

Brocade Fabric OS FCoE Configuration Guide, 8.2.0

Supporting Fabric OS 8.2.0

Copyright © 2017 Brocade Communications Systems LLC. All Rights Reserved.

Brocade and the stylized B logo are among the trademarks of Brocade Communications Systems LLC. Broadcom, the pulse logo, and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Limited and/or its subsidiaries

Brocade, a Broadcom Limited Company, reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Brocade is believed to be accurate and reliable. However, Brocade does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

Configuring FCoE.....	5
FCoE overview.....	5
Supported platforms.....	7
License requirements.....	7
FCoE port types.....	7
Supported FCoE topologies.....	7
FCoE-to-FC direct-attach topology.....	7
FCoE to FC through FSB topology.....	8
Configuring Ethernet ports in nonbreakout mode.....	9
Configuring Ethernet ports in breakout mode.....	10
FCoE over dynamic port channels.....	12
Configuring the LACP protocol and port channels.....	12
Configuring global LLDP parameters.....	12
Configuring and activating an LLDP profile for a group of ports.....	14
FCoE logical ports (E_Nodes).....	14
Configuring an FCoE fabric map.....	15
Enabling the FCoE protocol on Ethernet ports.....	16
Enabling the FCoE protocol on Ethernet port channels.....	16
Displaying FCoE fabric details.....	17
Resetting the configuration to the default.....	18
Duplicate port WWN support for FCoE devices.....	19
Configuring zoning with FCoE.....	19
Configuring Virtual Fabrics with FCoE.....	20
Configuring FCoE with FCR.....	20
Considerations and limitations.....	20
HA configuration considerations and limitations.....	21
Upgrade considerations and limitations.....	21
Downgrade considerations and limitations.....	21
FCoE scalability limits.....	21

Configuring FCoE


• FCoE overview.....	5
• Supported FCoE topologies.....	7
• Configuring Ethernet ports in nonbreakout mode.....	9
• Configuring Ethernet ports in breakout mode.....	10
• FCoE over dynamic port channels.....	12
• Configuring global LLDP parameters.....	12
• FCoE logical ports (E_Nodes).....	14
• Configuring an FCoE fabric map.....	15
• Duplicate port WWN support for FCoE devices.....	19
• Configuring zoning with FCoE.....	19
• Configuring Virtual Fabrics with FCoE.....	20
• Configuring FCoE with FCR.....	20
• Considerations and limitations.....	20
• FCoE scalability limits.....	21

FCoE overview

Fabric OS 8.2.0 and later allow you to connect FCoE-capable ports from hosts to the FCoE-enabled ports on the FC32-64 port blades in the Brocade X6 chassis as part of a unified FC fabric. FC and Ethernet storage traffic is allowed to share the same underlying network and thus extend to FCoE hosts and switches. FC frames are encapsulated in an Ethernet packet, and they are sent from one FCoE-aware device across a storage network to a second FC or FCoE-aware device. The FCoE-aware devices can be FC N_Port devices, such as servers and FCoE-aware switches providing FC fabric services and FCoE-to-FC bridging. The Brocade FC32-64 port blade acts as the FCoE host communication termination point to an FC target residing within the FC fabric.

The following are the supported FCoE-aware entities:

- E_Node—An E_Node or FCoE node is an FCoE device that supports FCoE VN_Ports such as hosts and target devices, which are referred to as E_Node devices.
- FCoE forwarder (FCF)—An FCoE forwarder is an FCoE switch that supports FCoE VF_Ports, which is the equivalent of an FC switch.
- FCoE Initialization Protocol (FIP)—FCoE-to-FC bridging allows FIP to be terminated at the first FCoE-capable switch that acts as an FCF. The connectivity between the host and the FCF is a CEE-capable network.
- A top of the rack (TOR) or an FIP snooping bridge (FSB) is an interconnect between the storage network and the Ethernet network or Ethernet hosts.
- Keep Alive Protocol (KAP)—KAP is used to track the following keep alive transmit intervals:
 - FIP default transmit interval: 8 seconds.
 - LLDP default transmit interval: 30 seconds.
 - LACP default longer transmit interval: 30 seconds.



NOTE

The intermediate switching devices in the Ethernet network (the Ethernet cloud in the diagram) need not be FCoE aware. The Ethernet devices merely route the FCoE traffic based on the Ethernet destination address in the FCoE packet. The only requirement is that the Ethernet network be CEE capable.

NOTE

Since the FC frames are larger than the default Ethernet frames, FCoE connectivity requires that any Ethernet network used for the transport of FCoE traffic support Ethernet jumbo frames.

NOTE

There is no STP support on FCoE-capable switches and blades. All FCoE-capable blades and switches act as FCoE termination points. Hence, ensure that there are no loops between FSB or TOR and FCF.

Supported platforms

The following ports support FCoE connectivity.

- All 16 QSFP ports on the Brocade FC32-64 blade in Brocade X6-4 and Brocade X6-8 chassis.
 - All the ports are FlexPorts, which you can configure as FC or Ethernet ports manually. The default port mode is FC.

NOTE

The Brocade FC32-64 does not interoperate with the older FCOE10-24 blade or Brocade 8000 switch.

License requirements

No special license is required to connect FCoE-capable storage devices, hosts, or switches.

FCoE port types

The following port types can be configured to connect to FCoE hosts and storage devices.

