Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide for Cisco Nexus 7000 Series Switches, Release 8.x

First Published: 2011-07-22

Last Modified: 2019-07-29

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

CONTENTS

PREFACE

Preface vii

Audience vii

Document Conventions vii

Related Documentation for Cisco Nexus 7000 Series NX-OS Software viii

Communications, Services, and Additional Information ix

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Overview 3

Information About the Cisco Nexus 2000 Series Fabric Extender 3

Fabric Extender Terminology 4

Fabric Interface Features 5

Host Interfaces 5

Layer 3 Host Interfaces 5

Layer 2 Host Interfaces 6

Host Interface Port Channels 6

Layer 3 Host Interface Port Channels 6

Layer 2 Host Interface Port Channels 7

Minimum Number of Links on a Fabric Port Channel 7

Load Balancing Using Host Interface Port Channels 7

VLANs 8

Protocol Offload 9

Quality of Service 9

Access Control Lists 9

IGMP Snooping 9

Switched Port Analyzer 10
Oversubscription 10
Management Model 11
Forwarding Model 11
Port Channel Fabric Interface Connection 12
Port Numbering Convention 13
Fabric Extender Image Management 14
Licensing Requirements for the Fabric Extender 14
Guidelines and Limitations for the Fabric Extender 14
Associating with F2-Series Modules 16
FEX Queuing Support 16
Configuration Limits 23
Default Settings 23

CHAPTER 3 **Configuring the Cisco Nexus 2000 Series Fabric Extender** 25
Information About Cisco Nexus 2000 Series Fabric Extender 25
FEX in Active-Active Mode 25
Managing the Fabric Extender Feature Set 26
Installing the Fabric Extender Feature Set 27
Enabling the Fabric Extender Feature Set 28
Disallowing the Fabric Extender Feature Set 29
Associating a Fabric Extender to a Fabric Interface 29
Associating a Fabric Extender to a Port Channel 30
Disassociating a Fabric Extender from an Interface 32
Associating a Fabric Extender to an F-Series Module 32
Configuring Fabric Extender Global Features 35
Enabling DSCP to Queue Mapping 37
Configuration Examples 38
Configuring a FEX with a Layer 3 Host Interface 38
Configuring a Host Interface in a vPC Topology Connected to Two FEXs 39
Dual-Homing of a Server to a FEX with FabricPath 39
Configuring a FEX in Active-Active Mode 41
Verifying the Status of DSCP-to-queue Mapping 41
Verifying the Configuration 42

Verifying the Fabric Extender Configuration 42
Verifying the Chassis Management Information 43
Additional References 47
Related Documents 47
Feature History 48

CHAPTER 4 vPC Configuration Synchronization 51
Prerequisites for vPC Configuration Synchronization 52
Guidelines and Limitations for Configuration Synchronization 52
Information about vPC Configuration Synchronization 52
Benefits of vPC Configuration Synchronization 52
Supported Commands for vPC Configuration Synchronization 52
How to configure vPC Config Sync 53
Enabling vPC Configuration Synchronization 53
Synchronizing Configuration for a Physical Port vPC 54
Synchronizing Configuration of vPC Member Port Channel 56
Verifying vPC Configuration Synchronization 57
Feature Information for vPC Configuration Synchronization 58

CHAPTER 5 FEX Queuing 59
Finding Feature Information 59
Information About FEX Queuing 59
Queuing Ethernet Frames Using COS 59
Queuing Ethernet Frames Using DSCP 60
Queueing FCoE Frames Using COS 3 60
How to Configure FEX Queuing 60
Changing COS2Q mapping 60
Changing DSCP2Q Mapping 61
Changing network-qos Template 62
Configuring FEX Queue Parameters 62
Verifying FEX Queuing 63
Feature Information for FEX Queuing 65

CHAPTER 6 Supported FEX Devices 67

Feature Information for Supported FEX Devices 67

Preface

This preface describes the audience, organization, and conventions of the *Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide*. It also provides information on how to obtain related documentation.

This chapter includes the following sections:

- Audience, on page vii
- Document Conventions, on page vii
- Related Documentation for Cisco Nexus 7000 Series NX-OS Software, on page viii
- Communications, Services, and Additional Information, on page ix

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco Nexus 2000 Series Fabric Extender.

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 7000 Series NX-OS Software

The documentation set for the Cisco Nexus 7000 Series Switches is available at the following URLs:

http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/tsd-products-support-series-home.html

Release Notes

The release notes are available at the following URL:

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

Installation and Upgrade Guides

The installation and upgrade guides are available at the following URL:

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

Configuration Guides

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

Command References

The command references are available at the following URL:

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

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.

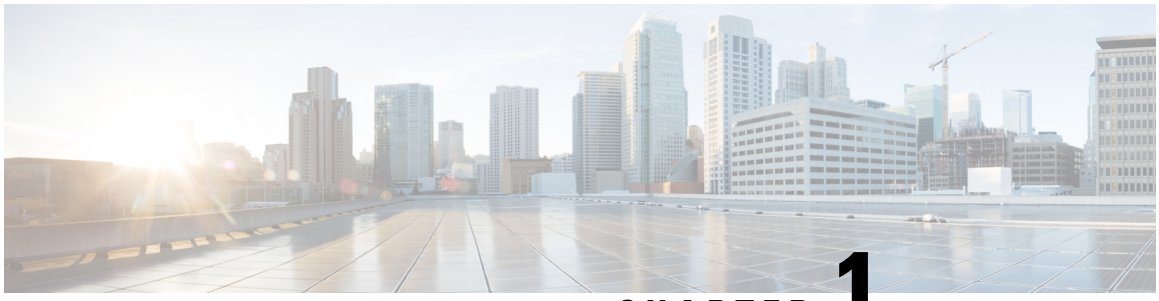

CHAPTER 1

New and Changed Information

- New and Changed Information, on page 1

New and Changed Information

The table below summarizes the new and changed features for this document and shows the releases in which each feature is supported. Your software release might not support all the features in this document. For the latest caveats and feature information, see the Bug Search Tool and the release notes for your software release.

Table 1: New and Changed Interfaces Features

Feature	Description	Changed in Release
F4-Series module support for FEX	Added F4 Series 30-port 100-Gigabit Ethernet I/O module (N77-F430CQ-36) support for FEX.	8.4(1)
Support added for B22 Dell FEX on F4-Series I/O modules	B22 Dell FEX is now supported with F4-Series I/O modules	8.4(1)
FEX support on F4-Series breakout (40G -> 4x10G) interfaces	FEX is now supported on F4-Series breakout (40G -> 4x10G) interfaces.	8.4(1)
Support added for B22 Dell FEX on F3-Series and M3-Series I/O modules	B22 Dell FEX is now supported with F3-Series and M3-Series I/O modules	8.2(1)
M3-Series Modules Support for FEX	Added M3 Series modules support for FEX. **Note** Not supported on M3-Series 12-port 100-Gigabit Ethernet I/O module (N77-M312CQ-26L)	8.1(1)

New and Changed Information

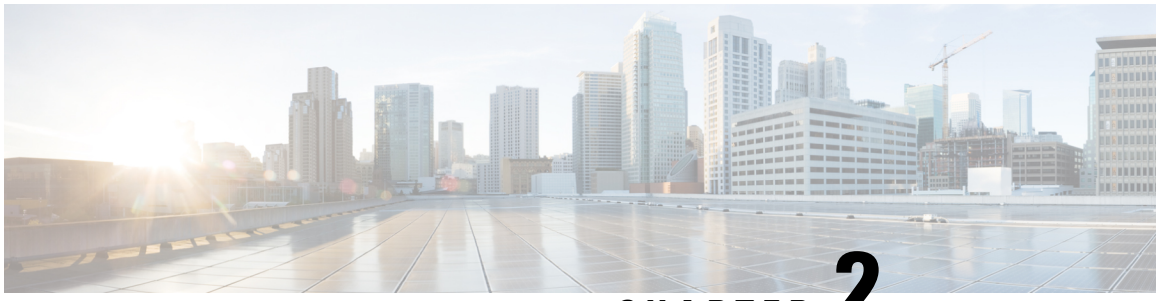

CHAPTER 2

Overview

This chapter provides an architectural overview of the Cisco Nexus 2000 Series Fabric Extender and includes the following sections:

- Information About the Cisco Nexus 2000 Series Fabric Extender, on page 3
- Fabric Extender Terminology, on page 4
- Fabric Interface Features , on page 5
- Host Interfaces, on page 5
- Host Interface Port Channels, on page 6
- VLANs, on page 8
- Protocol Offload, on page 9
- Quality of Service, on page 9
- Access Control Lists, on page 9
- IGMP Snooping, on page 9
- Switched Port Analyzer, on page 10
- Oversubscription, on page 10
- Management Model, on page 11
- Forwarding Model, on page 11
- Port Channel Fabric Interface Connection, on page 12
- Port Numbering Convention, on page 13
- Fabric Extender Image Management, on page 14
- Licensing Requirements for the Fabric Extender, on page 14
- Guidelines and Limitations for the Fabric Extender, on page 14
- Default Settings, on page 23

Information About the Cisco Nexus 2000 Series Fabric Extender

The Cisco Nexus 2000 Series Fabric Extender, also known as FEX, is a highly scalable and flexible server networking solution that works with Cisco Nexus Series devices to provide high-density, low-cost connectivity for server aggregation. Scaling across 1-Gigabit Ethernet, 10-Gigabit Ethernet, unified fabric, rack, and blade server environments, the Fabric Extender is designed to simplify data center architecture and operations.

The Fabric Extender integrates with its parent switch, which is a Cisco Nexus Series device, to allow automatic provisioning and configuration taken from the settings on the parent device. This integration allows large numbers of servers and hosts to be supported by using the same feature set as the parent device, including security and quality-of-service (QoS) configuration parameters, with a single management domain. The Fabric

Extender and its parent switch enable a large multipath, loop-free, active-active data center topology without the use of the Spanning Tree Protocol (STP).

The Cisco Nexus 2000 Series Fabric Extender forwards all traffic to its parent Cisco Nexus Series device over 10-Gigabit Ethernet fabric uplinks, which allows all traffic to be inspected by policies established on the Cisco Nexus Series device.

Note You must connect the Fabric Extender to its parent Cisco Nexus 7000 Series device with a 32-port, 10-Gigabit M1 module (N7K-M132XP-12), a 32-port 10-Gigabit M1-XL module (N7K-M132XP-12L), an M2 module, an M3 module, an F2 module, or an F3 module. ,or an F4 module.

No software is included with the Fabric Extender. The software is automatically downloaded and upgraded from its parent device.

Note When you configure a Cisco Nexus 2248 port to 100mbps speed (instead of auto-negotiation), FEX does not auto negotiate with the peer. You must manually set the peer not to auto negotiate and set the peer to 100mbps speed.

Fabric Extender Terminology

Some terms used in this document are as follows:

- Fabric interface—A 10-Gigabit Ethernet uplink port that is designated for connection from the Fabric Extender to its parent switch. A fabric interface cannot be used for any other purpose. It must be directly connected to the parent switch.

  

  Note A fabric interface includes the corresponding interface on the parent switch. This interface is enabled when you enter the **switchport mode fex-fabric** command.

- Port channel fabric interface—A port channel uplink connection from the Fabric Extender to its parent switch. This connection consists of fabric interfaces that are bundled into a single logical channel.
- Host interface—An Ethernet host interface for connection to a server or host system.

  

  Note Do not connect a bridge or switch to a host interface. These interfaces are designed to provide end host or server connectivity.

Note On Cisco Nexus 2348TQ and Nexus 2348UPQ FEX, if a port channel is used to connect a parent switch with a Fabric Extender device, the port channels can have maximum of 8 ports.

The Nexus 2348 FEX devices have a total of 6 * 40 Gigabit Ethernet uplink ports towards the parent switch. If these are used with native 40G uplinks port on a parent switch, then there is no limitation. All 6 ports can be used in either single homed or dual homed configuration. You can also use 40 Gigabit Ethernet uplink ports on the N2348 Fabric Extender device with 10 Gigabit Ethernet ports on the parent switch when used with the appropriate cabling. A maximum of 8 ports can be added to the port channel between the parent switch and Fabric Extender device. If it is a dual homed setup, VPC to the Fabric Extender device, only 4 ports per switch are allowed in the port channel.

- Port channel host interface—A port channel host interface for connection to a server or host system.

Fabric Interface Features

Note Flow control is not supported on Cisco Nexus 2348TQ fabric extender.

The FEX fabric interfaces support static port channels. During the initial discovery and association process, SFP+ validation and digital optical monitoring (DOM) are performed as follows:

- The FEX performs a local check on the uplink SFP+ transceiver. If it fails the security check, the LED flashes but the link is still allowed to come up.
- The FEX local check is bypassed if it is running its backup image.
- The parent switch performs SFP validation again when the fabric interface is brought up. It keeps the fabric interface down if SFP validation fails.

After an interface on the parent switch is configured in fex-fabric mode, all other features that were configured on that port and are not relevant to this mode are deactivated. If the interface is reconfigured to remove fex-fabric mode, the previous configurations are reactivated.

For more information about PFC, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide*.

Host Interfaces

Layer 3 Host Interfaces

Beginning with Cisco NX-OS Release 5.2, by default, all host interfaces on a Fabric Extender that are connected to a Cisco Nexus 7000 Series parent switch run in Layer 3 mode.

Note If you have updated the parent switch to Cisco Nexus Release 5.2, previously configured fabric Extender host interfaces retain their default port mode, Layer 2. You can change these ports to Layer 3 mode with the **no switchport** command.

