



Overview

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

- [Information About the Cisco Nexus 2000 Series Fabric Extender, page 2](#)
- [Fabric Extender Terminology, page 2](#)
- [Fabric Interface Features , page 3](#)
- [Host Interfaces, page 3](#)
- [Host Interface Port Channels, page 4](#)
- [VLANs, page 6](#)
- [Protocol Offload, page 6](#)
- [Quality of Service, page 6](#)
- [Access Control Lists, page 7](#)
- [IGMP Snooping, page 7](#)
- [Switched Port Analyzer, page 7](#)
- [Oversubscription, page 7](#)
- [Management Model, page 8](#)
- [Forwarding Model, page 9](#)
- [Port Channel Fabric Interface Connection, page 10](#)
- [Port Numbering Convention, page 10](#)
- [Fabric Extender Image Management, page 11](#)
- [Licensing Requirements for the Fabric Extender, page 11](#)
- [Guidelines and Limitations for the Fabric Extender, page 11](#)
- [Default Settings, page 19](#)

Information About the Cisco Nexus 2000 Series Fabric Extender

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

The Fabric Extender integrates with its parent switch, which is a Cisco Nexus Series device, to allow automatic provisioning and configuration taken from the settings on the parent device. This integration allows large numbers of servers and hosts to be supported by using the same feature set as the parent device, including security and quality-of-service (QoS) configuration parameters, with a single management domain. The Fabric Extender and its parent switch enable a large multipath, loop-free, active-active data center topology without the use of the Spanning Tree Protocol (STP).

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

**Note**

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

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 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.

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

Host Interfaces

Layer 3 Host Interfaces

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

**Note**

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

The host interfaces also support subinterfaces. You can create up to 32 subinterfaces on a Fabric Extender host interface.

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

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

Layer 2 Host Interfaces

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

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

The Fabric Extender provides connectivity for computer hosts and other edge devices in the network fabric. Follow these guidelines when connecting devices to Fabric Extender host interfaces:

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

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

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

Host Interface Port Channels

Layer 3 Host Interface Port Channels

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

**Note**

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

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

Layer 3 mode is supported on host interface port channels.

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

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

Layer 2 Host Interface Port Channels

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

**Note**

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

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

Layer 2 mode is supported on host interface port channels.

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

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

Minimum Number of Links on a Fabric Port Channel

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

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

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

Load Balancing Using Host Interface Port Channels

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

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

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



Note

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

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

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

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

**Note**

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

VLANs

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

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

**Note**

The Fabric Extender does not support private VLANs (PVLANS).

Protocol Offload

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

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

Quality of Service

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

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

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

Access Control Lists

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

For more information about ACLs, see the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

IGMP Snooping

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

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

**Note**

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

Switched Port Analyzer

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

**Note**

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

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

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

Oversubscription

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

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

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

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

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

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

The Cisco Nexus 2232PP Fabric Extender has 8 10-Gigabit Ethernet fabric interfaces and 32 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces. (Static pinning is not supported. Port-channel mode is supported only on fabric interfaces.) When all host interfaces are sending traffic to all fabric interfaces, the maximum oversubscription ratio for the Cisco Nexus 2232PP is 4:1.

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

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

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

Management Model

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

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

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


The switch assigns an IP address in the range of loopback addresses (127.15.1.0/24) to the Fabric Extender to avoid conflicts with IP addresses that might be in use on the network.

- 3 The switch pushes the configuration data to the Fabric Extender. The Fabric Extender does not store any configuration locally.
- 4 The Fabric Extender updates the switch with its operational status. All Fabric Extender information is displayed using the switch commands for monitoring and troubleshooting.

Forwarding Model

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

Figure 1: Forwarding Model



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



Note

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


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

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

Port Channel Fabric Interface Connection

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

Figure 2: Port Channel Fabric Interface Connection



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

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



Note

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

Port Numbering Convention

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

interface ethernet *chassis/slot/port*

where

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

The chassis ID ranges from 101 to 199.



