Cisco Nexus 5500 Series NX-OS Fibre Channel over Ethernet Configuration Guide, Release 7.x

First Published: 2013-01-29

Last Modified: 2019-08-06

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: www.cisco.com go trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2014 Cisco Systems, Inc. All rights reserved.

CONTENTS

PREFACE

CHAPTER 1

CHAPTER 2

CHAPTER 3

Preface

The preface contains the following sections:

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices and Cisco Nexus 2000 Series Fabric Extenders.

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.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.

Convention	Description
{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 for Cisco Nexus 5500 Series NX-OS Software

The entire Cisco NX-OS 5500 Series documentation set is available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-5000-series-switches/tsd-products-support-series-home.html

Release Notes

The release notes are available at the following URL:

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

Configuration Guides

These guides are available at the following URL:

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

The documents in this category include:

- *Cisco Nexus 5500 Series NX-OS Adapter-FEX Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS FabricPath Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS FCoE Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Fundamentals Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Interfaces Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Layer 2 Switching Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Multicast Routing Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Quality of Service Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS SAN Switching Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Security Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS System Management Configuration Guide*
- *Cisco Nexus 5500 Series NX-OS Unicast Routing Configuration Guide*

Installation and Upgrade Guides

These guides are available at the following URL:

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

The document in this category include:

- *Cisco Nexus 5500 Series NX-OS Software Upgrade and Downgrade Guides*

Licensing Guide

The *License and Copyright Information for Cisco NX-OS Software* is available at http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_0/nx-os/license_agreement/nx-ossw_lisns.html.

Command References

These guides are available at the following URL:

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

The documents in this category include:

- *Cisco Nexus 5500 Series NX-OS Fabric Extender Command Reference*
- *Cisco Nexus 5500 Series NX-OS FabricPath Command Reference*
- *Cisco Nexus 5500 Series NX-OS Fundamentals Command Reference*
- *Cisco Nexus 5500 Series NX-OS Interfaces Command Reference*
- *Cisco Nexus 5500 Series NX-OS Layer 2 Interfaces Command Reference*
- *Cisco Nexus 5500 Series NX-OS Multicast Routing Command Reference*
- *Cisco Nexus 5500 Series NX-OS Quality of Service Command Reference*
- *Cisco Nexus 5500 Series NX-OS Security Command Reference*
- *Cisco Nexus 5500 Series NX-OS System Management Command Reference*
- *Cisco Nexus 5500 Series NX-OS TrustSec Command Reference*
- *Cisco Nexus 5500 Series NX-OS Unicast Routing Command Reference*
- *Cisco Nexus 5500 Series NX-OS Virtual Port Channel Command Reference*

Technical References

The *Cisco Nexus 5500 Series NX-OS MIB Reference* is available at http://www.cisco.com/en/US/docs/switches/datacenter/nexus5500/sw/mib/reference/NX5500_MIBRef.html.

Error and System Messages

The *Cisco Nexus 5500 Series NX-OS System Message Guide* is available at http://www.cisco.com/en/US/docs/switches/datacenter/nexus5500/sw/system_messages/reference/sl_nxos_book.html.

Troubleshooting Guide

The *Cisco Nexus 5500 Series NX-OS Troubleshooting Guide* is available at http://www.cisco.com/en/US/docs/switches/datacenter/nexus5500/sw/troubleshooting/guide/N5K_Troubleshooting_Guide.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: nexus5k-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 Marketplace.
- 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

This chapter contains the following sections:

- New and Changed Information, on page 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 5500 Series NX-OS FCoE Configuration Guide, Release 7.x*.

Feature	Description	Changed in Release	Where Documented
Implicit vFC	This feature was introduced.	7.3(0)N1(1)	Creating an Implicit Virtual Fibre Channel Port Channel Interface, on page 26
Dynamic FCoE Using FabricPath	This feature was introduced.	7.0(1)N1(1)	Configuring Dynamic FCoE Using FabricPath, on page 43

CHAPTER 2

Overview

This chapter contains the following sections:

Overview

Fibre Channel over Ethernet (FCoE) allows Fibre Channel traffic to be encapsulated over a physical Ethernet link. FCoE 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; in the event of congestion, Ethernet will discard packets, relying 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.

For FCoE traffic, the Ethernet link must provide a lossless service. Ethernet links on Cisco Nexus devices provide two mechanisms to ensure lossless transport for FCoE traffic: link-level flow control (LL-FC) and priority flow control (PFC).

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 the traffic on the link.

The priority flow control feature 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. PFC can provide different levels of service to specific classes of Ethernet traffic (using IEEE 802.1p traffic classes).

Cisco Nexus switches support T11-compliant FCoE on all 10-Gigabit Ethernet interfaces.

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. Two versions of FIP are supported by the Cisco Nexus device:

- FIP—The Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) protocol supports T11-compliant Gen-2, Gen-3, and Gen-4 CNAs.
- Pre-FIP—The Cisco, Intel, Nuova Data Center Bridging Exchange (CIN-DCBX) protocol supports Gen-1 converged network adapters (CNAs).

The Cisco Nexus device detects the capabilities of the attached CNA and switches to the correct FIP mode.

FIP Virtual Link Instantiation

Cisco NX-OS support the T11-compliant FIP on Cisco Nexus devices.

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

- FIP Discovery—When a FCoE device is connected to the fabric, it sends out 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 parameters (ELP) frames 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 the connection is still valid.

FCoE Frame Format

FCoE is implemented by encapsulating a Fibre Channel frame in an Ethernet packet with a dedicated EtherType, 0x8906. That 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 might 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 Cisco Nexus device 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.

The switch expects frames from a pre-FIP CNA to be priority tagged with the FCoE CoS value. The switch will still accept untagged frames from the CNA.

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.

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 Keep Alive
	0x02	FIP Clear Virtual Links
0x0004	0x01	FIP VLAN Request
	0x02	FIP VLAN Notification

Pre-FIP Virtual Link Instantiation

Pre-FIP virtual link instantiation consists of two phases; link discovery using the Data Center Bridging Exchange protocol (DCBX), which is followed by Fabric Login.

The Cisco Nexus device is backward compatible with Gen-1 CNAs that operate in pre-FIP mode.

Note Pre-FIP is also known as the Cisco, Intel, Nuova Data Center Bridging Exchange (CIN-DCBX) protocol.

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.

The Cisco Nexus device supports two versions of DCBX:

- CEE-DCBX—The Converged Enhanced Ethernet DCBX is supported on all T11-compliant Gen-2, Gen-3, and Gen-4 CNAs.
- CIN-DCBX—The Cisco, Intel, Nuova DCBX is supported on Gen-1 converged network adapters (CNAs). CIN-DCBX is used to perform link detection in addition to other functions.

DCBX runs on the physical Ethernet link between the Cisco Nexus device 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. If the CNA supports both CIN and CEE mode, the switch and CNA will operate in CEE-DCBX mode.

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 Cisco Nexus device are programmed to accept the configuration values 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.

DCBX Feature Negotiation

The switch and CNA exchange capability information and configuration values. The Cisco Nexus devices support the following capabilities:

- FCoE—If the CNA supports FCoE capability, the switch sends the IEEE 802.1p CoS value to be used with FCoE packets.
- Priority Flow Control (PFC)—If the adapter supports PFC, the switch sends the IEEE 802.1p CoS values to be enabled with PFC.
- Priority group type-length-value (TLV).
- Ethernet logical link up and down signal.
- FCoE logical link up and down signal for pre-FIP CNAs.

The following rules determine whether the negotiation results in a capability being enabled:

- If a capability and its configuration values match between the switch and the CNA, the feature is enabled.
- If a capability matches, but the configuration values do not match, the following occurs:
    - If the CNA is configured to accept the switch configuration value, the capability is enabled using the switch value.
    - If the CNA is not configured to accept the switch configuration value, the capability remains disabled.
- If the CNA does not support a DCBX capability, that capability remains disabled.
- If the CNA does not implement DCBX, all capabilities remain disabled.

Note The Cisco Nexus device provides CLI commands to manually override the results of the PFC negotiation with the adapter. On a per-interface basis, you can force capabilities to be enabled or disabled.

Note The priority flow control (PFC) mode does not send PFC TLV and PFC will not negotiate between CNA and Cisco Nexus 5000 Series switches.

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

The following expansion modules provide native Fibre Channel ports to connect the Cisco Nexus 5000 Series switch to other Fibre Channel devices.

- N5K-M1404 Cisco Nexus 5000 1000 Series Module 4x10GE 4xFC 4/2/1
- N5K-M1008 Cisco Nexus 5000 1000 Series Module 8xFC 4/2/1
- N5K-M1060 Cisco Nexus 5000 1000 Series Module 6xFC 8/4/2/1

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.

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.

For pre-FIP CNAs, the switch sends a DCBX message to request the CNA to reset only the virtual Fibre Channel interface.

Note If the CNA does not support the logical link level up/down feature, the CNA resets the physical link. In this case, all traffic on the Ethernet interface is disrupted.

DCBX-based FC Logical Link Status signaling applies only to FCoE sessions to pre-FIP CNAs.

Converged Network Adapters

The following types of CNAs are available:

- 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.

Two generations of CNAs are supported by the Cisco Nexus device:

- A FIP adapter uses the FIP to exchange information about its available capabilities and to negotiate the configurable values with the switch.
- A pre-FIP adapter uses DCBX to exchange information about its available capabilities and to negotiate the configurable values with the switch.