The host interfaces also support subinterfaces. You can create up to 63 subinterfaces on a Fabric Extender host interface.

Beginning with Cisco NX-OS Release 6.2, port profiles are supported on the host interfaces of a Fabric Extender.

For information about interfaces, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*.

Layer 2 Host Interfaces

In Cisco NX-OS Release 5.1 and earlier releases, the default port mode is Layer 2.

To run a host interface in Layer 2 mode, use the **switchport** command. For Cisco NX-OS Release 5.2 and later releases, to change the port mode to Layer 3, use the **no switchport** command.

The Fabric Extender provides connectivity for computer hosts and other edge devices in the network fabric.

Follow these guidelines when connecting devices to Fabric Extender host interfaces:

- All Fabric Extender host interfaces run as spanning tree edge ports with BPDU Guard enabled and you cannot configure them as spanning tree network ports.
- You can connect servers that use active/standby teaming, 802.3ad port channels, or other host-based link redundancy mechanisms to Fabric Extender host interfaces.
- Any device that is running spanning tree connected to a Fabric Extender host interface results in that host interface being placed in an error-disabled state when a BPDU is received.
- You can connect only virtual switches that leverages a link redundancy mechanism not dependent on spanning tree such as Cisco FlexLink or vPC (with the BPDU Filter enabled) to a Fabric Extender host interface. Because spanning tree is not used to eliminate loops, you should ensure a loop-free topology below the Fabric Extender host interfaces.

Ingress and egress packet counters are provided on each host interface.

For more information about BPDU Guard, see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*.

Host Interface Port Channels

Layer 3 Host Interface Port Channels

The Fabric Extender (FEX) supports host interface port channel configurations. You can combine up to 8 interfaces in a standard mode port channel and 16 interfaces when configured with the Link Aggregation Control Protocol (LACP).

Note Port channel resources are allocated when the port channel has one or more members.

All members of the port channel must be FEX host interfaces and all host interfaces must be from the same FEX. You cannot mix interfaces from the FEX and the parent switch.

Layer 3 mode is supported on host interface port channels.

A host interface port channel also supports subinterfaces. You can create up to 1000 subinterfaces on a FEX host interface port channel.

For more information about port channels, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*.

Layer 2 Host Interface Port Channels

The Fabric Extender supports host interface port channel configurations. You can combine up to 8 interfaces in a standard mode port channel and 16 interfaces when configured with the Link Aggregation Control Protocol (LACP).

Note Port channel resources are allocated when the port channel has one or more members.

All members of the port channel must be Fabric Extender host interfaces and all host interfaces must be from the same Fabric Extender. You cannot mix interfaces from the Fabric Extender and the parent switch.

Layer 2 mode is supported on host interface port channels.

You can configure Layer 2 port channels as access or trunk ports.

Beginning with Cisco NX-OS Release 5.2(1), Fabric Extenders support the host vPC feature where a server can be dual-attached to two different FEXs through a port channel. You must configure parent switches that connect each Fabric Extender (one parent switch per FEX) in a vPC domain.

Minimum Number of Links on a Fabric Port Channel

In a network configuration of dual-homed hosts (active/standby), you can configure the Fabric Extender to support a minimum number of links for fabric port channels (FPCs) with the **port-channel min-links** command.

When the number of FPC links falls below the specified threshold, the host-facing Cisco Nexus 2000 interfaces are brought down. This process allows for a NIC switchover on the connection between the host and the FEX.

The automatic recovery of Cisco Nexus 2000 Series interfaces to the standby FEX is triggered when the number of FPC links reaches the specified threshold.

Load Balancing Using Host Interface Port Channels

The Cisco NX-OS software allows for load balancing traffic across all operational interfaces on a FEX host interface port-channel by hashing the addresses in the frame to a numerical value that selects one of the links in the channel. Port-channels provide load balancing by default.

You can configure the type of load-balancing algorithm used. You can choose the load-balancing algorithm that determines which member port to select for egress traffic by looking at the fields in the frame.

You can configure the load-balancing mode to apply to all Fabric Extenders or to specified ones. If load-balancing mode is not configured, Fabric Extenders use the default system configuration. The per-FEX configuration takes precedence over the load-balancing configuration for the entire system. You cannot configure the load-balancing method per port channel.

Note The default load-balancing mode for Layer 3 interfaces is the source and destination IP address, and the default load-balancing mode for non-IP interfaces is the source and destination MAC address. For more details, see the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x*.

You can configure the device to use one of the following methods to load balance across the port channel:

- Destination MAC address
- Source MAC address
- Source and destination MAC address
- Destination IP address
- Source IP address
- Source and destination IP address
- Source TCP/UDP port number
- Destination TCP/UDP port number
- Source and destination TCP/UDP port number
- Dot1Q VLAN number

Note You must be in the default virtual device context (VDC) to configure load-balancing method for FEX; if you attempt to configure this feature from another VDC, the system displays an error.

VLANs

The Fabric Extender supports Layer 2 VLAN trunks and IEEE 802.1Q VLAN encapsulation.

For more information about VLANs, see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide*.

Note The Fabric Extender does not support private VLANs (PVLANs).

Protocol Offload

To reduce the load on the control plane of the Cisco Nexus Series device, Cisco NX-OS allows you to offload link-level protocol processing to the Fabric Extender CPU. The following protocols are supported:

- Link Layer Discovery Protocol (LLDP)
- Cisco Discovery Protocol (CDP)
- Link Aggregation Control Protocol (LACP)

Quality of Service

The Fabric Extender uses IEEE 802.1p class of service (CoS) values to associate traffic with the appropriate class. Per-port quality of service (QoS) configuration is also supported.

Host interfaces support pause frames, which are implemented using IEEE 802.3x link-level flow control (LLC). By default, flow control send is on and flow control receive is off on all host interfaces. Autonegotiation is enabled on the host interfaces. Per-class flow control is set according to the QoS classes.

For more information about LLC and quality-of-service, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide*.

Access Control Lists

The Fabric Extender supports the full range of ingress access control lists (ACLs) that are available on its parent Cisco Nexus Series device.

For more information about ACLs, see the Security Configuration Guide for your device.

IGMP Snooping

IGMP snooping is supported on all host interfaces of the Fabric Extender.

The Fabric Extender and its parent switch support IGMPv2 and IGMPv3 snooping based only on the destination multicast MAC address. It does not support snooping that is based on the source MAC address or on proxy reports.

Note For more information about IGMP snooping, see http://tools.ietf.org/wg/magma/draft-ietf-magma-snoop/rfc4541.txt. Also, see the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide*.

Switched Port Analyzer

You can configure the host interfaces on the Fabric Extender as Switched Port Analyzer (SPAN) source ports. You cannot configure Fabric Extender ports as a SPAN destination. Only one SPAN session is supported for all the host interfaces on the same Fabric Extender. Ingress source (Rx), egress source (Tx), or both ingress and egress monitoring are supported.

Note All IP multicast traffic on the VLANs that a Fabric Extender host interface belongs to is captured in the SPAN session. You cannot separate the traffic by IP multicast group membership.

If you configure ingress monitoring and egress monitoring for host interfaces on the same Fabric Extender, you might see a packet twice: once as the packet ingresses on an interface with Rx configured, and again as the packet egresses on an interface with Tx configured.

For more information about SPAN, see the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide*.

Oversubscription

In a switching environment, oversubscription is the practice of connecting multiple devices to the same interface to optimize port usage. An interface can support a connection that runs at its maximum speed. Because most interfaces do not run at their maximum speeds, you can take advantage of unused bandwidth by sharing ports. Oversubscription, which is a function of the available fabric interfaces to active host interfaces, provides cost-effective scalability and flexibility for Ethernet environments.

The Cisco Nexus 2248TP Fabric Extender has 4 10-Gigabit Ethernet fabric interfaces and 48 100/1000BASE-T (100-Mb/1-Gigabit) Ethernet host interfaces. When its host interfaces are running in Gigabit Ethernet mode, it offers the following configurations:

- No oversubscription (40 host interfaces for four fabric interfaces)
- 1.2 to 1 oversubscription (48 host interfaces for four fabric interfaces)
- 4.8 to 1 oversubscription (48 host interfaces for one fabric interface)

The Cisco Nexus 2248TP can be run with no oversubscription when its host interfaces are running in 100-Mb mode.

The Cisco Nexus 2248TP-E Fabric Extender has 4 10-Gigabit Ethernet fabric interfaces and 48 100/1000BASE-T (100-Mb/1-Gigabit) Ethernet host interfaces. When its host interfaces are running in Gigabit Ethernet mode, it offers 1.2 to 1 oversubscription (48 host interfaces for four fabric interfaces).

The Cisco Nexus 2248PQ Fabric Extender has 16 10-Gigabit Ethernet fabric interfaces and 48 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus 2248PQ is 3:1.

The Cisco Nexus 2232PP Fabric Extender has 8 10-Gigabit Ethernet fabric interfaces and 32 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. (Static pinning is not

supported. Port-channel mode is supported only on fabric interfaces.) When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus 2232PP is 4:1.

The Cisco Nexus 2232TM and Cisco Nexus 2232TM-E Fabric Extenders have 8 10-Gigabit Ethernet fabric interfaces and 32 Gigabit and 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus 2232TM and Cisco Nexus 2232TM-E is 4:1.

The Cisco Nexus 2224TP Fabric Extender has 2 10-Gigabit Ethernet fabric interfaces and 24 100/1000BASE-T (100-Mb/1-Gigabit) Ethernet host interfaces. With this system, you can configure a 1.2 to 1 oversubscription (24 host interfaces for 2 fabric interfaces) or higher.

The Cisco Nexus B22 Fabric Extender for HP (NB22HP) has 8 10-Gigabit Ethernet fabric interfaces and 16 1G/10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus B22 Fabric Extender for HP (N2K-B22HP-P) is 2:1.

Management Model

The Cisco Nexus 2000 Series Fabric Extender is managed by its parent switch over the fabric interfaces through a zero-touch configuration model. The switch discovers the Fabric Extender by detecting the fabric interfaces of the Fabric Extender.

After discovery, if the Fabric Extender has been correctly associated with the parent switch, the following operations are performed:

1. The switch checks the software image compatibility and upgrades the Fabric Extender if necessary.
2. The switch and Fabric Extender establish in-band IP connectivity with each other.

   The switch assigns an IP address in the range of loopback addresses (127.15.1.0/24) to the Fabric Extender to avoid conflicts with IP addresses that might be in use on the network.
3. The switch pushes the configuration data to the Fabric Extender. The Fabric Extender does not store any configuration locally.
4. The Fabric Extender updates the switch with its operational status. All Fabric Extender information is displayed using the switch commands for monitoring and troubleshooting.

Forwarding Model

The Cisco Nexus 2000 Series Fabric Extender does not perform any local switching. All traffic is sent to the parent switch that provides central forwarding and policy enforcement, including host-to-host communications between two systems that are connected to the same Fabric Extender as shown in the following figure.

Figure 1: Forwarding Model

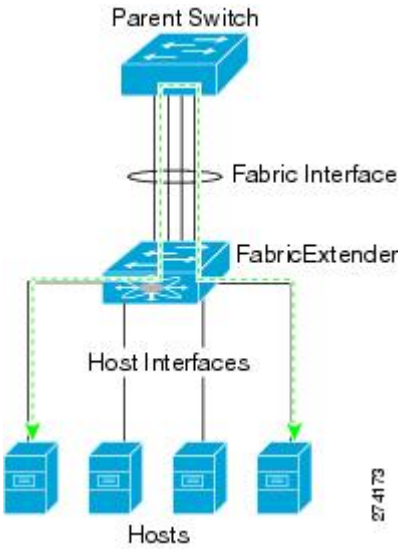

The forwarding model facilitates feature consistency between the Fabric Extender and its parent Cisco Nexus Series device.

Note The Fabric Extender provides end-host connectivity into the network fabric. As a result, BPDU Guard is enabled on all its host interfaces. If you connect a bridge or switch to a host interface, that interface is placed in an error-disabled state when a BPDU is received.

You cannot disable BPDU Guard on the host interfaces of the Fabric Extender.

The Fabric Extender supports egress multicast replication from the network to the host. Packets that are sent from the parent switch for multicast addresses attached to the Fabric Extender are replicated by the Fabric Extender ASICs and are then sent to corresponding hosts.

Port Channel Fabric Interface Connection

To provide load balancing between the host interfaces and the parent switch, you can configure the Fabric Extender to use a port channel fabric interface connection. This connection bundles 10-Gigabit Ethernet fabric interfaces into a single logical channel as shown in the following figure.

Figure 2: Port Channel Fabric Interface Connection

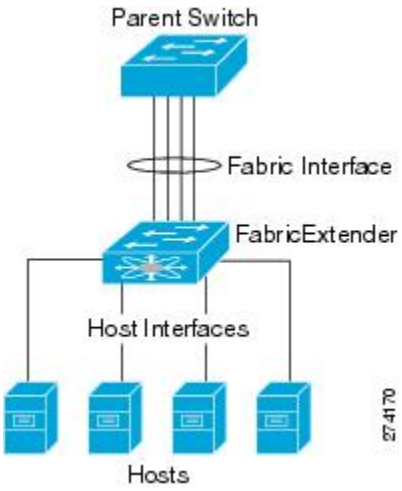

When you configure the Fabric Extender to use a port channel fabric interface connection to its parent switch, the switch load balances the traffic from the hosts that are connected to the host interface ports by using the following load-balancing criteria to select the link:

- For a Layer 2 frame, the switch uses the source and destination MAC addresses.
- For a Layer 3 frame, the switch uses the source and destination MAC addresses and the source and destination IP addresses.

