



Configuring the Fabric Extender

This chapter contains the following sections:


- [Information About the Cisco Nexus 2000 Series Fabric Extender, page 2](#)
- [Fabric Extender Terminology, page 3](#)
- [Fabric Extender Features, page 3](#)
- [Oversubscription, page 9](#)
- [Management Model, page 10](#)
- [Forwarding Model, page 11](#)
- [Connection Model, page 11](#)
- [Port Numbering Convention, page 14](#)
- [Fabric Extender Image Management, page 14](#)
- [Fabric Extender Hardware, page 14](#)
- [Speed and Duplex Mode, page 16](#)
- [Associating a Fabric Extender to a Fabric Interface, page 19](#)
- [Configuring Fabric Extender Global Features, page 22](#)
- [Enabling the Fabric Extender Locator LED, page 24](#)
- [Redistributing the Links, page 25](#)
- [Verifying the Fabric Extender Configuration, page 27](#)
- [Verifying the Chassis Management Information, page 28](#)
- [Configuring the Cisco Nexus N2248TP-E Fabric Extender, page 31](#)
- [Configuring the Cisco Nexus N2248PQ Fabric Extender, page 35](#)

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).

Figure 1: Single Management Domain



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



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

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



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



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

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

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

Fabric Extender Features

The Cisco Nexus 2000 Series Fabric Extender allows a single switch—and a single consistent set of switch features—to be supported across a large number of hosts and servers. By supporting a large server-domain under a single management entity, policies can be enforced more efficiently.

Some of the features of the parent switch cannot be extended onto the Fabric Extender.

Layer 2 Host Interfaces

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.

You can enable host interfaces to accept Cisco Discovery Protocol (CDP) packets. This protocol only works when it is enabled for both ends of a link.



Note CDP is not supported on fabric interfaces when the Fabric Extender is configured in a virtual port channel (vPC) topology.

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

For more information about BPDU Guard, see [Understanding BPDU Guard](#).

Host Port Channel

The following fabric extenders support port channel host interface configurations. Up to eight interfaces can be combined in a port channel. The port channel can be configured with or without Link Aggregation Control Protocol (LACP).

- Cisco Nexus 2248TP
- Cisco Nexus 2348UPQ
- Cisco Nexus 2348TQ
- Cisco Nexus 2232PP
- Cisco Nexus 2332TQ
- Cisco Nexus 2224TP
- Cisco Nexus 2248PQ
- Cisco Nexus B22 Fabric Extender for Fujitsu (N2K-B22FTS-P)
- Cisco Nexus B22 Fabric Extender for Dell (N2K-B22DELL-P)
- Cisco Nexus B22 Fabric Extender for HP (N2K-B22HP-P)

- Cisco Nexus B22 Fabric Extender for IBM (N2K-B22IBM-P)

VLANs and Private VLANs

The Fabric Extender supports Layer 2 VLAN trunks and IEEE 802.1Q VLAN encapsulation. Host interfaces can be members of private VLANs with the following restrictions:

- You can configure a host interface as an isolated or community access port only.
- You cannot configure a host interface as a promiscuous port.
- You cannot configure a host interface as a private VLAN trunk port.

For more information about VLANs, see the chapter in this guide on Configuring VLANs.

Virtual Port Channels

With a virtual port channel (vPC), you can configure topologies where a Cisco Nexus Fabric Extender is connected to a pair of parent switches or a pair of Fabric Extenders are connected to a single parent switch. The vPC can provide multipath connections, which allow you to create redundancy between the nodes on your network.



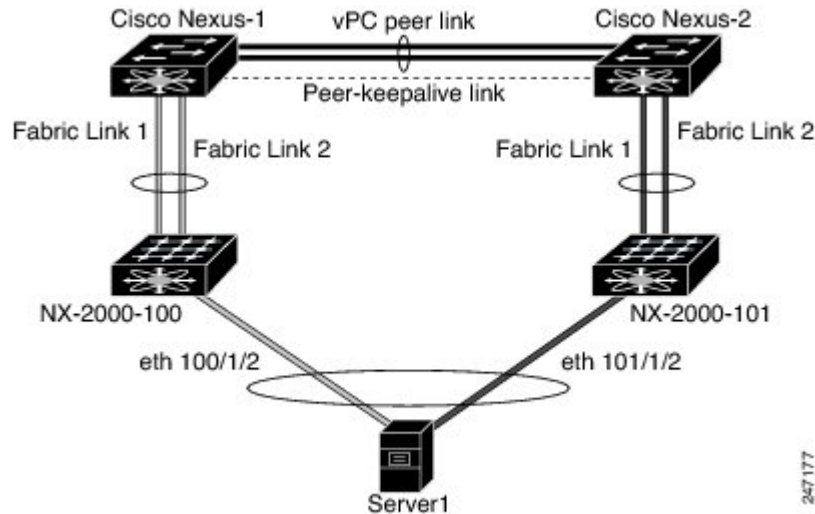
Note

A port channel between two FEXs that are connected to the same Cisco Nexus device is not supported. Virtual port channels (vPCs) cannot span two different FEXs when connected to the same Cisco Nexus device.

