



Cisco Nexus 2000 Series NX-OS Fabric Extender Configuration Guide for Cisco Nexus 9000 Series Switches, Release 10.2(x)

First Published: 2021-08-24

Last Modified: 2022-10-28

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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

© 2021–2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface	vii
Audience	vii
Document Conventions	vii
Related Documentation for Cisco Nexus 9000 Series Switches	viii
Documentation Feedback	viii
Communications, Services, and Additional Information	viii

CHAPTER 1

New and Changed Information	1
New and Changed Information	1

CHAPTER 2

Overview	3
Licensing Requirements	3
Supported Platforms	4
Information About the Cisco Nexus 2000 Series Fabric Extender	4
Fabric Extender Terminology	4
Fabric Interface Features	5
Host Interfaces	5
Layer 2 Host Interfaces	5
Host Interface Port Channels	6
Layer 2 Host Interface Port Channels	6
Load Balancing Using Host Interface Port Channels	6
VLANs	7
FCoE Over FEX	7
Protocol Offload	7
Access Control Lists	7
IGMP Snooping	8

- Switched Port Analyzer 8
- Oversubscription 8
- Management Model 11
- Forwarding Model 11
- Port Channel Fabric Interface Connection 12
- Port Numbering Convention 13
- Fabric Extender Image Management 13
- Guidelines and Limitations for the Fabric Extender 14
- Configuration Limits 19
- Default Settings 19

CHAPTER 3

Configuring the Fabric Extender 21

- Managing the Fabric Extender Feature Set 21
 - Installing the Fabric Extender Feature Set 21
 - Enabling the Fabric Extender Feature Set 22
 - Associating a Fabric Extender to a Fabric Interface 23
 - Associating a Fabric Extender to a Port Channel 23
 - Disassociating a Fabric Extender from an Fabric Interface 24
 - Configuring Fabric Extender Global Features 25
 - Configuration Examples 27
 - Configuring a Host Interface in a vPC Topology Connected to Two FEXs 27
 - Switch vPC Topology and Straight Through FEX Topologies (Host vPC) 27
 - Dual-Homed FEX Topology (Active-Active FEX Topology) 34
- Verifying the Configuration 42
 - Verifying the Fabric Extender Configuration 42
 - Verifying the Chassis Management Information 45
- Additional References 49
 - Related Documents 49

CHAPTER 4

Software FEX Mode Configuration 51

- Software FEX Mode Configuration 51

CHAPTER 5

Upgrading Procedures 53

- Upgrade Process for a vPC Topology on the Primary Switch 53

Upgrade Process for a vPC Topology on the Secondary Switch 54

CHAPTER 6**Minimizing the Impact of a Disruptive Upgrade 55**

Upgrading a Direct vPC or a Single-Homed FEX Access Layer 55

Upgrading a Dual-Homed FEX Access Layer 57



Preface

This preface includes the following sections:

- [Audience, on page vii](#)
- [Document Conventions, on page vii](#)
- [Related Documentation for Cisco Nexus 9000 Series Switches, on page viii](#)
- [Documentation Feedback, on page viii](#)
- [Communications, Services, and Additional Information, on page viii](#)

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which you supply the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments that are separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments that are 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.

Convention	Description
<i>variable</i>	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 includes the quotation marks.

Examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information that you must enter is in boldface screen font.
<i>italic screen font</i>	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.

Related Documentation for Cisco Nexus 9000 Series Switches

The entire Cisco Nexus 9000 Series switch documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps13386/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

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

Cisco Bug Search Tool

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



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

Table 1: New and Changed Features for Cisco NX-OS Release 10.2(x)

Feature	Description	Changed in Release	Where Documented
No new features are introduced for this release.			



CHAPTER 2

Overview

- This chapter provides an architectural overview of the Cisco Nexus 2000 Series Fabric Extender and includes the following sections:
 - [Licensing Requirements, on page 3](#)
 - [Supported Platforms, on page 4](#)
 - [Information About the Cisco Nexus 2000 Series Fabric Extender, on page 4](#)
 - [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 7](#)
 - [FCoE Over FEX, on page 7](#)
 - [Protocol Offload, on page 7](#)
 - [Access Control Lists, on page 7](#)
 - [IGMP Snooping, on page 8](#)
 - [Switched Port Analyzer, on page 8](#)
 - [Oversubscription, on page 8](#)
 - [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 13](#)
 - [Guidelines and Limitations for the Fabric Extender, on page 14](#)
 - [Configuration Limits, on page 19](#)
 - [Default Settings, on page 19](#)

Licensing Requirements

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

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the [Nexus Switch Platform Support Matrix](#) to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

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 with a single management domain. The Fabric Extender and its parent switch enable a large multipath, loop-free 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.

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

Fabric Extender Terminology

Some terms used in this document are as follows:

- Fabric interface—A 10-Gigabit/40-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.

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

Fabric Interface Features

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.

Host Interfaces

Layer 2 Host Interfaces

Beginning with Cisco NX-OS Release 10.2(1), the default port mode for FEX HIF is system default layer for ST mode. The default port mode for AA mode continues to be L2.

To run a host interface in Layer 2 mode, use the **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 any edge switch that leverages a link redundancy mechanism not dependent on spanning tree such as 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 9000 Series NX-OS Layer 2 Switching Configuration Guide*.

Host Interface Port Channels

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.

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.

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 non-IP interfaces is the source and destination MAC address. For more details, see the *Cisco Nexus 9000 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

VLANs

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

For more information about VLANs, see the *Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide*.



Note Configuring a native VLAN on a FEX fabric interface is not supported.

FCoE Over FEX

The Fabric Extender supports Fibre Channel over Ethernet (FCoE) over Fabric Extenders (FEX). This feature enables Fibre Channel traffic to be carried on a FEX port. The FEX is connected to a Cisco Nexus 9000 device that is in FCoE NPV mode through a Fabric Port Channel (FPC). FCoE over FEX enables the provisioning of FCoE on host connections.

For more information, see the *Cisco Nexus 9000 Series NX-OS FCoE Configuration Guide*.

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)

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.

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 IP address. It does not support snooping that is based on the MAC address.



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 9000 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. Up to four SPAN sessions for host interfaces are supported on the same or different Fabric Extenders. Ingress source (Rx) monitoring is 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.



Note Tx monitoring on the FEX host interface (HIF) source is supported only for known Layer2 unicast traffic.



Note An interface that has port ACLs or router ACLs (PACL/RACL) configured with **statistics per-entry** is not supported in a SPAN/ERSPAN session with a configured ACL filter.

For more information about SPAN, see the *Cisco Nexus 9000 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.

The Cisco Nexus B22 Fabric Extender for Dell (NB22DELL) 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 Dell (N2K-B22DELL-P) is 2:1.

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.

The Cisco Nexus B22 Fabric Extender for Dell (NB22DELL) 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 Dell (N2K-B22DELL-P) is 2:1.

Beginning with 7.0(3)I2(1), the Cisco Nexus B22 Fabric Extender for IBM (NB22IBM) has 8 10-Gigabit Ethernet fabric interfaces and 14 1G/10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. Static pinning and port-channel mode are supported. When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus B22 Fabric Extender for IBM (N2K-B22IBM-P) is 2:1.

Beginning with 7.0(3)I2(1), the Cisco Nexus B22 Fabric Extender for Fujitsu (NB22FSTS) has 8 10-Gigabit Ethernet fabric interfaces and 16 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. Static pinning and port-channel mode are supported. When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus B22 Fabric Extender for Fujitsu (N2K-B22FSTS-P) is 2:1.

Beginning with 7.0(3)I2(1), the Cisco Nexus 2348UPQ Fabric Extender has 6 40-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 2348UPQ is 2:1.

Beginning with 7.0(3)I2(1), the Cisco Nexus 2348TQ Fabric Extender has 6 40-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 2348TQ is 2:1.

Cisco Nexus 2348TQ-E Fabric Extender has 48 1-Gigabit Ethernet fabric interfaces 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 2348TQ-E is 2:1.

