



Preparing the Switch for Configuring FCoE

- [Information About FCoE, on page 1](#)
- [Default Settings for FCoE, on page 5](#)
- [Guidelines and Limitations, on page 5](#)
- [Enabling and Disabling FCoE, on page 5](#)
- [Configuring the FC-Map , on page 5](#)
- [Configuring the Fabric Priority, on page 6](#)
- [Setting the Advertisement Interval, on page 6](#)
- [Disabling LAN Traffic on an FCoE Link, on page 7](#)
- [Additional References for FCoE, on page 8](#)

Information About FCoE

FCoE provides a method of transporting Fibre Channel traffic over a physical Ethernet connection. FCoE requires the underlying Ethernet to be full duplex and to provide lossless behavior for Fibre Channel traffic.



Note Lossless behavior on Ethernet is provided by using a priority flow control (PFC) mechanism that prevents packet loss during congestion conditions.

The Cisco NX-OS software supports T11-compliant FCoE on all 10-Gigabit and 40-Gigabit Ethernet interfaces.

FCoE and FIP

FCoE Initiation Protocol

The FCoE Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN. Cisco NX-OS switches support the Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) protocol for T11-compliant Generation 2 CNAs.

The following switches do not support Pre-FIP on Generation 1 CNAs:

FIP Virtual Link Instantiation

FIP is used to perform device discovery, initialization, and link maintenance. FIP performs the following tasks:

- FIP VLAN discovery— Discovers the FCoE VLAN that will be used by all other FIP protocols as well as by the FCoE encapsulation for Fibre Channel payloads on the established virtual link. FIP VLAN discovery occurs in the native VLAN used by the initiator or target to exchange Ethernet traffic. The FIP VLAN discovery protocol is the only FIP protocol running on the native VLAN; all other FIP protocols run on the discovered FCoE VLANs.
- FIP FCF discovery—When a FCoE device is connected to the fabric, it sends a Discovery Solicitation message. A Fibre Channel Forwarder (FCF) or a switch responds to the message with a Solicited Advertisement that provides an FCF MAC address to use for subsequent logins.
- FCoE virtual link instantiation— FIP defines the encapsulation of fabric login (FLOGI), fabric discovery (FDISC), logout (LOGO), and exchange link parameter (ELP) frames along with the corresponding reply frames. The FCoE devices use these messages to perform a fabric login.
- FCoE virtual link maintenance— FIP periodically sends maintenance messages between the switch and the CNA to ensure that the connection is still valid.

FCoE Frame Format

FCoE is implemented when the switch encapsulates a Fibre Channel frame in an Ethernet packet with a dedicated Ethernet type, 0x8906. The packet has a 4-bit version field. The other header fields in the frame (the source and destination MAC addresses, VLAN tags, and frame markers) are all standard Ethernet fields. Reserved bits pad the FCoE frame to the IEEE 802.3 minimum packet length of 64 bytes.

A Fibre Channel frame consists of 36 bytes of headers and up to 2112 bytes of data for a total maximum size of 2148 bytes. The encapsulated Fibre Channel frame has all the standard headers, which allow it to be passed to the storage network without further modification. To accommodate the maximum Fibre Channel frame in an FCoE frame, the class-fcoe is defined with a default Maximum Transmission Unit (MTU) of 2240 bytes.

VLAN Tagging for FCoE Frames

The Ethernet frames that are sent by the switch to the adapter include the IEEE 802.1Q tag. This tag includes a field for the class of service (CoS) value used by the priority flow control (PFC). The IEEE 802.1Q tag also includes a VLAN field.

The switch expects frames from a FIP T11-compliant CNA to be tagged with the VLAN tag for the FCoE VLAN. Frames that are not correctly tagged are discarded.



Note You cannot map VLAN 1 or the native VLAN to an FCoE VSAN.