The following vPC topologies are possible with the Fabric Extender:


- The parent switches are connected single homed to Fabric Extenders that are subsequently connected to servers with dual interfaces (see the following figure).

Figure 2: Single Homed Fabric Extender vPC Topology



- The Fabric Extender is connected dual homed to two upstream parent switches and connected downstream to single homed servers (see the following figure).

Figure 3: Dual Homed Fabric Extender vPC Topology



This configuration is also called an Active-Active topology.



Note

Port channels between two Fabric Extenders connected to the same Cisco Nexus device is not supported. vPCs cannot span two different Fabric Extenders that are connected to the same physical Cisco Nexus device.

Fibre Channel over Ethernet Support

The Cisco Nexus 2232PP and Cisco Nexus 2248PQ support Fibre Channel over Ethernet (FCoE) with the following restrictions:

- Only FCoE Initialization Protocol (FIP) enabled converged network adapters (CNAs) are supported on the Fabric Extender.
- Binding to a port channel is limited to only one member in the port channel.

For configuration details, see the Fibre Channel over Ethernet Configuration Guide for the Nexus software release that you are using. The available versions of this document can be found at the following URL: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html.

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 provides two user queues for its quality-of-service (QoS) support, one for all no-drop classes and one for all drop classes. The classes configured on its parent switch are mapped to one of these two queues; traffic for no-drop classes is mapped to one queue and traffic for all drop classes is mapped to the other. Egress policies are also restricted to these two classes.

The Cisco Nexus Series device provides two predefined type qos class maps for matching broadcast or multicast traffic; class-all-flood and class-ip-multicast. These classes are ignored on the Fabric Extender.

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 and CoS-based egress queuing are 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.

Host interfaces support jumbo frames (up to 9216 bytes); however, a per-host interface maximum transmission unit (MTU) is not supported. Instead, MTU is set according to the QoS classes. You modify MTU by setting policy and class maps on the parent switch. Because the Fabric Extender has only two user queues, the MTU for the drop-queue is set to the maximum MTU of all drop classes and the MTU on the no-drop queue is set to the maximum MTU of all no-drop classes.

For more information about LLC and quality-of-service, see the Quality of Service Configuration Guide for your device.

Access Control Lists

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

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

IGMP Snooping

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

The Fabric Extender and its parent switch support 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 Multicast Routing Configuration Guide for the Nexus software release that you are using. The available versions of this document can be found at the following URL: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html.

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 System Management Configuration Guide for your device.

Fabric Interface Features

The FEX fabric interfaces support static port channels and priority flow control (PFC). PFC allows you to apply pause functionality to specific classes of traffic on an interface (instead of all the traffic on the interface). 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.

**Note**

Per class flow control mode is enabled by default on the fabric interfaces. When a fabric interface is configured on the parent switch, PFC mode is enabled by default and cannot be changed.

**Note**

For a 2248PQ, all fabric interfaces need to be bundled into a single fabric port channel. They cannot be used to connect the parent switch as individual ports.

For more information about PFC, see the Quality of Service Configuration Guide for your device.

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 2148T Fabric Extender has 4 10-Gigabit Ethernet fabric interfaces and 48 1000BASE-T (1-Gigabit) Ethernet host interfaces. With this system, you can have any number of configurations. For example, you can configure the following:

- 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 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 similar configurations to the Cisco Nexus 2148T.

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

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. (Static pinning is not supported. Port-channel mode is only supported on 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. 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 2232™ Fabric Extender has 8 10-Gigabit Ethernet fabric interfaces and 32 10G-BASE-T (10-Gigabit) Ethernet host interfaces. With this system, you can configure a 4 to 1 oversubscription (4 host interfaces for one fabric interface) or higher.

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 two 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. 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 HP (N2K-B22HP-P) is 2:1.

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.

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. 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 Dell (N2K-B22DELL-P) is 2:1.

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.

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 4: 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.

Connection Model

Two methods (the static pinning fabric interface connection and the Port Channel fabric interface connection) allow the traffic from an end host to the parent switch to be distributed when going through the Cisco Nexus 2000 Series Fabric Extender.


**Note**

The Cisco Nexus 2248PQ Fabric Extender does not support the static pinning fabric interface connection.

Static Pinning Fabric Interface Connection

To provide a deterministic relationship between the host interfaces and the parent switch, you can configure the Fabric Extender to use individual fabric interface connections. This configuration connects the 10-Gigabit Ethernet fabric interfaces as shown in the following figure. You can use any number of fabric interfaces up to the maximum available on the model of the Fabric Extender.

Figure 5: Static Pinning Fabric Interface Connections



When the Fabric Extender is brought up, its host interfaces are distributed equally among the available fabric interfaces. As a result, the bandwidth that is dedicated to each end host toward the parent switch is never changed by the switch but instead is always specified by you.

**Note**

If a fabric interface fails, all its associated host interfaces are brought down and remain down until the fabric interface is restored.