Cisco Nexus 2332TQ Fabric Extender has 32 1-Gigabit Ethernet fabric interfaces 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 2332TQ 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.
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.

Guidelines and Limitations for the Fabric Extender

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

- The FEX QoS system level queuing policy does not support WRED, queue-limit, shaping, or policing features.
- The FEX QoS system level queuing policy does not support multiple priority levels.
- Before converting a port from trunk to FEX fabric, remove/unconfigure any explicit native VLAN configuration.
- NAT is not supported on the FEX host interfaces.
- The FEX host interface is the system default layer.
- You cannot use N9K-C93108TC-FX3P ports 1–48(10GT) as a parent to connect FEX NIF ports. Instead you must use N9K-C93108TC-FX3P Ports 49–54(40G/100G) as a parent to connect FEX NIF ports.
- Only the 4Q queuing policy model is supported on FEX. When you try to bring up FEX in 8Q queuing policy mode, you get an error message.
- 10G GLC-T optics are not supported on FEX ports.
- Beginning with Cisco NX-OS Release 9.2(1), dual-homed FEX support is added to Cisco Nexus 93180YC-FX, and 93108TC-FX switches in addition to straight-through FEX support.
- Beginning with Cisco NX-OS Release 9.2(1), straight-through FEX support is added to Cisco Nexus 93240YC-FX2 and 9336C-FX2 switches.
- Beginning with Cisco NX-OS Release 9.3(1), straight-through FEX support is added to Cisco Nexus 93360YC-FX2.
- For FEX HIF port channels, enable the STP port type edge using the **spanning tree port type edge [trunk]** command.
- The Cisco Nexus 2248PQ, 2348TQ, 2348TQ-E, and 2348UPQ FEXs support connections to the Cisco Nexus 9300 or 9500 platform switches by using supported breakout cables to connect a QSFP+ uplink on the FEX and an SFP+ link on the parent switch (4x10 G links).
- Beginning with Cisco NX-OS Release 9.3(5), the Active-Active FEX topology is supported on the N9K-C9336C-FX2, N9K-C93240YC-FX2, N9K-C93360YC-FX2, N9K-C93216TC-FX2, N9K-C93108TC-FX3P, N9K-C93180YC-FX3S switches. The Cisco Nexus 9300-FX2 and FX3 switches are supported on the ST and the AA FEX modes.
- Beginning with Cisco NX-OS Release 9.3(1), all FEX types support for N9K-C93360YC-FX2 switch in straight-through mode.
- Beginning with Cisco NX-OS Release 9.2(3), FEX supports IEEE 802.1X port-based authentication on FEX-ST and host interface (HIF) ports. IEEE 802.1X port-based authentication support applies to both straight-through and dual-homed FEX.

For more information about configuring port-based authentication, see the *Cisco Nexus 9000 Series NX-OS Security Configuration Guide, Release 9.x*.

- Beginning with Cisco NX-OS Release 9.2(1), all FEX types support for the N9K-C93180YC-FX and N9K-C93108TC-FX switches in dual-homed mode.
- Beginning with Cisco NX-OS Release 9.2(1), all FEX types support for the N9K-C93240YC-FX2 and N9K-C9336C-FX2 switches in straight-through mode.
- Beginning with Cisco NX-OS Release 9.2(1), FCoE over FEX is supported on N9K-C93180YC-FX switches in both straight-through and dual-homed mode with N2K-C2348UPQ, N2K-C2232PP, N2K-B221BM-P and N2K-B22HP-P FEX models.
- Beginning with Cisco NX-OS Release 9.2(1), NetFlow for FEX Layer 3 ports is now supported on Cisco Nexus 9300-EX and 9300-FX platform switches.
- The configuration is purged when:
 - Straight-through FEXes are converted to dual-homed.
 - Dual-homed FEXes are converted to straight-through.
- Converting from an active-active to straight-through or a straight-through to active-active FEX topology with Cisco Nexus 9000 Series switches requires reloading the parent switch. See also: <https://bst.cloudapps.cisco.com/bugsearch/bug/CSCve15816>
 - While the FEX is online: the FEX goes down as a dual-homed FEX on conversion and comes back up a straight-through FEX. The configuration is purged on bring up.
 - While the FEX is offline: the FEX goes down as a dual-homed FEX, then the **no vpc id** command is entered on the fabric port channel. No configuration purge takes place. In this scenario, default the configuration on FEX interfaces while toggling the mode from active-active to straight-through.
- ASCII/POAP Replay are supported from Release 7.0(3)I7(1) onwards. Earlier releases require manually reapplying the FEX configuration after the FEX is online.
- An upgrade performed over **install all** command for Release 7.0(3)I2(2b) to Release 7.0(3)I6(2) or to Release 7.0(3)I7(x) and later may result in the VLANs being unable to be added to the existing FEX HIF trunk ports. To recover from this, the following steps should be performed after all FEXs have come online and the HIFs are operationally up:
 - Enter the **copy run bootflash:fex_config_restore.cfg** command at the prompt.
 - Enter the **copy bootflash:fex_config_restore.cfg running-config echo-commands** command at the prompt.
- The 2332TQ FEX now supports Cisco Nexus 9300, 9300-EX, and 9500 platform switches as the parent switch (on all FEX supported platforms).
- Beginning with Cisco NX-OS Release 9.2(1), FEX is supported on Cisco Nexus 9500 chassis with N9K-X9432PQ, N9K-X9536PQ, and N9K-X9636PQ line cards in breakout mode.
- Beginning with Cisco NX-OS Release 10.2(1), the default port mode for FEX HIF is the system default layer for ST mode. The default port mode for AA mode remains L2.
- You can configure a maximum of eight ports as part of a fabric port channel (the uplink from the Fabric Extender to the switch).
- The Fabric Extender is supported with Cisco Nexus 9500 Series switches and X9464PX and X9564PX line cards and with Cisco Nexus 9372PX and 9396PX switches.

- The Fabric Extender is supported on the N9K-C93108TC-EX, N9K-C93180LC-EX, N9K-C93180YC-EX, N9K-C93180YC-FX, and N9K-C93108TC-FX switches. Support includes straight-through and dual-homed (active-active) FEX topologies.
- The 2348TQ-E Fabric Extender is supported.
- 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.
- When you connect a FEX to a Cisco Nexus 9000 Series device, the queuing capability on the FEX host interface is limited. A router that is connected to a Layer 2 (using SVI interfaces) cannot participate in routing protocol adjacency. The FEX cannot be used as a peer because when congestion occurs on the FEX host interface, the control plane traffic is not prioritized. This limitation also applies to the FEX when it is connected to other Layer 3 devices, such as an ASA firewall, an ACE load balancer, or other Layer 3 networking devices that are running a dynamic routing protocol. Static routes to routers, ASA firewalls, ACE load balancers, and other Layer 3 network devices are supported.
- For FEX HIF port channels, Cisco recommends that you enable STP port type edge using the **spanning tree port type edge [trunk]** command.
- When you unconfigure a NIF port, you must use the **no fex associate fex-id** command with an existing and correct FEX ID to avoid any error.
- If you configure the FEX with **speed 100/full-duplex** and you do not explicitly configure the neighboring device with **speed 100/full-duplex**, the data packets might not pass properly although the link may appear as being "up".

Interface Configuration	Description
no speed	Autonegotiates and advertises all speeds (only full duplex).
speed 100	Does not autonegotiate; pause cannot be advertised. The peer must be set to not autonegotiate (only 100-Mbps full duplex supported).
speed 1000	Autonegotiates and advertises pause (advertises only for 1000 Mbps full duplex).