Note

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

interface ethernet *slot/port*

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

Fabric Extender Image Management

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

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

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

Licensing Requirements for the Fabric Extender

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

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

Guidelines and Limitations for the Fabric Extender

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

- Beginning with Cisco NX-OS Release 5.2(1), the default port mode is Layer 3. Before Cisco NX-OS Release 5.2(1), the default port mode was Layer 2.

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

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

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


Note

FEX interfaces do not support the spanning tree protocol.

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

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

- FEX HIF (FEX Host Interface) port to connect to a Protocol Independent Multicast (PIM) router.
- For Cisco NX-OS Release 6.2(2) and later releases, the Cisco Fabric Extender supports optimized multicast flooding (OMF) is available on FEX ports.
- The Cisco Fabric Extender does not support policy based routing (PBR).

Associating with F2-Series Modules

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

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

 - Supported: Port 1 of ASIC 1 and port 1 of ASIC 2 are added to a port channel.
 - Not supported: Port 1 of ASIC 1 and port 2 of ASIC 2 to form a port channel.

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

FEX Queuing Support

- **FEX QoS Queuing Support**

Fabric Extenders (FEXs) follow the network quality of service (QoS) queuing model for supporting queuing on FEX host interfaces, regardless of whether the FEX is connected to M-series or F-series fabric uplinks.

- Depending on the network-QoS template that is attached to the system QoS, the following parameters are inherited for queuing support on a FEX:
 - Number of queues
 - Class of service (CoS2q) mapping
 - Differentiated services code point (DSCP2q) mapping
 - Maximum transmission unit (MTU)

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

• Hardware Queue-limit Support

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

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

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

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

<snippet>
Queue limit: 66560 bytes
<snippet>
```

• Global Enable/Disable Control of DSCP2Q

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

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

• Show Command Support for FEX Host Interfaces

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



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

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

slot 1
=====

Interface is not in this module.

slot 2
=====

Interface is not in this module.

slot 4
=====

Interface is not in this module.

slot 6
=====

Interface is not in this module.

slot 9
=====

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

    Qos-group: 0 2 (shared)
    frh: 2
    drop-type: drop
    cos: 0 1 2 3 4 5 6
    xon      xoff      buffer-size
    -----+-----+-----
    34560     39680     48640

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

  Queue limit: 66560 bytes

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

```

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

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

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

FEX TCAM programmed successfully

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

slot 10
=====

slot 11
=====

Interface is not in this module.

slot 15
=====

Interface is not in this module.

slot 16
=====

Interface is not in this module.

slot 17
=====

Interface is not in this module.

slot 18
=====

Interface is not in this module.

```

• ISSU Behavior

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

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

```

Ethernet133/1/32 queuing information:
  Input buffer allocation:

```



```

Qos-group: ctrl
frh: 0
drop-type: drop
cos: 7
xon      xoff      buffer-size
-----+-----+-----
2560      7680      10240

```

```

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

```

```

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

```

Queue limit: 66560 bytes

```

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

```

```

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

```

```

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

```

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

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

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

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

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

```

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

Ethernet143/1/1 queuing information:
Network-QOS is disabled for N2248PQ <=====
Displaying the default configurations
Input buffer allocation:
Qos-group: ctrl

```

```

frh: 0
drop-type: drop
cos: 7
xon      xoff      buffer-size
-----+-----+-----
2560      7680      10240

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

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

Queue limit: 0 bytes

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

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

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

```

• Fabric Port Queuing Restrictions

- For FEXs that are connected to M-series uplinks, the queuing structure is different on FEX host interfaces and FEX fabric interfaces. The M series queuing policies must be consistent with the FEX queuing policies.
- DSCP-based queuing is not supported on M1 FEX fabric uplinks.
- You cannot modify the CoS 0-4 mappings in the default queue on M1 FEX fabric uplinks.

• MTU

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

Configuration Limits

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

Default Settings

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

Table 1: Default Cisco Nexus 2000 Series Fabric Extender Parameter Settings

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