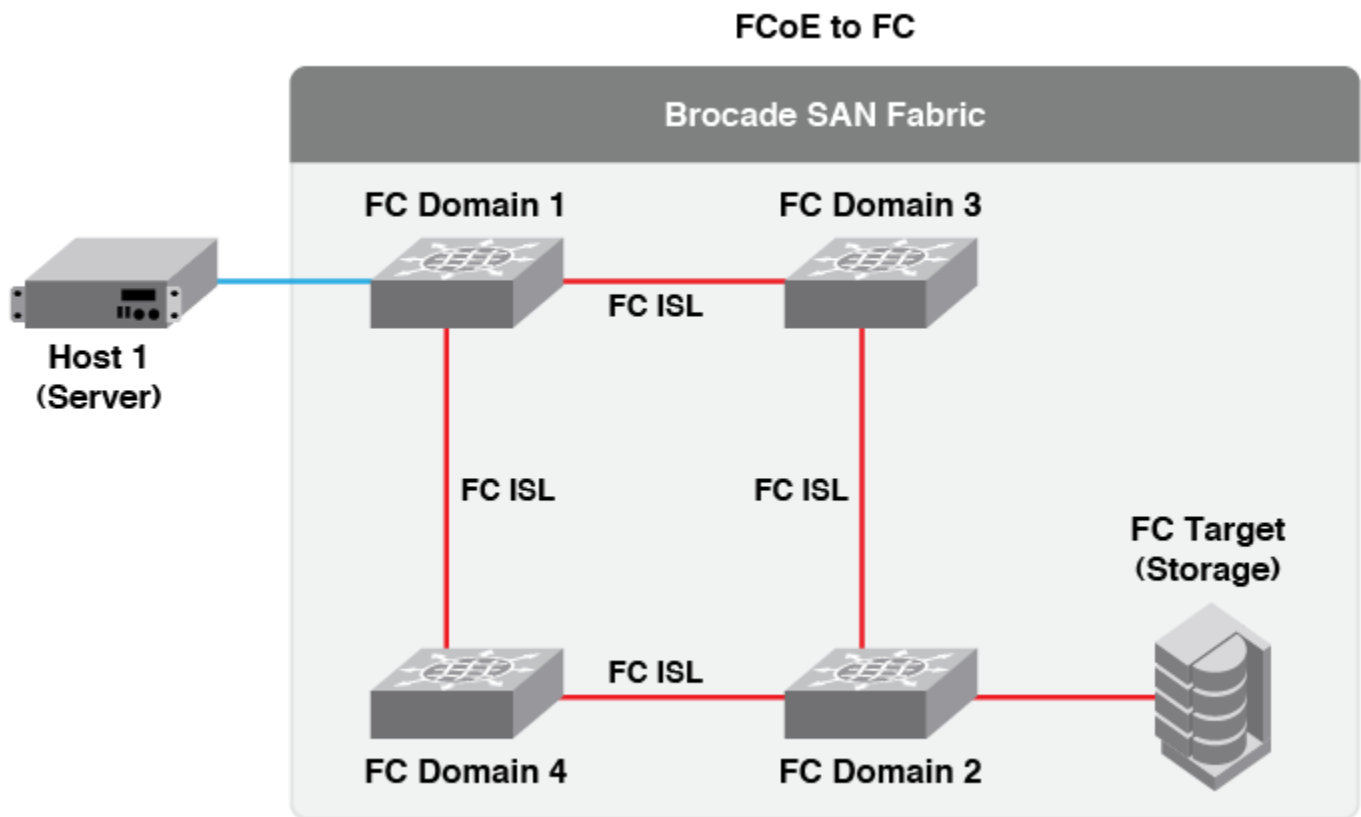
- VF_Port—The FCoE equivalent of a Fibre Channel F_Port.
- VN_Port—The FCoE equivalent of a Fibre Channel N_Port.

Supported FCoE topologies

FCoE-to-FC direct-attach topology

In this topology, there can be an FCoE host (for example, Host 1) connected via a rack server directly to an FCoE-capable switch and an FC target (for example, FC Target) connected to an FC switch within the FC fabric. This is FCoE-to-FC bridging where the FCF acts as an FCoE-to-FC bridge. All FIP communications are terminated at the connected FCF. FCoE data traffic is de-capsulated at the FCF and traverses the FC fabric as pure FC traffic to the FC target.