FIP Ethernet Frame Format

FIP is encapsulated in an Ethernet packet with a dedicated EtherType, 0x8914. The packet has a 4-bit version field. Along with the source and destination MAC addresses, the FIP packet also contains a FIP operation code and a FIP operation subcode. The following table describes the FIP operation codes and subcodes.

Table 1: FIP Operation Codes

FIP Operation Code	FIP Subcode	FIP Operation
0x0001	0x01	Discovery Solicitation
	0x02	Discovery Advertisement
0x0002	0x01	Virtual Link Instantiation Request
	0x02	Virtual Link Instantiation Reply
0x0003	0x01	FIP Keepalive
	0x02	FIP Clear Virtual Links
0x0004	0x01	FIP VLAN Request
	0x02	FIP VLAN Notification

DCBX

Data Center Bridging Exchange Protocol

The Data Center Bridging Exchange (DCBX) protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX end points exchange request and acknowledgment messages. For flexibility, parameters are coded in a type-length-value (TLV) format. Cisco NX-OS switches support the Converged Enhanced Ethernet Data Center Bridging Exchange (CEE-DCBX) is supported on all T11-compliant Generation 2 CNAs.

DCBX runs on the physical Ethernet link between the switch and the CNA. By default, DCBX is enabled on Ethernet interfaces. When an Ethernet interface is brought up, the switch automatically starts to communicate with the CNA.

During the normal operation of FCoE between the switch and the CNA, DCBX provides link-error detection.

DCBX is also used to negotiate capabilities between the switch and the CNA and to send configuration values to the CNA.

The CNAs that are connected to a switch are programmed to accept the configuration values that are sent by the switch, allowing the switch to distribute configuration values to all attached CNAs, which reduces the possibility of configuration errors and simplifies CNA administration.

Lossless Ethernet

Standard Ethernet is a best-effort medium which means that it lacks any form of flow control. In the event of congestion or collisions, Ethernet drops packets. The higher level protocols detect the missing data and retransmit the dropped packets.

To properly support Fibre Channel, Ethernet has been enhanced with a priority flow control (PFC) mechanism.

Logical Link Up/Down

On a native Fibre Channel link, some configuration actions (such as changing the VSAN) require that you reset the interface status. When you reset the interface status, the switch disables the interface and then immediately reenables the interface.

Caution:

If an Ethernet link provides FCoE service, do not reset the physical link because this action is disruptive to all traffic on the link.

The logical link up/down feature allows the switch to reset an individual virtual link. The logical link down is signaled with a FIP Clear Virtual Link message.

Caution:

If the CNA does not support the logical link level up/down feature, the CNA resets the physical link, which means that all traffic on the Ethernet interface is disrupted.

Converged Network Adapters

Cisco NX-OS switches support the following CNA types:

- Hardware adapter
 - Works with the existing Fibre Channel host bus adapter (HBA) driver and Ethernet Network Interface Card (NIC) driver in the server.
 - Server operating system view of the network is unchanged; the CNA presents a SAN interface and a LAN interface to the operating system.
- FCoE software stack
 - Runs on existing 10-Gigabit Ethernet adapters.

The following Cisco NX-OS series and platforms support Generation 2 CNAs that use the FIP to exchange information about its available capabilities and to negotiate the configurable values with the switch:

To reduce configuration errors and simplify administration, the switch distributes the configuration data to all the connected adapters.

STP Lite

FCoE does not require full Spanning Tree Protocol (STP) because FCoE has no bridging functionality, which means that no STP loops are created in the network. STP Lite on FCoE interfaces ensures rapid convergence across the network by sending an agreement Bridge Protocol Data Unit (BPDU) whenever it receives a proposal BPDU. The FCoE link sends the identical agreement BPDU in response to either an Multiple Spanning Tree (MST) or a Per VLAN Rapid Spanning Tree Plus (PVRST+) proposal BPDU. Additionally, STP Lite suppresses the MAC address flushing function for FCoE VLANs.