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

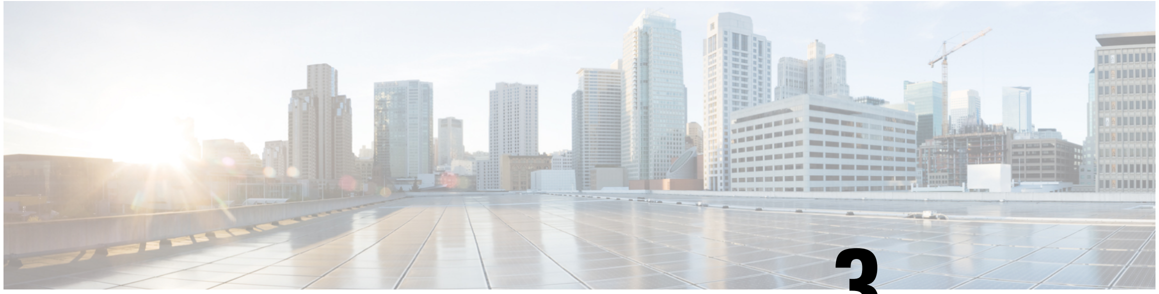

CHAPTER 3

Configuring FCoE

This chapter contains the following sections:

FCoE Topologies

Directly Connected CNA Topology

The Cisco Nexus device can be deployed as a Fibre Channel Forwarder (FCF) as shown in the following figure.

Figure 1: Directly Connected Fibre Channel Forwarder

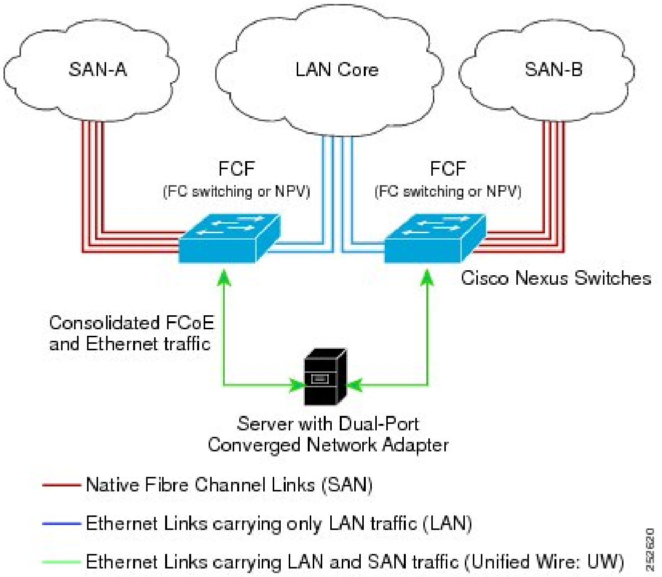
SAN-A
LAN Core
SAN-B
FCF
(FC switching or NPV)
FCF
(FC switching or NPV)
Cisco Nexus Switches
Consolidated FCoE
and Ethernet traffic
Server with Dual-Port
Converged Network Adapter
Native Fibre Channel Links (SAN)
Ethernet Links carrying only LAN traffic (LAN)
Ethernet Links carrying LAN and SAN traffic (Unified Wire: UW)
252620

The following rules are used to process FIP frames to avoid the FCF being used as a transit between an FCoE node (ENode) and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:
    - The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric).
    - The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric).

CNAs cannot discover or log in to FCFs that are reachable only through a transit Cisco Nexus FCF. The Cisco Nexus device cannot perform the FCoE transit function between a CNA and another FCF due to hardware limitations.

Because the Cisco Nexus FCF cannot perform the transit FCoE function, you must design your network topology so that the active Spanning Tree Protocol (STP) path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

Remotely Connected CNA Topology

The Cisco Nexus device can be deployed as a Fibre Channel Forwarder (FCF) for remotely connected CNAs, but not as a FIP snooping bridge, as shown in the following figure.

Figure 2: Remotely Connected Fibre Channel Forwarder

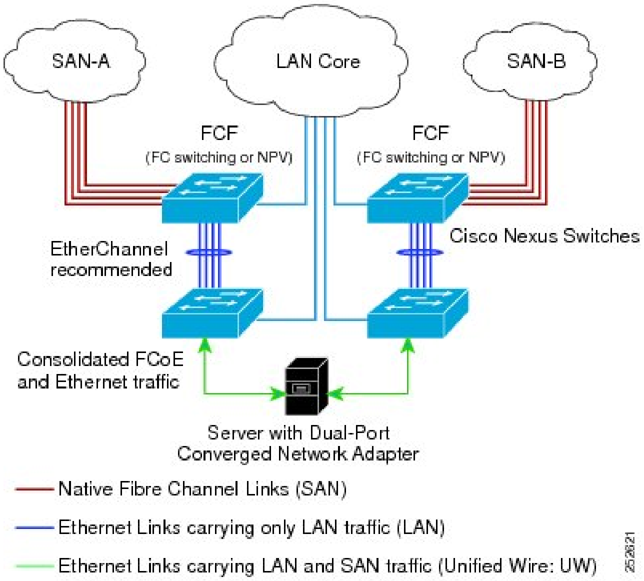
SAN-A
LAN Core
SAN-B
FCF
(FC switching or NPV)
FCF
(FC switching or NPV)
Cisco Nexus Switches
EtherChannel recommended
Consolidated FCoE and Ethernet traffic
Server with Dual-Port Converged Network Adapter
Native Fibre Channel Links (SAN)
Ethernet Links carrying only LAN traffic (LAN)
Ethernet Links carrying LAN and SAN traffic (Unified Wire: UW)
252621

The following rules are used to process FIP frames to avoid the FCF being used as a transit between an ENode and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:
  - The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric).
  - The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric).

Because the Cisco Nexus FCF cannot perform the transit FCoE function, you must design your network topology so that the active STP path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

FCoE Best Practices

Directly Connected CNA Best Practice

The following figure shows a best practices topology for an access network that is using directly connected CNAs with Cisco Nexus devices.

Figure 3: Directly Connected CNA

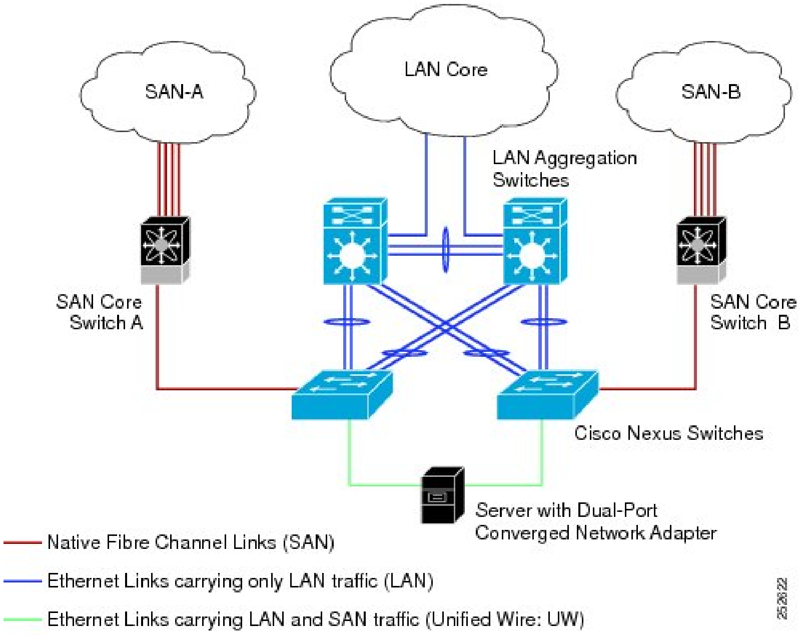
SAN-A
LAN Core
SAN-B
LAN Aggregation Switches
SAN Core Switch A
SAN Core Switch B
Cisco Nexus Switches
Server with Dual-Port Converged Network Adapter
Native Fibre Channel Links (SAN)
Ethernet Links carrying only LAN traffic (LAN)
Ethernet Links carrying LAN and SAN traffic (Unified Wire: UW)
252622

Follow these configuration best practices for the deployment topology in the preceding figure:

1. You must configure a 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 Multiple Spanning Tree (MST), you must use a separate MST instance for FCoE VLANs.

2. You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

Note A unified wire carries both Ethernet and FCoE traffic.

3. You must configure the UF links as spanning-tree edge ports.
4. You must not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure that the scope of the STP for the FCoE VLANs is limited to UF links only.
5. If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, you must explicitly configure such links to exclude all FCoE VLANs from membership. This action ensures that the scope of the STP for the FCoE VLANs is limited to UF links only.
6. You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

Note All Gen-1 (pre-FIP) and Gen-2, Gen-3, and Gen-4 (FIP) CNAs are supported in a directly connected topology.

Remotely Connected CNA Best Practice

The following figure shows a best practices topology for an access network using remotely connected CNAs with Cisco Nexus devices.

Figure 4: Remotely Connected CNAs

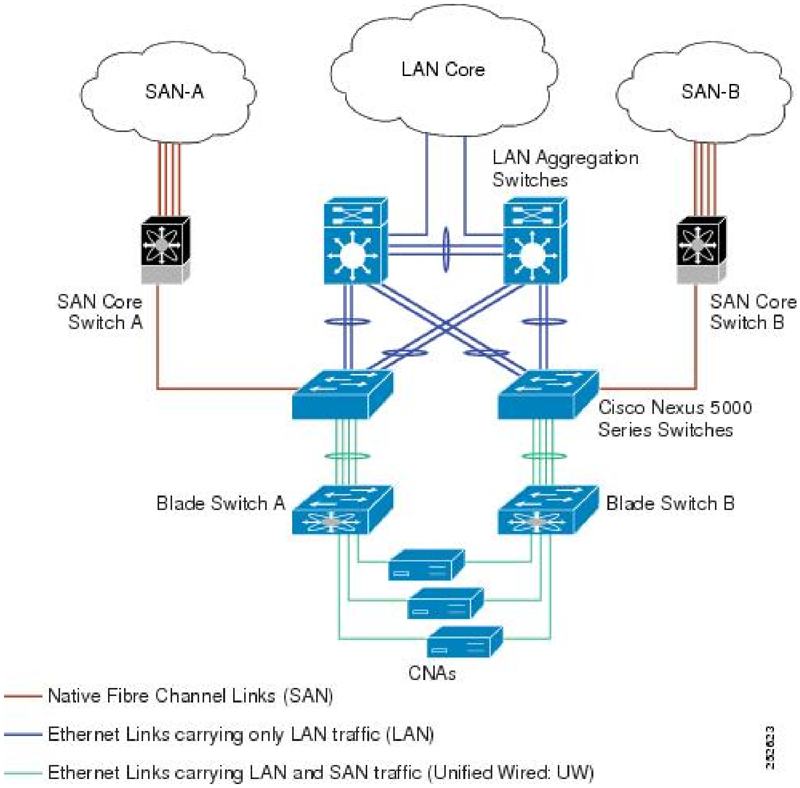
SAN-A
LAN Core
SAN-B
LAN Aggregation Switches
SAN Core Switch A
SAN Core Switch B
Cisco Nexus 5000 Series Switches
Blade Switch A
Blade Switch B
CNAs
Native Fibre Channel Links (SAN)
Ethernet Links carrying only LAN traffic (LAN)
Ethernet Links carrying LAN and SAN traffic (Unified Wired: UW)
252623

Follow these configuration best practices for the deployment topology in the preceding figure:

1. You must configure a 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.
2. You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

Note A unified fabric link carries both Ethernet and FCoE traffic.