- Cisco Nexus 2332TQ, 2348TQ, 2348TQ-E, and 2348UPQ support 40G connectivity or 4x10G breakout.
- Cisco Nexus 2248PQ, 2348TQ, 2348TQ-E, 2332TQ, and 2348UPQ support 4x10g breakout on N9K-C93180YC-EX, N9K-C93180YC-FX, and N9K-C93240YC-FX2.
- Cisco Nexus 2348TQ, 2332TQ, 2348TQ-E, and 2348UPQ support original 40G connectivity on N9K-C9332PQ, N9K-C93180YC-EX, N9K-C93108TC-EX, N9K-C93180YC-FX, N9K-C93108TC-FX, N9K-C93240YC-FX2, and N9K-C9336C-FX2.
- For FEX support on various hardware platforms, see the FEX matrix at the location: <https://www.cisco.com/c/dam/en/us/td/docs/Website/datacenter/fexmatrix/fexmatrix.html>
- FEX supports in-service software upgrades (ISSU) on the following switches when only in straight-through mode:
 - 9372PX

- 9396PX
- Cisco Nexus 9500 Series switches with a X9464PX or X9564PX line card.
- In-service software upgrade (ISSU) is not supported on Cisco Nexus 9000 Series switches with dual-homed FEX.
- Beginning Cisco NX-OS Release 9.3(1), MTU 9216 is made default for FEX fabric ports-channels. Only MTU 9216 is allowed to be configured on FEX fabric port-channels. Configuring any other value throws an error.
- If the MTU value on a FEX fabric port-channel was set to 9216 before upgrading to Cisco NX-OS Release 9.3(1), the show running config command will not display the MTU config as it is the new default in Cisco NX-OS Release 9.3(1). Due to this the show running-config diff command displays the difference which is expected.
- A jumbo ping (greater than MTU 2344) from a switch supervisor to a FEX host fails because the control queue on a FEX supports an MTU 2344 maximum.
- Layer 3 is supported on FEX interfaces and subinterfaces on Cisco Nexus 9300 Series switches.
- Layer 3 is supported on FEX port channel interfaces on Cisco Nexus 9300 Series switches.
- Layer 3 routing is supported on FEX interfaces, subinterfaces, and port-channel interfaces on Cisco Nexus 9300-EX, N9K-C93180YC-FX, and N9K-C93240YC-FX2 platform switches.
- The following limitations apply to Layer 3 support on FEX interfaces for all platforms:
 - FEX Layer 3 is not supported on Cisco Nexus 9500 Series switches.
 - FEX Layer 3 is not supported on port channel subinterfaces.
 - FEX Layer 3 does not support HSRP, ACLs, QoS, BFD, or DHCP on FEX Layer 3 ports and subinterfaces.
 - FEX Layer 3 supports only 1500-bytes MTU.
 - PIM is supported on FEX Layer 3 interfaces.
 - Layer 3 and FEX consistency-checker do not support Layer 3 HIF ports.
 - Layer 3 is not supported on AA FEX interfaces.
- Static routes and all routing protocols are supported on FEX Layer 3 ports.
- When connecting a FEX module to a 9300-EX Series switch, the switch queuing policy must be changed from 8Q to 4Q if QoS queuing is going to be used.

Configuration Example:

```
switch(config)# system qos
switch(config-sys-qos)# service-policy type queuing out default-out-policy
switch(config-sys-qos)# service-policy type network-qos default-nq-policy
```

- The FEX configuration is not supported on the Cisco Nexus 9348GC-FXP and N9K-C92348GC switches.

- Postrouted flood is not supported on FEX HIF interfaces for Cisco Nexus 31128PQ switch and 3100-V platform switches, Cisco Nexus 9300 platform switches and the Cisco Nexus 3164Q switches, and Cisco Nexus 9500 platform switches.
- IPSG (IP Source Guard) is not supported on FEX ports.
- URPF is not supported on the FEX host interfaces.
- VTEP connected to FEX host interface ports is not supported.
- Dual-homed and straight-through FEXes are not supported if the parent switch is 48x10GT + 6x40G/100G.
- Enhanced vPC is not supported.
- The **configure replace** feature is not supported with a FEX HIF configuration.
- The **show** commands with the **internal** keyword are not supported.
- First generation Cisco Nexus 9300 platform switches do not support FEXs on uplink module ports (ALE – Application Leaf Engine). First generation switches are those that do not have a suffix (such as -EX, -FX, or -FX2) in the model name.
- The following features are not supported on the Cisco Nexus 9364C switch:
 - 100 G Port cannot support breakout (HW limitation)
 - FEX
 - ISSU
 - Segment routing
 - Tetratation (HW limitation)
- FEX is supported only on the Cisco Nexus 9332PQ, 9372PX, 9372PX-E, 9396PX, 93180YC-EX, and 9500 platform switches (FEX is not supported on the N9K-X9732C-EX line card, and Cisco Nexus 9200 platforms).
- FEX vPC is not supported between any model of FEX and the Cisco Nexus 9500 platform switches as the parent switches.
- FEX Layer 3 is not supported on the Cisco Nexus 2348TQ-E fabric.
- Beginning with Cisco NX-OS Release 9.3(9), FX3 used in FEX-mode fails for the first 20 HIF ports. However, this issue does not affect the FX3 in TOR-mode.
- Beginning with Cisco NX-OS Release 9.3(9) and Release 10.2(2), the Cisco Nexus 2248PQ, 2348TQ, 2348TQ-E, and 2348UPQ FEXs support using a QSA Adapter on the FEX NIF to connect to an 10G/SFP+ link on the parent switch.
- Cisco Nexus switches must have auto-negotiation that are disabled when using 40G or 100G FEX NIF uplinks.
- Cisco Nexus 9300-FX3 series switches operating in FEX mode do not support autonegotiation on FEX NIF ports.
- If FEX is offline and binary configuration exists, some of the HIF configuration will not be displayed in show startup configuration, such as:

- Switch port configuration excludes allowed VLAN.
- Spanning-tree

Configuration Limits

The configuration limits are documented in the *Cisco Nexus 9000 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
<code>feature-set fex</code> command	Disabled
Port mode	Layer 2 (Till Cisco NX-OS Release 10.1x) Layer 3 (only for ST mode - from Cisco NX-OS Release 10.2(1))



CHAPTER 3

Configuring the Fabric Extender

This chapter describes how to configure a Cisco Nexus 2000 Series Fabric Extender using a parent device and includes the following sections:

- [Managing the Fabric Extender Feature Set, on page 21](#)
- [Verifying the Configuration, on page 42](#)
- [Additional References, on page 49](#)

Managing the Fabric Extender Feature Set

Installing the Fabric Extender Feature Set

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>configure terminal</code> Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	<code>install feature-set fex</code> Example: <pre>switch(config)# install feature-set fex</pre>	To uninstall the Fabric Extender feature set, use the no install feature-set fex command. Note Before you can uninstall the feature set, you must ensure that the feature set is installed in the default VDC and that the feature set is not enabled in any VDC.

	Command or Action	Purpose
Step 3	exit Example: <pre>switch(config)# exit switch#</pre>	Exits global configuration mode.

Enabling the Fabric Extender Feature Set

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

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	feature-set fex Example: <pre>switch(config)# feature-set fex</pre>	<p>Enables the Fabric Extender feature set. The feature set must be installed before it shows as an option to this command.</p> <p>To disable the Fabric Extender feature set, use the no feature-set fex command.</p> <p>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.</p>
Step 3	exit Example: <pre>switch(config)# exit switch#</pre>	Exits global configuration mode.

Associating a Fabric Extender to a Fabric Interface

Associating a Fabric Extender to a Port Channel

Before you begin

Ensure that you have enabled the Fabric Extender feature set.

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 <i>channel</i> 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.
Step 4	fex associate <i>FEX-number</i> Example: switch(config-if)# fex associate 101	Associates a FEX number to the Fabric Extender unit attached to the interface. The range is from 101 to 199.
Step 5	(Optional) show interface port-channel <i>channel fex-intf</i> 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

```

**Note**

- The **fex associate** command must be entered from the port channel interface, not from the physical interface.
- 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

```

Disassociating a Fabric Extender from an Fabric Interface

Before you begin

Ensure that you have enabled the Fabric Extender feature set.