FIGURE 1 FCoE-to-FC bridging



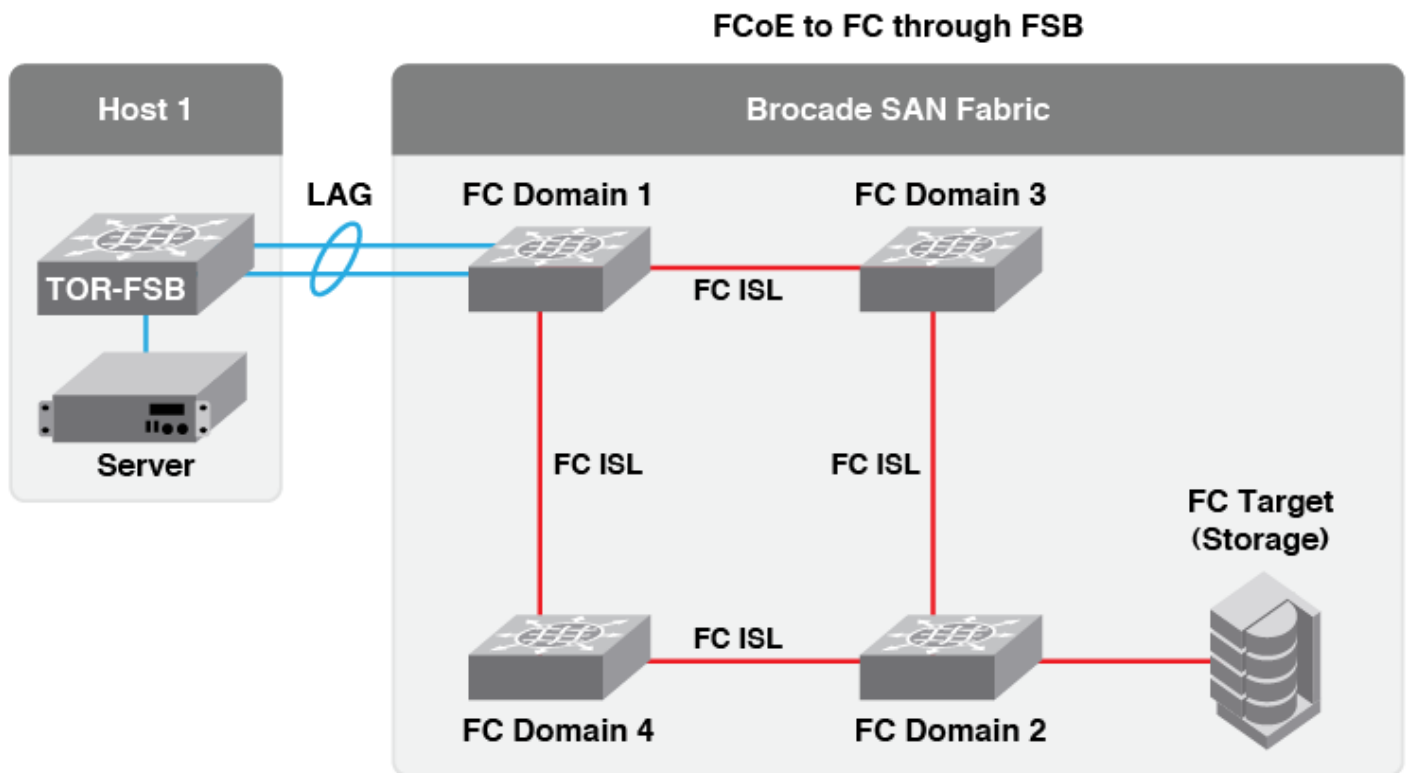
FCoE to FC through FSB topology

In this topology, there can be an FCoE host (for example, Host 1) connected via a rack server through a top of the rack (TOR) that acts as a FIP snooping bridge (FSB) to the FCoE-capable switch and an FC target (for example, FC Target) connected to an FC switch within the FC fabric. This is FCoE-to-FC bridging where the FCF acts as an FCoE-to-FC bridge. All FIP traffic is terminated at the connected FCF. FCoE data traffic is de-capsulated at the FCF and traverses the FC fabric as pure FC frames to the FC target. All security ACLs are installed by the FSB to provide a secure access to the storage.

NOTE

TOR-FSB to FC Domain 1 can either be a single link or a LAG.

FIGURE 2 FCoE to FC through FSB

**NOTE**

A FIP snooping bridge is not a strict requirement. The intermediate switch can be any DCB-capable switch that is capable of forwarding FIP frames to the FCF and back to the Converged Network Adapter (CNA).

Configuring Ethernet ports in nonbreakout mode

By default, all ports in the FC32-64 port blade function in FC mode. You must change the mode of the ports to ETH before you plug in FCoE-capable devices, hosts, or switches. To convert a port to ETH mode, follow these steps:

NOTE

Enabling a port, disabling a port, changing a port to ETH mode, and changing the port speed are applied to all four ports within a QSFP. All four ports within a quad in the ASIC must be in the same logical switch before the ports are changed to ETH mode or back to FC mode. They can be moved to different logical switches after converting to ETH mode or FC mode.

1. Install the Ethernet QSFP transceivers and the cables in nonbreakout mode.
2. Disable the ports that you want to configure in ETH mode. You can use the `portcfgflexport --quadshow` command to display the ports within a quad in the ASIC.

```
switch#admin> portcfgflexport --quadshow 9/60
```

UPort	slot	Port	Protocol	QuadNum
860	9	60	ETH	60

```

862    9    62    ETH    60
863    9    63    ETH    60
861    9    61    ETH    60

```

- Run the **portcfgflexport --proto eth** command to change to ETH mode. All four ports in the quad are changed to ETH mode.

```
switch#admin> portcfgflexPort --proto eth 9/60
```

If you need to change the ports back to FC mode:

```
switch#admin> portcfgflexPort --proto fc 9/60
```

NOTE

A port cannot be converted to ETH mode if FMS mode is enabled or if the port is in a FICON logical switch.

