



CHAPTER 1

Configuring FCoE

Fibre Channel over Ethernet (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.

This chapter describes how to configure FCoE and includes the following sections:

- [Information About FCoE, page 1-1](#)
- [Configuring FCoE, page 1-4](#)
- [Configuring LLDP, page 1-7](#)
- [Verifying FCoE Configuration, page 1-8](#)

Information About FCoE

In Cisco Nexus 5000 Series switches, FCoE is supported on all 10-Gigabit Ethernet interfaces.

To use FCoE, the switch must be directly connected to the server and the server port must terminate the Ethernet with a converged network adapter.

This section includes the following topics:

- [Licensing Requirements, page 1-1](#)
- [Converged Network Adapters, page 1-2](#)
- [DCBX Capabilities, page 1-2](#)
- [DCE Bridging Capability Exchange Protocol, page 1-3](#)
- [DCBX Feature Negotiation, page 1-3](#)
- [Ethernet Frame Formats, page 1-4](#)

Licensing Requirements

On Cisco Nexus 5000 Series switches, FCoE capability is included in the Storage Protocol Services License.

Before using FCoE capabilities, ensure that:

- The correct license is installed (N5010SS or N5020SS).
- FCoE is activated by entering the **feature fcoe** command in configuration mode.

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Converged Network Adapters

The following types of converged network adapters (CNAs) are available:

- Hardware adapter
 - Works with the existing FC HBA driver and LAN 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.

FCoE offers a number of optional capabilities. The available capabilities and their configurable values are negotiated between the switch and the adapter. To accomplish this, the adapter and the switch exchange information using the Data Center Bridging Exchange Protocol (DCBX).

To reduce configuration errors and simplify administration, you can configure the switch to distribute the configuration data to all the connected adapters.

DCBX Capabilities

The DCBX capabilities supported by Cisco Nexus 5000 Series switches are described in the following topics:

- [FCoE, page 1-2](#)
- [Priority Flow Control, page 1-2](#)
- [Logical Link Up/Down, page 1-3](#)

FCoE

By default, each Ethernet interface attempts to enable FCoE capability by advertising the capability to the adapter.

If the FCoE negotiation fails, you can configure the switch to disable FCoE or to force-enable FCoE for this interface.

Priority Flow Control

The priority flow control (PFC) capability allows you to apply pause functionality to specific classes of traffic. PFC decides whether to apply pause based on the IEEE 802.1p CoS value. When the switch enables PFC, it configures the connected adapter to apply the pause functionality to packets with specific CoS values.

By default, the switch negotiates to enable the PFC capability. If the negotiation succeeds, PFC is enabled and link-level flow control remains disabled (regardless of its configuration settings).

If the PFC negotiation fails, you can either force PFC to be enabled on the interface or you can enable IEEE 802.x link-level flow control.

Link-level flow control can be enabled on the interface only if PFC is not enabled. For additional information about link-level flow control, see the [“Link-Level Flow Control” section on page 1-3](#).

Send feedback to nx5000-docfeedback@cisco.com

Logical Link Up/Down

On a native Fibre Channel link, some configuration actions (such as changing the VSAN) require you to reset the interface status. The switch achieves the reset by disabling the interface, and then immediately reenabling the interface.

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

The logical link up/down feature provides the ability to reset an individual virtual link. The switch sends a DCBX message to request the adapter to reset only the virtual Fibre Channel interface.



Note

If the adapter does not support the logical link level up/down feature, the adapter resets the physical link. In this case, all traffic on the Ethernet interface is disrupted.

DCE Bridging Capability Exchange Protocol

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.

DCBX runs on the physical Ethernet link between the Cisco Nexus 5000 Series switch and the converged network adapter on the server. By default, DCBX is enabled on Ethernet interfaces. When an Ethernet interface is brought up, the switch automatically starts to communicate with the adapter.

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

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

Adapters connected to a Cisco Nexus 5000 Series switch are programmed to accept the configuration values sent by the switch, allowing the switch to distribute configuration values to all attached adapters. The capability reduces the possibility of configuration error and simplifies administration of the adapters.

The switch provides a manual configuration override for each parameter negotiated using DCBX. The override takes effect if the adapter does not support DCBX, or if the adapter does not support the specific capability requested by the switch.

