	Sets the interface type to be an uplink port for a Fabric Extender (FEX).
Step 10	**priority-flow-control mode** {**auto** | **off** | **on**} **Example:** `switch-ethernet-vdc(config-if)# priority-flow-control mode on`	Configure Priority Flow Control on a per-port basis and on the port channel. **Note** This command is enabled by default when you enable the **hardware qos fcoe-fex** command in the default VDC.
Step 11	**fex associate** *chassis-id* **Example:** `switch-ethernet-vdc(config-if)# fex associate 100`	Associates a Fabric Extender (FEX) to a fabric interface.
Step 12	**no shutdown** **Example:** `switch-ethernet-vdc(config-if)# no shutdown`	Administratively enables the Ethernet shared interface.

Step 13	**interface ethernet** *slot/port-list*
	Example:
	`switch-ethernet-vdc(config)# interface ethernet 1/37`
	`switch-ethernet-vdc(config-if)#`
		Configures an Ethernet interface and enter interface configuration mode.

Step 14	**switchport** **Example:**	Sets the interface as a Layer 2 switching port.

	Command or Action	Purpose
	`switch-ethernet-vdc(config-if)# switchport`	
Step 15	**switchport mode fex-fabric** **Example:** `switch-ethernet-vdc(config-if)# switchport mode fex-fabric`	Sets the interface type to be an uplink port for a Fabric Extender (FEX).
Step 16	**priority-flow-control mode** {**auto** | **off** | **on**} **Example:** `switch-ethernet-vdc(config-if)# priority-flow-control mode on`	Configure Priority Flow Control on a per-port basis. **Note** This command is enabled by default when you enable the **hardware qos fcoe-fex** command in the default VDC.
Step 17	**fex associate** *chassis-id* **Example:** `switch-ethernet-vdc(config-if)# fex associate 100`	Associates a Fabric Extender (FEX) to a fabric interface.
Step 18	**channel-group** *number* **Example:** `switch-ethernet-vdc(config-if)# channel-group 4`	Assigns and configures a physical interface to a port-channel group.
Step 19	**no shutdown** **Example:** `switch-ethernet-vdc(config-if)# no shutdown`	Administratively enables the Ethernet shared interface.

Step 20	**interface** *if-range*
	Example:
	`switch-ethernet-vdc(config-if)# interface ethernet 100/1/1`
	`switch-ethernet-vdc(config-if)#`
		Configures the interface and enters interface configuration mode in the ethernet VDC.

Step 21	**switchport** **Example:** `switch-ethernet-vdc(config-if)# switchport`	Sets the interface as a Layer 2 switching port.
Step 22	**switchport mode trunk** **Example:** `switch-ethernet-vdc(config-if)# switchport mode trunk`	Puts the Ethernet interface into trunk mode.
Step 23	**spanning tree port type edge trunk** **Example:** `switch-ethernet-vdc(config-if)# spanning tree port type edge trunk`	Configures an interface connected to a host as an edge port
Step 24	**no shutdown** **Example:**	Administratively enables the Ethernet shared interface.

	Command or Action	Purpose
	`switch-ethernet-vdc(config-if)# no shutdown`	
Step 25	**exit** **Example:** `switch-ethernet-vdc(config-if)# exit`	Exits the interface configuration mode.
Step 26	**feature lldp** **Example:** `switch-ethernet-vdc(config)# feature lldp`	Enables the LLDP feature in the ethernet VDC.

Step 27	**switchback**
	Example:
	`switch-ethernet-vdc(config)# switchback`
	`switch#`
		Go to default VDC.

Step 28	**switchto vdc** *storage-vdc-name*
	Example:
	`switch# switchto vdc storage-vdc`
	`switch-storage-vdc#`
		Switches to the storage VDC.

Step 29	**configure terminal**
	Example:
	`switch-storage-vdc# configure terminal`
	`switch-storage-vdc(config)#`
		Enters configuration mode.

Step 30	**feature-set fex** **Example:** `switch-storage-vdc(config)# feature-set fex`	Enables the Fabric Extender (FEX) feature set in the storage VDC.
Step 31	**feature lldp** **Example:** `switch-storage-vdc(config)# feature lldp`	Enables the LLDP feature in the storage VDC.

Step 32	**switchback**
	Example:
	`switch-storage-vdc(config)# switchback`
	`switch#`
		Go to default VDC.

Step 33	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode in the default VDC.