Note A fabric interface that fails in the port channel does not trigger a change to the host interfaces. Traffic is automatically redistributed across the remaining links in the port channel fabric interface. If all links in the fabric port channel go down, all host interfaces on the FEX are set to the down state.

Port Numbering Convention

The following port numbering convention is used for the Fabric Extender:

interface ethernet *chassis*/*slot*/*port*

where

- *chassis* is configured by the administrator. A Fabric Extender must be directly connected to its parent Cisco Nexus Series device via a port channel fabric interface. You configure a chassis ID on a port channel on the switch to identify the Fabric Extender that is discovered through those interfaces.

  The chassis ID ranges from 101 to 199.

Note The chassis ID is required only to access a host interface on the Fabric Extender. A value of less than 101 indicates a slot on the parent switch. The following port numbering convention is used for the interfaces on the switch:

interface ethernet *slot/port*

- *slot* identifies the slot number on the Fabric Extender.
- *port* identifies the port number on a specific slot and chassis ID.

Fabric Extender Image Management

No software ships with the Cisco Nexus 2000 Series Fabric Extender. The Fabric Extender image is bundled into the system image of the parent switch. The image is automatically verified and updated (if required) during the association process between the parent switch and the Fabric Extender.

When you enter the **install all** command, it upgrades the software on the parent Cisco Nexus Series switch and also upgrades the software on any attached Fabric Extender. To minimize downtime as much as possible, the Fabric Extender remains online while the installation process loads its new software image. Once the software image has successfully loaded, the parent switch and the Fabric Extender both automatically reboot.

This process is required to maintain version compatibility between the parent switch and the Fabric Extender.

Licensing Requirements for the Fabric Extender

The following table shows the licensing requirements for the Cisco Nexus 2000 Series Fabric Extender:

Product	License Requirement
Cisco NX-OS	The Cisco Nexus 2000 Series Fabric Extender requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For an explanation of the licensing scheme, see the *Cisco NX-OS Licensing Configuration Guide*.

Guidelines and Limitations for the Fabric Extender

The Cisco Nexus 2000 Series Fabric Extender (FEX) has the following configuration guidelines and limitations:

- Beginning with Cisco NX-OS Release 8.4(1), B22 Dell FEX is supported with F4-Series modules.
- Beginning with Cisco NX-OS Release 5.2(1), the default port mode is Layer 3. Before Cisco NX-OS Release 5.2(1), the default port mode was Layer 2.
- You must enable the Fabric Extender feature set in the default virtual device context (VDC). After you enable the feature set in the default VDC, the FEX can belong to any VDC and can be configured from those VDCs.

- Each Fabric Extender that is connected to a chassis must have a unique FEX ID. The same FEX ID cannot be configured for two or more Fabric Extenders even if the Fabric Extenders are in separate VDCs.
- The FEX ID for a Fabric Extender is persistent across a chassis. The FEX ID is not reset when used in a VDC.
- All the uplinks and host ports of a Fabric Extender belong to a single VDC. The ports cannot be allocated or split among multiple VDCs.
- You must connect the Fabric Extender to its parent Cisco Nexus 7000 Series device with a 32-port 10-Gigabit M1 module (N7K-M132XP-12), a 32-port, 10-Gigabit M1-XL module (N7K-M132XP-12L), an M2-Series module, an M3-Series module, an F2 module, or an F3 module.,or an F4 module.
- The Fabric Extender feature set operation might cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up. You can check whether the standby supervisor is stable by using the **show modules** command. When the standby supervisor is stable, it is indicated as ha-standby.
- You can configure the Fabric Extender host interfaces as edge ports only. The interface is placed in an error-disabled state if a downstream switch is detected.
- The Fabric Extender does not support PVLANs.
- For Cisco NX-OS Release 6.2(2) and later releases, the FEX supports queuing, which allows a router to be connected to a Layer 3 FEX interface or a router to be connected to a Layer 2 FEX interface (using SVI).

  Follow these guidelines for a router that is connected to a Layer 2 FEX interface (using SVI):

  - You can configure routing adjacency with Layer 3 on the peer router.
  - You can configure routing adjacency with SVI on the router using access/trunk interfaces.

  

  Note FEX interfaces do not support the spanning tree protocol.

  You must configure the network without the possibility of any loops.

- For Cisco NX-OS Release 6.2(2) and later releases, the Cisco Fabric Extender supports routing protocol adjacency. Before Cisco NX-OS Release 6.2(2), the Fabric Extender cannot participate in a routing protocol adjacency with a device attached to its port. Only a static direct route is supported. This restriction applies to both of the following supported connectivity cases:
  - An SVI with a FEX single port or portchannel in Layer 2 mode.
  - A FEX port or portchannel in Layer 3 mode.
- For Cisco NX-OS Release 6.2(2) and later releases, the Cisco Fabric Extender supports the following:
  - Queuing for Ethernet frames on a FEX-based CoS and DSCP values and support for queuing Fibre Channel over Ethernet (FCoE) frames on a FEX.
  - FEX HIF (FEX Host Interface) port to connect to a Protocol Independent Multicast (PIM) router.
- For Cisco NX-OS Release 6.2(2) and later releases, the Cisco Fabric Extender supports optimized multicast flooding (OMF) is available on FEX ports.

- The Cisco Fabric Extender does not support policy based routing (PBR).

  Beginning with Cisco NX-OS Release 6.2(2), the configured MTU for the FEX ports is controlled by the network QoS policy. To change the MTU that is configured on the FEX ports, modify the network QoS policy to change when the fabric port MTU is also changed.

- In Cisco NX-OS Release 8.2(4), when you use the **no negotiate auto** command (for a FEX Host Interface (HIF) without a transceiver) after setting the speed to 1000, you get an error message as given below. This is a known limitation.

  ```
  ERROR: Ethernet103/1/23: Configuration does not match the port capability.
  ```

- Starting from Cisco NX-OS Release 8.4(1), FEX is supported on F4-Series breakout (40G -> 4x10G) interfaces.

Associating with F2-Series Modules

- The following FEX devices support F2 modules:
  - 2248TP
  - 2248TP-E
  - 2248PQ
  - 2232TP
  - 2232PP
  - 2232TM
  - 2224TP
- Each port in the ASIC has an index. Allow only ports with similar indices across ASICs to be added to a port channel.

  For example, if port 1 has an index of 1 and port 2 has an index of 2, the following ports are supported and not supported:
  - Supported: Port 1 of ASIC 1 and port 1 of ASIC 2 are added to a port channel.
  - Not supported: Port 1 of ASIC 1 and port 2 of ASIC 2 to form a port channel.

  A set of ports from an ASIC that has an index sub-set S, such as {1,2,4}, is allowed to be added to a port channel only if the port channel has an equivalent or an empty set.

FEX Queuing Support

- **FEX QoS Queuing Support**

  Fabric Extenders (FEXs) follow the network quality of service (QoS) queuing model for supporting queuing on FEX host interfaces, regardless of whether the FEX is connected to M-series or F-series fabric uplinks.
  - Depending on the network-QoS template that is attached to the system QoS, the following parameters are inherited for queuing support on a FEX:

- Number of queues
- Class of service (CoS2q) mapping
- Differentiated services code point (DSCP2q) mapping
- Maximum transmission unit (MTU)

- For both ingress and egress queuing on the FEX host interfaces, all of the preceding parameters are derived from the ingress queuing parameters that are defined in the active network-QoS policy. The egress queuing parameters of the active network-QoS policy do not affect the FEX host-port queuing.
- Such parameters as the bandwidth, queue limit, priority, and set CoS in the network-QoS type queuing-policy maps are not supported for a FEX.

- **Hardware Queue-limit Support**

  The following example shows how to configure the queue limit for a FEX by using the **hardware** *fex-type* **queue-limit** command in the FEX configuration mode:

```
switch(config)# fex 101
switch(config-fex)# hardware ?
  B22DELL Fabric Extender 16x10G SFP+ 8x10G SFP+ Module
  B22HP Fabric Extender 16x10G SFP+ 8x10G SFP+ Module
  N2224TP Fabric Extender 24x1G 2x10G SFP+ Module
  N2232P Fabric Extender 32x10G SFP+ 8x10G SFP+ Module
  N2232TM Fabric Extender 32x10GBase-T 8x10G SFP+ Module
  N2232TM-E Fabric Extender 32x10GBase-T 8x10G SFP+ Module
  N2248T Fabric Extender 48x1G 4x10G SFP+ Module
  N2248TP-E Fabric Extender 48x1G 4x10G SFP+ Module
switch(config-fex)# hardware N2248T ?
  queue-limit Set queue-limit
switch(config-fex)# hardware N2248T queue-limit ?
  <5120-652800> Queue limit in bytes ======> Allowed range of values varies dependent
on the FEX type for which it is configured
switch(config-fex)# hardware N2248T queue-limit ======> Default configuration that sets
 queue-limit to default value of 66560 bytes
switch(config-fex)# hardware N2248T queue-limit 5120 ======> Set user defined queue-limit
 for FEX type N2248T associated on fex id 101
switch(config-fex)# no hardware N2248T queue-limit ======> Disable queue-limit for FEX
 type N2248T associated on fex id 101
switch(config-fex)# hardware N2248TP-E queue-limit ?
  <32768-33538048> Queue limit in Bytes
  rx Ingress direction
  tx Egress direction
switch(config-fex)# hardware N2248TP-E queue-limit 40000 rx
switch(config-fex)# hardware N2248TP-E queue-limit 80000 tx ======> For some FEX types,
 different queue-limit can be configured on ingress & egress directions
```

  The value of the queue limit that is displayed for a FEX interface is 0 bytes until after the first time the FEX interface is brought up. After the interface comes up, the output includes the default queue limit or the user-defined queue limit based on the hardware queue-limit configuration. If the hardware queue limit is unconfigured, “Queue limit: Disabled” is displayed in the command output. The following partial output of the **show queuing interface** *interface* command shows the queue limit that is enforced on a FEX:

```
switch# show queuing interface ethernet 101/1/48

<snippet>
Queue limit: 66560 bytes
```

```
<snippet>
```

- **Global Enable/Disable Control of DSCP2Q**

  In the following example, the **all** or the **f-series** keyword enables DSCP2q mapping for the FEX host interfaces, regardless of the module type to which the FEX is connected:

```
switch(config)# hardware QoS dscp-to-queue ingress module-type ?
  all       Enable dscp based queuing for all cards
  f-series  Enable dscp based queuing for f-series cards
  m-series  Enable dscp based queuing for m-series cards
```

- **Show Command Support for FEX Host Interfaces**

  The **show queuing interface** *interface* command is supported for FEX host interfaces. The following sample output of this command for FEX host interfaces includes the number of queues used, the mapping for each queue, the corresponding queue MTU, the enforced hardware queue limit, and the ingress and egress queue statistics.

Note There is no support to clear the queuing statistics shown in this output.

```
switch# show queuing interface ethernet 199/1/2

slot  1
=======

Interface is not in this module.

slot  2
=======

Interface is not in this module.

slot  4
=======

Interface is not in this module.

slot  6
=======

Interface is not in this module.

slot  9
=======

Ethernet199/1/2 queuing information:
  Input buffer allocation:
  Qos-group: ctrl
  frh: 0
  drop-type: drop
  cos: 7
  xon       xoff      buffer-size
  ---------+---------+-----------
  2560      7680      10240

  Qos-group: 0  2  (shared)
  frh: 2
  drop-type: drop
```

```
  cos: 0 1 2 3 4 5 6
  xon       xoff      buffer-size
  ---------+---------+-----------
  34560     39680     48640

  Queueing:
  queue    qos-group    cos             priority  bandwidth mtu
  --------+------------+---------------+---------+---------+----
  ctrl-hi  n/a          7                PRI         0      2400
  ctrl-lo  n/a          7                PRI         0      2400
  2        0            0 1 2 3 4        WRR        80      1600
  4        2            5 6              WRR        20      1600

  Queue limit: 66560 bytes

  Queue Statistics:
  queue  rx              tx              flags
  ------+---------------+---------------+-----
  0      0               0                ctrl
  1      0               0                ctrl
  2      0               0                data
  4      0               0                data

  Port Statistics:
  rx drop         rx mcast drop   rx error        tx drop         mux ovflow
  ---------------+---------------+---------------+---------------+--------------
  0               0               0               0               InActive

  Priority-flow-control enabled: no
  Flow-control status: rx 0x0, tx 0x0, rx_mask 0x0
  cos     qos-group   rx pause  tx pause  masked rx pause
  -------+-----------+---------+---------+---------------
  0              0    xon       xon       xon
  1              0    xon       xon       xon
  2              0    xon       xon       xon
  3              0    xon       xon       xon
  4              0    xon       xon       xon
  5              2    xon       xon       xon
  6              2    xon       xon       xon
  7            n/a    xon       xon       xon

  DSCP to Queue mapping on FEX
----+--+-----+-------+--+---

FEX TCAM programmed successfully

  queue         DSCPs
-----          -----
02             0-39,
04             40-63,
03             **EMPTY**
05             **EMPTY**


slot 10
=======


slot 11
=======

Interface is not in this module.

slot 15
```

```
=======

Interface is not in this module.

slot 16
=======

Interface is not in this module.

slot 17
=======

Interface is not in this module.

slot 18
=======

Interface is not in this module.
```