You must use the **pinning max-links** command to create a number of pinned fabric interface connections so that the parent switch can determine a distribution of host interfaces. The host interfaces are divided by the number of the max-links and distributed accordingly. The default value is max-links 1.

**Caution**

Changing the value of the **max-links** is disruptive; all the host interfaces on the Fabric Extender are brought down and back up as the parent switch reassigns its static pinning.

The pinning order of the host interfaces is initially determined by the order in which the fabric interfaces were configured. When the parent switch is restarted, the configured fabric interfaces are pinned to the host interfaces in an ascending order by the port number of the fabric interface.

To guarantee a deterministic and sticky association across a reboot, you can manually redistribute the pinning.




Note The redistribution of the host interfaces will always be in an ascending order by the port number of the fabric interface.

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 6: 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/QSFP-module/port*

where

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

The chassis ID ranges from 100 to 199.



Note The chassis ID is required only to access a host interface on the Fabric Extender. A value of less than 100 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.
- *QSFP-module* identifies the 10G breakout Linecard Expansion Module (LEM).
- *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.

Fabric Extender Hardware

The Cisco Nexus 2000 Series Fabric Extender architecture allows hardware configurations with various host interface counts and speeds.

Chassis

The Cisco Nexus 2000 Series Fabric Extender is a 1 RU chassis that is designed for rack mounting. The chassis supports redundant hot-swappable fans and power supplies.

Ethernet Interfaces

There are 8 models of the Cisco Nexus 2000 Series Fabric Extender:

- The Cisco Nexus 2148T has 48 100BASE-T Ethernet host interfaces for its downlink connection to servers or hosts and 4 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent switch.
- The Cisco Nexus 2224TP has 24 100BASE-T/1000Base-T Ethernet host interfaces for its downlink connection to servers or hosts and 2 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent switch.
- The Cisco Nexus 2248PQ has 48 10-Gigabit Ethernet host interfaces with SFP+ interface adapters and 16 10-Gigabit Ethernet fabric interfaces corresponding to 4 QSFP interface adapters for its uplink connection to the parent switch.
- The Cisco Nexus 2232PP has 32 10-Gigabit Ethernet host interfaces with SFP+ interface adapters and 8 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent switch.
- Cisco Nexus N2332TQ—FEX supporting 32 10GBaseT host ports and 4 QSFP+ network ports.
- Cisco Nexus 2348UPQ—FEX for QSA (FET-10G, SFP-10G-SR, SFP-10G-ER).
- Cisco Nexus 2348TQ FEX (N2K-C2348TQ-10GE)
- The Cisco Nexus 2248TP has 48 100BASE-T/1000Base-T Ethernet host interfaces for its downlink connection to servers or hosts and 4 10-Gigabit Ethernet fabric interfaces with SFP+ interface adapters for its uplink connection to the parent switch.

The Cisco Nexus 2248TP-E has all the features of the Cisco Nexus 2248TP with these additional features:

- A larger buffer to absorb large bursts.
- Support for an ingress and egress queue-limit per port.
- Support for debug counters.
- Support for pause no-drop behavior over a cable distance of 3000 meters between the Fabric Extender and switch.
- Support for a user configurable shared-buffer.
- The Cisco Nexus B22 Fabric Extender for HP (NB22HP) has 16 1G/10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces.
- The Cisco Nexus B22 Fabric Extender for Fujitsu (NB22FTS) has 16 10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces.
- The Cisco Nexus B22 Fabric Extender for Dell (NB22DELL) has 16 1G/10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces.

- The Cisco Nexus B22 Fabric Extender for IBM (NB22IBM) has 14 1G/10-Gigabit Ethernet host interfaces. All host interfaces use all of the available fabric interfaces.

Speed and Duplex Mode

The table listed below shows the autonegotiation matrix for N2K-C2348TQ-10GE and N2K-C2332TQ-10GE fabric extenders.

Configuration N2K-C2348TQ-10GE and N2K-C2332TQ-10GE HIF (Speed/Duplex)	Configuration NIC (Speed/Duplex)	Resulting HIF status (Speed/Duplex)	Resulting NIC status (Speed/Duplex)	Comments
Table for 100 Mbps				
AUTO	AUTO	UP, 100 Mbps, Full duplex	UP, 100 Mbps, Full duplex	Assuming maximum capability of NIC is 100 Mbps, Full duplex
100 Mbps, Full duplex (force mode)	100 Mbps, Full duplex	UP, 100 Mbps, Full duplex	UP, 100 Mbps, Full duplex	Link DOWN in some scenarios. 1
100 Mbps, Full duplex (force mode)	100 Mbps, Half duplex	UP, 100 Mbps, Full duplex	UP, 100 Mbps, Half duplex	Duplex mismatch as per standard; results in collision errors. Hence, not a functional scenario.
100 Mbps, Full duplex (force mode)	AUTO	UP, 100 Mbps, Full duplex	UP, 100 Mbps, Half Duplex	Duplex mismatch as per standard; results in collision errors. Hence, not a functional scenario. (Link will be DOWN if NIC is not 100 Mbps, Half Duplex capable)
AUTO	100 Mbps, Full duplex	DOWN	DOWN	—
AUTO	100 Mbps, Half duplex	DOWN	DOWN	—
Table for 1G (1000 Mbps)				