DCBX Feature Negotiation

The switch and adapter exchange capability information and configuration values. The Cisco Nexus 5000 Series switches support the following capabilities:

- FCoE

If the adapter supports FCoE capability, the switch sends the IEEE 802.1p CoS value to be used with FCoE packets.

- Priority Flow Control (PFC)

If the adapter supports PFC, the switch sends the IEEE 802.1p CoS values to be enabled with PFC.

- Ethernet logical link up and down signal
- FCoE logical link up and down signal

Send feedback to nx5000-docfeedback@cisco.com

The following rules determine whether the negotiation results in a capability being enabled:

- If a capability and its configuration values match between the switch and the adapter, the feature is enabled.
- If the capability matches, but the configuration values do not match:
 - If the adapter is configured to accept the switch configuration value, the capability is enabled using the switch value.
 - If the adapter is not configured to accept the switch configuration value, the capability remains disabled.
- If the adapter does not support the DCBX capability, the capability remains disabled.
- If the adapter does not implement DCBX, all capabilities remain disabled.

**Note**

The Cisco Nexus 5000 Series switch provides CLI commands to manually override the results of the negotiation with the adapter. On a per-interface basis, you can force capabilities to be enabled or disabled. For additional information, see the [“Configuring FCoE” section on page 1-4](#).

Ethernet Frame Formats

Ethernet frames sent by the switch to the adapter may include the IEEE 802.1Q tag. This tag includes a field for the CoS value used by PFC. The IEEE 802.1Q tag also includes a VLAN field, currently not used by the Cisco Nexus 5000 Series switch.

The switch will always accept tagged or untagged frames from the adapter.

If the FCoE capability is enabled between the switch and the adapter:

- All Ethernet frames sent to the adapter include IEEE 802.1Q tags.

The CoS field is set to the CoS value of the associated system class. For additional information, see the [“System Classes” section on page 1-2](#).

- The switch accepts tagged or untagged frames from the adapter.

If the FCoE capability is disabled between the switch and the adapter:

- Ethernet frames sent to the adapter do not include IEEE 802.1Q tags.
- The switch accepts tagged or untagged frames from the adapter.

Configuring FCoE

This section contains the following topics:

- [Enabling FCoE, page 1-5](#)
- [Enabling FCoE on Ethernet Interfaces, page 1-5](#)
- [Configuring Priority Flow Control, page 1-6](#)
- [Configuring IEEE 802.3x Link-Level Flow Control, page 1-6](#)

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Enabling FCoE

You need to enable the FCoE capability after the FC_FEATURES_PKG is installed. To enable FCoE, perform this task:

	Command	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# feature fcoe	Enables the FCoE capability.

The following example shows how to enable FCoE on the switch:

```
switch# configure terminal
switch(config)# feature fcoe
switch(config)# 2008 Nov 11 20:43:38 switch %$ VDC-1 %$ %PFMA-2-FC_LICENSE_DESIRE:
FCoE/FC feature will be enabled after the configuration is saved followed by a reboot
```



Note

After you enable the FCoE capability, you must reboot the Cisco Nexus 5000 Series switch before you can use the features.

Enabling FCoE on Ethernet Interfaces

By default, 10-Gigabit Ethernet interfaces negotiate FCoE capability with the adapter. You can override the negotiation result by force-enabling the FCoE capability. To force-enable the FCoE capability, perform this task:

	Command	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode for the specified interface.
Step 3	switch(config-if)# fcoe mode {auto on}	When you set the FCoE parameter to auto , the interface negotiates the setting with the connected adapter. FCoE will not be enabled on the interface if the adapter does not support FCoE. If the mode is set it to on , the negotiation result is ignored and FCoE is set to enabled on the interface.

The following example shows how to force-enable FCoE for an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# fcoe mode on
```

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

To disable the FCoE capability, perform this task:

Command	Purpose
switch(config-if)# no fcoe mode [auto on]	Disables FCoE capability for this interface.

This example shows how to disable FCoE for an Ethernet interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/4
switch(config-if)# no fcoe mode auto
```

The **fcoe** command can only be applied to a physical Ethernet interface.