- **ISSU Behavior**

  In Cisco NX-OS Release 6.2(2) and later releases, FEX queuing is disabled by default on all existing FEXs after an in-service software upgrade (ISSU). FEX queuing is enabled upon flapping the FEX. You can reload the FEX to enable queuing on any FEX after an ISSU. A message is displayed in the output of the **show queuing interface** *interface* command for the FEX host interface after an ISSU.

```
switch# show queuing interface ethernet 133/1/32 module 9

Ethernet133/1/32 queuing information:
  Input buffer allocation:
  Qos-group: ctrl
  frh: 0
  drop-type: drop
  cos: 7
  xon       xoff      buffer-size
  ---------+---------+-----------
  2560      7680      10240

  Qos-group: 0
  frh: 8
  drop-type: drop
  cos: 0 1 2 3 4 5 6
  xon       xoff      buffer-size
  ---------+---------+-----------
  0         126720    151040

  Queueing:
  queue    qos-group    cos            priority  bandwidth mtu
  --------+------------+--------------+---------+---------+----
  ctrl-hi  n/a          7               PRI         0      2400
  ctrl-lo  n/a          7               PRI         0      2400
  2        0            0 1 2 3 4 5 6   WRR       100      9440

  Queue limit: 66560 bytes

  Queue Statistics:
  queue  rx              tx              flags
  ------+---------------+---------------+-----
  0      0               0                ctrl
  1      0               0                ctrl
  2      0               0                data

  Port Statistics:
```

```
rx drop          rx mcast drop   rx error        tx drop          mux ovflow
---------------+---------------+---------------+---------------+--------------
0                0               0               0                InActive

Priority-flow-control enabled: no
Flow-control status: rx 0x0, tx 0x0, rx_mask 0x0
cos     qos-group   rx pause  tx pause  masked rx pause
-------+----------+---------+---------+---------------
0              0    xon       xon       xon
1              0    xon       xon       xon
2              0    xon       xon       xon
3              0    xon       xon       xon
4              0    xon       xon       xon
5              0    xon       xon       xon
6              0    xon       xon       xon
7            n/a    xon       xon       xon

***FEX queuing disabled on fex 133. Reload the fex to enable queuing.<======
```

For any new FEXs brought online after an ISSU, queuing is enabled by default.

The queue limit is enabled by default for all FEXs, regardless of whether queuing is enabled or disabled for the FEX. In Cisco NX-OS Release 6.2(2), all FEXs come up with the default hardware queue-limit value. Any user-defined queue limit that is configured after an ISSU by using the **hardware queue-limit** command takes effect even if queuing is not enabled for the FEX.

- **No Support on the Cisco Nexus 2248PQ 10-Gigabit Ethernet Fabric Extender**

  The following sample output shows that FEX queuing is not supported for the Cisco Nexus 2248PQ 10-Gigabit Ethernet Fabric Extender (FEX2248PQ):

```
switch# show queuing interface ethernet 143/1/1 module 5

Ethernet143/1/1 queuing information:
Network-QOS is disabled for N2248PQ <=======
Displaying the default configurations
  Input buffer allocation:
  Qos-group: ctrl
  frh: 0
  drop-type: drop
  cos: 7
  xon       xoff      buffer-size
  ---------+---------+-----------
  2560      7680      10240

  Qos-group: 0
  frh: 8
  drop-type: drop
  cos: 0 1 2 3 4 5 6
  xon       xoff      buffer-size
  ---------+---------+-----------
  0         126720    151040

  Queueing:
  queue    qos-group    cos            priority  bandwidth mtu
  --------+------------+--------------+---------+---------+----
  ctrl-hi  n/a          7               PRI         0      2400
  ctrl-lo  n/a          7               PRI         0      2400
  2        0            0 1 2 3 4 5 6   WRR       100      9440

  Queue limit: 0 bytes

  Queue Statistics:
```

```
queue  rx              tx              flags
------+---------------+---------------+-----
0      0               0                ctrl
1      0               0                ctrl
2      0               0                data

Port Statistics:
rx drop         rx mcast drop   rx error        tx drop         mux ovflow
---------------+---------------+---------------+---------------+--------------
0               0               0               0                InActive

Priority-flow-control enabled: no
Flow-control status: rx 0x0, tx 0x0, rx_mask 0x0
cos     qos-group   rx pause  tx pause  masked rx pause
-------+----------+---------+---------+---------------
0              0    xon       xon       xon
1              0    xon       xon       xon
2              0    xon       xon       xon
3              0    xon       xon       xon
4              0    xon       xon       xon
5              0    xon       xon       xon
6              0    xon       xon       xon
7            n/a    xon       xon       xon
```

- **Fabric Port Queuing Restrictions**
  - For FEXs that are connected to M-series uplinks, the queuing structure is different on FEX host interfaces and FEX fabric interfaces. The M series queuing policies must be consistent with the FEX queuing policies.
  - DSCP-based queuing is not supported on the FEX fabric uplinks of the Cisco M1 series 1 Gigabit Ethernet modules. However, DSCP-based queuing is supported on the FEX fabric uplinks of the Cisco M1 Series 10 Gigabit Ethernet modules, Cisco M2 Series modules, Cisco M3 Series modules, Cisco F2 Series modules, Cisco F3 Series modules, and Cisco F4 Series modules.

- **MTU**
  - FEX queue MTU configurations are derived from type network-QoS policy-map templates. MTU changes are applied on cloned network-QoS policy maps. The MTU that is configured on a FEX port must match the MTU in the network-QoS policy map so that the FEX MTU can be applied to the FEX host interfaces. For more information, see the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide*.

Note Starting with Cisco NX-OS Release 6.2(2), the configured MTU for the FEX ports is controlled by the network QoS policy. To change the MTU that is configured on the FEX ports, modify the network QoS policy to change when the fabric port MTU is also changed.

If you change the FEX fabric port MTU on a version prior to Cisco NX-OS Release 6.2(x), and then upgrade via ISSU to Cisco NX-OS Release 6.2(x) or a later version, you will not get any issues until either a FEX or switch is reloaded. It is recommended that post-upgrade, the FEX HIF MTU be changed via the network QoS policy as described above.

Qos policy changes affects only F series cards and M series cards.

Configuration Limits

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

Default Settings

This table lists the default settings for the Fabric Extender parameters.

Table 2: Default Cisco Nexus 2000 Series Fabric Extender Parameter Settings

Parameters	Default
feature-set fex command	Disabled
Port mode	Layer 3 (Cisco NX-OS Release 5.2 and later releases). Layer 2 (Cisco NX-OS Release 5.1 and earlier releases).

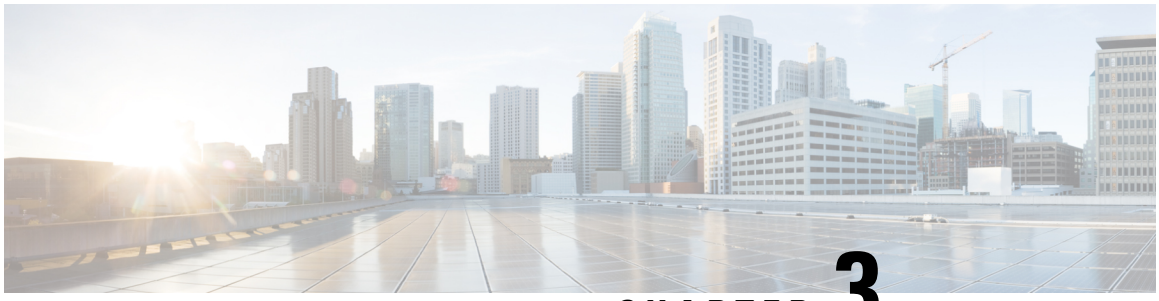

CHAPTER 3

Configuring the Cisco Nexus 2000 Series Fabric Extender

This section describes how to configure the Fabric Extender.

- Information About Cisco Nexus 2000 Series Fabric Extender, on page 25
- Managing the Fabric Extender Feature Set, on page 26
- Associating a Fabric Extender to a Fabric Interface, on page 29
- Configuring Fabric Extender Global Features, on page 35
- Enabling DSCP to Queue Mapping, on page 37
- Configuration Examples, on page 38
- Verifying the Configuration, on page 42
- Additional References, on page 47

Information About Cisco Nexus 2000 Series Fabric Extender

FEX in Active-Active Mode

Beginning with Cisco NX-OS Release 7.2(0)D1(1) , a Fabric Extender can support connections to two Cisco Nexus 7000 Series switches in active-active mode using a vPC.

A Fabric Extender in active-active (FEX-AA) mode:

- Uses a vPC to provide a seamless fail-over and fast convergence when one of the switches fail.
- Supports traffic across both switches to maintain efficiency.

Figure 3: FEX Active-Active

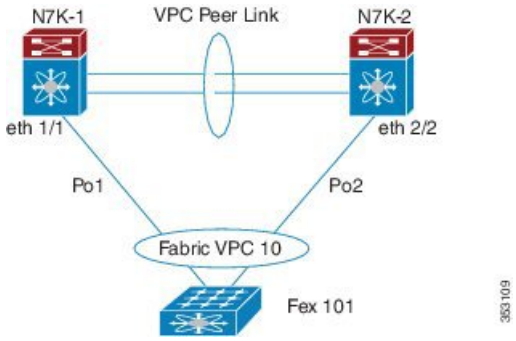

Configuration Synchronization and FEX-AA

The vPC configuration synchronization feature can be used for FEX-AA. Configuration synchronization allows you to synchronize the configuration between a pair of switches in a network. You use a switch profile to create a configuration file that is applied locally and used to synchronize the configuration to the peer.

Guidelines and Limitations for FEX-AA

The following are guidelines and limitations for FEX-AA:

- FEX-AA is not supported on vPC+ deployments.
- Only F-Series (except F1 series) modules support FEX-AA. M1, M2, and M3 Series modules do not support FEX-AA.
- Configuring FEX-AA across two VDCs on the same chassis is not supported.
- FEXs configured in AA mode cannot have host interfaces configured in L3 mode.
- Straight-Through FEX and Active-Active FEX cannot exist on a same ASIC instance.
- Both Cisco Nexus 7000 Series switches can configure the FEX.
- Both Cisco Nexus 7000 Series switches must configure the FEX in the same way so that the fex-id is the same for each.
- The configuration of host ports and host port-channels behind the FEX in AA mode must be the same on both the Cisco Nexus 7000 switches .
- The FEX image can be downloaded from either Cisco Nexus 7000 Series switch.

Managing the Fabric Extender Feature Set

You can install and manage the Fabric Extender feature set.

SUMMARY STEPS

1. Installing the Fabric Extender Feature Set, on page 27
2. Enabling the Fabric Extender Feature Set, on page 28
3. (Optional) Disallowing the Fabric Extender Feature Set, on page 29

DETAILED STEPS

Step 1 Installing the Fabric Extender Feature Set, on page 27

Step 2 Enabling the Fabric Extender Feature Set, on page 28

Step 3 (Optional) Disallowing the Fabric Extender Feature Set, on page 29

Installing the Fabric Extender Feature Set

Note You must enable the Fabric Extender feature set in the default virtual device context (VDC). Once enabled in the default VDC, the FEX can belong to any VDC and can be configured from those VDCs. A single Fabric Extender belongs exclusively to a single VDC.

Before you begin

Ensure that you are in the default VDC.

Ensure that you have disabled the Intrusion Detection System (IDS) reserved addresses check (it is disabled by default). Use the **show hardware ip verify** command and look for the string "address reserved" in the output. If the IDS reserved addresses check is enabled, disable it with the **no hardware ip verify address reserved** command.

SUMMARY STEPS

1. **configure terminal**
2. **install feature-set fex**
3. **exit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**install feature-set fex** **Example:** `switch(config)# install feature-set fex`	Installs the Fabric Extender feature set in the default VDC. To uninstall the Fabric Extender feature set, use the **no install feature-set fex** command. Before you can uninstall the feature set, you must ensure the following: • The feature set must be installed in the default VDC. • The feature set must not be enabled in any VDC.

	Command or Action	Purpose

Step 3	**exit**
	Example:
	`switch(config)# exit`
	`switch#`
		Exits global configuration mode.

Enabling the Fabric Extender Feature Set

You can enable the installed Fabric Extender feature set in any VDC on the device.

Before you begin

Ensure that you have installed the Fabric Extender feature set in the default VDC.

Ensure that you are in the correct VDC or use the **switchto vdc** command.

SUMMARY STEPS

1. **configure terminal**
2. **feature-set fex**
3. **exit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**feature-set fex** **Example:** `switch(config)# feature-set fex`	Enables the Fabric Extender feature set. The feature set must be installed before it shows as an option to this command. To disable the Fabric Extender feature set, use the **no feature-set fex** command. Before you can disable a feature set, you must install the feature set in the default VDC. **Note** The **no feature-set fex** command might take some time to complete if the size of the configuration is very large. The command cleans up all of the configurations associated with the Fabric Extender feature set.

Step 3	**exit**
	Example:
	`switch(config)# exit`
	`switch#`
		Exits global configuration mode.

Disallowing the Fabric Extender Feature Set

By default, when you install the Fabric Extender feature set, it is allowed in all VDCs. You can disallow the installed Fabric Extender feature set in a specific VDC on the device.

Before you begin

Ensure that you have installed the feature set in the default VDC.

Ensure that you have not enabled the feature set in the specified VDC.