Configuration N2K-C2348TQ-10GE and N2K-C2332TQ-10GE HIF (Speed/Duplex)	Configuration NIC (Speed/Duplex)	Resulting HIF status (Speed/Duplex)	Resulting NIC status (Speed/Duplex)	Comments
AUTO	AUTO	UP, 1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	Assuming maximum capability of NIC is 1000 Mbps, Full duplex
1000 Mbps, Full duplex	AUTO	UP, 1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	Assuming maximum capability of NIC is 1000 Mbps, Full duplex
AUTO	1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	—
1000 Mbps, Full duplex	1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	UP, 1000 Mbps, Full duplex	—
100 Mbps, Full duplex	1000 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch
1000 Mbps, Full duplex	100 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch
1000 Mbps, Full duplex	100 Mbps, Half duplex	DOWN	DOWN	Speed Mismatch
10000 Mbps, Full duplex	1000 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch
10000 Mbps, Full duplex	100 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch
10000 Mbps, Full duplex	100 Mbps, Half duplex	DOWN	DOWN	Speed Mismatch
Table for 10G (10000 Mbps)				
AUTO	AUTO	UP, 10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	Assuming maximum capability of NIC is 10000 Mbps, Full duplex
10000 Mbps, Full duplex	AUTO	UP, 10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	Assuming maximum capability of NIC is 10000 Mbps, Full duplex
AUTO	10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	—

Configuration N2K-C2348TQ-10GE and N2K-C2332TQ-10GE HIF (Speed/Duplex)	Configuration NIC (Speed/Duplex)	Resulting HIF status (Speed/Duplex)	Resulting NIC status (Speed/Duplex)	Comments
10000 Mbps, Full duplex	10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	UP, 10000 Mbps, Full duplex	—
100 Mbps, Full duplex	10000 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch
1000 Mbps, Full duplex	10000 Mbps, Full duplex	DOWN	DOWN	Speed Mismatch

¹ Refer to [CSCut35369](#) for more details.

Example: Configuring the Interface Speed Parameters

Configuration: AUTO

Configuring speed as AUTO (advertises all speeds and Full Duplex only)

```
switch(config)# interface ethernet 101/1/1
switch(config-if)# speed auto
```

Configuration: 100 (Force Mode)

Configuring speed as 100 Mbps

```
switch(config)# interface ethernet 101/1/1
switch(config-if)# speed 100
```

Configuration: 1G

Configuring speed as 1000 Mbps (This has autoneg enabled with 1000 Mbps)

```
switch(config)# interface ethernet 101/1/1
switch(config-if)# speed 1000
```

Configuration: 10G

Configuring speed as 10000 Mbps (This has autoneg enabled with 10000 Mbps)

```
switch(config)# interface ethernet 101/1/1
switch(config-if)# speed 10000
```

Associating a Fabric Extender to a Fabric Interface

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



Note The Fabric Extender may connect to the switch through a number of separate physical Ethernet interfaces or one port channel interface.



Note You must enable the Fabric Extender features with the **feature fex** command before you can configure and use a Fabric Extender that is connected to the parent switch.

Associating a Fabric Extender to an Ethernet Interface

Before You Begin

Ensure that you have enabled the Fabric Extender feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface ethernet slot/port Example: switch(config)# interface ethernet 1/40 switch(config)#	Specifies an Ethernet interface to configure. Note If this is a 10G breakout port, the <i>slot/port</i> syntax is <i>slot/QSFP-module/port</i> .
Step 3	switchport mode fex-fabric Example: switch(config-if)# switchport mode fex-fabric switch(config-if)#	Sets the interface to support an external Fabric Extender.
Step 4	fex associate FEX-number Example: switch(config-if)# fex associate 101 switch#	Associates the FEX number to the Fabric Extender unit attached to the interface. The range of the FEX number is from 100 to 199.

	Command or Action	Purpose
Step 5	show interface ethernet <i>port/slot</i> fex-intf Example: <pre>switch# show interface ethernet 1/40 fex-intf switch#</pre>	(Optional) Displays the association of a Fabric Extender to an Ethernet interface.

This example shows how to associate the Fabric Extender to an Ethernet interface on the parent device:

```
switch# configure terminal
switch(config)# interface ethernet 1/40
switch(config-if)# switchport mode fex-fabric
switch(config-if)# fex associate 101
switch(config)#
```