- By default, the port speed is set to 40G for QSFP ports in nonbreakout mode. If required, you can run the **portcfgspeed** command to configure the Ethernet port speed.

```
switch#admin> portcfgspeed 9/60 40
```

- Enable the ports. For example, run the following commands to enable a quad port and display the port status.

```

switch#admin> portenable 9/60
switch#admin> switchshow
<output truncated>
 344  9  60  01d0c0  id 40G  Online  ETH
 345  9  61  01d1c0  id 40G  In_Sync ETH  Disabled (Persistent) (QSFP Secondary port)
 346  9  62  01d2c0  id 40G  In_Sync ETH  Disabled (Persistent) (QSFP Secondary port)
 347  9  63  01d3c0  id 40G  In_Sync ETH  Disabled (Persistent) (QSFP Secondary port)
<output truncated>

```

- Run the **portcfgshow** command to display the port status.

```

switch#admin> portcfgshow 9/40
Area Number:          252
Octet Speed Combo:   1 (32G|16G|8G|4G)
Eth Speed Level:     10G
Encryption:          OFF
<output truncated>
10G/16G FEC:         ON
16G FEC via TTS:     OFF
Flex Port:           ETH
QSFP Breakout:       NO

```

- Now, you can run any other port configuration commands.

NOTE

Not all port configuration commands supported in FC mode are supported in ETH mode. Refer to the *Brocade Fabric OS Command Reference* for more details.

Configuring Ethernet ports in breakout mode

By default, the QSFP ports that are configured in ETH mode operate in nonbreakout mode. Whenever a QSFP Ethernet port is configured in nonbreakout mode, only the primary port is enabled and the remaining three ports in the QSFP are persistently disabled. To convert a QSFP Ethernet port from nonbreakout mode to breakout mode, follow these steps:

NOTE

A QSFP port in FC mode does not require breakout mode configuration to operate in breakout mode. All QSFP ports in FC mode can automatically recognize the breakout mode depending on the transceivers plugged in.

1. Insert the Ethernet QSFP transceiver in the QSFP port and install the breakout cables.
2. Disable the QSFP ports that you want to configure in breakout mode. You can use the **portcfgflexport --quadshow** command to display the ports within a QSFP.

```
switch#admin> portcfgflexport --quadshow 9/60
9/60 (Primay)
9/61
9/62
9/63
```

3. Convert the port to Ethernet mode using the **portcfgflexport --proto eth** command.

```
switch#admin> portcfgflexPort --proto eth 9/60
```

If you need to change the ports back to FC mode,

```
switch#admin> portcfgflexPort --proto fc 9/60
```

4. Convert the port to breakout mode using the **portcfgbreakout --enable** command, and then enable the ports. You can run the breakout mode command on any one port from the group of four ports within a QSFP.

```
switch#admin> portcfgbreakout --enable 9/60
switch#admin> portenable 9/60
switch#admin> switchshow
<output truncated>
 252  9  60  02fc00  id  10G      Online  ETH
 253  9  61  02fd00  id  10G      Online  ETH
 254  9  62  02fe00  id  10G      Online  ETH
 255  9  63  02ff00  id  10G      Online  ETH
<output truncated>
```

If you have installed a 100G QSFP transceiver, you must enable breakout mode. Nonbreakout mode is not supported with this transceiver. If you need to convert the QSFP Ethernet port back to nonbreakout mode, run the **portcfgbreakout --disable** command.

5. By default, the port speed is set to 10G for each port within a QSFP in breakout mode. If you have installed a 100G QSFP transceiver, you can set the individual port speed to 25G using the **portcfgspeed** command in breakout mode.

```
switch#admin> portcfgspeed 9/60 25
```

6. Run the **portcfgshow** command to display the status of the ports within the QSFP.

```
switch#admin> portcfgshow 9/60
Area Number:          253
Octet Speed Combo:    1 (32G|16G|8G|4G)
Eth Speed Level:      25G
Speed Level:          AUTO (HW)
<output truncated>
Encryption:           OFF
10G/16G FEC:          ON
16G FEC via TTS:      OFF
Flex Port:             ETH
QSFP Breakout:        YES
```

NOTE

Configuration download fails if the ports within a QSFP are not in the same logical switch.

FCoE over dynamic port channels

Link Aggregation Control Protocol (LACP) manages the dynamic port-channel creation, port-channel member addition, and port channel status. LACP dynamic port channels can be created with member ports spanning multiple blades. However, the port channels cannot have member ports spanning multiple logical switches. FCoE is not supported on static port channels. Port channels cannot have members with mixed port speeds. All member ports must be of the same port speed. A port channel supports only a single FCoE VLAN at any given point in time based on the FCoE configuration on the port channel.

Configuring the LACP protocol and port channels

To configure the LACP protocol and port channels, follow these steps:

NOTE

For FCoE connectivity, only dynamic port channels are supported.