SUMMARY STEPS

1. **configure terminal**
2. **vdc** *vdc_ID*
3. **no allow feature-set fex**
4. **end**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**vdc** *vdc_ID* **Example:** `switch(config)# vdc 1`	Specifies a VDC and enters VDC configuration mode.
Step 3	**no allow feature-set fex** **Example:** `switch(config-vdc)# no allow feature-set fex`	Disallows the feature set in the VDC. You cannot disallow a feature set that is enabled in the specified VDC. By default, the installed Fabric Extender feature set is allowed in all VDCs on the device. You can disallow a feature set in a specific VDC. Subsequently, you can change the status back to allowed with the **allow feature-set fex** command.

Step 4	**end**
	Example:
	`switch(config-vdc)# end`
	`switch#`
		Exits VDC configuration mode and returns to EXEC mode.

Associating a Fabric Extender to a Fabric Interface

A FEX is connected to its parent device through a port channel. By default, the parent device does not allow the attached Fabric Extender to connect until it has been assigned a FEX number and is associated with the connected interface.

Note You must have installed and enabled the Fabric Extender features with the **install feature-set fex** and **feature-set fex** commands before you can configure and use a Fabric Extender that is connected to the parent switch.

Note Nexus 7000 doesn't support static pinning configuration on FEX.

Note Associating a Fabric Extender to a fabric interface enables the specified port to operate only in shared rate-mode. The port does not support dedicated rate-mode when associated with a Fabric Extender.

Associating a Fabric Extender to a Port Channel

Before you begin

Ensure that you have enabled the Fabric Extender feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *channel*
3. **switchport mode fex-fabric**
4. **fex associate** *FEX-number*
5. (Optional) **show interface port-channel** *channel* **fex-intf**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface port-channel** *channel*
	Example:
	`switch(config)# interface port-channel 4`
	`switch(config-if)#`
		Specifies a port channel to configure.

Step 3	**switchport mode fex-fabric** **Example:** `switch(config-if)# switchport mode fex-fabric`	Sets the port channel to support an external Fabric Extender.

	Command or Action	Purpose
Step 4	**fex associate** *FEX-number* **Example:** `switch(config-if)# fex associate 101`	Associates a FEX number to the Fabric Extender unit attached to the interface. The range is from 100 to 199.
Step 5	(Optional) **show interface port-channel** *channel* **fex-intf** **Example:** `switch# show interface port-channel 4 fex-intf`	Displays the association of a Fabric Extender to a port channel interface.

Example

This example shows how to associate the Fabric Extender to a port channel interface on the parent device:

```
switch# configure terminal
switch(config)# interface ethernet 1/28
switch(config-if)# channel-group 4
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# interface ethernet 1/29
switch(config-if)# channel-group 4
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# interface ethernet 1/30
switch(config-if)# channel-group 4
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# interface ethernet 1/31
switch(config-if)# channel-group 4
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# interface port-channel 4
switch(config-if)# switchport
switch(config-if)# switchport mode fex-fabric
switch(config-if)# fex associate 101
```

Tip As a best practice, only enter the **fex associate** command from the port channel interface, not from the physical interface.

Note When adding physical interfaces to port channels, all configurations on the port channel and physical interface must match.

This example shows how to display the association of the Fabric Extender and the parent device:

```
switch# show interface port-channel 4 fex-intf
Fabric           FEX
Interface        Interfaces
---------------------------------------------------
Po4              Eth101/1/48   Eth101/1/47   Eth101/1/46   Eth101/1/45
                 Eth101/1/44   Eth101/1/43   Eth101/1/42   Eth101/1/41
```

```
Eth101/1/40   Eth101/1/39   Eth101/1/38   Eth101/1/37
Eth101/1/36   Eth101/1/35   Eth101/1/34   Eth101/1/33
Eth101/1/32   Eth101/1/31   Eth101/1/30   Eth101/1/29
Eth101/1/28   Eth101/1/27   Eth101/1/26   Eth101/1/25
Eth101/1/24   Eth101/1/23   Eth101/1/22   Eth101/1/21
Eth101/1/20   Eth101/1/19   Eth101/1/18   Eth101/1/17
Eth101/1/16   Eth101/1/15   Eth101/1/14   Eth101/1/13
Eth101/1/12   Eth101/1/11   Eth101/1/10   Eth101/1/9
Eth101/1/8    Eth101/1/7    Eth101/1/6    Eth101/1/5
Eth101/1/4    Eth101/1/3    Eth101/1/2    Eth101/1/1
```

Disassociating a Fabric Extender from an Interface

Before you begin

Ensure that you have enabled the Fabric Extender feature.

SUMMARY STEPS

1. **configure terminal**
2. **interface** {**ethernet** *slot*/*port* | **port-channel** *channel*}
3. **no fex associate**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface** {**ethernet** *slot*/*port* | **port-channel** *channel*}
	Example:
	`switch(config)# interface port-channel 4`
	`switch(config-if)#`
		Specifies the interface to configure. The interface can be an Ethernet interface or a port channel.

Step 3	**no fex associate** **Example:** `switch(config-if)# no fex associate`	Disassociates the Fabric Extender unit attached to the interface.

Associating a Fabric Extender to an F-Series Module

Before you begin

Ensure that you have installed and enabled the Fabric Extender feature set.

- switch# **configure terminal**
- switch(config)# **install feature-set fex**

- switch(config)# **feature-set fex**

Note F-Series modules (except F1 Series) are supported only by the following FEX devices:

- 2248TP
- 2248TP-E
- 2248PQ
- 2348UPQ
- 2348TQ
- 2232TP
- 2232PP
- 2232TM
- 2224TP
- B22-IBM
- B22-HP
- B22-DELL

Note Restriction for F2-Series Linecards in an F2 VDC:

Each port in the ASIC has an index. Allow only ports with similar indices across ASICs to be added to a port channel.

For example, if port 1 has an index of 1 and port 2 has an index of 2, the following ports are supported and not supported:

- Supported: Port 1 of ASIC 1 and port 1 of ASIC 2 are added to a port channel.
- Not supported (For F2-series): Port 1 of ASIC 1 and port 2 of ASIC 2 to form a port channel.

A set of ports from an ASIC that has an index sub-set S, such as {1,2,4}, is allowed to be added to a port channel only if the port channel has an equivalent or an empty set.

SUMMARY STEPS

1. **vdc switch**
2. **limit-resource module-type** [**f2e** | **f3**]
3. **interface ethernet** *mod-number port-range*
4. **allocate interface ethernet** *slot-port*
5. **switchport mode fex -fabric**
6. **fex associate** *FEX chassis number*

7. **channel-group** *port-channel number*
8. **no shutdown**
9. **int port** *port-channel number*
10. **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**vdc switch** **Example:** `switch(config)# vdc switch`	Specifies the VDC.
Step 2	**limit-resource module-type** [**f2e** | **f3**] **Example:** `switch(config)# limit-resource module-type f2e`	Specifies the module type.
Step 3	**interface ethernet** *mod-number port-range* **Example:** `switch(config)# interface ethernet 1/1`	Specifies the interface.
Step 4	**allocate interface ethernet** *slot-port* **Example:** `switch(config)# allocate interface ethernet 1`	Allocates the F2-Series interfaces to VDC.
Step 5	**switchport mode fex -fabric** **Example:** `switch(config-if)# switchport mode fex -fabric`	Specifies the FEX.
Step 6	**fex associate** *FEX chassis number* **Example:** `switch(config-if)# fex associate 101`	Specifies the chassis.
Step 7	**channel-group** *port-channel number* **Example:** `switch(config-if)# channel-group 1`	Specifies the port channel number.
Step 8	**no shutdown** **Example:** `switch(config-if)# no shutdown`	Brings up the port.
Step 9	**int port** *port-channel number* **Example:** `switch(config-if)# int port 1`	Specifies the port channel.
Step 10	**no shutdown** **Example:**	Brings up the port channel.

	Command or Action	Purpose
	`switch(config-if)# no shutdown`	

Configuring Fabric Extender Global Features

You can configure global features on the Fabric Extender.

Before you begin

Ensure that you have enabled the Fabric Extender feature set.

SUMMARY STEPS

1. **configure terminal**
2. **fex** *FEX-number*
3. (Optional) **description** *desc*
4. (Optional) **no description**
5. (Optional) **type** *FEX-type*
6. (Optional) **type** *FEX-type*
7. (Optional) **no type**
8. (Optional) **serial** *serial*
9. (Optional) **no serial**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fex** *FEX-number*
	Example:
	`switch(config)# fex 101`
	`switch(config-fex)#`
		Enters FEX configuration mode for the specified Fabric Extender. The range of the *FEX-number* is from to 199.

Step 3	(Optional) **description** *desc* **Example:** `switch(config-fex)# description Rack7A-N2K`	Specifies the description. The default is the string FEX*xxxx* where *xxxx* is the FEX number. If the FEX number is 123, the description is FEX0123.
Step 4	(Optional) **no description** **Example:** `switch(config-fex)# no description`	Deletes the description.
Step 5	(Optional) **type** *FEX-type* **Example:**	Specifies the type of Fabric Extender. The *FEX-type* is one of the following:

	Command or Action	Purpose
	`switch(config-fex)# type N2248T`	• N2224TP—24 100 BASE-T/1000 BASE-T Ethernet host interfaces and 2 10-Gigabit SFP+ Ethernet fabric interfaces module • N2232P and N2232TM—32 10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • N2248T and N2248TP-E—48 100BASE-T/1000 BASE-T Ethernet host interfaces and 4 10-Gigabit SFP+ Ethernet fabric interfaces module • NB22HP—16 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • NB22FJ—16 10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • NB22DELL—16 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module The parent Cisco Nexus Series device remembers the type of the Fabric Extender in its binary configuration. When this feature is configured, the Fabric Extender is allowed to come online only if its type matches the configured FEX type.
Step 6	(Optional) **type** *FEX-type* **Example:** `switch(config-fex)# type N2248T`	Specifies the type of Fabric Extender. The *FEX-type* is one of the following: • N2224TP—24 100 Base-T/1000 Base-T Ethernet host interfaces and 2 10-Gigabit SFP+ Ethernet fabric interfaces module • N2232P and N2232TM—32 10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • N2248T and N2248TP-E—48 100 Base-T/1000 Base-T Ethernet host interfaces and 4 10-Gigabit SFP+ Ethernet fabric interfaces module. • N2248PQ—48 10-Gigabit SFP+ Ethernet host interfaces and 16 10-Gigabit SFP+ Ethernet fabric interfaces module. • NB22HP—16 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module

	Command or Action	Purpose
		The parent Cisco Nexus Series device remembers the type of the Fabric Extender in its binary configuration. When this feature is configured, the Fabric Extender is allowed only to come online if its type matches the configured FEX type.
Step 7	(Optional) **no type** **Example:** `switch(config-fex)# no type`	Deletes the FEX type. When a Fabric Extender is connected to the fabric interfaces and does not match the configured type that is saved in the binary configuration on the parent switch, all configurations for all interfaces on the Fabric Extender are deleted.
Step 8	(Optional) **serial** *serial* **Example:** `switch(config-fex)# serial JAF1339BDSK`	Defines a serial number string. If this command is configured, a switch allows the corresponding chassis ID to associate (using the **fex associate** command) only if the Fabric Extender reports a matching serial number string. **Caution** Configuring a serial number that does not match the specified Fabric Extender forces the Fabric Extender offline.
Step 9	(Optional) **no serial** **Example:** `switch(config-fex)# no serial`	Deletes the serial number string.

Enabling DSCP to Queue Mapping

For Cisco NX-OS Release 6.2(2) and later releases, the Cisco Fabric Extenders support Layer 3 protocol adjacencies on host interfaces (HIFs) and DSCP to queue mapping. Before Cisco NX-OS Release 6.2(2), you can configure a Fabric Extender (FEX) port as a Layer 3 interface for host connectivity, but not for routing. See the Configuring the Cisco Nexus 2000 Series Fabric Extender for more information about fabric extenders.

Before Cisco NX-OS 6.2(2), the Fabric Extender cannot participate in a routing protocol adjacency with a device attached to its port. Only a static direct route is supported. This restriction applies to both of the following supported connectivity cases:

- Switch virtual interfaces (SVI) with Fabric Extender single port or portchannel in Layer 2 mode.
- Fabric Extender port or portchannel in Layer 3 mode.

SUMMARY STEPS

1. **configure terminal**
2. **hardware qos dscp-to-queue ingress module type** {**all** | **f-series** | **m-series**}
3. (Optional) **show hardware qos dscp-to-queue ingress**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch(config)# configure terminal`	Enters global configuration mode.
Step 2	**hardware qos dscp-to-queue ingress module type** {**all** | **f-series** | **m-series**} **Example:** `switch(config)# hardware qos dscp-to-queue ingress module type m-series`	Enables the dscp-to-queue mapping on the specified module(s).
Step 3	(Optional) **show hardware qos dscp-to-queue ingress** **Example:** `switch(config)# show hardware qos dscp-to-queue ingress`	Displays information about the status of dscp-to-queue mapping in ingress direction.
Step 4	(Optional) **copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	Saves the running configuration to the startup configuration.

Configuration Examples

This section contains examples of FEX configurations.