Step 34	**vdc** *storage-vdc-name* **Example:** `switch(config)# vdc storage-vdc`	Specifies the storage VDC in the default VDC.

	Command or Action	Purpose
Step 35	**allocate fcoe-vlan-range** *vlan-range* **from vdc** *ethernet-vdc-name* **Example:** `switch(config-vdc)# allocate fcoe-vlan-range 2-8 from vdc ethernet-vdc`	Allocates Fibre Channel over Ethernet (FCoE) VLANs to a virtual device context (VDC),

Step 36	[**no**] **allocate shared interface** *if-range*
	Example:
	`switch(config-vdc)# allocate shared interface ethernet 100/1/1`
	`Ports that share the port group of the interfaces you have specified will be affected as well. Continue (y/n)? y`
		Shares FEX interface from the Ethernet VDC to the storage VDC.
		Note It is expected that, if the feature-set FEX is not enabled, then the **allocate shared interface** *fex_interfaces* command fails on VDC write-erase/reload. To resolve, un-configure the shared interfaces and configure it back.

Example

This example shows how to configure FCoE over FEX:

```
!Default VDC
switch# configure terminal
switch(config)# hardware qos fcoe-fex
switch(config)# license fcoe module 3
switch(config)# exit
switch# switchto vdc ethernet-vdc
!Ethernet VDC
switch-ethernet-vdc# configure terminal
switch-ethernet-vdc(config)# interface port-channel 4
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode fex-fabric
switch-ethernet-vdc(config-if)# priority-flow-control mode on
switch-ethernet-vdc(config-if)# fex associate 100
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config)# interface ethernet 1/37
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode fex-fabric
switch-ethernet-vdc(config-if)# priority-flow-control mode on
switch-ethernet-vdc(config-if)# fex associate 100
switch-ethernet-vdc(config-if)# channel-group 4
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config-if)# interface ethernet 100/1/1
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode trunk
switch-ethernet-vdc(config-if)# spanning tree port type edge trunk
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config-if)# exit
switch-ethernet-vdc(config)# feature lldp
switch-ethernet-vdc(config)# switchback
!Default VDC
switch# switchto vdc storage-vdc
!Storage VDC
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# feature-set fex
switch-storage-vdc(config)# feature lldp
switch-storage-vdc(config)# switchback
!Default VDC
```

```
switch# configure terminal
switch(config)# vdc storage-vdc
switch(config-vdc)# allocate fcoe-vlan-range 2-8 from vdc ethernet-vdc
switch(config-vdc)# allocate shared interface ethernet 100/1/1
```

Binding VFC to Shared FEX Interface Explicitly

SUMMARY STEPS

1. **configure terminal**
2. **interface vfc** *vfc-id*
3. **bind** {**interface** {**ethernet** [*chassid-id*]/ *slot*/ *port*}}
4. **no shutdown**
5. (Optional) **show interface vfc**
6. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch-storage-vdc# configure terminal`
	`switch-storage-vdc(config)#`
		Enters configuration mode.

Step 2	**interface vfc** *vfc-id*
	Example:
	`switch-storage-vdc(config)# interface vfc 100`
	`switch-storage-vdc(config-if)#`
		Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode. The *chassis-id* range is from 100 to 199.

Step 3	Required: **bind** {**interface** {**ethernet** [*chassid-id*]/ *slot*/ *port*}} **Example:** `switch-storage-vdc(config-if)# bind interface ethernet 100/1/1`	Explicitly binds the virtual Fibre Channel interface to the specified interface. Use the **no** form of this command to unbind the virtual Fibre Channel interface from the specified interface. The *chassis-id* range is from 100 to 199.
Step 4	**no shutdown** **Example:** `switch-storage-vdc(config-if)# no shutdown`	Administratively enables the Ethernet shared interface.
Step 5	(Optional) **show interface vfc** **Example:** `switch-storage-vdc(config-if)# show interface vfc`	Displays information about the virtual Fibre Channel interfaces.

	Command or Action	Purpose
Step 6	(Optional) **copy running-config startup-config** **Example:** `switch-storage-vdc(config)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Example

This example shows how to bind a virtual Fibre Channel interface to a FEX interface:

```
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# interface vfc 100
switch-storage-vdc(config-if)# bind interface ethernet 100/1/1
switch-storage-vdc(config-if)# no shutdown
```

Binding VFC to Shared FEX Interface Implicitly

SUMMARY STEPS

1. **configure terminal**
2. **interface vfc** [*chassid-id*]/ *slot*/ *port*
3. **no shutdown**
4. (Optional) **show interface vfc**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch-storage-vdc# configure terminal`
	`switch-storage-vdc(config)#`
		Enters configuration mode.