1. (Optional) Configure the LACP protocol system priority using the **lACP -config -sysprio** command.

```
switch#admin> lACP --config -sysprio 100
```

2. Create dynamic port channels using the **portchannel --create** command.

```
switch#admin> portchannel --create po10 -key 10 -type dynamic -speed 10G
```

3. Add the required ports as members to the port channel using the **portchannel --add** command.

```
switch#admin> portchannel --add po10 -port 4/23-26 -timeout S
```

4. The port channel is enabled by default. If required, you can enable the port channels using the **portchannel --enable** command.

```
switch#admin> portchannel --enable po10
switch#admin> portchannel --disable po10
```

- a. You can use the **portchannel --show -dynamic** command to display the port channels.

```
switch#admin> portchannel --show -dynamic
Name                Type           Oper-State   Port-Count   Member Ports
-----
LAG_0               Dynamic        Offline      0
LAG_1               Dynamic        Offline      0
LAG_2               Dynamic        Offline      0
LAG_3               Dynamic        Online       4             3/15*, 3/16,
                                     3/17, 3/18
LAG_4               Dynamic        Offline      0
LAG_5               Dynamic        Offline      0
LAG_6               Dynamic        Offline      0
LAG_7               Dynamic        Offline      0
LAG_8               Dynamic        Offline      0
LAG_9               Dynamic        Offline      0
LAG_10              Dynamic        Offline      0
```

Configuring global LLDP parameters

The Link Level Discovery Protocol (LLDP) discovers the Ethernet neighbor at the link level, negotiates with the peer for configured parameters, negotiates DCBX capabilities for CEE, and manages the CEE status of the port. To enable FCoE connectivity, LLDP primarily maintains the FCoE priority and logical link status, and negotiates them with the peer device. The default transmit interval is 30 seconds.

NOTE

LLDP is enabled by default when a port is converted to ETH mode and the global parameters are applied to all ports in ETH mode.

To configure LLDP global parameters, follow these steps:

1. LLDP is enabled by default. If LLDP is not enabled, enable the LLDP protocol on the switch using the **lldp --enable** command.

```
switch#admin> lldp --enable
```

2. Configure the LLDP global parameters using the **lldp --config** command.

```
switch#admin> lldp --config -sysname Anakin_19
switch#admin> lldp --config -sysdesc venator_2_on_default_vf
switch#admin> lldp --config -mx 5
switch#admin> lldp --config -txintvl 40
```

3. Set the DCBX version as either auto, pre-CEE, or CEE. The default version is auto.

```
switch#admin> lldp --enable -port 3/4-8 -dcbxver cee
```

4. Enable the required TLVs globally using the **lldp --enable -tlv** command.

```
switch#admin> lldp --enable -tlv fcoe-app
```

NOTE

The **dcbx**, **fcoe-app**, **fcoe-lls**, and **sys-name** TLVs are enabled by default globally.

The following TLVs are supported:

- **dcbx**—IEEE Data Center Bridging eXchange.
- **fcoe-app**—IEEE DCBX FCoE Application
- **fcoe-lls**—IEEE DCBX FCoE Logical Link
- **dot1**—IEEE 802.1 Organizationally Specific
- **dot3**—IEEE 802.3 Organizationally Specific
- **mgmt-addr**—Management Address
- **port-desc**—Port Description
- **sys-cap**—System Capabilities
- **sys-desc**—System Description
- **sys-name**—System Name

5. You can use the following commands to display the LLDP neighbors and statistics.

```
switch#admin> lldp --show -nbr
Local Intf Dead Interval Remaining Life Remote Intf          Chassis ID      Tx  Rx  System
Name
8/4          120          104          port1              0005.1e8f.fd04 183 203 (null)
8/6          120          110          Eth1/3             8c60.4ff7.edc6 174 199 FCoEFI-A
8/7          120          110          Eth1/4             8c60.4ff7.edc7 174 199 FCoEFI-A
8/24         120          116          FortyGigabitEthernet 1/0/49 0005.3365.0bc4 174 198 sw0
```

Configuring and activating an LLDP profile for a group of ports

LLDP is enabled by default when a port is converted to ETH mode. You can also create multiple customized LLDP profiles with different parameters and apply them to specific groups of ports. If no specific LLDP profile is applied on a port, the global parameters are applied by default.

1. Create an LLDP profile using the **lldp --create -profile** command.

```
switch#admin> lldp --create -profile lldp_profile_1
```

2. Configure the LLDP profile parameters using the **lldp --config** command.

```
switch#admin> lldp --config -mx 4 -profile lldp_profile_1
switch#admin> lldp --config -txintvl 40 -profile lldp_profile_1
```

3. Enable the required TLVs on the LLDP profile using the **lldp --enable -tlv** command.

```
switch#admin> lldp --enable -tlv fcoe-app -profile lldp_profile_1
```

4. Use the **lldp --show -profile** command to display the configured LLDP profile parameters.

```
switch#admin> lldp --show -profile lldp_profile_1
```

5. Enable the LLDP profile on a group of ports using the following command.

```
switch#admin> lldp --enable -port 3/40-56 -profile lldp_profile_1
```

6. Verify the LLDP profile details.

```
Switch#admin> lldp --show -profile

Profile-name:brocade_san
  Enabled TLVs:dcbx;fcoe-app;mgmt-addr;
  Profile ports: 8/4;8/5
=====

Profile-name:brocade_fabric
  Enabled TLVs:dot1;fcoe-lls;sys-desc;mgmt-addr;
  Profile ports:8/6
=====

Number of profile entries = 2
```

FCoE logical ports (E_Nodes)

A logical port is dynamically bound to a physical Ethernet port when a FLOGI is received on that port. You can create and delete the FCOE logical ports.