STP Lite is enabled automatically by default across the entire device for FCoE VLANs as soon as the first FCoE VLAN comes up. At the same time, the system automatically converts all FCoE links as the STP-type normal ports. This feature runs only in FCoE VLANs.

Default Settings for FCoE

This table lists the default settings for FCoE parameters.

Table 2: Default FCoE Parameter Settings

Parameters	Default
FCoE feature	Not installed, disabled
FC-Map	0E.FC.00
Fabric priority	128
Advertisement interval	8 seconds

Guidelines and Limitations

Enabling and Disabling FCoE

Configuring the FC-Map

You can prevent data corruption due to cross-fabric talk by configuring an FC-Map that identifies the Fibre Channel fabric for this switch. When the FC-Map is configured, the switch discards the MAC addresses that are not part of the current fabric. An FCF can assign Fabric Provide MAC Addresses (FPMA) to the CNAs consisting of the FC-Map Value for the Fabric and the Fibre Channel ID (FCID) assigned during Fabric Login



Note In most deployments, changing the fc-map is not needed. We recommend that you use VSAN to VLAN mapping. For more information, see the [Configuring FCoE VLANs and Virtual Fibre Channel Interfaces](#) chapter.

Step 1 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 2 **fcoe fcmap fabric-map**

Example:

```
fcoe(config)# fcoe fcmap 0x0efc2a
```

Configures the global FC-Map. The default value is 0x0EFC00. The range is from 0x0EFC00 to 0x0EFCFF. Use the **no fcoe map** command to reset to the default value.

Example

Configuring the Fabric Priority

The FCoE switch advertises its priority. The priority is used by the CNAs in the fabric to determine the best switch to connect to.

Step 1 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 2 **fcoe fcf-priority fabric-priority**

Example:

```
fcoe(config)# fcoe fcf-priority 42
```

Configures the global fabric priority. The default value is 128. The range is from 0 (higher) to 255 (lower). Use the **no fcoe fcf-priority** command to reset the global fabric priority to the default value.

Example

Setting the Advertisement Interval

You can configure the interval for Fibre Channel fabric advertisement on the switch.

Step 1 **configure terminal**

Example:

```
fcoe# configure terminal
fcoe(config)#
```

Enters configuration mode.

Step 2 **fcoe fka-adv-period interval**

Example:

```
fcoe(config)# fcoe fka-adv-period 8
fcoe#
```

Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds.

Example

Disabling LAN Traffic on an FCoE Link

You can disable LAN traffic on an FCoE link.

DCBX allows the switch to send a LAN Logical Link Status (LLS) message to a directly-connected CNA. Enter the **shutdown lan** command to send an LLS-Down message to the CNA. This command causes all VLANs on the interface that are not enabled for FCoE to be brought down. If a VLAN on the interface is enabled for FCoE, it continues to carry SAN traffic without any interruption.

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **interface ethernet slot/port**

Example:

```
switch(config)# interface e 2/1
switch(config-if)#
```

Specifies an interface to configure, and enters interface configuration mode. Use **?** to view a list of supported interfaces.

Step 3 **shutdown lan**

Example:

```
switch(config-if)# shutdown lan
```

Shuts down Ethernet traffic on the interface. If the interface is part of an FCoE VLAN, the shutdown has no impact on the FCoE traffic. Use **no shutdown lan** to reenable Ethernet traffic on this interface.

Step 4 (Optional) **show interface**

Example:

```
switch(config-if)# show interface
```

Displays information about the interface.

Step 5 (Optional) **copy running-config startup-config**

Example:

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

Copies the running configuration to the startup configuration.

Additional References for FCoE

Related Documents

Related Topic	Document Title
Command reference	Cisco NX-OS FCoE Command Reference Guide, Nexus 7000 and MDS 9500
Configuration guide	
Cisco NX-OS licensing	

Standards and RFCs

Standard/RFC	Title
T11 FC BB-5	Fibre Channel Backbone 5

MIBs

MB	MIBs Link
	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/support