Step 2	**interface vfc** [*chassid-id*]/ *slot*/ *port*
	Example:
	`switch-storage-vdc(config)# interface vfc 100/1/1`
	`switch-storage-vdc(config-if)#`
		Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode. This command binds the VFC to the shared FEX interface implicitly. The *chassis-id* range is from 100 to 199. This command binds the Ethernet interface 100/1/1 to the shared FEX interface.

Step 3	**no shutdown** **Example:** `switch-storage-vdc(config-if)# no shutdown`	Administratively enables the Ethernet shared interface.

	Command or Action	Purpose
Step 4	(Optional) **show interface vfc** **Example:** `switch-storage-vdc(config-if)# show interface vfc`	Displays information about the virtual Fibre Channel interfaces.
Step 5	(Optional) **copy running-config startup-config** **Example:** `switch-storage-vdc(config)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Example

This example shows how to bind a virtual Fibre Channel interface to a FEX interface implicitly:

```
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# interface vfc 100/1/1
switch-storage-vdc(config)# no shutdown
```

Verifying the FCoE over FEX Configuration

To display FCoE over FEX configuration information, perform one of these tasks:

Command	Purpose
show fex	Displays information about all Fabric Extenders (FEX) available in a specific VDC. Use this command in storage VDC and ethernet VDC.
show fex detail	Displays detailed information about all Fabric Extenders (FEX) available in a specific VDC. Use this command in storage VDC and ethernet VDC.
show fex *chassis-id*	Displays information about a specific FEX. Use this command in storage VDC and ethernet VDC.
show vdc shared membership	Displays the FEX interfaces shared with the storage virtual device context (VDC). Use this command in the default VDC.
show vdc shared membership both	Displays the FPC ports that are implicitly shared as a result of the shared FEX interface. Use this command in the default VDC.

Command	Purpose
show flogi database	Lists all the FLOGI sessions. Use this command in the storage VDC.

This example shows how to display information about all attached Fabric Extender (FEX) chassis:

```
switch-storage-vdc# show fex
FEX             FEX                 FEX                             FEX

Number        Description            State                 Model                Serial
-------------------------------------------------------------------------------------------
100           FEX0100                Online                N2K-C2232PP-10GE     SSI181202V9
```

This example shows how to display detailed information about all attached FEX chassis:

```
switch-ethernet-vdc# show fex detail
 FEX: 101 Description: FEX0100   state: Online
  FEX version: 7.2(0)D1(1) [Switch version: 7.2(0)D1(1)]
  FEX Interim version: (1)FIP_0_28
  Switch Interim version: 7.2(0)D1(1)
  Extender Serial: SSI181202V9
  Extender Model: N2K-C2232PP-10GE,  Part No: 73-12533-06
  Card Id: 82, Mac Addr: 58:0a:20:37:e4:02, Num Macs: 64
  Keepalive pending for 1 intervals
  Module Sw Gen: 12594  [Switch Sw Gen: 21]
  Pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Eth1/19
  FCoE Admin: false
  FCoE Oper: true
  FCoE FEX AA Configured: false
  Fabric interface state:
    Po100 - Interface Up. State: Active
    Eth1/37 - Interface Up. State: Active
  Fex Port        State  Fabric Port
       Eth100/1/1  Up         Po4
       Eth100/1/2 Down        Po4
       Eth100/1/3  Down       Po4
       Eth100/1/4  Down       Po4
       Eth100/1/5  Down       Po4
       Eth100/1/6  Down       Po4
       Eth100/1/7  Down       Po4
       Eth100/1/8  Down       Po4
       Eth100/1/9  Down       Po4
      Eth100/1/10  Down       Po4
      Eth100/1/11  Down       Po4
      Eth100/1/12  Down       Po4
      Eth100/1/13  Down       Po4
      Eth100/1/14  Down       Po4
      Eth100/1/15  Down       Po4
      Eth100/1/16  Down       Po4
      Eth100/1/17  Down       Po4
      Eth100/1/18  Down       Po4
      Eth100/1/19  Down       Po4
      Eth100/1/20  Down       Po4
      Eth100/1/21  Down       Po4
      Eth100/1/22  Down       Po4
      Eth100/1/23  Down       Po4
      Eth100/1/24  Down       Po4
      Eth100/1/25  Down       Po4
      Eth100/1/26  Down       Po4
      Eth100/1/27  Down       Po4
      Eth100/1/28  Down       Po4
```

```
      Eth100/1/29  Down       Po4
      Eth100/1/30  Down       Po4
      Eth100/1/31  Down       Po4
      Eth100/1/32  Down       Po4
<---output truncated--->
```