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *channel*
3. **no fex associate** *<FEX-id>*
4. **default interface ethernet** *<interface>*
5. **no fex** *<FEX-id>*

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface port-channel <i>channel</i> Example: <pre>switch(config)# interface port-channel 4 switch(config-if)#</pre>	Specifies a port channel to configure.
Step 3	no fex associate <FEX-id> Example: <pre>switch(config-if)# no fex associate 101</pre>	Disassociates the Fabric Extender unit attached to the interface.
Step 4	default interface ethernet <interface> Example: <pre>switch(config)# default interface ethernet 1/1</pre>	Sets the default settings for the member interface of the fabric port channel.
Step 5	no fex <FEX-id> Example: <pre>switch(config)# no fex 101</pre>	Removes the FEX configuration.

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) **locator-led fex *FEX-number***
4. (Optional) **description *desc***
5. (Optional) **no description**
6. (Optional) **no type**
7. (Optional) **serial *serial***
8. (Optional) **no serial**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fex FEX-number Example: <pre>switch(config)# fex 101 switch(config-fex)#</pre>	Enters FEX configuration mode for the specified Fabric Extender. The range of the <i>FEX-number</i> argument is from 101 to 199. Note If you need to change the FEX number, you must remove the previous configuration (using the no fex FEX-number command) before applying a new FEX number.
Step 3	(Optional) locator-led fex FEX-number Example: <pre>switch(config)# locator-led fex 101</pre>	Turns on the locator LED of a Fabric Extender. The range of the <i>FEX-number</i> argument is from 101 to 199.
Step 4	(Optional) description desc Example: <pre>switch(config-fex)# description Rack7A-N2K</pre>	Specifies the description. The default is the string FEXxxxx where <i>xxxx</i> is the FEX number. If the FEX number is 123, the description is FEX0123.
Step 5	(Optional) no description Example: <pre>switch(config-fex)# no description</pre>	Deletes the description.
Step 6	(Optional) no type Example: <pre>switch(config-fex)# no type</pre>	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 7	(Optional) serial serial Example: <pre>switch(config-fex)# serial JAF1339BDSK</pre>	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 8	(Optional) no serial Example: <pre>switch(config-fex)# no serial</pre>	Deletes the serial number string.

Configuration Examples

This section contains examples of FEX configurations.

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):

Switch 1 Configuration	Switch 2 Configuration
<pre> config t feature lacp int e101/1/1-2 channel-group 10 mode active no shutdown Int port-channel10 switchport switchport mode trunk switchport trunk allowed vlan 1-20 vpc 10 </pre>	<pre> config t feature lacp int e101/1/1-2 channel-group 10 mode active no shutdown Int port-channel10 switchport switchport mode trunk switchport trunk allowed vlan 1-20 vpc 10 </pre>


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

The Straight through FEX is supported on the Cisco Nexus 9300-FX2 platform. Beginning with Cisco NX-OS Release 9.3(5) it is supported on the Cisco Nexus 9300-FX3 platform also.

In Cisco NX-OS Release 9.3(5), the following third-party equipment is not supported for straight-through FEX for Cisco Nexus 9300-FX3 platform switches: B22-HP, B22-IBM, B22-Dell, and B22-Fujitsu.

The following is an example of a straight through FEX topology:


Figure 3: Switch vPC Topology



Note that this topology shows:

- The Cisco Nexus 9000 Series switch ports are directly connected to another switch or host and are configured as part of a port channel that becomes part of a vPC.
- That vPC 20 is configured on port channel 20, which has Eth1/10 on N9k-1 and Eth2/1 on N9k-2 as members.

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



Note that this topology shows:

- Each FEX is single-homed (straight-through FEX topology) with a Cisco Nexus 9000 Series switch. The host interfaces on this FEX are configured as port channels and those port channels are configured as vPCs.
- Eth100/1/1 on N9k-1 and Eth102/1/5 on N9k-2 are configured as members of PO200 and PO200 is configured for vPC 200.

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

The following table compares the sample running configuration that must be configured for the peer switches shown in the Switch vPC Topology and the FEX Straight-Through Topology (Host vPC).

Table 3: Running Configuration Example for a vPC Straight-Through Topology

Basic Configuration—No Port Profile	Port Profile Configuration
vlan 1-10	vlan 1-10
interface port-channel 20 switchport mode trunk vpc 20 switchport trunk allowed vlan 1-5	port-profile type port-channel pc-profile switchport mode trunk state enabled
interface port-channel 200 switchport mode trunk vpc 200 switchport trunk allowed vlan 1-5	interface port-channel 20 inherit port-profile pc-profile vpc 20 switchport trunk allowed vlan 1-5

Basic Configuration—No Port Profile	Port Profile Configuration
vlan 1-10	vlan 1-10
interface Ethernet1/10 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 20	interface port-channel 20 inherit port-profile pc-profile vpc 20 switchport trunk allowed vlan 1-5
interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200	interface Ethernet1/10 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 20
	interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

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

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

The following example shows how to configure the peer switches in the Switch vPC Topology and the FEX Straight-Through Topology (Host vPC):

- Configure the CFS over IPV4 distribution to change the multicast address.

```
N9k-1(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network
Are you sure? (y/n) [n] y
```

```
N9k-2(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network.
Are you sure? (y/n) [n] y
```

- Configure the CFSoIP multicast address on each peer switch

```
N9k-1(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

```
N9k-2# config terminal
N9k-2(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

- Enable CFSoIP on both switches.

```
N9k-1(config)# cfs ipv4 distribute
```

```
N9k-2(config)# cfs ipv4 distribute
```

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

```

N9k-1# config sync
N9k-1(config-sync)# switch-profile Test
N9k-1(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N9k-2>
N9k-2# config sync
N9k-2(config-sync)# switch-profile Test
N9k-2(config-sync-sp)# sync-peers destination <out of band mgmt0 IP address of peer switch
N9k-1>

```

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

```

N9k-1(config-sync-sp)# port-profile type port-channel pc-profile
N9k-1(config-sync-port-prof)# switchport mode trunk
N9k-1(config-sync-port-prof)# state enabled

```

- Create port-channel interfaces inside the switch profile.



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

```

N9k-1(config-sync-sp)# interface port-channel 20
N9k-1(config-sync-sp)# interface port-channel 200

```

- Commit the configuration in the switch profile.

```

N9k-1(config-sync-sp)# commit

```

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



Note In this topology, port-channel members must not be identical on the peer switches.

```

N9k-1(config)# interface Ethernet1/10
N9k-1(config-if)# channel-group 20 force
N9k-1(config)# interface Ethernet100/1/1
N9k-1(config-if)# channel-group 200 force

```

```

N9k-2(config)# interface Ethernet2/1
N9k-2(config-if)# channel-group 20 force
N9k-2(config)# interface Ethernet102/1/5
N9k-2(config-if)# channel-group 200 force

```

```

N9k-1(config)# interface Ethernet100/1/1
N9k-1(config-if)# switchport mode trunk
N9k-1(config-if)# switchport trunk allowed vlan 1-5

```

```

N9k-2(config)# interface Ethernet2/1
N9k-2(config-if)# switchport mode trunk

```



```
N9k-2(config-if)# switchport trunk allowed vlan 1-5

N9k-2(config)# interface Ethernet102/1/5
N9k-2(config-if)# switchport mode trunk
N9k-2(config-if)# switchport trunk allowed vlan 1-5
```

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

```
N9k-1(config-sync-sp)# interface port-channel 20
N9k-1(config-sync-sp-if)# inherit port-profile pc-profile
N9k-1(config-sync-sp-if)# vpc 20
N9k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
N9k-1(config-sync-sp)# interface port-channel 200
N9k-1(config-sync-sp-if)# inherit port-profile pc-profile
N9k-1(config-sync-sp-if)# vpc 200
N9k-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
```

- Commit the configuration in the switch profile.

```
N9k-1(config-sync-sp)# commit
```

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

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

The following example shows how to configure the peer switches in the Switch vPC Topology and the FEX Straight-Through Topology (Host vPC):

- Configure the CFS over IPV4 distribution to change the multicast address

```
N9k-1(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network
Are you sure? (y/n) [n] y
```

```
N9k-2(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network.
Are you sure? (y/n) [n] y
```