3. You must configure the CNAs and the blade switches as spanning-tree edge ports.
4. A blade switch must connect to exactly one Cisco Nexus device converged access switch, preferably over an EtherChannel, to avoid disruption due to STP reconvergence on events such as provisioning new links or blade switches.
5. You must configure the Cisco Nexus device converged access switch with a better STP priority than the blade switches that are connected to it. This requirement allows you to create an island of FCoE VLANs where the converged access switch is the spanning-tree root and all the blade switches connected to it become downstream nodes.

6. Do not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure that the scope of the STP for the FCoE VLANs is limited to UF links only.
7. If the converged access switches and/or the blade switches need to be connected to each over Ethernet links for the purposes of LAN alternate pathing, you must explicitly configure such links to exclude all FCoE VLANs from membership. This action ensures the scope of the STP for FCoE VLANs is limited to UF links only.
8. You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

Note A remotely connected topology is supported only with Gen-2, Gen-3, and Gen-4 (FIP) CNAs.

Guidelines and Limitations

FCoE has the following guidelines and limitations:

- FCoE on Cisco Nexus devices support the Gen-1 (pre-FIP) and Gen-2, Gen-3, and Gen-4 (FIP) CNAs. FCoE on the Cisco Nexus 2232PP fabric extender (FEX) supports Gen-2 CNAs only.
- Enabling FCoE on VLAN 1 is not supported.
- A combination of straight-through and active-active topologies is not supported on the same FEX.
- Direct connect FCoE (that is, a direct connect to CNAs through a bind interface) is not supported on a port channel of a Cisco Nexus device or FEX interface if it is configured to have more than one interface. Direct connect FCoE is supported on port channels with a single link to allow for FCoE from a CNA connected through a vPC with one 10 GB link to each upstream switch/FEX.

Note For a description of the default quality of service (QoS) policies for FCoE, see the Quality of Service guide for your device. for the Nexus software release that you are using. The available versions of this document can be found at the following URL: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html..

Configuring FCoE

Configuring QoS

You need to attach the system service policy to configure QoS. The **service-policy** command specifies the system class policy map as the service policy for the system.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **system qos**	Enters system qos configuration mode.
Step 3	switch(config-sys-qos)# **service-policy type** {**network-qos** | **qos** | **queuing**} [**input** | **output**] *fcoe default policy-name*	Specifies the default FCoE policy map to use as the service policy for the system. There are four pre-defined policy-maps for FCoE: • service-policy type queuing input fcoe-default-in-policy • service-policy type queuing output fcoe-default-out-policy • service-policy type qos input fcoe-default-in-policy • service-policy type network-qos fcoe-default-nq-policy **Note** Before enabling FCoE on a Cisco Nexus device, you must attach the pre-defined FCoE policy maps to the type qos, type network-qos, and type queuing policy maps.

Enabling FCoE

You can enable FCoE on the switch; however, enabling FCoE on VLAN 1 is not supported.

Note All the Fibre Channel features of the Cisco Nexus device are packaged in the FC Plugin. When you enable FCoE, the switch software checks for the FC_FEATURES_PKG license. If it finds the license, the software loads the plugin. If the license is not found, the software loads the plugin with a grace period of 180 days.

After the FC Plugin is loaded, the following occurs:

- All Fibre Channel and FCoE-related CLI are available
- The Fibre Channel interfaces of any installed expansion modules are available

If after 180 days, a valid license is not found, the FC Plugin is disabled. At the next switch reboot, all FCoE commands are removed from the CLI and the FCoE configuration is deleted.

Before you begin

You must have the FC_FEATURES_PKG (N5010SS or N5020SS) license installed.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature fcoe**	Enables the FCoE capability.

Example

This example shows how to enable FCoE on the switch:

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

Disabling FCoE

After you disable FCoE, all FCoE commands are removed from the CLI and the FCoE configuration is deleted.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **no feature fcoe**	Disables the FCoE capability.

Example

This example shows how to disable FCoE on the switch:

```
switch# configure terminal

switch(config)# no feature fcoe
```

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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **interface ethernet** *slot*/*port*	Specifies an interface to configure, and enters interface configuration mode.
Step 3	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.
Step 4	(Optional) switch(config-if)# **no shutdown lan**	Reenables Ethernet traffic on the interface.

Configuring the FC-Map

Note We recommend using the "Mapping a VSAN to a VLAN " method for preserving fabric isolation and leaving the FC-MAP default.

You can prevent data corruption due to cross-fabric talk by configuring an FC-Map that identifies the Fibre Channel fabric for this Cisco Nexus device. When the FC-Map is configured, the switch discards the MAC addresses that are not part of the current fabric.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fcmap** *fabric-map*	Configures the global FC-Map. The default value is 0E.FC.00. The range is from 0E.FC.00 to 0E.FC.FF.
Step 3	(Optional) switch(config)# **no fcoe fcmap** *fabric-map*	Resets the global FC-Map to the default value of 0E.FC.00.

Example

This example shows how to configure the global FC-Map:

```
switch# configure terminal
```

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

Configuring the Fabric Priority

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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fcf-priority** *fabric-priority*	Configures the global fabric priority. The default value is 128. The range is from 0 (higher) to 255 (lower).
Step 3	(Optional) switch(config)# **no fcoe fcf-priority** *fabric-priority*	Resets the global fabric priority to the default value of 128.

Example

This example shows how to configure the global fabric priority:

```
switch# configure terminal
```

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

Configuring Jumbo MTU

This example shows how to configure the default Ethernet system class to support the jumbo MTU:

```
switch(config)# policy-map type network-qos jumbo
switch(config-pmap-nq)# class type network-qos class-fcoe
switch(config-pmap-c-nq)# pause no-drop
switch(config-pmap-c-nq)# mtu 2158
switch(config-pmap-nq)# class type network-qos class-default
switch(config-pmap-c-nq)# mtu 9216
switch(config-pmap-c-nq)# exit
switch(config-pmap-nq)# exit
switch(config)# system qos
switch(config-sys-qos)# service-policy type qos input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing input fcoe-default-in-policy
switch(config-sys-qos)# service-policy type queuing output fcoe-default-out-policy
switch(config-sys-qos)# service-policy type network-qos jumbo
```

Setting the Advertisment Interval

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

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fka-adv-period** *inverval*	Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds.

	Command or Action	Purpose
Step 3	(Optional) switch(config)# **no fcoe fka-adv-period** *interval*	Resets the advertisement interval for the fabric to its default value of 8 seconds.

Example

This example shows how to configure the advertisement interval for the fabric:

```
switch# configure terminal

switch(config)# fcoe fka-adv-period 42
```

Verifying the FCoE Configuration

To verify FCoE configuration information, perform one of these tasks:

Command	Purpose
switch# **show fcoe**	Displays whether FCoE is enabled on the switch.
switch# **show fcoe database**	Displays the contents of the FCoE database.
switch# **show interface** [*interface number*] **fcoe**	Displays the FCoE settings for an interface or all interfaces.
switch# **show queuing interface**[*interface slot/port*]	Displays the queue configuration and statistics.
switch# **show policy-map interface**[*interface number*]	Displays the policy map 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
```

CHAPTER 4

Configuring FCoE VLANs and Virtual Interfaces

This chapter contains the following sections:

Information About Virtual Interfaces

Cisco Nexus devices support Fibre Channel over Ethernet (FCoE), which 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 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 Cisco Nexus device), a MAC address (when the CNA is remotely connected over a Layer 2 bridge), or an EtherChannel when the CNA connects to the Fibre Channel Forwarder (FCF) over a virtual port channel (vPC).

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.

On the Cisco Nexus 5000 and 6000 Series switches, traffic across members of a port channel that a VE_Port is bound to is load balanced based on SID, DID, and OXID.

In order to enable all links to be used in the port-channel for FCoE traffic, enter the **port-channel load-balance ethernet** *source-dest-port* command to configure 'port-channel load balancing' to 'source-dest-port'. The configuration 'source-destination-oxid' load balancing is used for FCoE traffic.

Guidelines and Limitations for FCoE VLANs and Virtual Interfaces

FCoE VLANs and Virtual Fiber Channel (vFC) interfaces have these guidelines and limitations:

- Each vFC interface must be bound to an FCoE-enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter. FCoE is supported on 10-Gigabit Ethernet interfaces.

  The Ethernet or EtherChannel interface that you bind to the vFC interface must be configured as follows:

  - The Ethernet or EtherChannel interface must be a trunk port (use the **switchport mode trunk** command).
  - The FCoE VLAN that corresponds to a vFC's VSAN must be in the allowed VLAN list.
  - You must not configure an FCoE VLAN as the native VLAN of the trunk port.

    Note The native VLAN is the default VLAN on a trunk. Any untagged frames transit the trunk as native VLAN traffic.

  - You should use an FCoE VLAN only for FCoE.
  - Do not use the default VLAN, VLAN1, as an FCoE VLAN.
  - You must configure the Ethernet interface as PortFast (use the **spanning-tree port type edge trunk** command).

    Note You are not required to configure trunking on the server interface even if the switch interface is configured with trunking enabled. All non-FCoE traffic from the server is passed on the native VLAN.

- The vFC interface can be bound to Ethernet port channels with multiple member ports connected to FCoE Initialization Protocol (FIP) snooping bridges.
- Each vFC interface is associated with only one VSAN.
- You must map any VSAN with associated vFC interfaces to a dedicated FCoE-enabled VLAN.
- FCoE is not supported on private VLANs.
- If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, then you must explicitly configure such links to exclude all FCoE VLANs from membership.
- You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B fabrics.
- FCoE connectivity to pre-FIP CNAs over virtual port channels (vPCs) is not supported.
- The maximum number of vFCs that can be bound to a port-channel is 48.

Note Virtual interfaces are created with the administrative state set to down. You must explicitly configure the administrative state to bring the virtual interface into operation.

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 VSAN in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If Multiple Spanning Tree (MST) is enabled, a separate MST instance must be used for FCoE VLANs.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan** *vlan-id*	Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 3	switch(config-vlan)# **fcoe** [**vsan** *vsan-id*]	Enables FCoE for the specified VLAN. If you do not specify a VSAN number, a mapping is created from this VLAN to the VSAN with the same number. Configures the mapping from this VLAN to the specified VSAN.
Step 4	switch(config-vlan)# **exit**	Exits VLAN configuration mode. You must exit this mode to execute the configured commands on your Cisco Nexus device.
Step 5	(Optional) switch(config)# **show vlan fcoe**	Displays information about the FCoE configuration for a VLAN.
Step 6	(Optional) switch(config-if)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to map VLAN 200 to VSAN 2:

```
switch(config)# vlan 200
```

```
switch(config-vlan)# fcoe vsan 2
```

Creating a Virtual Fibre Channel Interface

You can create a virtual Fibre Channel interface. You must bind the virtual Fibre Channel interface to a physical interface before it can be used.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode. The virtual Fibre Channel interface ID range is from 1 to 8192.
Step 3	switch(config-if)# **bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}	Binds the virtual Fibre Channel interface to the specified interface.
Step 4	(Optional) switch(config-if)# **no bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}	Unbinds the virtual Fibre Channel interface from the specified interface.
Step 5	(Optional) switch(config)# **no interface vfc** *vfc-id*	Deletes a virtual Fibre Channel interface.

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 bind a virtual Fibre Channel interface to a Cisco Nexus 2232PP Fabric Extender (FEX) Ethernet interface:

```
switch# configure terminal
switch(config)# interface vfc 1001
switch(config-if)# bind interface ethernet 100/1/1
```

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

```
switch# configure terminal
switch(config)# interface vfc 3
switch(config-if)# bind interface port-channel 1
```

This example shows how to bind a virtual Fibre Channel interface on a Cisco Nexus 2232PP FEX to create a vPC:

```
switch# configure terminal
switch(config)# interface vfc 1001
switch(config-if)# bind interface ethernet 100/1/1
```

Note An error message is displayed if you attempt to bind the interface to a Cisco Nexus FEX that does not support FCoE.

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

```
switch# configure terminal
switch(config)# interface vfc 2
switch(config-if)# bind mac-address 00:0a:00:00:00:36
```

This example shows how to bind a virtual Fibre Channel interface to a Cisco Nexus 2232PP FEX MAC address:

```
switch# configure terminal
switch(config)# interface vfc 1001
switch(config-if)# bind mac-address 00:01:0b:00:00:02
```

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

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

Associating a Virtual Fibre Channel Interface to a VSAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If MST is enabled, a separate MST instance must be used for FCoE VLANs.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vsan database**	Enters VSAN configuration mode.
Step 3	switch(config-vsan)# **vsan** *vsan-id* **interface vfc** *vfc-id*	Configures the association between the VSAN and virtual Fibre Channel interface. The VSAN number must map to a VLAN on the physical Ethernet interface that is bound to the virtual Fibre Channel interface.
Step 4	(Optional) switch(config-vsan)# **no vsan** *vsan-id* **interface vfc** *vfc-id*	Disassociates the connection between the VSAN and virtual Fibre Channel interface.

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)# vsan 2 interface vfc 4
```

Creating an Implicit Virtual Fibre Channel Port Channel Interface

You can create a virtual Fibre Channel (vFC), and implicitly bind it to an Ethernet interface or a port-channel using a single command. For this, the vFC identifier must match the Ethernet interface or port-channel identifier. The Ethernet interface can be a module (slot or port) or a Fabric Extender (FEX) interface (chassis, slot or port).

Note You cannot create an implicit vFC in a breakout port.

Configuring virtual Fibre Channel Interface

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you have enabled FCoE.

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Create a VFC (if it does not already exist):

Additionally, *vfc slot/port* binds the vFC to an Ethernet *slot/port* interface. *vfc chassis/slot/port* binds the vFC to a FEX interface.

switch(config) # **interface vfc** {id | *slot/port* | *chassis/slot/port*

Step 3 Bring up the vFC interface:

switch(config-if) # **no shutdown**

Step 4 Required: Exit the interface configuration mode:

switch(config-if) # **exit**

Configuring virtual Fibre Channel Interface

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

```
switch# configure terminal
switch(config)# interface vfc 101/1/1
switch(config-if)# no shutdown
```

```
.
.
.
.
Switch# show interface vfc 101/1/1
    vfc101/1/1 is trunking
    Bound interface is Ethernet101/1/1
    Hardware is Ethernet
    Port WWN is 20:00:00:2a:6a:15:d2:7b
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 600
    Trunk vsans (admin allowed and active) (1,500,600)
    Trunk vsans (up)                       (600)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1,500)
    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
      49 frames input, 5516 bytes
        0 discards, 0 errors
      49 frames output, 5772 bytes
        0 discards, 0 errors
    last clearing of "show interface" counters Fri Oct 30 06:19:33 2015
    Interface last changed at Fri Oct 30 06:19:33 2015
```

Configuring virtual Fibre Channel – Port Channel Interface

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Create a vFC that implicitly binds to the Ethernet port-channel based on its number:

The port number range is from 257 to 4096.

switch(config) # **interface vfc-port-channel** *port number*

Step 3 Bring up the vFC port:

switch(config-if) # **no shutdown**

Step 4 Required: Exit from the current interface configuration mode:

switch(config-if) # **exit**

Configuring virtual Fibre Channel - Port Channel Interface

The example shows how you can create a vFC-port-channel that implicitly binds to Ethernet port-channel:

```
switch# configure terminal
switch(config)# interface vfc-port-channel 300
switch(config-if)# no shutdown
.
.
.
.
switch# show interface vfc-port-channel 258

    vfc-po258 is trunking
    Bound interface is port-channel258
    Hardware is Ethernet
    Port WWN is 21:01:8c:60:4f:59:31:3f
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 600
    Trunk vsans (admin allowed and active) (1,100,500,600)
    Trunk vsans (up)                       (600)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1,100,500)
    1 minute input rate 3448 bits/sec, 431 bytes/sec, 4 frames/sec
    1 minute output rate 9064 bits/sec, 1133 bytes/sec, 4 frames/sec
      977735 frames input, 77172556 bytes
        0 discards, 0 errors
      977733 frames output, 205924892 bytes
        0 discards, 0 errors
    last clearing of "show interface" counters Thu Oct 29 06:35:41 2015
    Interface last changed at Thu Oct 29 06:35:41 2015
```

Verifying the Virtual Interface

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

Command	Purpose
switch# **show interface vfc** *vfc-id*	Displays the detailed configuration of the specified Fibre Channel interface.
switch# **show interface brief**	Displays the status of all interfaces.
switch# **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 a virtual Fibre Channel interface bound to a MAC address:

```
switch# show interface vfc 1001
vfc1001 is down
    Bound MAC is 00:0a:00:00:00:01
    Hardware is Virtual Fibre Channel
    Port WWN is 23:e8:00:0d:ec:6d:95:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 901
    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
```

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
```

Mapping VSANs to VLANs Example Configuration

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.

```
switch(config)# vlan 200
switch(config-vlan)# fcoe vsan 2
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.

```
switch(config)# interface vfc 4

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

switch(config-if)# exit
```

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.

```
switch(config)# vsan database

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

switch(config-vsan)# exit
```

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

```
switch# show vsan 2 membership
vsan 2 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 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 Mar 11 04:44:42 2010
```

FCoE over Enhanced vPC

Although Ethernet traffic is dual homed between a FEX and a switch pair in an enhanced vPC topology, FCoE traffic must be single-homed to maintain SAN isolation. Therefore, while enhanced vPC supports FCoE, a single homed FEX topology can be a better choice when SAN isolation and high FCoE bandwidth are required.

Consider the following disadvantages of enhanced vPC for a single-homed topology:

- A typical SAN network maintains two fabrics, SAN A and SAN B, with traffic isolated between the two. In an enhanced vPC topology, each switch must be paired (single homed) with a FEX to ensure that FCoE traffic from one FEX is sent to only one switch, while Ethernet traffic is dual homed between each FEX and both switches. Because FCoE traffic from the FEX flows to only one switch while Ethernet traffic flows to both, the traffic load for the FEX uplinks is not evenly balanced.
- In a FEX with eight uplink ports, Ethernet traffic can use all eight ports, while the single-homed FCoE traffic is limited by this topology to using only four of those ports, restricting the maximum bandwidth available for FCoE. As a further restriction, the default QoS template for the shared link allocates only half the link bandwidth to FCoE traffic, with the other half allocated to Ethernet traffic.
- In an enhanced vPC topology with FCoE, the host vPC is limited to two ports, one to each FEX.