Configuring Priority Flow Control

By default, the Ethernet interfaces negotiate PFC capability with the adapter. You can override the negotiation result by force-enabling the PFC capability.

To force-enable the PFC capability, perform this task:

	Command	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Specifies the interface to be changed.
Step 3	switch(config-if)# priority-flow-control mode {auto on }	Sets PFC mode for the selected interface. Specify auto to negotiate PFC capability. Specify on to force-enable PFC.

The following example shows how to force-enable PFC on an interface:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# priority-flow-control mode on
```

To disable PFC capability for this interface, perform this task:

Command	Purpose
switch(config-if)# no priority-flow-control mode [auto on]	Disables the PFC capability for this interface.

Configuring IEEE 802.3x Link-Level Flow Control

By default, link-level flow control capability on Ethernet interfaces is disabled. Only enable the link-level flow control capability if PFC is disabled on the interface.

To configure link-level flow control, see the [“Configuring IEEE 802.3x Link-Level Flow Control” section on page 1-8](#).

Send feedback to nx5000-docfeedback@cisco.com

Configuring LLDP

This section shows how to configure LLDP both globally and on individual interfaces. This section includes the following topics:

- [Configuring Global LLDP Commands, page 1-7](#)
- [Configuring Interface LLDP Commands, page 1-8](#)

Configuring Global LLDP Commands

You can set global LLDP settings. These settings include the length of time before discarding LLDP information received from peers, the length of time to wait before performing LLDP initialization on any interface, and the rate at which LLDP packets are sent.

To configure LLDP settings, perform this task:

	Command	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# lldp { holdtime <i>seconds</i> reinit <i>seconds</i> timer <i>seconds</i> }	Configures LLDP options. Use the holdtime option to set the length of time (10 to 255 seconds, default 120 seconds) that a device should save LLDP information received before discarding it. Use the reinit option to set the length of time (1 to 10 seconds, default 2 seconds) to wait before performing LLDP initialization on any interface. Use the timer option to set the rate (5 to 254 seconds, default 30 seconds) at which LLDP packets are sent.

The following example shows how to set LLDP timer option to 15 seconds:

```
switch# configure terminal
switch(config)# lldp timer 15
```

To reset LLDP settings, perform this task:

	Command	Purpose
	switch(config)# no lldp { holdtime reinit timer }	Reset the LLDP values to their defaults.

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Configuring Interface LLDP Commands

To configure the LLDP feature for a physical Ethernet interface, perform this task:

	Command	Purpose
Step 1	switch# configure terminal	Enters configuration mode.
Step 2	switch(config)# interface <i>type slot/port</i>	Selects the interface to change.
Step 3	switch(config-if)# [no] lldp { receive transmit }	Sets the selected interface to either receive or transmit. The no form of the command disables the LLDP transmit or receive.

The following example shows how to set an interface to transmit LLDP packets:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# lldp transmit
```

The following example shows how to configure an interface to disable LLDP:

```
switch# configure terminal
switch(config)# interface ethernet 1/2
switch(config-if)# no lldp transmit
switch(config-if)# no lldp receive
```

You can only apply this command to a physical Ethernet interface.

Verifying FCoE Configuration

To verify FCoE configuration information, perform one of these tasks:

Command	Purpose
switch# show fcoe	Displays whether FCoE is enabled on the switch.
switch# show lldp	Displays LLDP configuration.

The following example shows how to verify that the FCoE capability is enabled:

```
switch# show fcoe
FCoE/FC feature is desired.
```

The following example shows how to display LLDP interface information:

```
switch# show lldp interface ethernet 1/2
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE

Port MAC address: 00:0d:ec:a3:5f:48

Remote Peers Information
```


Send feedback to nx5000-docfeedback@cisco.com

No remote peers exist

The following example shows how to display LLDP neighbor information:

```
switch# show lldp neighbors
```

The following example shows how to display LLDP timer information:

```
switch# show lldp timers
```

```
LLDP Timers
```

```
holdtime 120 seconds
```

```
reinit 2 seconds
```

```
msg_tx_interval 30 seconds
```

The following example shows how to display LLDP counters:

```
switch# show lldp traffic
```

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)