- Configure the CFSoIP multicast address on each peer switch

```
N9k-1(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

```
N9k-2# config terminal
N9k-2(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

- Enable CFSoIP on both switches.

```
N9k-1(config)# cfs ipv4 distribute
```

```
N9k-2(config)# cfs ipv4 distribute
```

- Create a switch profile on both switches.

```
N9k-1# config sync
N9k-1(config-sync)# switch-profile Test
```

```
N9k-2# config sync
N9k-2(config-sync)# switch-profile Test
```

- Import the running configuration.

```
N9k-1(config-sync-sp)# import running-config
N9k-1(config-sync-sp-import)# show switch-profile Test buffer
```

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

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

The following shows the command sequence to import the running configuration:

Table 4: Command Sequence to Import the Running Configuration

Buffer Sequence Number	Command
1	vlan 1-10
2 2.1 2.2 2.3	interface port-channel20 switchport mode trunk vpc 20 switchport trunk allowed vlan 1-5
3 3.1 3.2	interface port-channel100 switchport mode fex-fabric fex associate 101
4 4.1 4.2 4.3	Note Switch vPC Topology interface port-channel20 switchport mode trunk vpc 20 switchport trunk allowed vlan 1-5
4 4.1 4.2 4.3	Note FEX Straight-Through Topology (Host vPC) interface port-channel200 switchport mode trunk vpc 200 switchport trunk allowed vlan 1-5

Buffer Sequence Number	Command
5 5.1 5.2 5.3	interface Ethernet1/1 fex associate 101 switchport mode fex-fabric channel-group 100
6 6.1 6.2 6.3	interface Ethernet1/2 fex associate 101 switchport mode fex-fabric channel-group 100
7 7.1 7.2 7.3	interface Ethernet1/10 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 20
8 8.1 8.2 8.3	interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

```
N9k-2(config-sync-sp) # import running-config
```

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

```
N9k-1(config-sync-sp-import) # buffer-delete 3,5,6-8
```

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

- Commit the configuration in the switch profile on both switches.

```
N9k-1(config-sync-sp-import) # commit
```

```
N9k-2(config-sync-sp-import) # commit
```

- Add the sync peer on both switches.



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

```
N9k-1# config sync
N9k-1(config-sync) # switch-profile Test
N9k-1(config-sync-sp) # sync-peers destination <out of band mgmt0 IP address of peer switch
N9k-2>
```

```
N9k-2# config sync
N9k-2(config-sync) # switch-profile Test
N9k-2(config-sync-sp) # sync-peers destination <out of band mgmt0 IP address of peer
```

```

switch
N9k-1>

N9K-1# configure sync
Enter configuration commands, one per line. End with CNTL/Z.
N9K-1(config-sync)# no switch-profile SP ?
    all-config    Deletion of profile, local and peer configurations
    local-config  Deletion of profile and local configuration
    profile-only  Deletion of profile only and no other configuration

N9396PX-1(config-sync)# no switch-profile SP

```



Caution When you remove a switch profile using the **no switch-profile name [all-config | local-config]** command, the configuration in the switch profile is immediately removed from the running configuration. This disrupts the configurations that were present in the switch profile, such as port channel and vPC configurations.

When you remove a switch profile using the **no switch-profile name [profile-only]** command, the configuration in the switch profile is immediately removed from the switch profile only. This does not disrupt the configurations that were present in running config.

It is recommended to execute the CLI **resync-database** on both peer switches before deleting a large configuration in the switch-profile.

Perform the following action if you received the "Deletion of switch profile failed" error message when attempting to delete switch-profile:

```

N9K-1(config-sync)# resync-database
Re-synchronization of switch-profile db takes a few minutes...
Re-synchronize switch-profile db completed successfully.
N9K-1(config-sync)#
N9K-2(config-sync)# resync-database
Re-synchronization of switch-profile db takes a few minutes...
Re-synchronize switch-profile db completed successfully.
N9K-2(config-sync)#

```

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

The dual-homed FEX (active-active) topology is supported beginning with Cisco NX-OS Release 7.0(3)I5(2) for Cisco Nexus 9300 and 9300-EX Series switches.

Beginning with Cisco NX-OS Release 9.3(5), the Dual-Homed FEX is supported on the N9K-C9336C-FX2, N9K-C93240YC-FX2, N9K-C93360YC-FX2, and N9K-C93216TC-FX2 switches. The Cisco Nexus 9300-FX2 and FX3 switches are supported on the ST and the AA FEX modes.


In Cisco NX-OS Release 9.3(5), the following third-party equipment is not supported for Dual-Homed FEX for Cisco Nexus 9300-FX2/FX3 platform switches and straight-through FEX for Cisco Nexus 9300-FX3 platform switches: B22-HP, B22-IBM, B22-Dell, and B22-Fujitsu.

The following topology shows that each FEX is dual-homed with two Cisco Nexus 9300 Series switches. The FEX-fabric interfaces for each FEX are configured as a vPC on both peer switches. The host interfaces on the FEX appear on both peer switches.



Note The port configuration should be the same on both switches.

Figure 5: Dual-Homed FEX Topology



In the dual-homed FEX topology, the vPC is already operational. FEX 101 is dual-homed to both parent switches: N9k-1 and N9k-2 on FEX-fabric interfaces Ethernet 1/1.



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

FEX 101 is configured to have two types of host interfaces. One interface is Ethernet100/1/1, which is singly attached to a server (nonport-channel member), and the other interface is Ethernet 100/1/2-3, which is configured in a port channel to the server (port-channel member).

The following table shows the sample running configuration for the peer switches. Two types of configurations are shown:

- Basic Configuration.
- Port profile configuration.

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



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

Table 5: Running Configuration of a FEX in a Dual-Homed Topology for the Peer Switches

Basic Configuration—No Port Profile	Port Profile Configuration
vlan 1-10	vlan 1-10
interface port-channel100 switchport mode fex fabric vpc 100 fex associate 101	port-profile type ethernet eth-profile switchport mode trunk state enabled
interface port-channel 200 switchport mode trunk switchport trunk allowed vlan 1-5	port-profile type port-channel pc-profile switchport mode trunk state enabled
interface Ethernet1/1 fex associate 101 switchport mode fex fabric channel-group 100	interface port-channel100 switchport mode fex fabric vpc 100 fex associate 101
interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-10	interface port-channel 200 inherit port-profile pc-profile switchport trunk allowed vlan 1-5
interface Ethernet100/1/2 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200	interface Ethernet1/1 fex associate 101 switchport mode fex fabric channel-group 100
interface Ethernet100/1/3 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200	interface Ethernet100/1/1 inherit port-profile eth-profile switchport trunk allowed vlan 1-10
	interface Ethernet100/1/2 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200
	interface Ethernet100/1/3 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

New Deployments in a Dual-Homed FEX Topology

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

The following example shows how to configure the dual-homed FEX (active-active) topology:

- Configure the CFS over IPV4 distribution to change the multicast address.

```
N9K-1(config)# no cfs ipv4 distribute
```

```
This will prevent CFS from distributing over IPv4 network
Are you sure? (y/n) [n] y
```

```
N9K-2(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network.
Are you sure? (y/n) [n] y
```

- Configure the CFSoIP multicast address on each peer switch.

```
N9K-1(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

```
N9K-2# config terminal
N9K-2(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

- Enable CFSoIP on both switches.

```
N9K-1(config)# cfs ipv4 distribute
```

```
N9K-2(config)# cfs ipv4 distribute
```

- Create a switch profile on both switches.

```
N9K-1# config sync
N9K-1(config-sync)# switch-profile Test
N9K-1(config-sync-sp)# sync-peers destination /***out of band mgmt0 IP address of peer
switch***/
N9K-2>
```

```
N9K-2# config sync
N9K-2(config-sync)# switch-profile Test
N9K-2(config-sync-sp)# sync-peers destination /***out of band mgmt0 IP address of peer
switch***/
N9K-1>
```

- Add referred global configuration to the switch profile.



