



## Overview

---

- This chapter provides an architectural overview of the Cisco Nexus 2000 Series Fabric Extender and includes the following sections:
  - [Licensing Requirements, on page 1](#)
  - [Supported Platforms, on page 2](#)
  - [Information About the Cisco Nexus 2000 Series Fabric Extender, on page 2](#)
  - [Fabric Extender Terminology, on page 2](#)
  - [Fabric Interface Features , on page 3](#)
  - [Host Interfaces, on page 3](#)
  - [Host Interface Port Channels, on page 4](#)
  - [VLANs, on page 5](#)
  - [FCoE Over FEX, on page 5](#)
  - [Protocol Offload, on page 5](#)
  - [Access Control Lists, on page 5](#)
  - [IGMP Snooping, on page 6](#)
  - [Switched Port Analyzer, on page 6](#)
  - [Oversubscription, on page 6](#)
  - [Management Model, on page 9](#)
  - [Forwarding Model, on page 9](#)
  - [Port Channel Fabric Interface Connection, on page 10](#)
  - [Port Numbering Convention, on page 11](#)
  - [Fabric Extender Image Management, on page 11](#)
  - [Guidelines and Limitations for the Fabric Extender, on page 12](#)
  - [Configuration Limits, on page 17](#)
  - [Default Settings, on page 17](#)

## 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 (NB22FTS) 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-B22FTS-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
<b>no speed</b>	Autonegotiates and advertises all speeds (only full duplex).
<b>speed 100</b>	Does not autonegotiate; pause cannot be advertised. The peer must be set to not autonegotiate (only 100-Mbps full duplex supported).
<b>speed 1000</b>	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.
  - 93180YC-FX
- 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, 93180YC-FX, 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.
- Beginning with Cisco NX-OS Release 10.3(1)F, FCoE over FEX is supported on N9K-C93180YC-FX switches in straight-through mode with N9K-C93180YC-FX3 FEX models.



- When upgrading from an earlier release to 10.3(1)F, to enable FCoE on N9K-C93180YC-FX3, reload the FEX after the upgrade.
- On flapping FX3 port-channel, member PO takes approximately 3–5 minutes to come up.
- N9K-C93180YC-FX3 FEX supports 10/25G HIF ports and 10/25/40/100G NIF ports.
- For LACP port priority configuration of an AA FEX interface, the configuration can be given from either of the VPC switches to reflect the configuration changes done on an AA FEX interface.
- 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 1: Default Cisco Nexus 2000 Series Fabric Extender Parameter Settings**

Parameters	Default
<b>feature-set fex</b> 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))