The following figure shows the FCoE traffic flow in a system with two Cisco Nexus 2000 FEXs, each associated with a different Cisco Nexus device.

Figure 5: FCoE over Enhanced vPC

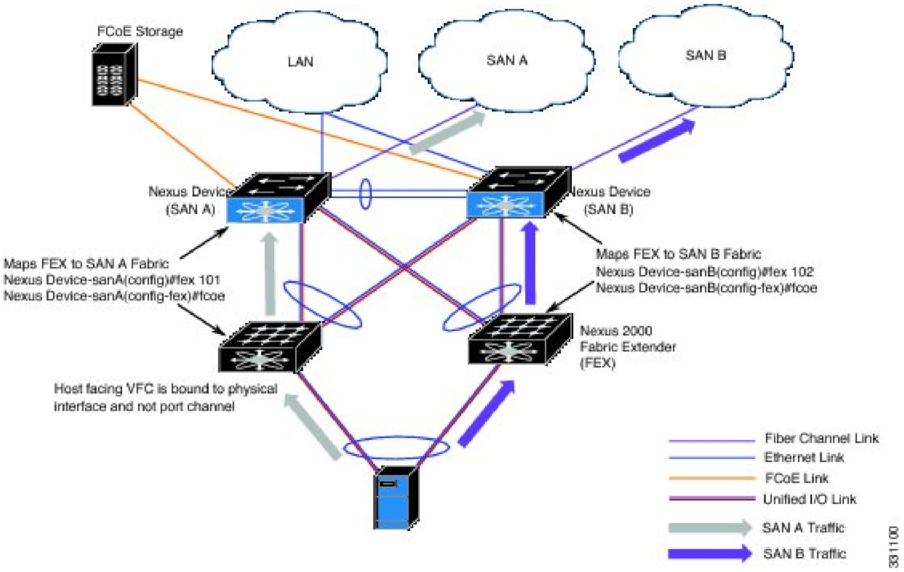
FCoE Storage
LAN
SAN A
SAN B
Nexus Device (SAN A)
Nexus Device (SAN B)
Maps FEX to SAN A Fabric
Nexus Device-sanA(config)#fex 101
Nexus Device-sanA(config-fex)#fcoe
Maps FEX to SAN B Fabric
Nexus Device-sanB(config)#fex 102
Nexus Device-sanB(config-fex)#fcoe
Nexus 2000 Fabric Extender (FEX)
Host facing VFC is bound to physical interface and not port channel
Fiber Channel Link
Ethernet Link
FCoE Link
Unified I/O Link
SAN A Traffic
SAN B Traffic
331100

Configuring FCoE over Enhanced vPC

FCoE traffic must be single homed to maintain SAN isolation. You must first associate a FEX with only one switch. When the FEX and switch are associated, you can then create a virtual Fibre Channel (vFC) interface and bind it to a port.

After pairing the FEX and switch on the first peer, you repeat the configuration on the second peer using a different port number to ensure SAN traffic isolation. The different configuration will not cause a consistency error because the FCoE portion of the enhanced vPC configuration is not subject to the vPC consistency check.

Before you begin

Review the limitations in FCoE over Enhanced vPC, on page 32.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config) # **fex** *fex-chassis_ID*	Enters configuration mode for the specified FEX. The range for *fex-chassis_ID* is from 100 to 199.
Step 3	switch(config-fex) # **fcoe**	Configures the FEX to send FCoE traffic only to this switch.
Step 4	switch(config-fex) # **interface vfc** *vfc-id*	Enters configuration mode for the virtual Fibre Channel interface. If the interface does not already exist, this command also creates that interface. The range of *vfc-id* is from 1 to 8192.
Step 5	switch(config-if) # **bind interface ethernet** [*fex-chassis-ID/*]*slot/port*	Binds the vFC interface to the specified physical Ethernet interface. The range for *fex-chassis_ID* is from 100 to 199. The *slot* must be 1.For FCoE, the range for *port* is from 1 to 32.
Step 6	switch(config-if) # **no shutdown**	Returns the interface to its default operational state.
Step 7	(Optional) switch(config-if) # **end**	Return to privileged EXEC mode.
Step 8	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to pair each FEX to a switch for FCoE traffic:

```
nexus5000-sanA# configure terminal
nexus5000-sanA(config) # fex 101
nexus5000-sanA(config-fex) # fcoe
nexus5000-sanA(config-fex) # interface vfc 1
nexus5000-sanA(config-if) # bind interface ethernet 101/1/1
```

```
nexus5000-sanA(config-if) # no shutdown
nexus5000-sanA(config-if) # end
nexus5000-sanA# copy running-config startup-config
nexus5000-sanA#

nexus5000-sanB# configure terminal
nexus5000-sanB(config) # fex 102
nexus5000-sanB(config-fex) # fcoe
nexus5000-sanB(config-fex) # interface vfc 1
nexus5000-sanB(config-if) # bind interface ethernet 102/1/1
nexus5000-sanB(config-if) # no shutdown
nexus5000-sanB(config-if) # end
nexus5000-sanB# copy running-config startup-config
nexus5000-sanB#


nexus5500-sanA# configure terminal
nexus5500-sanA(config) # fex 101
nexus5500-sanA(config-fex) # fcoe
nexus5500-sanA(config-fex) # interface vfc 1
nexus5500-sanA(config-if) # bind interface ethernet 101/1/1
nexus5500-sanA(config-if) # no shutdown
nexus5500-sanA(config-if) # end
nexus5500-sanA# copy running-config startup-config
nexus5500-sanA#

nexus5500-sanB# configure terminal
nexus5500-sanB(config) # fex 102
nexus5500-sanB(config-fex) # fcoe
nexus5500-sanB(config-fex) # interface vfc 1
nexus5500-sanB(config-if) # bind interface ethernet 102/1/1
nexus5500-sanB(config-if) # no shutdown
nexus5500-sanB(config-if) # end
nexus5500-sanB# copy running-config startup-config
nexus5500-sanB#
```

SAN Boot with vPC

A Cisco Nexus Series switch can use SAN boot if the following conditions are met:

- The FEX that contains the port assigned to the vPC must be associated with the Cisco Nexus switch.
- Only one VFC interface is bound to a vPC member. You cannot bind multiple interfaces to multiple members.

Note If you want to ensure backward compatibility for all previous configurations and supported topologies, you must configure the FEX in a straight-through FEX topology that does not use Enhanced vPC.

SAN Boot with vPC Configuration Example

In this example, virtual Fibre Channel interface 1 is bound to physical Ethernet interface 101/1/1 on Fabric A and on interface 102/1/1 on Fabric B. The interface is also associated with virtual port channel 1 on both fabrics.

```
nexus5000-sanA(config) # interface vfc 1
nexus5000-sanA(config-if) # bind interface eth 101/1/1
nexus5000-sanA(config) # interface eth 101/1/1
nexus5000-sanA(config-if) # channel-group 1 mode active
nexus5000-sanA(config-if) # interface port-channel 1
nexus5000-sanA(config-if) # vpc 1
nexus5000-sanA(config-if) #

nexus5000-sanB(config) # interface vfc 1
nexus5000-sanB(config-if) # bind interface eth 102/1/1
nexus5000-sanB(config) # interface eth 102/1/1
nexus5000-sanB(config-if) # channel-group 1 mode active
nexus5000-sanB(config-if) # interface port-channel 1
nexus5000-sanB(config-if) # vpc 1
nexus5000-sanB(config-if) #
```

```
nexus5500-sanA(config) # interface vfc 1
nexus5500-sanA(config-if) # bind interface eth 101/1/1
nexus5500-sanA(config) # interface eth 101/1/1
nexus5500-sanA(config-if) # channel-group 1 mode active
nexus5500-sanA(config-if) # interface port-channel 1
nexus5500-sanA(config-if) # vpc 1
nexus5500-sanA(config-if) #

nexus5500-sanB(config) # interface vfc 1
nexus5500-sanB(config-if) # bind interface eth 102/1/1
nexus5500-sanB(config) # interface eth 102/1/1
nexus5500-sanB(config-if) # channel-group 1 mode active
nexus5500-sanB(config-if) # interface port-channel 1
nexus5500-sanB(config-if) # vpc 1
nexus5500-sanB(config-if) #
```

CHAPTER 5

Configuring Cisco Adapter FEX with FCoE

This chapter contains the following sections:

Overview

The Cisco Adapter Fabric Extender (FEX) feature allows you to create an FCoE connection to a FEX so that you can establish an FCoE connection to a server with a virtual interface card (VIC) adapter.

For example, you could use this feature to connect your Nexus switch to a Cisco UCS C-Series Rack-Mount Server that contains a Cisco UCS P81E Virtual Interface Card, or you could connect it to a third-party server that has a Broadcom BCM57712 Convergence Network Interface Card (C-NIC) installed.

The switch connects to the FEX through a virtual port channel (vPC) while the FEX connects to the server using a standard FCoE link between the FEX and the VIC adapter.

Guidelines and Limitations

If you are using Enhanced vPC, the FEX can be associated with one and only one Cisco Nexus fabric for FCoE forwarding.

If you are using FabricPath, you must use a dedicated link for FCoE traffic.

If you are using a Cisco UCS C-Series Rack-Mount Server with a Cisco UCS P81E Virtual Interface Card (VIC), you must do the following:

- Configure the VIC in Network Interface Virtualization (NIV) mode, which makes the two unified ports appear to the system as virtual Host Bus Adapters (vHBAs).
- You cannot connect to the FEX through a VNP port. If this type of connection is used, NIV mode cannot be enabled on the VIC.
- You must set the NIC mode on the Cisco UCS C-Series Rack-Mount Server to **active-standby**.