1. Configure the number of VF_Ports or E_Nodes that you require using the **fcoe --config -enodes** command. You can have a maximum of 1600 E_Nodes. By default, each logical switch comes up with zero E_Nodes.

```
switch:admin> fcoe --config -enodes 300
```

2. You can check the port status using the **switchshow** command.

```
switch:admin> switchshow
switchName:    sw0
switchType:    165.0
switchState:   Online
```

```

switchMode:    Native
switchRole:    Principal
switchDomain:   1
switchId:      fffc01
switchWwn:     10:00:c4:f5:7c:64:29:00
zoning:        OFF
switchBeacon:  OFF
FC Router:     OFF
HIF Mode:      OFF
Allow XISL Use: OFF
LS Attributes: [FID: 128, Base Switch: No, Default Switch: Yes, Ficon Switch: No, Address Mode 0]

```

Index	Slot	Port	Address	Media	Speed	State	Proto
192	8	0	01c000	id	40G	Online	ETH
196	8	4	01c400	id	40G	No_Sync	ETH
200	8	8	01c800	id	40G	Online	ETH
204	8	12	01cc00	id	40G	No_Sync	ETH
216	8	24	01d800	--	40G	No_Module	ETH
203	8	27	01cb00	id	16G	In_Sync	FC
220	8	28	01dc00	--	40G	No_Module	ETH
...							
900	-1	900	01a400	--	--	Offline	FCoE
901	-1	901	01a500	--	--	Offline	FCoE
902	-1	902	01a600	--	--	Offline	FCoE
903	-1	903	01a700	--	--	Offline	FCoE
904	-1	904	01a800	--	--	Offline	FCoE
905	-1	905	01a900	--	--	Offline	FCoE
906	-1	906	01aa00	--	--	Offline	FCoE
907	-1	907	01ab00	--	--	Offline	FCoE
908	-1	908	01ac00	--	--	Offline	FCoE
909	-1	909	01ad00	--	--	Offline	FCoE

Configuring an FCoE fabric map

An FCoE fabric map is created whenever a logical switch is created. Each logical switch has an FCoE fabric map with a set of attributes that you can configure using the **fcoe --config** commands. To configure a fabric map, follow these steps:

1. Configure the FCoE FC map using the **fcoe --config -fcmmap** command. The map can be from 0xefc00 to 0xefcff.

```
switch#admin> fcoe --config -fcmmap 0x0efc02
```

NOTE

The default FCoE fabric map is 0x0efc00

2. Configure the FCoE VLAN using the **fcoe --config -vlan** command. The VLAN can be any user VLAN supported in the Ethernet network. The default FCoE VLAN is 1002.

```
switch#admin> fcoe --config -vlan 2010
```

3. Configure the FCoE priority using the **fcoe --config -priority** command. The priority can be between 0 and 6. The default priority is 3.

```
switch#admin> fcoe --config -priority 4
```

4. Configure the FKA interval in milliseconds using the **fcoe --config -fka** command. The interval can be between 250-90000 ms. The default FKA interval is 8000 ms.

```
switch#admin> fcoe --config -fka 15000
```

- Configure the keepalive timer ON or OFF using the **fcoe --config -katimeout** command. The default value is ON. When the timer is on, the default keepalive timeout value in the FCoE protocol is used.

```
switch#admin> fcoe --config -katimeout off
```

Enabling the FCoE protocol on Ethernet ports

- Check the LLDP neighbor status to make sure that the Ethernet ports are up before enabling the FCoE protocol on the ports.

```
switch#admin> lldp --show -nbr
Local Port   Dead Interval  Remaining Life  Remote Port                               Chassis ID
Tx          Rx            System Name
8/0         120          94             FortyGigabitEthernet 3/0/52    0005.3365.0bc4
3          71          sw0
8/8         120          95             FortyGigabitEthernet 3/0/49    0005.3365.0bc4
3          71          sw0
```

- Enable the FCoE protocol on the Ethernet ports using the **fcoe --enable -port** command. You can specify a port number or a port range.

```
switch#admin> fcoe --enable -port 9/4-20
```

NOTE

You must set the UCS mode on the port or port channel connected to the Cisco UCS devices.

- Check if the port is provisioned for FCoE using the **fcoe --show -provision** command.
- By default, the VF_Ports are enabled. However, if you need to enable a VF_Port, use the **portenable** command.

```
switch#admin> portenable -i 776
```

Enabling the FCoE protocol on Ethernet port channels

- Check the port channel status before enabling the FCoE protocol on the port channels.

```
switch#admin> portchannel --show
Name          Type          Oper-State  Port-Count  Member Ports
-----
LAG_0         Dynamic       Offline     0
LAG_1         Dynamic       Offline     0
LAG_2         Dynamic       Offline     0
LAG_3         Dynamic       Online      3           3/15*, 3/16,
                                     3/17 ,3/18
LAG_4         Dynamic       Offline     0
LAG_5         Dynamic       Offline     0
<output truncated>
```

- Enable the FCoE protocol on the Ethernet port channels using the **fcoe --enable -portchannel** command.

```
switch#admin> fcoe --enable -portchannel port_ch_1
```

- By default, the port channel is enabled. You can also use the **portchannel --enable** command to enable it.

```
switch#admin> portchannel --enable port_ch_1
```

- Check if the port channel is provisioned for FCoE using the **fcoe --show -provision** command.
- You can check the status of the port channel using the **fcoe --show -portchannel** command.

Displaying FCoE fabric details

You can run the `fcoe --show` command with the following options to display the configured FCoE fabric map attributes and current status.