This example shows how to display information about a specific FEX:

```
switch-storage-vdc# show fex 100
FEX: 101 Description: FEX0100   state: Online
  FEX version: 7.2(0)D1(1) [Switch version: 7.2(0)D1(1)]
  Extender Serial: SSI181202V9
  Extender Model: N2K-C2232PP-10GE,  Part No: 73-12533-06
  Keepalive pending for 1 intervals
  Pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Eth1/37
  FCoE Admin: false
  FCoE Oper: false
  FCoE FEX AA Configured: false
  Fabric interface state:
    Po101 - Interface Up. State: Active
    Eth3/1 - Interface Up. State: Active
    Eth3/2 - Interface Up. State: Active
```

This example shows how to display the FEX interfaces shared with the storage virtual device context (VDC):

```
switch# show vdc shared membersip
vdc_id: 3 vdc_name: fcoe-vdc interfaces:
        Ethernet100/1/1
```

This example shows how to display the FPC ports that are implicitly shared as a result of the shared FEX interface:

```
switch# show vdc shared membership both
vdc_id: 3 vdc_name: fcoe-vdc interfaces:
        Ethernet100/1/1
Implicitly shared interfaces:
        Ethernet100/1/2
```

This example shows how to list the FLOGI sessions:

```
switch-storage-vdc# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
vfc100/1/1       100   0x800020  10:00:00:00:06:67:e9:00 20:00:00:00:06:67:e9:00

Total number of flogi = 1.
```

Example: FCoE Over FEX Configuration

```
!Configuring FCoE over FEX:
!Default VDC
switch# configure terminal
switch(config)# hardware qos fcoe-fex
switch(config)# license fcoe module 3
switch(config)# exit
switch# switchto vdc ethernet-vdc
!Ethernet VDC
switch-ethernet-vdc# configure terminal
switch-ethernet-vdc(config)# interface port-channel 4
```

```
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode fex-fabric
switch-ethernet-vdc(config-if)# priority-flow-control mode on
switch-ethernet-vdc(config-if)# fex associate 100
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config)# interface ethernet 1/37
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode fex-fabric
switch-ethernet-vdc(config-if)# priority-flow-control mode on
switch-ethernet-vdc(config-if)# fex associate 100
switch-ethernet-vdc(config-if)# channel-group 4
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config-if)# interface ethernet 100/1/1
switch-ethernet-vdc(config-if)# switchport
switch-ethernet-vdc(config-if)# switchport mode trunk
switch-ethernet-vdc(config-if)# spanning tree port type edge trunk
switch-ethernet-vdc(config-if)# no shutdown
switch-ethernet-vdc(config-if)# exit
switch-ethernet-vdc(config)# feature lldp
switch-ethernet-vdc(config)# switchback
!Default VDC
switch# switchto vdc storage-vdc
!Storage VDC
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# feature-set fex
switch-storage-vdc(config)# feature lldp
switch-storage-vdc(config)# switchback
!Default VDC
switch# configure terminal
switch(config)# vdc storage-vdc
switch(config-vdc)# allocate fcoe-vlan-range 2-8 from vdc ethernet-vdc
switch(config-vdc)# allocate shared interface ethernet 100/1/1

!Bind VFC to Shared FEX Interface Explicitly:
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# interface vfc 100
switch-storage-vdc(config-if)# bind interface ethernet 100/1/1
switch-storage-vdc(config-if)# no shutdown

!Bind VFC to Shared FEX Interface Implicitly:
switch-storage-vdc# configure terminal
switch-storage-vdc(config)# interface vfc 100/1/1
switch-storage-vdc(config)# no shutdown
```

Additional References for FCoE over FEX

Related Documents

Related Topic	Document Title
Command reference	*Cisco NX-OS FCoE Command Reference for Cisco Nexus 7000*
Cisco NX-OS licensing	*Cisco NX-OS Licensing Guide*
Configuring vPCs	*Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*