Note The HP chassis has internal HP-branded VirtualConnect cards (the specific model is HP VC FlexFabric 10Gb/24-Port Module), which runs its own internal host-hiding NPV processes for both FC and Ethernet sides of things. Then the external links of that card go upstream to regular Cisco 2232PP FEX's, which are then attached to our Nexus core.

Configuring Cisco Adapter FEX with FCoE

Procedure

	Command or Action	Purpose

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

Step 2	**install feature-set virtualization**
	Example:
	`switch(config) # install feature-set virtualization`
	`switch(config) #`
		Installs the virtualization feature set.

Step 3	**feature-set virtualization**
	Example:
	`switch(config) #` **`feature-set virtualization`**
	`switch(config)#`
		Enables the virtualization feature.

Step 4	**fex** *fex-chassis-ID*
	Example:
	`switch(config) #` **`fex 101`**
	`switch(config-fex) #`
		Enters configuration mode for the specified FEX.
		The range for *fex-chassis_ID* is from 100 to 199.

Step 5	**fcoe**
	Example:
	`switch(config-fex) #` **`fcoe`**
	`switch(config-fex) #`
		Enables Fibre Channel over Ethernet traffic on the FEX.

Step 6	**interface ethernet** [*fex-chassis-ID*/*slot*/*port*]
	Example:
	`switch(config-fex)#` **`interface ethernet 101/1/1`**
	`switch(config-if)#`
		Enters configuration mode for the specified Ethernet interface.
		The range for *fex-chassis-ID* is from 100 to 199. The *slot* For FCoE, the range for *port* is from 1 to 32.

	Command or Action	Purpose

Step 7	**switchport mode vntag**
	Example:
	`switch(config-if) # switchport mode vntag`
	`switch(config-if) #`
		Configures the interface in port mode.

Step 8	**interface vethernet** *veth-id*
	Example:
	`switch(config-if) # interface vethernet 2`
	`switch(config-if) #`
		Creates a virtual Ethernet interface and enters configuration mode for that interface.
		The range of *veth-id* is from 1 to 1,048,575.
		Note If you have two Cisco Nexus Series switches configured for redundancy, the virtual Ethernet interface ID must be unique on each switch.

Step 9	**bind interface ethernet** [*fex-chassis-ID/*]*slot*/*port* **channel** *channel-no*
	Example:
	`switch(config-if) # bind interface ethernet 101/1/1 channel 1`
	`switch(config-if) #`
		Binds the specified Ethernet interface to the specified port channel.
		The range for *fex-chassis-ID* is from 100 to 199. The *slot* must be 1. For FCoE, the range for *port* is from 1 to 32. The range for *channel-no* is from 1 to 4096.

Step 10	**switchport mode {trunk|access}**
	Example:
	`switch(config-if) # switchport mode trunk`
	`switch(config-if) #`
		Configures the interface as a trunk port or an access port.

Step 11	(Optional) **switchport trunk allowed vlan** *vlan-ID*
	Example:
	`switch(config-if) # switchport trunk allowed vlan 33`
	`switch(config-if) #`
		If you configured the interface as a trunk port, use this command to specify the VLAN for FCoE traffic.
		The range for *vlan-ID* is from 1 to 4094, except for the VLANs reserved for internal use.

Step 12	(Optional) **switchport access vlan** *vlan-ID*
	Example:
	`switch(config-if) # switchport access vlan 33`
	`switch(config-if) #`
		If you configured the interface as an access port, use this command to specify the VLAN for FCoE traffic.

Step 13	**interface vfc** *vfc-id*
	Example:
	`switch(config-if) # interface vfc 4`
	`switch(config-if) #`
		Creates a virtual Fibre Channel interface on the switch and enters configuration mode.
		The range of *vfc-id* is from 1 to 8192.

	Command or Action	Purpose

Step 14	**bind interface vethernet** *veth-num*
	Example:
	`switch(config-if) # bind interface veth 2`
	`switch(config-if) #`
		Binds the virtual Fibre Channel interface to the specified virtual Ethernet interface.
		The range of *veth-num* is from 1 to 1048575.

Step 15	**no shutdown**
	Example:
	`switch(config-if) # no shutdown`
	`switch(config-if) #`
		Returns the interface to its default operational state.

Example

This example show how to configure Cisco Adapter FEX with FCoE on SAN fabric A using FEX 101 and the Ethernet interface on channel 1 configured as a trunk port.

```
nexus5000-sanA(config)#configure terminal
nexus5000-sanA(config)# install feature-set virtualization
nexus5000-sanA(config)# feature-set virtualization
nexus5000-sanA(config)# fex 101
nexus5000-sanA(config-fex)# fcoe
nexus5000-sanA(config-fex)# interface ethernet 101/1/1
nexus5000-sanA(config-if)# switchport mode vntag
nexus5000-sanA(config-if)# interface veth 2
nexus5000-sanA(config-if)# bind interface eth 101/1/1 channel 1
nexus5000-sanA(config-if)# switchport mode trunk
nexus5000-sanA(config-if)# switchport trunk allowed vlan 33
nexus5000-sanA(config-if)# interface vfc 4
nexus5000-sanA(config-if)# bind interface veth 2
nexus5000-sanA(config-if)# no shutdown
```

```
nexus5500-sanA(config)#configure terminal
nexus5500-sanA(config)# install feature-set virtualization
nexus5500-sanA(config)# feature-set virtualization
nexus5500-sanA(config)# fex 101
nexus5500-sanA(config-fex)# fcoe
nexus5500-sanA(config-fex)# interface ethernet 101/1/1
nexus5500-sanA(config-if)# switchport mode vntag
nexus5500-sanA(config-if)# interface veth 2
nexus5500-sanA(config-if)# bind interface eth 101/1/1 channel 1
nexus5500-sanA(config-if)# switchport mode trunk
nexus5500-sanA(config-if)# switchport trunk allowed vlan 33
nexus5500-sanA(config-if)# interface vfc 4
nexus5500-sanA(config-if)# bind interface veth 2
nexus5500-sanA(config-if)# no shutdown
```

This example shows how to configure Cisco Adapter FEX with FCoE on SAN fabric B using FEX 102 and Ethernet interface on channel 2 as an access port:

```
nexus5000-sanB(config)#configure terminal
nexus5000-sanB(config)# install feature-set virtualization
nexus5000-sanB(config)# feature-set virtualization
nexus5000-sanB(config)# fex 102
nexus5000-sanB(config-fex)# fcoe
```

```
nexus5000-sanB(config-fex)# interface ethernet 102/1/1
nexus5000-sanB(config-if)# switchport mode vntag
nexus5000-sanB(config-if)# interface veth 5
nexus5000-sanB(config-if)# bind interface eth 102/1/1 channel 2
nexus5000-sanB(config-if)# switchport mode access
nexus5000-sanB(config-if)# switchport access vlan 40
nexus5000-sanB(config-if)# interface vfc 6
nexus5000-sanB(config-if)# bind interface veth 5
nexus5000-sanB(config-if)# no shutdown
```

```
nexus5500-sanB(config)#configure terminal
nexus5500-sanB(config)# install feature-set virtualization
nexus5500-sanB(config)# feature-set virtualization
nexus5500-sanB(config)# fex 102
nexus5500-sanB(config-fex)# fcoe
nexus5500-sanB(config-fex)# interface ethernet 102/1/1
nexus5500-sanB(config-if)# switchport mode vntag
nexus5500-sanB(config-if)# interface veth 5
nexus5500-sanB(config-if)# bind interface eth 102/1/1 channel 2
nexus5500-sanB(config-if)# switchport mode access
nexus5500-sanB(config-if)# switchport access vlan 40
nexus5500-sanB(config-if)# interface vfc 6
nexus5500-sanB(config-if)# bind interface veth 5
nexus5500-sanB(config-if)# no shutdown
```

CHAPTER 6

Configuring Dynamic FCoE Using FabricPath

This chapter contains the following sections:

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 6: FabricPath Topology

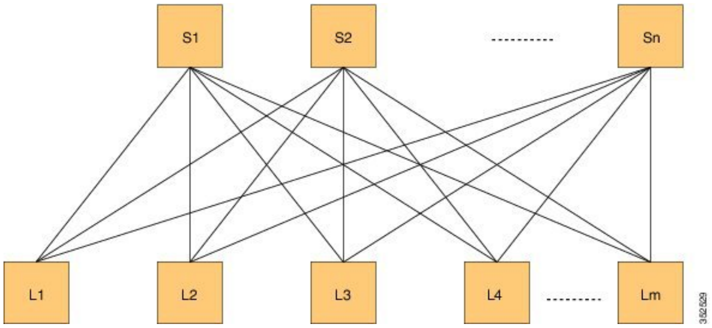

FCoE creates an overlay of FCoE virtual links on top of the underlying Ethernet topology, irrespective of how that Ethernet topology is constructed and which protocol is used to compute the MAC address routes.