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

```
switch# show interface ethernet 1/40 fex-intf
Fabric          FEX
Interface      Interfaces
-----
Eth1/40        Eth101/1/48  Eth101/1/47  Eth101/1/46  Eth101/1/45
                Eth101/1/44  Eth101/1/43  Eth101/1/42  Eth101/1/41
                Eth101/1/40  Eth101/1/39  Eth101/1/38  Eth101/1/37
                Eth101/1/36  Eth101/1/35  Eth101/1/34  Eth101/1/33
                Eth101/1/32  Eth101/1/31  Eth101/1/30  Eth101/1/29
                Eth101/1/28  Eth101/1/27  Eth101/1/26  Eth101/1/25
                Eth101/1/24  Eth101/1/23  Eth101/1/22  Eth101/1/21
                Eth101/1/20  Eth101/1/19  Eth101/1/18  Eth101/1/17
                Eth101/1/16  Eth101/1/15  Eth101/1/14  Eth101/1/13
                Eth101/1/12  Eth101/1/11  Eth101/1/10  Eth101/1/9
                Eth101/1/8   Eth101/1/7   Eth101/1/6   Eth101/1/5
                Eth101/1/4   Eth101/1/3   Eth101/1/2
```

Associating a Fabric Extender to a Port Channel

Before You Begin

Ensure that you have enabled the Fabric Extender feature.

Procedure

	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.

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

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

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



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



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

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

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

```

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

```

Disassociating a Fabric Extender from an Interface

Before You Begin

Ensure that you have enabled the Fabric Extender feature.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	interface {ethernet slot/port port-channel channel} Example: <pre>switch(config)# interface port-channel 4 switch(config-if)#</pre>	Specifies the interface to configure. The interface can be an Ethernet interface or a port channel. Note If this is a 10G breakout port, the <i>slot/port</i> syntax is <i>slot/QSFP-module/port</i> .
Step 3	no fex associate Example: <pre>switch(config-if)# no fex associate</pre>	Disassociates the Fabric Extender unit attached to the interface.

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.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fex <i>FEX-number</i> 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> is from 1 to 199.
Step 3	description <i>desc</i> Example: <pre>switch(config-fex)# description Rack7A-N2K</pre>	(Optional) Specifies the description. The default is the string FEXxxxx where xxx is the FEX number. If the FEX number is 123, the description is FEX0123.
Step 4	no description Example: <pre>switch(config-fex)# no description</pre>	(Optional) Deletes the description.
Step 5	type <i>FEX-type</i> Example: <pre>switch(config-fex)# type N2248T</pre>	(Optional) Specifies the type of Fabric Extender. The <i>FEX-type</i> is one of the following: <ul style="list-style-type: none"> • N2148T—48 1000 Base-T Ethernet host interfaces and 4 10-Gigabit SFP+ Ethernet fabric interfaces module • N2224TP—24 100 Base-T/1000 Base-T Ethernet host interfaces and 2 10-Gigabit SFP+ Ethernet fabric interfaces module • N2232P and N2232TM—32 10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • N2232TP—32 10-Gigabit Base-T Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module • N2232TT—32 10-Gigabit Base-T Ethernet host interfaces and 8 10-Gigabit Base-T Ethernet fabric interfaces module • N2248T and N2248TP-E—48 100 Base-T/1000 Base-T Ethernet host interfaces and 4 10-Gigabit SFP+ Ethernet fabric interfaces module. • N2248PQ—48 10-Gigabit SFP+ Ethernet host interfaces and 16 10-Gigabit SFP+ Ethernet fabric interfaces module. • NB22HP—16 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module

	Command or Action	Purpose
		<ul style="list-style-type: none"> NB22FTS—16 10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module NB22DELL—16 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module NB22IBM—14 1G/10-Gigabit SFP+ Ethernet host interfaces and 8 10-Gigabit SFP+ Ethernet fabric interfaces module <p>The parent Cisco Nexus Series device remembers the type of the Fabric Extender in its binary configuration. When this feature is configured, the Fabric Extender is allowed to come online only if its type matches the configured FEX type.</p>
Step 6	no type Example: <pre>switch(config-fex)# no type</pre>	(Optional) 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	serial serial Example: <pre>switch(config-fex)# serial JAF1339BDSK</pre>	(Optional) 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	no serial Example: <pre>switch(config-fex)# no serial</pre>	(Optional) Deletes the serial number string.

Enabling the Fabric Extender Locator LED

The locator beacon LED on the Fabric Extender allows you to locate a specific Fabric Extender in a rack.

Procedure

	Command or Action	Purpose
Step 1	locator-led fex FEX-number Example: <pre>switch# locator-led fex 101</pre>	Turns on the locator beacon LED for a specific Fabric Extender.

	Command or Action	Purpose
Step 2	no locator-led fex <i>FEX-number</i> Example: switch# no locator-led fex 101	(Optional) Turns off the locator beacon LED for a specific Fabric Extender.

Redistributing the Links

When you provision the Fabric Extender with statically pinned interfaces, the downlink host interfaces on the Fabric Extender are pinned to the fabric interfaces in the order that they were initially configured. If you want to maintain a specific relationship of host interfaces to fabric interface across reboots, you should repin the links.

You may want to perform this function in these two situations:

- A change in the max-links configuration.
- If you need to maintain the pinning order of host interfaces to fabric interfaces.



Note

The Cisco Nexus 2248PQ Fabric Extender does not support the static pinning fabric interface connection.