MIBs

MIB	MIBs Link
	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/support

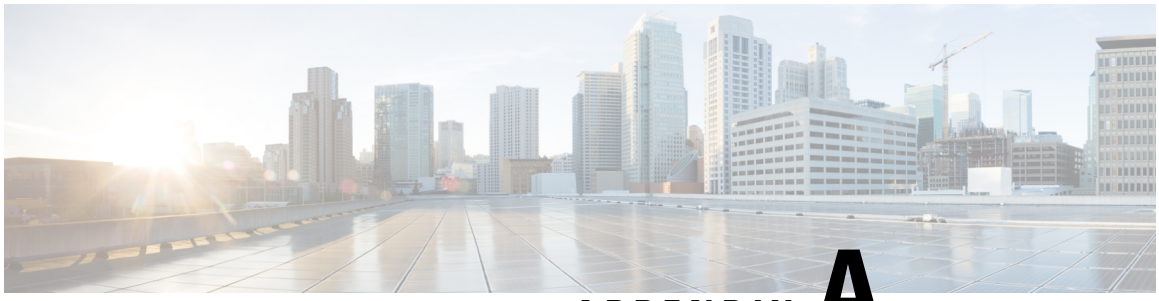

APPENDIX A

Feature History by Platform

This chapter contains the following sections:

- Supported Features, on page 75

Supported Features

Feature	Release	Feature Information
FCoE	7.2(0)D1(1)	Supports F3 Series modules.
FCoE over FEX	7.2(0)D1(1)	This feature was introduced.
Configuring Dynamic FCoE Using FabricPath	7.2(0)D1(1)	This feature was introduced.
FCoE over physical port vPC	6.2(6)	This feature was introduced. For information, see the *Cisco 7000 Series NX-OS Interfaces Configuration Guide.*
FIP	5.2(1)	Supports the T11-compliant FIP.
FCoE	5.2(1)	This feature was introduced.
FCoE	6.1(1)	Supports F2 Series modules.
FCoE	6.1(1)	Supports F2e Series modules (SFP+ only).

Feature History by Platform

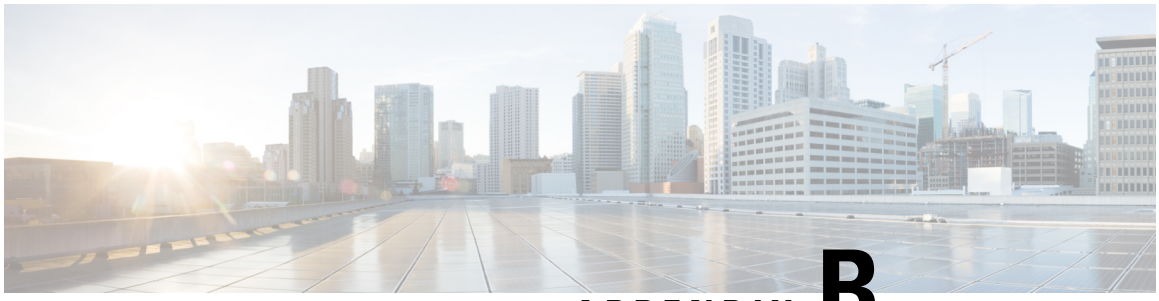

APPENDIX B

Guidelines and Limitations

This chapter includes the following sections:

- Verified Limits for FCoE, on page 77
- FCoE, on page 77

Verified Limits for FCoE

The configuration limits are documented in the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

FCoE

- VDCs apply only to Cisco Nexus 7000 Series Switches.
- You cannot enable FCoE on default VLAN.
- The QoS policy must be the same on all Cisco FCoE switches in the network.
- Beginning with Cisco NX-OS Release 6.1, FCoE is supported on F2 and F2e Series modules. F3 Series modules are supported from Cisco NX-OS Release 6.2(6) onwards.
  - FCoE supports only F2e (SFP+) modules.
  - FCoE does not support F2e (Copper) modules.

FCoE VDC

FCoE in a dedicated storage VDC has the following guidelines:

- Enable the FCoE feature set in only one VDC.
- Create VLANs in the FCoE allocated VLAN range.
- Do not enable any other features other than storage-related features in the dedicated FCoE VDC.
- Allocate resources for the dedicated FCoE VDC from an F Series module, such as the 32-port 10-Gigabit Ethernet I/O module (PID N7K-F132XP-15) .
- Rollback is not supported in a storage VDC.