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

```
N9K-1(config-sync-sp)# port-profile type ethernet eth-profile
N9K-1(config-sync-port-prof)# switchport mode trunk
N9K-1(config-sync-port-prof)# state enabled
```

```
N9K-1(config-sync-sp)# port-profile type port-channel pc-profile
N9K-1(config-sync-port-prof)# switchport mode trunk
N9K-1(config-sync-port-prof)# state enabled
```

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

```
N9K-1(config-sync-sp) # interface Ethernet100/1/1
N9K-1(config-sync-sp-if) # inherit port-profile eth-profile
N9K-1(config-sync-sp-if) # switchport trunk allowed vlan 1-10
```

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



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

This example shows that port channel 100 (vPC 100) is the EtherChannel from N9k to N2k:

```
N9K-1(config-sync-sp) # interface Port-channel100
```

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

```
N9K-1(config-sync-sp) # interface Port-channel200
```

- Commit the configuration inside the switch profile.

```
N9K-1(config-sync-sp) # commit
```

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



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

This example shows that N9k-1- Ethernet1/1 is a FEX-fabric member of port channel 100:

```
N9K-1(config) # int ether1/1
N9K-1(config-if) # channel-group 100 force
```

This example shows that N9K-1- Ethernet1/100/2-3 are members of port channel 200:

```
N9K-1(config) # interface Ethernet100/1/2-3
N9K-1(config-if-range) # channel-group 200 force
```

This example shows that N9K-2- Ethernet1/1 is a FEX-fabric interface that is in port channel 100:

```
N9K-2(config) # int ether1/1
N9K-2(config-if) # channel-group 100 force
```

This example shows that N9K-2- Ethernet1/100/2-3 are members of port channel 200:

```
N9K-2(config) # interface Ethernet100/1/2-3
N9K-2(config-if-range) # channel-group 200 force
```

```
N9K-1(config) # interface Ethernet100/1/2-3
N9K-1(config-if-range) # switchport mode trunk
N9K-1(config-if-range) # switchport trunk allowed vlan 1-5
```

```
N9K-2(config) # interface Ethernet100/1/2-3
```



```
N9K-2(config-if-range)# switchport mode trunk
N9K-2(config-if-range)# switchport trunk allowed vlan 1-5
```

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

```
N9K-1(config-sync-sp-if)# interface Port-Channel100
N9K-1(config-sync-sp-if)# switchport mode fex-fabric
N9K-1(config-sync-sp-if)# fex associate 101
N9K-1(config-sync-sp-if)# vpc 100

N9K-1(config-sync-sp)# interface Port-channel200
N9K-1(config-sync-sp-if)# inherit port-profile pc-profile
N9K-1(config-sync-sp-if)# switchport trunk allowed vlan 1-5
```

- Commit the configuration in the switch profile.

```
N9K-1(config-sync-sp)# commit
```

Existing Deployment with a Dual-Homed FEX Topology

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

The following example shows how to configure the peer switches in the vPC topology for the dual-homed FEX (active-active) topology:

- Configure the CFS over IPV4 distribution to change the multicast address.

```
N9K-1(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network
Are you sure? (y/n) [n] y
N9K-2(config)# no cfs ipv4 distribute
This will prevent CFS from distributing over IPv4 network.
Are you sure? (y/n) [n] y
```

- Configure the CFSoIP multicast address on each peer switch.

```
N9K-1(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y

N9K-2# config terminal
N9K-2(config)# cfs ipv4 mcast-address 239.255.1.1
Distribution over this IP type will be affected
Change multicast address for CFS-IP?
Are you sure? (y/n) [n] y
```

- Enable CFSoIP on both switches.

```
N9K-1(config)# cfs ipv4 distribute
N9K-2(config)# cfs ipv4 distribute
```

- Create a switch profile on both switches.

```
N9K-1# config sync
N9K-1(config-sync)# switch-profile Test
```

```
N9K-2# config sync
N9K-2(config-sync)# switch-profile Test
```

- Commit the configuration in the switch profile on both switches.

```
N9K-1(config-sync-sp)# commit
N9K-2(config-sync-sp)# commit
```

- Import the running configuration.

```
N9K-1(config-sync-sp)# import running-config
N9K-1(config-sync-sp-import)# show switch-profile Test buffer
```

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

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

The following shows the command sequence to import the running configuration:

Table 6: Command Sequence to Import the Running Configuration

Buffer Sequence Number	Command
1	vlan 1-10
2 2.1 2.2 2.3	interface port-channel100 switchport mode fex-fabric vpc 100 fex associate 101
3 3.1 3.2	interface port-channel200 switchport mode trunk switchport trunk allowed vlan 1-5
4 4.1 4.2 4.3	interface Ethernet1/1 fex associate 101 switchport mode fex-fabric channel-group 100
5 5.1 5.2 5.3	interface Ethernet100/1/1 switchport mode trunk switchport trunk allowed vlan 1-10

Buffer Sequence Number	Command
6 6.1 6.2 6.3	interface Ethernet100/1/2 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200
7 7.1 7.2 7.3	interface Ethernet100/1/3 switchport mode trunk switchport trunk allowed vlan 1-5 channel-group 200

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

```
N9K-1(config-sync-sp-import)# buffer-delete 4, 6, 7
```

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

- Commit the configuration in the switch profile on both switches.

```
N9K-1(config-sync-sp-import)# commit
```

```
N9K-2(config-sync-sp-import)# commit
```

- Add the sync peer on both switches.



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

```
N9K-1# config sync
N9K-1(config-sync)# switch-profile sp
N9K-1(config-sync-sp)# sync-peers destination /***out of band mgmt0 IP address of peer
switch***/
N9K-2>
```

```
N9K-2# config sync
N9K-2(config-sync)# switch-profile sp
N9K-2(config-sync-sp)# sync-peers destination /***out of band mgmt0 IP address of peer
switch***/
N9K-1>
```

```
N9K-1# configure sync
Enter configuration commands, one per line. End with CNTL/Z.
N9K-1(config-sync)# no switch-profile SP ?
  all-config    Deletion of profile, local and peer configurations
  local-config  Deletion of profile and local configuration
  profile-only  Deletion of profile only and no other configuration
```

```
N9396PX-1(config-sync)# no switch-profile SP
```



Caution When you remove a switch profile using the **no switch-profile name [all-config | local-config]** command, the configuration in the switch profile is immediately removed from the running configuration. This disrupts the configurations that were present in the switch profile, such as port channel and vPC configurations.

When you remove a switch profile using the **no switch-profile name [profile-only]** command, the configuration in the switch profile is immediately removed from the switch profile only. This does not disrupt the configurations that were present in running config.

It is recommended to execute the CLI **resync-database** on both peer switches before deleting a large configuration in the switch-profile.

Perform the following action if you received the "Deletion of switch profile failed" error message when attempting to delete switch-profile:

```
N9K-1(config-sync)# resync-database
Re-synchronization of switch-profile db takes a few minutes...
Re-synchronize switch-profile db completed successfully.
N9K-1(config-sync)#
N9K-2(config-sync)# resync-database
Re-synchronization of switch-profile db takes a few minutes...
Re-synchronize switch-profile db completed successfully.
N9K-2(config-sync)#
```

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 [<i>FEX-number</i>] [detail]	Displays information about a specific Fabric Extender or all attached units.
show interface <i>type number fex-intf</i>	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 <i>number transceiver</i> [fex-fabric]	Displays the SFP+ transceiver and diagnostic optical monitoring (DOM) information for the Fabric Extender uplinks.