- The `-fabric` option displays the global fabric-wide FCoE attributes that you have configured, essentially the FCoE fabric map.

```
switch#admin> fcoe --show -fabric
=====
VLAN      VFID      Pri      FCMAP      FKA      Timeout
=====
1008      128[D]    3[D]     0xefc00[D] 8000[D]   Enabled[D]
```

- The `-fcf` option displays the details of attributes applicable to the local FCoE forwarder (FCF) switch such as domain ID, number of E_Nodes configured, and FCF MAC address.

```
switch#admin> fcoe --show -fcf
Domain-id : 1
Number of Enodes : 10
Global FCF Mac
=====
      c4:f5:7c:64:29:02
Per Port FCF Mac
=====
8/0:      c4:f5:7c:00:80:f0
8/1:      c4:f5:7c:00:80:f1
8/2:      c4:f5:7c:00:80:f2
8/3:      c4:f5:7c:00:80:f3
8/4:      c4:f5:7c:00:80:f4
8/5:      c4:f5:7c:00:80:f5
8/6:      c4:f5:7c:00:80:f6
8/7:      c4:f5:7c:00:80:f7
8/24:     c4:f5:7c:00:80:08
8/25:     c4:f5:7c:00:80:09
8/26:     c4:f5:7c:00:80:0a
8/27:     c4:f5:7c:00:80:0b
p1:       c4:f5:7c:64:29:22
```

- The `-provision` option displays the details of ports and port channels configured as FCoE ports.

```
switch#admin> fcoe --show -provision
=====
Domain      Port(s)/Portchannel(s)      Mode
=====
1           8/4
1           p1
Total number of port(s) = 2
```

- The `-brief` option displays brief details of FCoE devices currently logged in.

```
switch#admin> fcoe --show -login brief
=====
FCOE VF-Port      Eth-port/LAG      #VN-Ports
=====
1800              8/4                3
Total number of Port(s) = 1
```

- The `-login` option displays the details of all FCoE devices currently logged in to the local switch.

```
switch#admin> fcoe --show -login
=====
FCOE VF-Port      Eth-port/LAG      Device WWN      Device MAC      Session MAC
=====
1800              8/4                10:00:00:05:1e:8f:fd:02  00:05:1e:8f:fd:02  0e:fc:00:01:90:40
```

```

1800          8/4          10:00:00:05:1e:8f:fd:00  00:05:1e:8f:fd:02  0e:fc:00:01:90:41
1800          8/4          10:00:00:05:1e:8f:fd:10  00:05:1e:8f:fd:02  0e:fc:00:01:90:42

Total number of Login(s) = 3

```

- The **-stats** option displays the FCoE protocol statistics of ports and port channels.

```

switch#admin> fcoe --show -stats
Interface: 3/0
-----
RX Statistics:
Num of FIP VLAN Discovery Requests : 3
Num of FIP Discovery Solicitations : 3
Num of FIP FLOGIs                  : 7
Num of FIP NPIV FDISCs             : 439
Num of FIP LOGOs                   : 0
Num of FIP ENode Keep Alives       : 302
Num of FIP VN_Port Keep Alives     : 3036
Errors                              : 7
TX Statistics:
Num of FIP VLAN Discovery Responses : 3
Num of FIP Discovery SA             : 3
Num of FIP Discovery UA            : 0
Num of FLOGI ACCs                  : 3
Num of FDISC ACCs                  : 390
Num of LS_RJT (FLOGI, FDISC, LOGO) : 53
Num of CVLs                        : 62
Errors                              : 0
<output truncated>

```

NOTE

You can use the **fcoe --clear -stats** command and the **fcoe --clear -login** command to clear the FCoE device login details, and the **lldp --clear** command to clear the discovered neighbors and statistics.

Resetting the configuration to the default

NOTE

All configurations of LLDP, LACP, and FCoE are effective on the currently selected logical switch. They can be reset to the default within the logical switch using the following commands.

NOTE

After running the **config default** command, the FCoE feature gets blocked. You must run the **dlisset --enable -lossless** to unblock the FCoE feature. FCoE E_Nodes can be created only when the "Lossless" feature is enabled in the default switch.

- The **fcoe --default** command resets the switch to the default FCoE configuration.
- The **lldp --default** command resets the switch to the default LLDP configuration.
- The **lACP --default** command resets the switch to the default LACP configuration.
- The **configdefault** command is supported only in switch disabled state. The FCoE fabric map configuration is retained even after the **configdefault** command is issued.
- The **portcfgdefault** command is not supported on the ports in ETH mode. To reset to the default port configuration in ETH mode, you must convert the port to FC mode and then back to ETH mode.

Duplicate port WWN support for FCoE devices

Only limited support is available for duplicates between FCoE and FCoE devices on the same switch. Duplicates between FC and FCoE devices are fully supported. FCoE uses the existing FC device policies for duplicate port WWN detection and action. Use the **configure** command for both FC and FCoE devices.

```
switch#admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no] y
  Enforce FLOGI/FDISC login: (0..2) [0] 1
<output truncated>
```

Currently, there are three policies for duplicate port WWN handling:

1. Policy 0 (default) : Retains the old login.
2. Policy 1: Retains the new login.
3. Policy 2: Retains the old login for FLOGI and the new login for FDISC.