- For Cisco NX-OS Release 7.2(0)D1(1), ports from only 24 FEXes can be shared to storage VDC. System will not restrict the user to go beyond 24 but, more than 24 is not tested and not supported.
- FCoE on F2, F2e, and F3 Series modules is supported with the Supervisor 2 module (N7K-SUP2 for Cisco Nexus 7000 Series devices) and the Supervisor 2E module (N77-SUP2E for Cisco Nexus 7700 Series devices and N7K-SUP2E for Cisco Nexus 7000 Series devices).
- In order to enable FCoE over FEX on the storage VDC, you must execute the **allow feature-set FEX** command from the Admin or default VDC beforehand for storage VDC. FCoE over FEX is available from 7.2(0)D1(1) and onwards.
- IVR (Inter VSAN route) zone configuration is not supported for FCoE over FEX.
- F3 Fiber Channel over Ethernet (FCoE) feature licensing is supported from 7.2.0 release onwards. To downgrade to the older version of the image 6.2.x, first uninstall the F3 FCoE license and then proceed. For more information about licensing, refer Cisco NX-OS Licensing Guide.
- F2, F2e, and F3 Series modules can co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F3 Series modules cannot co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F2 series modules cannot exist in the same VDC. This applies to both LAN and storage VDCs.
- Use the **limit-resource module-type** command in the admin or default VDC to assign module resources such as F1, F2, F2e and F3 to a storage VDC. The supported line card modules are F1, F2, F2e and F3.
- When you configure a multi-hop FCoE, ensure that you use the same no-drop classes on both sides. Priority flow control does not work when you use different no-drop classes. Use the **show interface priority-flow-control** command to verify the priority flow control operation.

Shared Interfaces

Note

- Any change in protocol state that flaps the parent port of a shared interface because of any port feature also affects the FCoE traffic on the storage vdc.
- 1500 MTU do not carry FCoE traffic in all FCoE supported platforms.

The following interface config modes are not allowed while sharing an interface from Ethernet vdc to a storage vdc:

- SPAN destination
- Private VLAN mode
- Port-channel interfaces
- Access mode
- mac-packet-classify
- Interfaces that are part of a VLAN that has an associated QoS policy

Shared Ethernet interfaces must be in trunk mode and only shared with one other VDC.

Storage VDC

Configuring a VDC for the Out-Of-Band (OOB) management interface mgmt0 is accomplished with the **vrf context management** command. However, a storage VDC does not support VRF, so configuring mgmt0 requires a different approach.

The following table shows how to configure mgmt 0 for a VDC and for a storage VDC:

Configuring mgmt 0 for VDC	Configuring mgmt 0 for storage VDC

`vrf context management`
`ip route 0.0.0.0/0 default_gateway`
	`interface mgmt 0`
	`ip address mgmt0_ip_address mgmt0_subnet_mask`
	`no shut`
	`ip route 0.0.0.0/0 default_gateway`
	Note The **ip route** command specifies the default route that points to the default gateway.

where

- *mgmt0_ip_address* is the mgmt0 IPv4 address.
- *mgmt0_subnet_mask* is the mgmt0 IPv4 netmask.
- *default_gateway* is the IPv4 address of the default-gateway.

For more information about VDC, see the Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide.

Multi-Hop FCoE Distance Configuration

In Multi-Hop FCoE, when a device sends a pause, the interface that generates the pause frame must have an ingress queue with a buffer space large enough to buffer twice the link distance. This is because, when the pause is generated the wire might get congested. By the time the adjacent device receives or processes the generated pause frame, the wire might get congested again. Therefore, the device that generates the pause must have the ability to buffer twice the link distance.

As per calculations, there can be more than 100 packets traveling on the 10 kilometer link. Due to an ASIC limitation, the F1 series line card does not support lossless FCoE on a link greater than or equal to 10 kilometers.

For more information about Multi-Hop FCoE distance limitations, see http://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/117785-probsol-nexus7000-00.html

The F3 line cards support long haul lossless distance of up to 40 kilometers. In Cisco NX-OS Release 7.2(0) and later, you can change the ingress queuing buffer configuration.

Table 6: Buffer Tuning Table for FCoE Long Distance on F2, F2E, and F3 Line Cards

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
< 5 km	F2/F2e	LR	60% no-drop and 40% drop queue
> 5 km - 10 km	F2/F2e	LR	70% no-drop and 30% drop queue

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
> 10 km - 40 km	F2/F2e	ER	80% no-drop and 20% drop queue
< 10 km	F3	LR	90% no-drop and 10% drop queue
< 40 km	F3	ER	90% no-drop and 10% drop queue