Changing the Number of Links

If you initially configured a specific port on the parent switch, for example port 33, as your only fabric interface, all 48 host interfaces are pinned to this port. If you provision another port, for example 35, you must enter the **pinning max-links 2** command to redistribute the host interfaces. All host interfaces are brought down and host interfaces 1 to 24 are pinned to fabric interface 33 and host interfaces 25 to 48 are pinned to fabric interface 35.

Maintaining the Pinning Order

The pinning order of the host interfaces is initially determined by the order in which the fabric interfaces were configured. In this example, four fabric interfaces were configured in the following order:

```

switch# show interface ethernet 1/35 fex-intf
Fabric          FEX
Interface       Interfaces
-----
Eth1/35         Eth101/1/12  Eth101/1/11  Eth101/1/10  Eth101/1/9
                  Eth101/1/8   Eth101/1/7   Eth101/1/6   Eth101/1/5
                  Eth101/1/4   Eth101/1/3   Eth101/1/2   Eth101/1/1

switch# show interface ethernet 1/33 fex-intf
Fabric          FEX
Interface       Interfaces
-----
Eth1/33         Eth101/1/24  Eth101/1/23  Eth101/1/22  Eth101/1/21
    
```

```

Eth101/1/20 Eth101/1/19 Eth101/1/18 Eth101/1/17
Eth101/1/16 Eth101/1/15 Eth101/1/14 Eth101/1/13

switch# show interface ethernet 1/38 fex-intf
Fabric          FEX
Interface       Interfaces
-----
Eth1/38         Eth101/1/36 Eth101/1/35 Eth101/1/34 Eth101/1/33
                Eth101/1/32 Eth101/1/31 Eth101/1/30 Eth101/1/29
                Eth101/1/28 Eth101/1/27 Eth101/1/26 Eth101/1/25

switch# show interface ethernet 1/40 fex-intf
Fabric          FEX
Interface       Interfaces
-----
Eth1/40         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

```

The next time that you reboot the Fabric Extender, the configured fabric interfaces are pinned to the host interfaces in an ascending order by port number of the fabric interface. If you want to configure the same fixed distribution of host interfaces without restarting the Fabric Extender, enter the **fex pinning redistribute** command.

**Note**

It is a misconfiguration to have more fabric ports than pinning number even if the extra port is in DOWN state.

Redistributing Host Interfaces

**Caution**

This command disrupts all the host interface ports of the Fabric Extender.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fex pinning redistribute <i>FEX-number</i> Example: switch(config) # fex pinning redistribute 101 switch(config) #	Redistributes the host connections. The range of the <i>FEX-number</i> is from 100 to 199.

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</i> transceiver [fex-fabric]	Displays the SFP+ transceiver and diagnostic optical monitoring (DOM) information for the Fabric Extender uplinks.
show feature-set	Displays the status of the feature sets on the device.

Configuration Examples for the Fabric Extender

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

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

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

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

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

```
switch# show interface ethernet 1/40 transceiver
Ethernet1/40
  sfp is present
  name is CISCO-MOLEX INC
  part number is 74752-9026
  revision is A0
  serial number is MOC13321057
  nominal bitrate is 12000 Mbits/sec
  Link length supported for copper is 3 m(s)
  cisco id is --
  cisco extended id number is 4
```

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

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

```
Link length supported for 62.5/125mm fiber is 0 m(s)
cisco id is --
cisco extended id number is 4
```

Verifying the Chassis Management Information

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

Command or Action	Purpose
show environment fex {all <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  MinorThres  CurTemp     Status
(Celsius) (Celsius)
-----
1        Outlet-1    60           50          33          ok
1        Outlet-2    60           50          38          ok
1        Inlet-1     50           40          35          ok
1        Die-1      100          90          44          ok

Fan Fex: 101:
-----
Fan      Model              Hw      Status
-----
Chassis  N2K-C2148-FAN     --      ok
```

```
PS-1          --          --          absent
PS-2          NXK-PAC-400W --          ok
```

Power Supply Fex 101:

Voltage: 12 Volts

PS	Model	Power (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  : 0x1ale
EEPROM Size     : 65535
Block Count     : 3
FRU Major Type  : 0x6002
FRU Minor Type  : 0x0
OEM String      : Cisco Systems, Inc.
Product Number  : N2K-C2248TP-1GE
Serial Number   : JAF1339BDSK
Part Number     : 73-12748-01
Part Revision   : 11
Mfg Deviation   : 0
H/W Version    : 0.103
Mfg Bits       : 0
Engineer Use    : 0
snmpOID        : 9.12.3.1.9.78.3.0
Power Consump  : 1666
RMA Code       : 0-0-0-0
CLEI Code      : XXXXXXXXXTBDV00
VID            : V00
Supervisor Module specific block:
Block Signature : 0x6002
Block Version   : 2
Block Length    : 103
Block Checksum  : 0x2686
Feature Bits    : 0x0
HW Changes Bits : 0x0
Card Index     : 11016
MAC Addresses   : 00-00-00-00-00-00
Number of MACs  : 0
Number of EPLD : 0
```

```
Port Type-Num   : 1-48;2-4
Sensor #1       : 60,50
Sensor #2       : 60,50
Sensor #3       : -128,-128
Sensor #4       : -128,-128
Sensor #5       : 50,40
Sensor #6       : -128,-128
Sensor #7       : -128,-128
Sensor #8       : -128,-128
Max Connector Power: 4000
Cooling Requirement: 65
Ambient Temperature: 40
```