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 7: FCoE Overlay of VE_Port to VE_Port Virtual Links

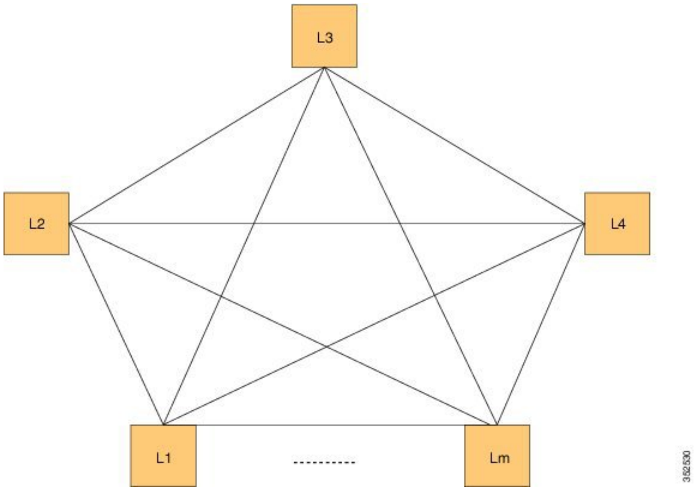

Only the FCFs, that are implemented by the leaf switches are part of this overlay topology. This topology is seen by Fabric Shortest Path First (FSPF), for each FCoE VLAN. FSPF computes over which virtual link to forward an FCoE frame based on its DomainID (D_ID). A virtual link is uniquely identified 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 8: Physical Topology Diagram

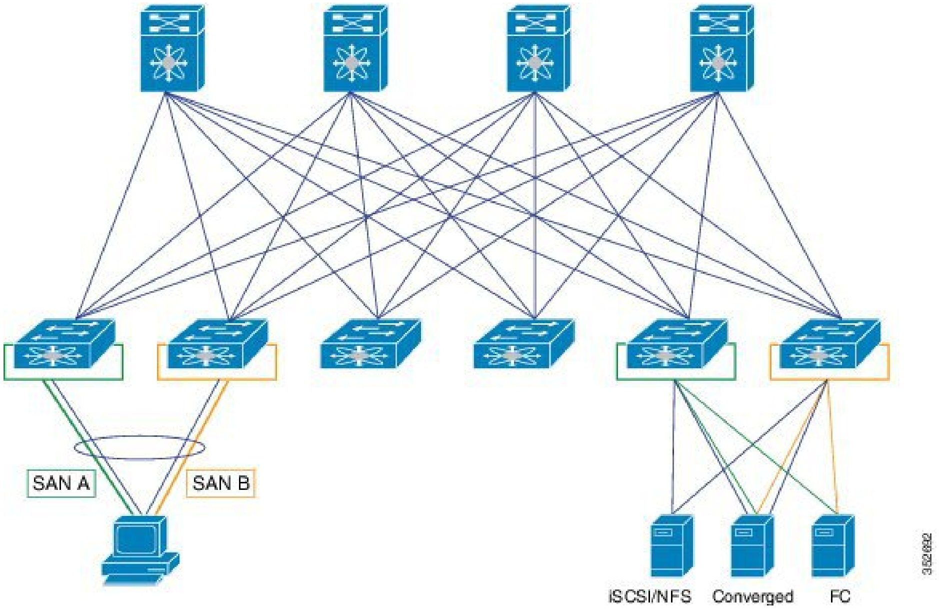

In the previous figure, the physical connectivity for the topology follows typical leaf/spine CLOS architectural best practices. Logically, SAN A and SAN B are isolated at the Top of Rack (ToR) switches physically. Once the traffic enters the FabricPath network, the storage traffic is logically separated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

Figure 9: Logical Topology Diagram

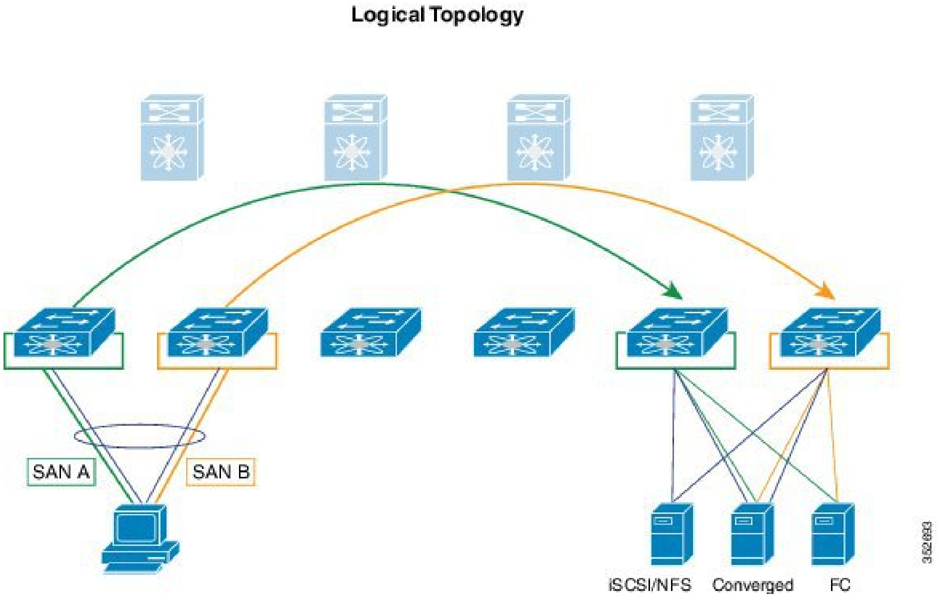

Dynamic FCoE gains the additional redundancy that is inherent in the FabricPath network by using the increased spine connectivity. A larger network with a large number of spines means increased reliability and stability for the storage network. This is achieved while retaining the best practices requirements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

Licensing Requirements for Dynamic FCoE Using FabricPath

The following table shows the licensing requirements for this feature:

Product	License Requirement
Fibre Channel over Ethernet	Feature FCoE license and feature FabricPath license for the leaf role.
FabricPath	Feature FabricPath for leaf.

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 leaf nodes.

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 10: Configuration Example

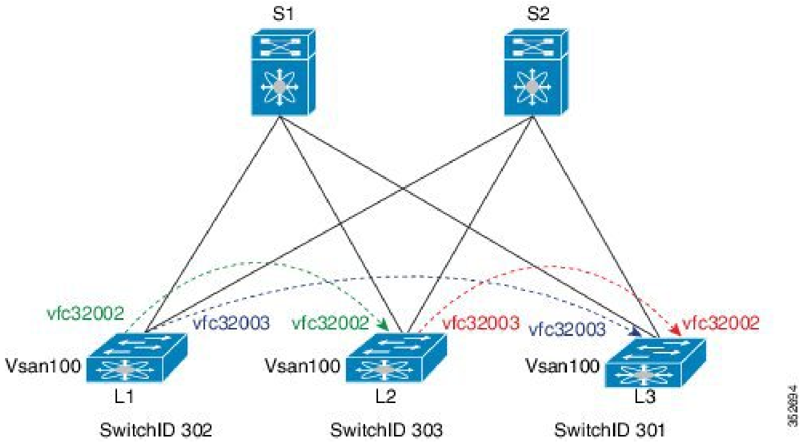

Note The component labels in the previous diagram are for illustrative purposes only.

Configuring Dynamic FCoE Using FabricPath

Procedure

Step 1 Establish the FabricPath infrastructure.

All spines and leafs must have FabricPath infrastructure configured. See Configuring All Leafs in the FabricPath Topology, on page 50.

Step 2 Configure spines for FCoE traffic.

See Configuring All Leafs in the FabricPath Topology, on page 50

Step 3 Configure non-FCoE leafs for FCoE traffic.

A leaf needs this configuration for failover cases. See Configuring All Leafs in the FabricPath Topology, on page 50.

Step 4 Configure leafs for FCoE (FCF) processing.

A leaf needs this configuration for failover cases.

1. See Configuring FCF Leafs, on page 51.
2. See Configuring FCoE and FabricPath-Enabled VLANs, on page 52.

Note If vPC or vPC+ was enabled on the leaf, follow the steps at Increasing the FabricPath Cost for a vPC+ Peer Link for FCF Leafs, on page 53.

Step 5 Configure ports on leafs for FC/FCoE.

If vPC or vPC+ is enabled, follow the steps at Increasing the FabricPath Cost for a vPC+ Peer Link for FCF Leafs, on page 53.

Related Topics

Enabling FabricPath, on page 50
Configuring All Leafs in the FabricPath Topology, on page 50
Configuring FCF Leafs, on page 51
Configuring FCoE and FabricPath-Enabled VLANs, on page 52
Defining FabricPath VLANs, on page 53
Increasing the FabricPath Cost for a vPC+ Peer Link for FCF Leafs, on page 53

Enabling FabricPath

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **install feature-set fabricpath**	Installs the FabricPath feature set on the switch.
Step 3	switch(config)# **feature-set fabricpath**	Enables the FabricPath feature set on the switch.
Step 4	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable FabricPath:

```
switch# configure terminal
switch(config)# install feature-set fabricpath
switch(config)# feature-set fabricpath
switch(config)# copy running-config startup-config
```

Configuring All Leafs in the FabricPath Topology

Quality of Service (QoS) settings are enabled on the spine. FCFs are not being established.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **install feature-set fabricpath**	Installs the FabricPath feature set on the switch.
Step 3	switch(config)# **feature-set fabricpath**	Enables the FabricPath feature set on the switch.
Step 4	switch(config)# **system qos**	Enters system class configuration mode.
Step 5	switch(config-sys-qos)# **service-policy type {network-qos | qos | queuing} [input | output]** *fcoe-default-policy-name*	Sets up the service policy for the system to specify the default policy map for FCoE traffic. Four predefined policy-maps for FCoE are as follows: • service-policy type qos input fcoe-default-in-policy • service-policy type queuing input fcoe-default-in-policy

	Command or Action	Purpose
		• service-policy type queuing output fcoe-default-out-policy • service-policy type network-qos fcoe-default-nq-policy **Note** Before enabling FCoE on a Cisco Nexus device, you must attach the predefined FCoE policy maps to the type qos, type network-qos, and type queuing policy maps.
Step 6	switch(config-sys-qos)# **vlan** *vlan-id*	Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 7	switch(config-vlan)# **mode fabricpath**	Configures the VLANs as FabricPath VLANs.
Step 8	switch(config-vlan)# **interface** [**ethernet** *slot/port* | **port-channel** *channel-no*]	Enters interface configuration mode and specifies the interfaces that you want to configure as FabricPath. The port number within a particular slot can be from 1 to 128. The port channel number assigned to the EtherChannel logical interface can be from 1 to 4096.
Step 9	switch(config-if)# **switchport mode fabricpath**	Specifies interfaces as FabricPath ports.
Step 10	(Optional) switch(config-if)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring FCF Leafs

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature fcoe**	Enables the FCoE capability.
Step 3	switch(config)# **fcoe fka-adv-period** *interval*	Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds.
Step 4	switch(config)# **fabricpath switch-id** *switch-id-value*	Configures the switch ID. The range is from 1 to 4094.