Configuring a FEX with a Layer 3 Host Interface

This example shows how to configure a Fabric Extender with a Layer 3 host interface (at the interface level, subinterface level, port channel level, and port channel subinterface level):

Layer 3 Configuration	Layer 3 Subinterface Configuration

`config t`
`interface ethernet 101/1/1`
`no switchport`
`ip address 192.0.1.1/24`
`Mtu 9000`
`no shutdown`
	`config t`
	`interface ethernet 101/1/1.12`
	`ip address 192.0.2.1/24`
	`encapsulation dot1Q 12`
	`mtu 850`
	`no shutdown`

Layer 3 Host Interface Port Channel Configuration	Layer 3 Host Interface Port Channel Subinterface Configuration
config t interface ethernet 101/1/1-2 no switchport channel-group 12 no shutdown interface port-channel 12 ip address 192.0.3.1/24 mtu 2000 no shutdown	config t interface ethernet 101/1/1-2 no switchport channel-group 12 no shutdown interface port-channel 12.14 ip address 192.0.4.1/24 encapsulation dot1Q 14 mtu 1700 no shutdown

Note The VLAN used in the Layer 3 host interface (HIF) or host interface port channel (HIFPC) subinterface has only local significance to its parent interface. The same VLAN ID can be reused between Layer 3 subinterfaces in the same switch or VDC.

Configuring a Host Interface in a vPC Topology Connected to Two FEXs

This example shows how to configure a host vPC with a FEX (host vPC attached to two different FEXs):

Note The **switchport trunk allowed vlan** *vlan-list* command is not supported on FEX fabric interfaces.

Switch 1 Configuration	Switch 2 Configuration

config t
feature lacp
int e101/1/1-2
channel-group 12 mode active
no shutdown
 	
int port-channel12
switchport
  switchport mode trunk
  switchport trunk allowed vlan 1-20
vpc 10
	config t
	feature lacp
	int e101/1/1-2
	channel-group 12 mode active
	no shutdown
	
	int port-channel12
	switchport
	  switchport mode trunk
	  switchport trunk allowed vlan 1-20
	vpc 10

Dual-Homing of a Server to a FEX with FabricPath

Note Cisco Nexus 7000 Series switches do not support dual-homing.

To configure FabricPath interfaces for two switches that are connected with FabricPath, perform the following tasks on each switch:

- Enable FabricPath on each switch.

- Configure the interfaces that you want to designate as FabricPath interfaces.
- Set the STP priority device to 8192 on all FabricPath Layer 2 gateway devices.
- (Optional) Set the STP domain ID for each of the separate STP domains that are connected to the FabricPath network.
- (Optional) Configure a FEX switch ID.

To configure FabricPath interfaces, follow these steps:

1. (Optional) Enable FabricPath on each switch.

```
switch# config terminal
switch(config)# feature fabricpath
switch(config-lldp)# exit
switch(config)#
```

2. After you enable FabricPath on the switch, configure the specified interface as FabricPath interfaces.

```
switch(config)# interface ethernet 1/2
switch(config-if)# switchport mode fabricpath
switch(config-if)# exit
switch(config)#
```

3. Configure the STP priority for all Rapid PVST+ VLANs as 8192.

```
switch# config terminal
switch(config)# spanning-tree vlan 11-20 priority 8192
switch(config)#
```

4. Configure the STP priority for all MST instances as 8192.

```
switch# config terminal
switch(config)# spanning-tree mst 1-5 priority 8192
switch(config)#
```

5. (Optional) Configure the STP domain ID on each FabricPath Layer 2 gateway switch attached to the FabricPath network.

```
switch# config terminal
switch(config)# spanning-tree domain 5
switch(config)
```

6. (Optional) Configure the FEX switch ID.

Note See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for information on configuring FEX.

Note FEX VPC+ configurations are supported only on F2-Series modules.

If you are setting up an initial FEX VPC+ configuration on an F2-Series module, follow these steps:

1. In the VPC domain configuration mode, enable partial DF mode with the **fabricpath multicast load-balance** command.
2. If disabled, enable TRILL style mac-address learning with the **mac address-table core-port-learning** command.

3. In the VPC domain configuration mode, configure the emulated switch ID with the **fabricpath switchid** *switchid#* command.
4. On each of the VPC/VPC+ peer link interfaces in interface configuration mode, enter the**switchport mode fabricpath** command.
5. On each VPC/VPC+ peer link port channel, enter the **VPC peer-link** command.
6. Configure the VPC ID with the **vpc vpcid** command.

If you are changing an existing FEX VPC configuration to a FEX VPC+ configuration on an F2-Series module, follow these steps:

1. In the VPC domain configuration mode, enable partial DF mode with the **fabricpath multicast load-balance** command.
2. If disabled, enable trill style mac-address learning with the **mac address-table core-port-learning** command.
3. In the VPC domain configuration mode, configure the emulated switch ID with the **fabricpath switchid** *switchid#* command.

7. Copy the configuration.

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

Configuring a FEX in Active-Active Mode

This example shows how to configure a FEX in active-active (FEX-AA) mode:

Switch 1 Configuration	Switch 2 Configuration
int eth 1/1 channel-group 1 int po1 switchport switchport mode fex-fabric fex associate 101 vpc 10	int eth 2/2 channel-group 2 int po2 switchport switchport mode fex-fabric fex associate 101 vpc 10

Note

- The chassis number (fex-id) for both switches must be the same.
- Configuring FEX-AA across two VDCs on the same chassis is not supported.
- FEX-AA does not support Layer 3 interfaces.

Verifying the Status of DSCP-to-queue Mapping

The following sample output from the **show hardware qos dscp-to-queue ingress** command displays the status of DSCP-to-queue mapping enabled in ingress direction on F-series modules:

```
Switch# show hardware qos dscp-to-queue ingress

status: Enabled
module_type : f-series
```

Verifying the Configuration

This section describes how to display the configuration of the Fabric Extender and verify the chassis hardware status.

Verifying the Fabric Extender Configuration

Use the following commands to display configuration information about the defined interfaces on a Fabric Extender:

Command or Action	Purpose
show fex [*FEX-number*] [**detail**]	Displays information about a specific Fabric Extender or all attached units.
show interface *type number* **fex-intf**	Displays the Fabric Extender ports that are pinned to a specific switch interface.
show interface fex-fabric	Displays the switch interfaces that have detected a Fabric Extender uplink.
show interface ethernet *number* **transceiver** [**fex-fabric**]	Displays the SFP+ transceiver and diagnostic optical monitoring (DOM) information for the Fabric Extender uplinks.
show feature-set	Displays the status of the feature sets on the device.

Configuration Examples for the Fabric Extender

This example shows how to display all the attached Fabric Extender units:

This example shows how to display the detailed status of a specific Fabric Extender:

This example shows how to display the Fabric Extender interfaces pinned to a specific switch interface:

This example shows how to display the switch interfaces that are connected to a Fabric Extender uplink:

This example shows how to display the SFP+ transceiver and diagnostic optical monitoring (DOM) information for Fabric Extender uplinks for an SFP+ transceiver that is plugged into the parent switch interface:

```
switch# show interface ethernet 1/40 transceiver
Ethernet1/40
    sfp is present
    name is CISCO-MOLEX INC
    part number is 74752-9026
    revision is A0
```

```
    serial number is MOC13321057
    nominal bitrate is 12000 MBits/sec
    Link length supported for copper is 3 m(s)
    cisco id is --
    cisco extended id number is 4
```

This example shows how to display the SFP+ transceiver and DOM information for Fabric Extender uplinks for an SFP+ transceiver that is plugged into the uplink port on the Fabric Extender:

```
switch# show interface ethernet 1/40 transceiver fex-fabric
Ethernet1/40
    sfp is present
    name is CISCO-MOLEX INC
    part number is 74752-9026
    revision is A0
    serial number is MOC13321057
    nominal bitrate is 12000 MBits/sec
    Link length supported for 50/125mm fiber is 0 m(s)
    Link length supported for 62.5/125mm fiber is 0 m(s)
    cisco id is --
    cisco extended id number is 4
```

Verifying the Chassis Management Information

Use the following to display configuration information used on the switch supervisor to manage the Fabric Extender.

Command or Action	Purpose
show environment fex {**all** | *FEX-number*} [**temperature** | **power** | **fan**]	Displays the environmental sensor status.
show inventory fex *FEX-number*	Displays inventory information for a Fabric Extender.
show module fex [*FEX-number*]	Displays module information about a Fabric Extender.
show sprom fex *FEX-number* {**all** | **backplane** | **powersupply** *ps-num*} | **all**	Displays the contents of the serial PROM (SPROM) on the Fabric Extender. The unit of the power for the **show sprom** command is displayed in centi-amperes.

Configuration Examples for Chassis Management

This example shows how to display the module information about all connected Fabric Extender units:

This example shows how to display the inventory information about a specific Fabric Extender:

```
switch# show inventory fex 101
NAME: "FEX 101 CHASSIS", DESCR: "N2K-C2248TP-1GE  CHASSIS"
PID: N2K-C2248TP-1GE   , VID: V00 , SN: SSI13380FSM

NAME: "FEX 101 Module 1", DESCR: "Fabric Extender Module: 48x1GE, 4x10GE Supervisor"
PID: N2K-C2248TP-1GE   , VID: V00 , SN: JAF1339BDSK

NAME: "FEX 101 Fan 1", DESCR: "Fabric Extender Fan module"
PID: N2K-C2248-FAN     , VID: N/A , SN: N/A
```

```
NAME: "FEX 101 Power Supply 2", DESCR: "Fabric Extender AC power supply"
PID: NXK-PAC-400W      , VID:  000, SN:      LIT13370QD6
```

This example shows how to display the environment status for a specific Fabric Extender:

```
switch# show environment fex 101


Temperature Fex 101:
-----------------------------------------------------------------------
Module   Sensor     MajorThresh   MinorThres   CurTemp     Status
                    (Celsius)     (Celsius)    (Celsius)
-----------------------------------------------------------------------
1        Outlet-1   60            50           33          ok
1        Outlet-2   60            50           38          ok
1        Inlet-1    50            40           35          ok
1        Die-1      100           90           44          ok


Fan Fex: 101:
-----------------------------------------------------
Fan             Model                Hw         Status
-----------------------------------------------------
Chassis         N2K-C2148-FAN        --         ok
PS-1            --                   --         absent
PS-2            NXK-PAC-400W         --         ok


Power Supply Fex 101:
-------------------------------------------------------------------------------
Voltage: 12 Volts
-------------------------------------------------------
PS  Model                Power       Power     Status
                         (Watts)     (Amp)
-------------------------------------------------------
1   --                       --        --     --
2   NXK-PAC-400W           4.32      0.36     ok


Mod Model                Power     Power       Power     Power       Status
                         Requested Requested   Allocated Allocated
                         (Watts)   (Amp)       (Watts)   (Amp)
--- -------------------  -------   ----------  --------- ----------  ----------
1    N2K-C2248TP-1GE     0.00      0.00        0.00      0.00        powered-up


Power Usage Summary:
--------------------
Power Supply redundancy mode:                redundant

Total Power Capacity                              4.32 W

Power reserved for Supervisor(s)                  0.00 W
Power currently used by Modules                   0.00 W

                                                -------------
Total Power Available                             4.32 W
                                                -------------
```

This example shows how to display the SPROM for a specific Fabric Extender:

```
switch# show sprom fex 101 all
DISPLAY FEX 101 SUP sprom contents
Common block:
 Block Signature : 0xabab
 Block Version   : 3
 Block Length    : 160
 Block Checksum  : 0x1a1e
 EEPROM Size     : 65535
 Block Count     : 3
 FRU Major Type  : 0x6002
 FRU Minor Type  : 0x0
 OEM String      : Cisco Systems, Inc.
 Product Number  : N2K-C2248TP-1GE
 Serial Number   : JAF1339BDSK
 Part Number     : 73-12748-01
 Part Revision   : 11
 Mfg Deviation   : 0
 H/W Version     : 0.103
 Mfg Bits        : 0
 Engineer Use    : 0
 snmpOID         : 9.12.3.1.9.78.3.0
 Power Consump   : 1666
 RMA Code        : 0-0-0-0
 CLEI Code       : XXXXXXXXXTBDV00
 VID             : V00
Supervisor Module specific block:
 Block Signature : 0x6002
 Block Version   : 2
 Block Length    : 103
 Block Checksum  : 0x2686
 Feature Bits    : 0x0
 HW Changes Bits : 0x0
 Card Index      : 11016
 MAC Addresses   : 00-00-00-00-00-00
 Number of MACs  : 0
 Number of EPLD  : 0
 Port Type-Num   : 1-48;2-4
 Sensor #1       : 60,50
 Sensor #2       : 60,50
 Sensor #3       : -128,-128
 Sensor #4       : -128,-128
 Sensor #5       : 50,40
 Sensor #6       : -128,-128
 Sensor #7       : -128,-128
 Sensor #8       : -128,-128
 Max Connector Power: 4000
 Cooling Requirement: 65
 Ambient Temperature: 40

DISPLAY FEX 101 backplane sprom contents:
Common block:
 Block Signature : 0xabab
 Block Version   : 3
 Block Length    : 160
 Block Checksum  : 0x1947
 EEPROM Size     : 65535
 Block Count     : 5
 FRU Major Type  : 0x6001
 FRU Minor Type  : 0x0
 OEM String      : Cisco Systems, Inc.
 Product Number  : N2K-C2248TP-1GE
 Serial Number   : SSI13380FSM
 Part Number     : 68-3601-01
 Part Revision   : 03
```

```
 Mfg Deviation   : 0
 H/W Version     : 1.0
 Mfg Bits        : 0
 Engineer Use    : 0
 snmpOID         : 9.12.3.1.3.914.0.0
 Power Consump   : 0
 RMA Code        : 0-0-0-0
 CLEI Code       : XXXXXXXXXTDBV00
 VID             : V00
Chassis specific block:
 Block Signature : 0x6001
 Block Version   : 3
 Block Length    : 39
 Block Checksum  : 0x2cf
 Feature Bits    : 0x0
 HW Changes Bits : 0x0
 Stackmib OID    : 0
 MAC Addresses   : 00-0d-ec-e3-28-00
 Number of MACs  : 64
 OEM Enterprise  : 0
 OEM MIB Offset  : 0
 MAX Connector Power: 0
WWN software-module specific block:
 Block Signature : 0x6005
 Block Version   : 1
 Block Length    : 0
 Block Checksum  : 0x66
wwn usage bits:
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00 00 00 00 00 00 00
 00 00
License software-module specific block:
 Block Signature : 0x6006
 Block Version   : 1
 Block Length    : 16
```

```
 Block Checksum  : 0x86f
lic usage bits:
 ff ff ff ff ff ff ff ff

DISPLAY FEX 101 power-supply 2 sprom contents:
Common block:
 Block Signature : 0xabab
 Block Version   : 3
 Block Length    : 160
 Block Checksum  : 0x1673
 EEPROM Size     : 65535
 Block Count     : 2
 FRU Major Type  : 0xab01
 FRU Minor Type  : 0x0
 OEM String      : Cisco Systems Inc   NXK-PAC-400W
 Product Number  : NXK-PAC-400W
 Serial Number   :      LIT13370QD6
 Part Number     :         341
 Part Revision   : -037
 CLEI Code       : 5-01     01  000
 VID             :  000
 snmpOID         : 12336.12336.12336.12336.12336.12336.12374.12336
 H/W Version     : 43777.2
 Current         : 36
 RMA Code        : 200-32-32-32
Power supply specific block:
 Block Signature : 0x0
 Block Version   : 0
 Block Length    : 0
 Block Checksum  : 0x0
 Feature Bits    : 0x0
 Current 110v    : 36
 Current 220v    : 36
 Stackmib OID    : 0
```