DISPLAY FEX 101 backplane sprom contents:

```
Common block:
Block Signature : 0xabab
Block Version   : 3
Block Length    : 160
Block Checksum  : 0x1947
EEPROM Size     : 65535
Block Count     : 5
FRU Major Type  : 0x6001
FRU Minor Type  : 0x0
OEM String      : Cisco Systems, Inc.
Product Number  : N2K-C2248TP-1GE
Serial Number   : SSI13380FSM
Part Number     : 68-3601-01
Part Revision   : 03
Mfg Deviation   : 0
H/W Version     : 1.0
Mfg Bits        : 0
Engineer Use    : 0
snmpOID        : 9.12.3.1.3.914.0.0
Power Consump   : 0
RMA Code        : 0-0-0-0
CLEI Code       : XXXXXXXXXXXTDBV00
VID             : V00
Chassis specific block:
Block Signature : 0x6001
Block Version   : 3
Block Length    : 39
Block Checksum  : 0x2cf
Feature Bits    : 0x0
HW Changes Bits : 0x0
Stackmib OID    : 0
MAC Addresses   : 00-0d-ec-e3-28-00
Number of MACs  : 64
OEM Enterprise  : 0
OEM MIB Offset  : 0
MAX Connector Power: 0
WWN software-module specific block:
Block Signature : 0x6005
Block Version   : 1
Block Length    : 0
Block Checksum  : 0x66
wnw usage bits:
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
```

```

00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00
License software-module specific block:
Block Signature : 0x6006
Block Version   : 1
Block Length    : 16
Block Checksum  : 0x86f
lic usage bits:
ff ff ff ff ff ff ff ff

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

```

Configuring the Cisco Nexus N2248TP-E Fabric Extender

The Cisco Nexus 2248TP-E Fabric Extender supports all of the CLI commands of the Cisco Nexus 2248TP Fabric Extender with additional commands to configure the following:

- Shared buffer (FEX global level)
- Queue limit in ingress direction (FEX global level and interface level)
- Queue limit in egress direction (FEX global level and interface level)
- No drop class over a distance of 3000 meters between the FEX and switch (FEX global level)

Configuring the Shared Buffer

The following are guidelines for the configuration of the shared buffer:

- Configuring the shared buffer is done at the FEX global level.
- The total available buffer is 32 MB which is shared in both the ingress and egress directions.
- The default size of the shared buffer is 25392KB.

However, when configuring an Ethernet-based pause no-drop class, the shared buffer size changes to 10800 KB. This change is required to increase the dedicated buffer that supports the pause no-drop class. The pause no-drop class does not use buffer space from the shared-pool.



Note Performing these commands might result in traffic disruption on all ports.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fex chassis_id Example: switch(config)# fex 100 switch(config-fex)#	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248TP-E shared-buffer-size buffer-size Example: switch(config-fex)# hardware N2248TP-E shared-buffer-size 25000	Specifies the shared buffer size (KB). The range of the <i>buffer-size</i> value is 10800 KB to 25392 KB. Note The hardware N2248TP-E shared-buffer-size command specifies the default shared buffer size of 25392 KB.

This example shows how to configure the shared buffer.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248TP-E shared-buffer-size 25000
switch(config-fex)#
```

Configuring the Queue Limit at the Global Level

The following are guidelines for the configuration of the queue limit:

- The tx queue limit specifies the buffer size used for each queue in the egress (n2h) direction.
- The rx queue limit specifies the buffer size used for each port in the ingress (h2n) direction.
- You can adjust the ingress queue limit when the FEX uplink experiences temporary congestion.
- You can adjust the egress queue limit for improved burst absorption or in a situation where there is a many to one traffic pattern.
- When you disable the tx queue limit, any output port is able to use the entire shared buffer.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fex chassis_id Example: <pre>switch(config)# fex 100 switch(config)#</pre>	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248TP-E queue-limit queue-limit tx rx Example: <pre>switch(config-fex)# hardware N2248TP-E queue-limit 83000 tx</pre>	Controls the egress (tx) or ingress (rx) queue tail drop threshold level on a FEX. <ul style="list-style-type: none"> • The default queue limit for tx (egress) is 4 MB. <p>Note The hardware N2248TP-E queue-limit command specifies the default tx queue limit.</p> <ul style="list-style-type: none"> • The default queue-limit for rx (ingress) is 1 MB. <p>Note The hardware N2248TP-E queue-limit rx command specifies the default rx queue limit.</p>

This example shows how to configure the queue limit.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248TP-E queue-limit 83000 tx
switch(config-fex)#
```

Configuring the Queue Limit at the Port Level

You can overwrite the global level configuration by configuring the queue limit at the port level.