Command or Action	Purpose
<code>show feature-set</code>	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:

```
switch# show fex
      FEX          FEX          FEX          FEX
Number  Description      State      Model          Serial
-----
101     FEX0101             Online     N2K-C2248TP-1GE  JAF1418AARL
```

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

```
switch# show fex 101 detail
FEX: 101 Description: FEX0101 state: Online
  FEX version: 5.1(1) [Switch version: 5.1(1)]
  FEX Interim version: 5.1(0.159.6)
  Switch Interim version: 5.1(1)
  Extender Model: N2K-C2248TP-1GE, Extender Serial: JAF1418AARL
  Part No: 73-12748-05
  Card Id: 99, Mac Addr: 54:75:d0:a9:49:42, Num Macs: 64
  Module Sw Gen: 21 [Switch Sw Gen: 21]
  pinning-mode: static Max-links: 1
  Fabric port for control traffic: Po101
  Fabric interface state:
    Po101 - Interface Up. State: Active
    Eth2/1 - Interface Up. State: Active
    Eth2/2 - Interface Up. State: Active
    Eth4/1 - Interface Up. State: Active
    Eth4/2 - Interface Up. State: Active
  Fex Port      State  Fabric Port  Primary Fabric
    Eth101/1/1  Up    Po101        Po101
    Eth101/1/2  Up    Po101        Po101
    Eth101/1/3  Down  Po101        Po101
    Eth101/1/4  Down  Po101        Po101
    Eth101/1/5  Down  Po101        Po101
    Eth101/1/6  Down  Po101        Po101
    Eth101/1/7  Down  Po101        Po101
    Eth101/1/8  Down  Po101        Po101
    Eth101/1/9  Down  Po101        Po101
    Eth101/1/10 Down  Po101        Po101
    Eth101/1/11 Down  Po101        Po101
    Eth101/1/12 Down  Po101        Po101
    Eth101/1/13 Down  Po101        Po101
    Eth101/1/14 Down  Po101        Po101
    Eth101/1/15 Down  Po101        Po101
    Eth101/1/16 Down  Po101        Po101
    Eth101/1/17 Down  Po101        Po101
    Eth101/1/18 Down  Po101        Po101
    Eth101/1/19 Down  Po101        Po101
    Eth101/1/20 Down  Po101        Po101
    Eth101/1/21 Down  Po101        Po101
    Eth101/1/22 Down  Po101        Po101
    Eth101/1/23 Down  Po101        Po101
    Eth101/1/24 Down  Po101        Po101
    Eth101/1/25 Down  Po101        Po101
    Eth101/1/26 Down  Po101        Po101
    Eth101/1/27 Down  Po101        Po101
```

```

Eth101/1/28 Down Po101 Po101
Eth101/1/29 Down Po101 Po101
Eth101/1/30 Down Po101 Po101
Eth101/1/31 Down Po101 Po101
Eth101/1/32 Down Po101 Po101
Eth101/1/33 Down Po101 Po101
Eth101/1/34 Down Po101 Po101
Eth101/1/35 Down Po101 Po101
Eth101/1/36 Down Po101 Po101
Eth101/1/37 Down Po101 Po101
Eth101/1/38 Down Po101 Po101
Eth101/1/39 Down Po101 Po101
Eth101/1/40 Down Po101 Po101
Eth101/1/41 Down Po101 Po101
Eth101/1/42 Down Po101 Po101
Eth101/1/43 Down Po101 Po101
Eth101/1/44 Down Po101 Po101
Eth101/1/45 Down Po101 Po101
Eth101/1/46 Down Po101 Po101
Eth101/1/47 Down Po101 Po101
Eth101/1/48 Down Po101 Po101

```

Logs:

```

09/21/2010 21:14:26.843850: Module register received
09/21/2010 21:14:26.845778: Registration response sent
09/21/2010 21:14:27.663073: Module Online Sequence
09/21/2010 21:14:30.191121: Module Online

```

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

```

switch# show interface port-channel 101 fex-intf
Fabric          FEX
Interface       Interfaces
-----
Po101           Eth101/1/2   Eth101/1/1

```

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

```

switch# show interface fex-fabric
Fabric      Fabric      Fex      FEX
Fex  Port    Port State  Uplink    Model      Serial
-----
101  Eth2/1    Active     1         N2K-C2248TP-1GE  JAF1418AARL
101  Eth2/2    Active     2         N2K-C2248TP-1GE  JAF1418AARL
101  Eth4/1    Active     3         N2K-C2248TP-1GE  JAF1418AARL
101  Eth4/2    Active     4         N2K-C2248TP-1GE  JAF1418AARL

```

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 <i>FEX-number</i> } [temperature power fan]	Displays the environmental sensor status.
show inventory fex <i>FEX-number</i>	Displays inventory information for a Fabric Extender.
show module fex [<i>FEX-number</i>]	Displays module information about a Fabric Extender.
show sprom fex <i>FEX-number</i> {all backplane powersupply <i>ps-num</i> } all	Displays the contents of the serial PROM (SPROM) on the Fabric Extender.

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 (Celsius)	MinorThres (Celsius)	CurTemp (Celsius)	Status
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 (Watts)	Power (Amp)	Status
1	--	--	--	--
2	NXK-PAC-400W	4.32	0.36	ok

Mod	Model	Power Requested (Watts)	Power Requested (Amp)	Power Allocated (Watts)	Power Allocated (Amp)	Status
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       : XXXXXXXXXXXTBDV00
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 srom 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       : XXXXXXXXXXXTDBV00

```



```

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	<i>Cisco NX-OS Licensing Guide</i>
Interface configuration	<i>Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide</i>
Command reference	<i>Cisco Nexus 9000 Series Command References</i>



CHAPTER 4

Software FEX Mode Configuration

- [Software FEX Mode Configuration, on page 51](#)

Software FEX Mode Configuration

By default, Cisco Nexus switches operate in the switch mode. However, Cisco Nexus 9300-FX3 platform switches also support the FEX mode. This FEX mode allows a switch to operate like a Cisco Nexus 2000 Series Fabric Extender. As a result, the switch will not require any independent software upgrades, configuration backups, or other maintenance tasks.

Cisco Nexus switches in software FEX mode support 25G FEX connectivity to the host for a single point of management use cases.

See [Platform Support Matrix](#), to check for the supported switches.

TOR/Switch to FEX Conversion

This section describes how to convert the switch usage from TOR/switch mode to FEX mode.

- Configure the switch in a way that it does not boot from Cisco NX-OS mode.
- Run the **copy running-config startup-config** command before booting the FEX image.
- Run the **boot fex** command. This command sets the FEX as the boot variable.
- Reload the switch.



Note Do not run the **copy running-config startup-config** command after you run the **boot fex** command.

A sample ToR to FEX conversion is provided below.

```
switch(config)# write erase
switch(config)# no boot nxos
switch(config)# copy running-config startup-config
switch(config)# boot fex
switch(config)# reload
```

FEX to TOR/Switch Conversion

This section describes how to convert the switch usage from FEX to switch/TOR mode.

- Run the conversion command, `boot nx-os bootflash:<nxos image>` from the FEX terminal.
- You must upload a Cisco NX-OS image when you use this conversion command.
- This conversion command verifies the Cisco NX-OS image and sets the boot variable. Hence the FEX boots with the specified Cisco NX-OS image on the reload.



Note FEX does not have or save any configuration. Hence you must save the running configuration as the startup configuration.

A sample FEX to ToR conversion is provided below.

```
fex-1(config)# boot nxos bootflash:<nxos image>
fex-1(config)# reload
```

You can use the following commands to configure management IP and copy the NX-OS image to FEX.

Commands	Uses
dir	Lists all files in bootflash.
delete <i>file-name</i>	Removes file in bootflash.
interface mgmt <i>0</i> ip address <i>ip address network mask</i> ip route <i>network_gateway</i>	Configures management IP to FEX. The management IP must be connected physically.
show interface mgmt <i>0</i> brief	Verifies the configured management IP.
copy scp: [[/username@]server][[/path] bootflash: [filename]	Copies the Cisco NX-OS image file to FEX for converting it to ToR mode.



CHAPTER 5

Upgrading Procedures

The section includes the following topics:

- [Upgrade Process for a vPC Topology on the Primary Switch, on page 53](#)
- [Upgrade Process for a vPC Topology on the Secondary Switch, on page 54](#)

Upgrade Process for a vPC Topology on the Primary Switch

The following list summarizes the upgrade process on a primary switch in a vPC topology. Steps that differ from a switch upgrade in a non-vPC topology are in bold.