TABLE 1 Duplicate WWN support matrix for FC and FCoE devices

Device type	Duplicates found on the same switch (local duplicates)		Duplicates found on remote switches (remote duplicates)	
	FC	FCoE	FC	FCoE
FC	Yes	Yes	Yes	Yes
FCoE	Yes	Yes (always default policy behavior)	Yes	Yes

NOTE

Unlike FC F_Ports, FCoE VF_Ports are not persistently disabled on duplicate port WWN detection.

Configuring zoning with FCoE

Once the FCoE port login process to the fabric is complete, the FCoE devices are no different than FC devices in the FC fabric. All functionality supported for the FC devices in terms of zoning is also supported for the FCoE devices. There is no special configuration that is required for FCoE devices.

NOTE

Zoning is enforced on the target port as no CAM is installed on the Ethernet ports.

- Soft zoning isolation is done using the name server defined by the zoning configuration. Different from session-based zoning, soft zoning does not intercept the initiator in the hardware; that is no PLOGI trapping is done.
- Hard zoning isolation is done using the TCAM zoning ACLs for FCoE devices.
- Hard port zoning, that is "D,I" zoning, is not supported for FCoE devices. The following table lists the zoning scenarios where there is no connectivity.

TABLE 2 FCoE zoning enforcement and connectivity support

Zone type	-->	WWN		[D,I]	
		Device type	FC	FCoE	FC
WWN	FC	Hard-zoned WWN	Hard-zoned WWN***	Session-based; PLOGI trap installed	-

TABLE 2 FCoE zoning enforcement and connectivity support (continued)

Zone type	-->	WWN		[D,I]	
	FCoE	Hard-zoned WWN***	Soft-zoned; No PLOGI trap	Session-based: PLOGI trap installed at FC Port; Soft-zoned at FCoE Port	-
[D,I]	FC	Session-based; PLOGI trap installed	Session-based: PLOGI trap installed at FC Port; Soft-zoned at FCoE Port	Hard-zoned port	-
	FCoE	-	-	-	-

- A "-" indicates the cases where the combination of such zoning results in "no connectivity"
- A "****" indicates the cases where it is hard zoned at the FC port and soft zoned at the FCoE port.
- TI zoning is not supported with FCoE devices.
- QoS zoning is not supported with FCoE devices.

Configuring Virtual Fabrics with FCoE

FCoE is fully supported with Virtual Fabrics.

- No special configuration is required to enable FCoE with Virtual Fabrics. The FCoE, LLDP, and LACP configuration must be done on each logical switch.
- Every time a new logical switch is created, the default fabric map attributes are automatically created.
- Each logical fabric is represented by an FCoE fabric map, and each logical fabric has a only one FCoE VLAN mapping to the VFID.
- Each logical fabric has a unique FCoE forwarder (FCF) MAC address.
- All FCoE features supported on the default VF are supported on the nondefault VFs as well.
- FCoE is not supported in the base switch.
- FCoE is not supported on the logical switch where FICON is enabled.

Configuring FCoE with FCR

- FCoE is fully supported with an FC router where FCoE hosts can communicate with remote targets across an FC router backbone fabric in both edge-to-edge and backbone-to-edge topologies.
- In the backbone-to-edge topology, the backbone switch **nsshow** command does not display the device type of FCoE for the FCoE proxy devices.

Considerations and limitations

The following must be considered when configuring FCoE-capable switches and blades:

- Use the same **configupload** and **configdownload** commands that are used for FC switches. The FCoE configuration is retained even after reboot.
- CSCTL mode is not supported with FCoE devices.

FCoE-capable switches and blades have the following general limitations:

- FCoE is not supported in Access Gateway mode.
- FCoE does not support SNMP, Flow Vision, and MAPS features.
- FCoE connectivity to the Brocade Analytics Monitoring Platform is not supported.
- There is no provision to configure security ACLs for directly connected FCoE devices. Connectivity through TOR is secure as FSB takes care of the security ACLs.
- The D_Port feature is not supported on FCoE or Ethernet ports.
- FCoE target connectivity is not supported.
- Only a single lossless priority is supported for FCoE. Multiple lossless priorities are not supported.
- Ethernet SPAN and sFLOW are not supported.
- MAC learning, Ethernet-based forwarding, and flooding are not supported.

HA configuration considerations and limitations

HA is supported with FCoE.

Upgrade considerations and limitations

Upgrading to Fabric OS 8.2.0 on the DCX 8510 chassis is nondisruptive but FCOE10-24 port blades are not recognized after the upgrade.

Downgrade considerations and limitations

Downgrading from Fabric OS 8.2.0 is not supported if a FC32-64 port blade is present in the chassis and FCoE, LLDP, or LACP are configured. To remove the configurations, run the following commands.

- `fcoe --default`
- `lldp --default`
- `lACP --default`

FCoE scalability limits

TABLE 3 FCoE scalability limits

Supported entity	Tested limit	Maximum limit	Notes
Number of FCoE devices (FLOGI + FDISC)	512	4000	Across chassis
Number of FCoE E_Nodes/logical ports	512	1600	Across chassis
Number of NPIVs per E_Node	256	3999	Per VF_Port
Number of logical switches	16	16	
Number of FCoE VLANs	16	16	One VLAN per VF
Number of LACP port channels	64	64	Per chassis
Number of ports in a port channel	16	16	Per port channel

TABLE 3 FCoE scalability limits (continued)

Supported entity	Tested limit	Maximum limit	Notes
Number of LLDP profiles	64 across chassis	512 (64 per VF)	Across chassis