Additional References

This section includes additional information that is related to configuring the Cisco Nexus 2000 Series Fabric Extender.

Related Documents

Related Topic	Document Title
Cisco NX-OS Licensing	*Cisco NX-OS Licensing Guide*
Virtual device contexts (VDC)	*Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide*
Interface configuration	*Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*
Command reference	*Cisco Nexus 7000 Series Command References* available at the following URL: http://www.cisco.com/en/US/products/ps9402/prod_command_reference_list.html

Feature History

This table lists the release history for this feature.

Table 3: Feature History for the Cisco Nexus 2000 Series Fabric Extender

Feature Name	Releases	Feature Information
vPC configuration synchronization	7.2(0)D1(1)	Added support for vPC configuration synchronization for FEX
Cisco Nexus 7000 Series FEX Fabric Active-Active	7.2(0)D1(1)	Added support for FEX Active-Active.
Feature set commands	5.1(1)	The Fabric Extender is enabled on the Cisco Nexus 7000 Series device with the **install feature-set** and **feature-set** commands.
Port channel and vPC support	5.2(1)	Support was added for port channels and the integration of vPC on the host interfaces.
Layer 3 support	5.2(1)	Layer 3 capability was added to the Fabric Extender host interface ports including IPv4 and IPv6 and IGMP snooping.

This table lists the supported line cards and FEXs.

Table 4: Supported Line cards and FEXs History

FEX and Line card	Releases
Cisco Nexus B22 Fabric Extender for Dell is now supported with F4-Series modules.	8.4(1)
Cisco Nexus B22 Fabric Extender for Dell is now supported with F3-Series and M3-Series (N77-M348XP-23L and N77-M324FQ-25) modules.	8.2(1)
Cisco Nexus B22 Fabric Extender for IBM Cisco Nexus 2348UPQ Fabric Extender Cisco Nexus 2348TQ Fabric Extender	7.2(0)D1(1)
Cisco Nexus B22 Fabric Extender for HP Cisco Nexus 2248PQ Fabric Extender	6.2(2)

FEX and Line card	Releases
48-port, 100/1000 BASE-T (100-Mb/1-Gigabit) N2248TP-E 32-port 10-Gigabit SFP+ N2248TP-E Support for M2 series modules Cisco Nexus 2224TP Fabric Extender Cisco Nexus 2232TM Fabric Extender	6.1(1)
48-port, 1/10-Gigabit F2-Series module (N7K-F248XP-25).	6.0(1)
Cisco Nexus 2248TP Fabric Extender Cisco Nexus 2232PP Fabric Extender	5.2(1)
32-port, 10-Gigabit M1 module (N7K-M132XP-12) 32-port, 10-Gigabit M1 XL module (N7K-M132XP-12L).	5.1(1)

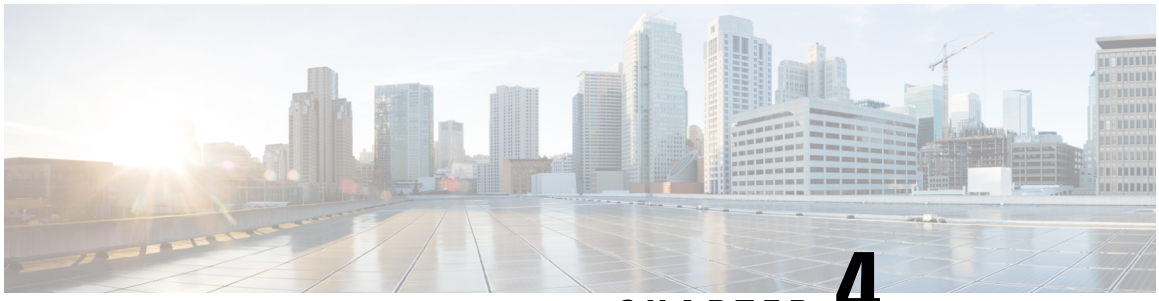

CHAPTER 4

vPC Configuration Synchronization

Virtual port channels (vPC) topologies require identical configurations on peer switches. As a result, you must repeat configurations on both peer switches. This process, which can cause errors due to misconfigurations or omissions, can result in additional service disruptions because of mismatched configurations. Configuration synchronization eliminates these problems by allowing you to configure one switch and automatically synchronize the configuration on the peer switch.

In a vPC topology, each Cisco Nexus 7000 Series switch must have some matching parameters. You can use a vPC consistency check to verify that both Cisco Nexus 7000 Series switches have the same configuration (Type 1 or Type 2). If they do not match, depending on whether it is a global (for example, spanning-tree port mode), a port-level (for example, speed, duplex, or channel-group type), or even a port-channel interface, the vPC can go into a suspended state or a VLAN can go into a blocking state on both peer switches. As a result, you must ensure that the configuration from one switch is copied identically to the peer switch.

Configuration synchronization allows you to synchronize the configuration between a pair of switches in a network. Configuration synchronization and vPCs are two independent features and configuration synchronization does not eliminate vPC consistency checks. The checks will continue. If there is a configuration mismatch, the vPC can still go into a suspended state.

In a FEX Active-Active setup:

- All the Host Interfaces (HIFs) ports are mapped to the internal vPC.
- The vPC Config-Sync feature listens to the internal vPC creation notification and triggers a merge of the HIF port configuration.
- All the future HIF configuration are synchronized with the peer switch, if the merge is successful.
- The status of HIF is marked as "peer out of synchronization" and the configuration of the interface is not synchronized, if the merge fails.
- We recommend that you disable **vpc-config-sync** command before starting ASCII configuration. After the ASCII configuration is completed, enable **config-sync** command for regular operation.

Note

- vPC peer-link should be configured and up state.
- You cannot chose which commands are synchronized.

- Prerequisites for vPC Configuration Synchronization, on page 52

• Guidelines and Limitations for Configuration Synchronization, on page 52
• Information about vPC Configuration Synchronization, on page 52
• How to configure vPC Config Sync, on page 53
• Feature Information for vPC Configuration Synchronization, on page 58

Prerequisites for vPC Configuration Synchronization

- vPC peer-link should be configured and up state.

Guidelines and Limitations for Configuration Synchronization

- You cannot chose which commands are synchronized.

Information about vPC Configuration Synchronization

Benefits of vPC Configuration Synchronization

Configuration synchronization benefits are as follows:

- Provides a mechanism to synchronize configuration from one switch to another switch.
- Merges configurations when connectivity is established between peers.
- Provides mutual exclusion for commands.
- Supports existing session and port profile functionality.
- Provides minimal user intervention.
- Minimizes the possibility of user error.

Supported Commands for vPC Configuration Synchronization

The following types of commands are enabled for configuration synchronization:

Note The **show vpc config-sync cli syntax** command lists all the commands that are enabled for configuration synchronization. You cannot choose which commands are synchronized. For more information, see the Cisco Nexus 7000 Series NX-OS Interfaces Command Reference.

- Type-1 configurations:
  - Global configurations
  - vPC member port-channel configurations

- vPC configurations.

Note The configurations can be given on either of the vPC peer switches.

How to configure vPC Config Sync

Enabling vPC Configuration Synchronization

Before you begin

- You must create identical vPC domain IDs on both vPC peer switches.
- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vpc domain** *domain-id*
3. switch(config-vpc-domain)# **config-sync**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **config-sync**	Enables vPC configuration synchronization. **Note** This command must be configured on both the primary and secondary switch.

The table below shows the process of configuration synchronization on switch 1 and switch 2:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# vpc domain 300`
`switch-1(config-vpc-domain)# config-sync`
	`switch-2# configure terminal`
	`switch-2(config)# vpc domain 300`
	`switch-2(config-vpc-domain)# config-sync`

Configuration synchronization is enabled on both switches in the same vPC domain.	

Primary Switch	Secondary Switch
`switch-1# configure terminal` `switch-1(config)# spanning-tree mode mst`	
The above configuration is applied on the primary switch and is configuration synchronized to the secondary switch. The configuration is either successfully applied to both switches or will be failed on both.	

`switch-1# show running-config`
`...`
`spanning-tree mode mst`
`...`
	`switch-2# show running-config`
	`...`
	`spanning-tree mode mst`
	`...`

	`switch-2# configure terminal` `switch-2(config)# spanning-tree port type switch-2 default`
The configuration is applied on the secondary switch and is configuration synchronized to the primary switch. **Note** The configuration can be applied to either switch.	

`switch-1# show running-config`
`...`
`spanning-tree port type network default`
`...`
	`switch-2# show running-config`
	`...`
	`spanning-tree port type network default`
	`...`

Synchronizing Configuration for a Physical Port vPC

Before you begin

Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot/port*
3. switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]
4. (Optional) switch(config-if)# **show running-config interface ethernet** *slot/port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot/port*	Specifies the vPC physical port, and enters interface configuration mode.

	Command or Action	Purpose
Step 3	switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]	Moves port channel into a vPC and enters interface vPC configuration mode. The range is from 1 to 4096. • **sync export** enables the primary switch configuration to be exported to the secondary switch. • **sync import** enables the secondary switch configuration to be imported to primary switch.
Step 4	(Optional) switch(config-if)# **show running-config interface ethernet** *slot/port*	Displays the running configuration for the physical port.

Asymmetric Mapping

The table below shows the process of enabling configuration synchronization (asymmetric mapping) on the vPC physical port on the primary and the secondary switch:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# interface eth1/1`
`switch-1(config-if)# vpc 100`	

The physical port (ethernet1/1) is added to the vPC 100 domain on the primary switch. vPC 100 is not configured on the secondary switch. The configuration will not be synchronized until vPC 100 is added to the secondary switch.	

	`switch-2# configure terminal`
	`switch-2(config)# interface eth2/3`
	`switch-2(config-if)# vpc 100`

Following the configuration of vPC 100 to the secondary switch, the physical ports (interface ethernet2/3 on the secondary switch and interface ethernet1/1 on the primary switch) will be configuration synchronized.	

Symmetric Mapping

The table below shows the process of enabling configuration synchronization (symmetric mapping) on the vPC physical port on the primary and the secondary switch:

Primary switch	Secondary switch

`switch-1# configure terminal`
`switch-1(config)# interface eth1/1`
`switch-1(config-if)# vpc 100 symmetric`
	`switch-2# configure terminal`
	`switch-2(config)# interface eth1/1`

The physical port (ethernet1/1) is added to the vPC 100 domain on the primary switch. The physical port (ethernet 1/1) is also present on the secondary switch. The configuration of the physical port on both the primary and secondary switch will be kept in synchronization.	

Primary switch	Secondary switch

`switch-1# show running-config interface eth1/10`
 	
`interface ethernet1/1`
`  switchport`
`  switchport mode trunk`
`  vpc 100`
	`switch-2# show running-config interface eth1/10`
	
	`interface ethernet1/1`
	`  switchport`
	`  switchport mode trunk`
	`  vpc 100`

Synchronizing Configuration of vPC Member Port Channel

Before you begin

Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface port-channel** *channel-number*
3. switch(config-if)# **switchport**
4. switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]
5. (Optional) switch(config-if)# **show running-config interface port-channel** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	Selects the port channel that you want to use as the vPC peer link for this device, and enters interface configuration mode.
Step 3	switch(config-if)# **switchport**	Configures the interface as a Layer 2 switching port.
Step 4	switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]	Moves port channel into a vPC and enters interface vPC configuration mode. The range is from 1 to 4096. • **sync export** enables the primary switch configuration to be exported to the secondary switch. • **sync import** enables the secondary switch configuration to be imported to primary switch.
Step 5	(Optional) switch(config-if)# **show running-config interface port-channel** *channel-number*	Displays the running configuration for the port channel.