You can also disable the queue limit at the port level.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	interface ethernet <i>chassis_id / slot/port</i> Example: switch(config)# interface ethernet 100/1/1	Enters interface configuration mode. Note If this is a 10G breakout port, the <i>slot/port</i> syntax is <i>slot/QSFP-module/port</i> .
Step 3	hardware N2248TP-E queue-limit queue-limit tx rx Example: switch(config-if)# hardware N2248TP-E queue-limit 83000 tx	Controls the egress (tx) or ingress (rx) queue tail drop threshold level on a FEX. <ul style="list-style-type: none"> • The default queue limit for tx (egress) is 4 MB. • The default queue limit for rx (ingress) is 1 MB.

This example shows how to configure the queue limit.

```
switch# configure terminal
switch(config)# interface ethernet 100/1/1
switch(config-if)# hardware N2248TP-E queue-limit 83000 tx
switch(config-if)#
```

Configuring the Uplink Distance

The Cisco Nexus N2248TP-E FEX supports a pause no-drop class up to a distance of 3000 meters between the FEX and the switch.

The default cable length between the FEX and the switch is 300 meters.

**Note**

When the pause no-drop class is not configured, the uplink distance configuration has no effect.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	fex chassis_id Example: switch(config)# fex 100 switch(config-fex)#	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248TP-E uplink-pause-no-drop distance distance-value Example: switch(config-fex)# hardware N2248TP-E uplink-pause-no-drop distance 3000	Specifies the no-drop distance between the FEX and the switch. The maximum distance is 3000 meters. Note The hardware N2248TP-E uplink-pause-no-drop distance command specifies the default 300 meter cable length.

This example shows how to configure the uplink distance.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248TP-E uplink-pause-no-drop distance 3000
switch(config-fex)#
```

Configuring the Cisco Nexus N2248PQ Fabric Extender

The Cisco Nexus 2248PQ Fabric Extender supports all of the CLI commands of the Cisco Nexus 2248TP Fabric Extender with additional commands to configure the following:

- Shared buffer (FEX global level)
- Load-balancing queues (FEX global level)
- No drop class over a distance of 3000 meters between the FEX and switch (FEX global level)

Configuring the Shared Buffer

The following are guidelines for the configuration of the shared buffer:

- Configuring the shared buffer is done at the FEX global level.
- The total available buffer is 16 MB which is shared in both the ingress and egress directions.
- The default size of the shared buffer is 10240KB.



Note

Performing these commands might result in traffic disruption on all ports.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fex chassis_id Example: switch(config)# fex 100 switch(config-fex)#	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248PQ shared-buffer-size buffer-size Example: switch(config-fex)# hardware N2248PQ shared-buffer-size 8096	Specifies the shared buffer size (KB). The range of the <i>buffer-size</i> value is 3072 KB to 10240 KB. Note The hardware N2248PQ shared-buffer-size command specifies the default shared buffer size of 10240 KB.

This example shows how to configure the shared buffer.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248PQ shared-buffer-size 8096
switch(config-fex)#
```

Configuring the Uplink Distance

The Cisco Nexus N2248PQ FEX supports a pause no-drop class up to a distance of 3000 meters between the FEX and the switch.

The default cable length between the FEX and the switch is 300 meters.

**Note**

When the pause no-drop class is not configured, the uplink distance configuration has no effect.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fex chassis_id	Enters configuration mode for the specified FEX.

	Command or Action	Purpose
	Example: switch(config)# fex 100 switch(config-fex)#	The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248PQ uplink-pause-no-drop distance distance-value Example: switch(config-fex)# hardware N2248PQ uplink-pause-no-drop distance 3000	Specifies the no-drop distance between the FEX and the switch. The maximum distance is 3000 meters. Note The hardware N2248PQ uplink-pause-no-drop distance command specifies the default 300 meter cable length.

This example shows how to configure the uplink distance.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248PQ uplink-pause-no-drop distance 3000
switch(config-fex)#
```

Configuring Slow Drain

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# fex chassis_id	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	switch(config-fex)# hardware fex slow-port-error-disable-time val	Specifies the FEX and the time threshold. The value of <i>fex</i> is the PID of the configured FEX. The range of <i>val</i> is from 200ms to 1000ms. The default value is 1000 ms.

This example shows how to configure the slow drain feature on the N2232P FEX:

```
switch# configure terminal
switch(config)# fex N2232P
switch(config-fex)# hardware N2232P slow-port-error-disable-time 500
```

Load-balancing queues at the FEX global level

The Cisco Nexus 2248PQ provides 8 load balancing queues. These load balancing queues are designed to resolve port congestion.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fex chassis_id Example: switch(config)# fex 100 switch(config)#	Enters configuration mode for the specified FEX. The range of the <i>chassis_id</i> value is 100 to 199.
Step 3	hardware N2248PQ uplink-load-balance-mode Example: switch(config-fex)# hardware N2248PQ uplink-load-balance-mode	Enables and disables load balancing queues at the FEX global level.

This example shows how to configure the load balance queues.

```
switch# configure terminal
switch(config)# fex 100
switch(config-fex)# hardware N2248PQ uplink-load-balance-mode
switch(config-fex)#
```