	Command or Action	Purpose
Step 5	switch(config)# **vlan** *vlan-id*	Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 6	switch(config)# **vsan database**	Enters VSAN configuration mode.
Step 7	switch(config-vsan-db)# **vsan** *vsan-id*	Configures VSAN.
Step 8	switch(config-vsan-db)# **show vpc**	Displays information about the vPC. **Note** If vPC is enabled, perform the following procedure at Increasing the FabricPath Cost for a vPC+ Peer Link for FCF Leafs, on page 53.
Step 9	(Optional) switch(config)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure FCF leafs:

```
switch# configure terminal
switch(config)# feature fcoe
switch(config)# fcoe fka-adv-period 20
switch(config)# fabricpath switch-id 5
switch(config)# vlan 100

switch(config-vsan-db)# vsan database
switch(config-vsan-db)# vsan 100
switch(config)# show vpc
```

Configuring FCoE and FabricPath-Enabled VLANs

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan** *vlan-id*	Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 3	switch(config-vlan)# **fcoe** [**vsan** *vsan-id*]	Enables FCoE for the specified VLAN. If you do not specify a VSAN number, a mapping is created from this VLAN to the VSAN with the same number. Configures the mapping from this VLAN to the specified VSAN.

Example

This example shows how to configure FCoE and FabricPath-Enabled VLANs:

```
switch# configure terminal
switch(config)# vlan 100
switch(config-vlan)# fcoe vsan 10
```

Defining FabricPath VLANs

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan** *vlan-id*	Enter VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 3	switch(config-vlan)# **mode fabricpath**	Configures the operational mode of the VLAN.

Example

This example shows how to define a FabricPath VLAN:

```
switch# configure terminal
switch(config)# vlan 100
switch(config-vlan)# mode fabricpath
switch(config-vlan)# copy running-config startup-config
```

Increasing the FabricPath Cost for a vPC+ Peer Link for FCF Leafs

Procedure

	Command or Action	Purpose
Step 1	switch# **show vpc**	Based on the output of the **show vpc** command, you have three options: • If the **show vpc** command is not available, do not continue with this procedure. • If vPC+ is not in the command output, do not continue with this procedure. • If vPC+ is in the command output, perform the remaining steps in this procedure.
Step 2	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 3	switch(config)# **interface** [**ethernet** *slot/port* | **port-channel** *channel-no*]	Enters interface configuration mode and specifies the interfaces that you want to configure as FabricPath. The port number within a particular slot can be from 1 to 128. The port channel number assigned to the EtherChannel logical interface can be from 1 to 4096.
Step 4	switch(config-if)# **fabricpath isis metric** *default-metric*	Configures the metric for the MCT interface. You must set the *default-metric* to 16777215.
Step 5	(Optional) switch(config-if)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to increase the FabricPath for a vPC+ peer link:

```
switch# show vpc
switch# configure terminal
switch(config)# interface port-channel 93
switch(config-if)# fabricpath isis metric 16777215
```

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:
    1. Feature FCoE is enabled on the leaf.
    2. 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

To display Dynamic FCoE Using FabricPath configuration information, perform one of the following tasks:

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 running-config Command

```
switch# show running-config

!Command: show running-config
!Time: Mon Feb 14 19:58:47 2011

version 7.0(3)N1(1)
feature fcoe
```

```
install feature-set fabricpath
feature-set fabricpath

feature telnet
feature lldp

username admin password 5 $1$1dLADwhf$7Ip2IYSMp/0nsII8rU5qh/  role network-admin
no password strength-check
ip domain-lookup
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos fcoe-default-nq-policy
snmp-server user admin network-admin auth md5 0x95d13d5b1da2ee92b77769b4c177a94b
 priv 0x95d13d5b1da2ee92b77769b4c177a94b localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vlan 1
vlan 100
  fcoe vsan 100
  mode fabricpath
vrf context management
  ip route 0.0.0.0/0 10.193.48.1
vsan database
  vsan 100


interface vfc32002
  bind mac-address 54:7f:ee:b1:8a:00
  dce
  switchport mode E
  no shutdown

interface vfc32003
  bind mac-address 54:7f:ee:73:e8:00
  dce
  switchport mode E
  no shutdown

interface Ethernet1/1

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10
interface Ethernet1/11
  switchport mode fabricpath

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface mgmt0
  vrf member management
  ip address 10.193.52.117/21
line console
line vty
fabricpath domain default
fabricpath switch-id 302
```

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 11: Sample Dynamic FCoE Configuration

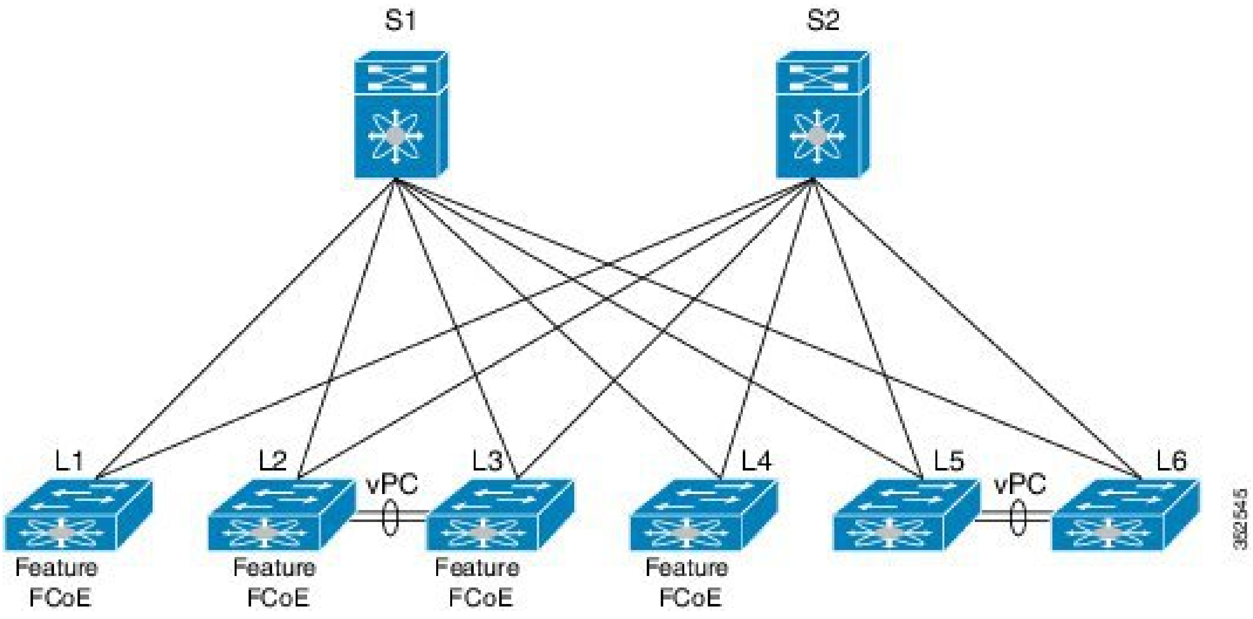

This example shows the configuration on S1 and S2:

```
switch# show running-config
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos fcoe-default-nq-policy
vlan 100
    mode fabric path
vlan 200
```

```
     mode fabric path
```

This example shows the configuration on L5 and L6 non-FCoE leafs:

```
switch# show running-config
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos fcoe-default-nq-policy
vlan 100
    mode fabric path
vlan 200
     mode fabric path
```

This example shows the configuration on L1 - FCF leaf (VSAN 100)

```
switch# show running-config
feature fcoe
vlan 100
    mode fabric path
    fcoe vsan 100
vlan 200
     mode fabric path

vsan database
     vsan 100

fabricpath switch-id 301

fcoe fka-adv-period  20
```

This example shows the configuration on the L4 FCF leaf (VSAN 100, VSAN 200):

```
switch# show running-config
feature fcoe
vlan 100
    mode fabric path
    fcoe vsan 100
vlan 200
     mode fabric path
     fcoe vsan 200
vsan database
     vsan 100
     vsan 200

fabricpath switch-id  304

fcoe fka-adv-period 20
```

This example shows the configuration on the L2 FCF leaf (VSAN 100):

```
switch# show running-config
feature fcoe

vlan 100
    mode fabric path
    fcoe vsan 100
vlan 200
     mode fabric path

vsan database
```

```
     vsan 100

fabricpath switch-id 302

fcoe fka-adv-period 20

switch# show vpc
vPC domain id                     : 1
vPC+ switch id                    : 123

:
vPC Peer-link status
-----------------------------------------------------------
id Port Status Active vlans
-- ---- ------ --------------------------------------------
1  Po93 up     1,10,20,30,101,201,500
interface port-channel93
           fabricpath isis metric 16777215
```

This example shows the configuration on the L3 FCF leaf (VSAN 200):

```
switch# show running-config
feature fcoe

vlan 100
    mode fabric path

vlan 200
     mode fabric path
     fcoe vsan 200
vsan database
     vsan 200

fabricpath switch-id 303

fcoe fka-adv-period 20

switch# show vpc
vPC domain id        : 1
vPC+ switch id       : 123

:
-----------------------------------------------------------
id Port Status Active vlans
-- ---- ------ --------------------------------------------
1  Po93 up     1,10,20,30,101,201,500
interface port-channel93
  fabricpath isis metric 16777215
```

For additional configuration output examples, refer Configuration Output Examples for Dynamic FCoE Using FabricPath.

INDEX