Note In vPC topologies, the two peer switches must be upgraded individually. An upgrade on one peer switch does not automatically update the vPC peer switch.

-
- | | |
|----------------|--|
| Step 1 | The install all command issued on the vPC primary switch triggers the installation upgrade. |
| Step 2 | The compatibility checks display the impact of the upgrade. |
| Step 3 | The installation proceeds or not based on the upgrade impact. |
| Step 4 | The configuration is locked on both vPC peer switches. |
| Step 5 | The current state is saved. |
| Step 6 | The system unloads and runs the new image. |
| Step 7 | The stateful restart of the system software and application occurs. |
| Step 8 | The installer resumes with the new image. |
| Step 9 | The FEXs are upgraded sequentially. |
| Step 10 | The installation is complete. |
-

What to do next

When the installation is complete, the vPC primary switch and the FEXs that are connected to the primary switch are upgraded. The single-homed FEXs and the dual-homed FEXs are now running the upgraded software.



Note The dual-homed FEXs are now connected to the primary and secondary switches that are running two different versions of the Cisco NX-OS software. The vPC primary switch is running the upgraded version and the vPC secondary switch is running the original software version.

Upgrade Process for a vPC Topology on the Secondary Switch

The following list summarizes the upgrade process on a secondary switch in a vPC topology. Steps that differ from a switch upgrade in a non-vPC topology are in bold.

-
- Step 1** **The install all command issued on the vPC second switch triggers the installation upgrade.**
- Step 2** The compatibility checks display the impact of the upgrade.
- Step 3** The installation proceeds or not based on the upgrade impact.
- Step 4** The current state is saved.
- Step 5** The system unloads and runs the new image.
- Step 6** The stateful restart of the system software and application occurs.
- Step 7** The installer resumes with the new image.
- Step 8** **The FEXs are upgraded sequentially. The upgrade completes on the single-homed FEXs and a sanity check is performed on the dual-homed FEXs.**
- Note** The dual-homed FEXs were upgraded by the primary switch.
- Step 9** **The configuration is unlocked on the primary and secondary switches.**
- Step 10** The installation is complete.
-



CHAPTER 6

Minimizing the Impact of a Disruptive Upgrade


A non-ISSU upgrade is a disruptive upgrade that results in the reload of the Cisco Nexus device and the Fabric Extenders. The reload is a cold reboot that brings down the control plane and the data plane. The reload causes disruptions to the connected servers and hosts. When a vPC is deployed in the access layer, it is possible to minimize the impact of a non-ISSU upgrade. When one of the vPC switches is being reset during the upgrade process, all the server traffic can flow through its vPC peer.

- [Upgrading a Direct vPC or a Single-Homed FEX Access Layer, on page 55](#)
- [Upgrading a Dual-Homed FEX Access Layer, on page 57](#)

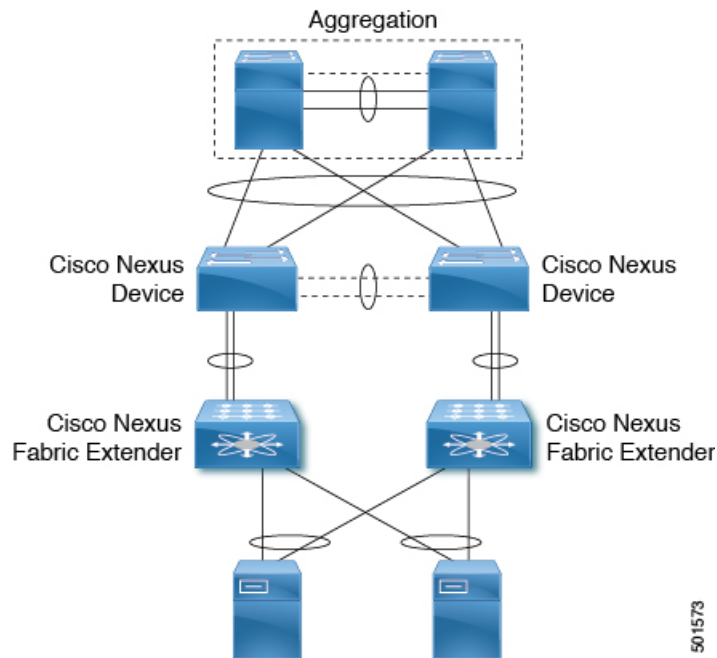
Upgrading a Direct vPC or a Single-Homed FEX Access Layer

The following figures show topologies in which the access layer includes a vPC configuration to hosts or downstream switches.

Hosts Directly Connected to vPC Peers




vPC Peered Dual-Supervisor Virtual Modular System Single-Homed FEXes



501573

Cisco Nexus Device Connected to Downstream Switches



501574

To upgrade the access layer without a disruption to hosts, follow these tasks:

- Upgrade the first vPC switch (vPC primary switch). During this upgrade, the switch is reloaded. When the switch is reloaded, the servers or the downstream switch detects a loss of connectivity to the first switch and starts forwarding traffic to the second (vPC secondary) switch.
- Verify that the upgrade of the switch has completed successfully. At the completion of the upgrade, the switch restores vPC peering, connected Nexus 2000 FEXes, and all the links.
- Upgrade the second switch. Repeating the same process on the second switch causes the second switch to reload during the upgrade process. During this reload, the first (upgraded) switch forwards all the traffic to/from servers.
- Verify that the upgrade of the second switch has completed successfully.




Note Flows that are forwarded to a switch during an upgrade on the switch will failover to the second switch. Also, flows are redistributed when vPC peers are active. The traffic disruption is limited to the time required for the server or host to detect the link-down and link-up events and to redistribute the flows.

Upgrading a Dual-Homed FEX Access Layer

A disruptive upgrade causes a switch and connected Fabric Extenders (FEX) to reload. The time required for a FEX to reload is less than the time required for a switch to reload. When hosts are connected to a dual-homed FEX, it is possible to keep the traffic disruption of the hosts to the same time as required by the FEX to download the image and reload (depending on the hardware platform it can be anywhere between 10 to 20 minutes), instead of the time required for an upgrade of the entire access layer.

The following figure shows a dual-homed FEX topology in which the access layer includes a vPC configuration to hosts or downstream switches.

vPC-Peered Dual-Supervisor Virtual Modular System Dual-Homed FEXes



The following dual-homed FEX procedure is supported only for an upgrade and not for a downgrade.

-
- Step 1** Upgrade the vPC primary switch with the new image. During the upgrade process, the switch is reloaded. When the switch is reloaded, only single-homed FEXes connected to the switch are reloaded and dual-homed FEXes are not reloaded. Servers connected to the dual-homed FEXes retain network connectivity through the vPC secondary switch.
- Step 2** Verify that the upgrade of the vPC primary switch is completed successfully. At the completion of the upgrade, the vPC primary switch restores vPC peering. However, dual-homed FEXes are connected only to the secondary vPC switch.
- Note**
- The FEX remains online on the vPC secondary switch while the vPC primary switch is reloaded.
 - On the vPC primary switch after the upgrade, the FEXes connected to the switch are in active-active mismatch state.
- Step 3** On the vPC secondary switch, shut the NIF (FEX uplink). The FEX downloads the new image from the vPC primary switch and it comes online on the newly upgraded switch. The servers connected to the dual-homed FEXes lose connectivity. Bring up the NIF (FEX uplink) on the vPC secondary.

Note Only the vPC primary switch displays that the FEX is online because the vPC secondary switch does not have the new image. The secondary switch displays the FEX in an active-active version mismatch state.

Step 4 Upgrade the vPC secondary switch with the new image. During the upgrade process, the switch is reloaded. When the switch is reloaded, only singled-homed FEXes connected to the switch are reloaded and dual-homed FEXes are not reloaded.

Step 5 Verify that the upgrade of the vPC secondary switch is completed successfully. At the completion of the upgrade, the vPC secondary switch restores vPC peering. Dual-homed FEXes connect to both the peer switches and start forwarding traffic.