The table below shows the process of enabling configuration synchronization under port channel 10 on the primary and the secondary switch:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# interface port-channel 10`
`switch-1(config-if)# switchport`
`switch-1(config-if)# vpc 10`	

The configuration under port-channel 10 is configuration synchronized to the secondary switch. **Note** The **vpc** *number* command can be given first on either the primary or secondary switch.	

	`switch-2# show running-config interface po10`
	
	`interface port-channel10`
	`  switchport`
	`  vpc 10`

The configuration is applied on the secondary switch and is configuration synchronized to the primary switch. **Note** The configuration can be applied to either switch.	

	`switch-2# configure terminal`
	`switch-2(config)# interface port-channel 10`
	`switch-2(config-if)# switchport mode trunk`

The **show running-config interface port-channel** *channel-number* command shows that the configuration synchronization for port channel 10 is successful:	

`switch-1# show running-config interface port-channel 10`
 	
`interface port-channel10`
`  switchport`
`  switchport mode trunk`
`  vpc 10`
	`switch-2# show running-config interface port-channel 10`
	
	`interface port-channel10`
	`  switchport`
	`  switchport mode trunk`
	`  vpc 10`

Verifying vPC Configuration Synchronization

To verify vPC configuration synchronization, perform one of the following tasks:

Command	Purpose
show running-config vpc-config-sync	Displays whether config-sync is available or not.
show vpc config-sync cli syntax	Displays the list of commands that are able to be configuration synchronized.
show vpc config-sync database	Displays the configuration synchronization database.
show vpc config-sync merge status	Displays the merge status of the switch and of each vPC interface.

Command	Purpose
show vpc config-sync status	Displays the status of the last 10 operations of the vPC configuration synchronization process. • Displays merge status (success/failure). • Displays the last action done by the vPC configuration synchronization process and the result of that action.

Feature Information for vPC Configuration Synchronization

Table 5: Feature Information for vPC Configuration Synchronization

Feature Name	Releases	Feature Information
vPC Configuration Synchronization	7.2(0)D1(1)	vPC Configuration Synchronization feature synchronizes the configurations of one switch automatically to other similar switches. .

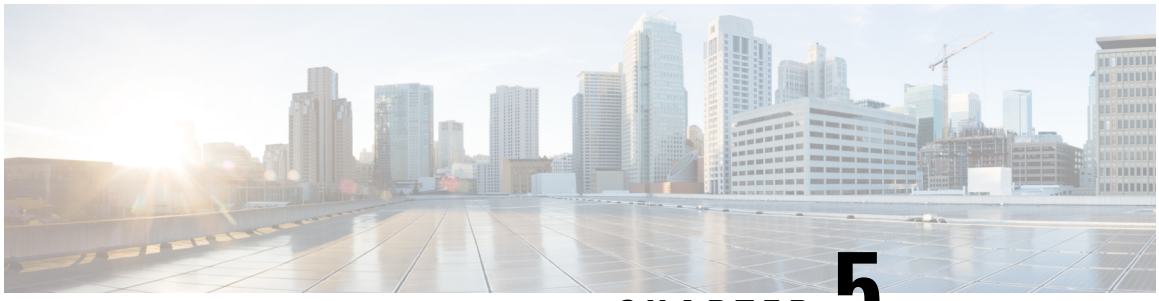

CHAPTER 5

FEX Queuing

FEX Queuing provides a mechanism to support queuing on host interfaces (HIF) of a FEX device. Queuing of data traffic is based on the COS or DSCP values of an Ethernet packet. Traffic that is not marked with these values are dropped to a default queue.

- Finding Feature Information, on page 59
- Information About FEX Queuing, on page 59
- How to Configure FEX Queuing, on page 60
- Feature Information for FEX Queuing, on page 65

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About FEX Queuing

Queuing Ethernet Frames Using COS

With this feature, data and control Ethernet frames can be prioritized based on the COS value of an Ethernet header and sent to separate queues of a FEX host interface. This ensures that control frames are not dropped during a traffic burst.

Every FEX interface has a COS2Q map associated with it. And depending on the COS2Q map, the ethernet frame is enqueued.

Queuing Ethernet Frames Using DSCP

With this feature, data and control Ethernet frames can be prioritized based on the DSCP value of an IP header and sent to separate queues of a FEX host interface. This ensures that control frames are not dropped during a traffic burst.

Queuing is based on the DSCP to Queue mapping configuration on the network-qos template

Queueing FCoE Frames Using COS 3

With this feature, FCoE and Ethernet frames can be prioritized based on the COS 3 of FCoE and COS value of Ethernet header and sent to separate queues of a FEX host interface. This ensures that control frames are not dropped during a traffic burst.

This is driven by the COS2Q mapping and the network-qos template on the Nexus 7000.

How to Configure FEX Queuing

Changing COS2Q mapping

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **class-map type queuing match-any** *queue*
4. **match cos 4-7**
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**class-map type queuing match-any** *queue* **Example:** `Device(config)# class-map type queuing match-any 2q4t-8e-in-q1`	Configures the global queuing class map.

	Command or Action	Purpose
Step 4	**match cos 4-7** **Example:** `Device(config-cmap-que)# match cos 4-7`	Configures COS2Q value
Step 5	**end** **Example:** `Device(config-cmap-que)# end`	Exits to previliged EXEC mode.

Changing DSCP2Q Mapping

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **class-map type queuing match-any 2q4t-8e-in-q-default**
4. **match dscp** *value-range*
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**class-map type queuing match-any 2q4t-8e-in-q-default** **Example:** `Device(config)# class-map type queuing match-any 2q4t-8e-in-q-default`	Configures the global queuing class map.
Step 4	**match dscp** *value-range* **Example:** `Device(config-cmap-que)# match dscp 1-10`	Configures DSCP2Q value
Step 5	**end** **Example:** `Device(config-cmap-que)# end`	Exits to previliged EXEC mode.

Changing network-qos Template

The cos2q map, bandwidth, mtu and priority for the queues are defined for each network-qos templates. The FEX is configured based on the default values for a given network-qos template. To change these default values, you should clone a template and change the values. The policy can then be applied to the system qos.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **system qos**
4. **service-policy type network-qos** *policy*
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**system qos** **Example:** `Device(config)# system qos`	Enters system class configuration mode.
Step 4	**service-policy type network-qos** *policy* **Example:** `Device(config-sys-qos)# service-policy type network-qos default-nq-7e-4q8q-policy`	Configures the policy map that is to be used as the network-wide service policy, and enters Network-wide (system qos) mode.
Step 5	**end** **Example:** `Device(config-sys-qos)# end`	Exits to previliged EXEC mode.

Configuring FEX Queue Parameters

SUMMARY STEPS

1. **hardware** *fex-card-type* **shared-buffer-size** *size*
2. **hardware** *fex-card-type* **queue-limit** *queue-limit*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**hardware** *fex-card-type* **shared-buffer-size** *size* **Example:** `Device# hardware N2248PQ shared-buffer-size 3072`	This command is applicable for N2248PQ only. The range is from 3072 to 10240. The command is hosted on default-vdc and admin-vdc.
Step 2	**hardware** *fex-card-type* **queue-limit** *queue-limit* **Example:** `Device# hardware N2248T queue-limit 327680`	The range is from 81920 to 652800 for a Cisco Nexus 2148T Fabric Extender and from 2560 to 652800 for all other supported Fabric Extenders. The command is hosted on default-vdc and admin-vdc.

Verifying FEX Queuing

SUMMARY STEPS

1. **show queuing interface** *fex-interface value/slot*

DETAILED STEPS

show queuing interface *fex-interface value/slot*

Example:

```
Device# show queuing interface ethernet 101/1/1

    invalid interface

    slot  4
    =======

    Interface is not in this module.

    slot  9
    =======

    Ethernet101/1/1 queuing information:
      Input buffer allocation:
      Qos-group: ctrl
      frh: 0
      drop-type: drop
      cos: 7
      xon       xoff
      buffer-size

      ---------+---------+-----------
      2560      7680      10240

      Qos-group: 0  2
      (shared)
      frh: 8
      drop-type: drop
      cos: 0 1 2 3 4 5 6
      xon       xoff
```

```
buffer-size

---------+---------+-----------
0         142080    151040

Queueing:
queue    qos-group    cos            priority  bandwidth mtu

--------+------------+--------------+---------+---------+----
ctrl-hi  n/a          7               PRI         0
2400
ctrl-lo  n/a          7               PRI         0
2400
2        0            0 1 2
  3         WRR        80
1600
4        2            4 5
  6           WRR        20
1600

Queue limit: 66560 bytes

Queue Statistics:
queue
  rx              tx              flags

------+---------------+---------------+-----
0      0
  0
ctrl
1      0
  0
ctrl
2      0
  0
data
4      6               0               data

Priority-flow-control enabled:
no
Flow-control status: rx 0x0, tx 0x0,
rx_mask 0x0
cos     qos-group   rx pause  tx pause  masked
rx pause

-------+-----------+---------+---------+---------------
0
             0    xon
  xon       xon
1              0    xon       xon
xon
2              0    xon       xon
xon
3              0    xon       xon
xon
4              2    xon       xon
xon
5              2    xon       xon       xon
6              2    xon       xon
xon
7            n/a    xon       xon
xon

DSCP to Queue mapping on
```

```
 FEX
----+--+-----+-------+--+---

queue           DSCPs
-----           -----
02              0-9,14-25,27-38,40-51,53-63,
04              10-13,26,39,52,
03              **EMPTY**
05              **EMPTY**


slot 10
=======
```

Feature Information for FEX Queuing

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 6: Feature Information for FEX Queuing

Feature Name	Releases	Feature Information
FEX Queuing	7.2(0)D1(1)	The FEX Queuing The following commands were introduced by this feature: **hardware shared-buffer-size, hardware queue-limit**.

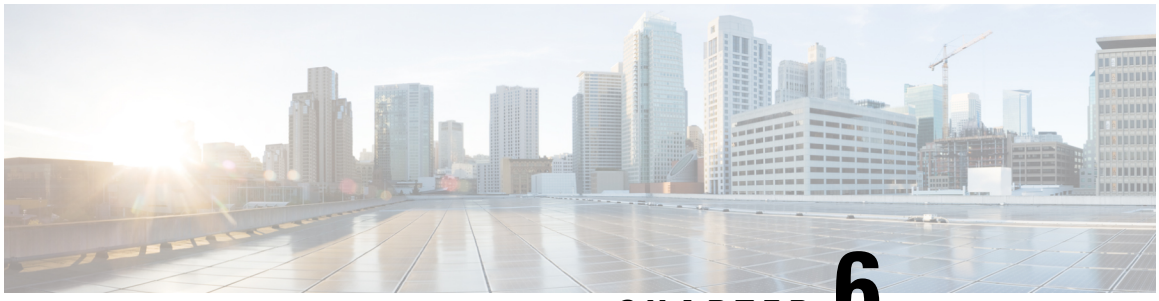

CHAPTER 6

Supported FEX Devices

Support has been added for the following FEX devices:

- 2348UPQ
- 2348TQ
- B22 IBM
- B22 DELL

• Feature Information for Supported FEX Devices, on page 67

Feature Information for Supported FEX Devices

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 7: Feature Information for Supported FEX Devices

Feature Name	Releases	Feature Information
Support added for B22 Dell FEX on F4-Series modules	8.4(1)	B22 Dell FEX is now supported with F4-Series modules
Support added for B22 Dell FEX on F3-Series and M3-Series I/O modules	8.2(1)	B22 Dell FEX is now supported with F3-Series and M3-Series I/O modules
Support added for 2348UPQ, 2348TQ and B22 IBM devices	7.2(0)D1(1)	Support has been added for 2348UPQ, 2348TQ and B22 IBM devices

INDEX

A

ACL support 9
associating fabric extender 29

B

BPDU Guard 6, 11

C

chassis configuration mode 35
chassis ID 13
class of service, *See* CoS
configuration data 11
CoS 9

D

Data Center Bridging Exchange, *See* DCBX
DCBX 9
description 35
digital optical monitoring, *See* DOM
documentation viii
    additional publications viii
DOM 5

E

edge port (PortFast) 6
Ethernet fabric interface 4

F

fabric extender 14
    guidelines 14
    licensing requirements 14
    limitations 14
fabric interface 4
fabric interface Port Channel 12
fail-over load balancing 12
FEX-number 13

G

guidelines 14

H

host interface 4
host interface autonegotiation 9

I

IEEE 802.1p 9
IEEE 802.3x 9
IGMP snooping 9
image management 14

L

LACP 9
layer 3 ports 5
licensing 14
    fabric extender 14
limitations 14
Link Aggregation Control Protocol, *See* LACP
Link Layer Discovery Protocol, *See* LLDP
LLDP 9
local switching 11
loopback address assignment 11
loopback address range 11

M

multicast replication 11

O

oversubscription 10
oversubscription ratio 10

P

packet counter 6
PFC 5
port channel 12

port channel fabric interface 4, 5
port channel host interface 4
port numbering 13
priority flow control, *See* PFC

Q

QoS 9
quality-of-service, *See* QoS

R

related documents viii, 47

S

serial number 35
SFP+ validation 5
show diagnostics 43
show environment 43
show fex 42
show inventory 43
show modules 43
show SPROM 43
show transceiver status 42
SPAN restrictions 10
SPAN source ports 10
subinterfaces 5
switchport fex-fabric mode 5
switchport saved configuration 5

T

type 35

V

version compatibility 14
VLANs 8