Fibre Channel over Ethernet (FCoE)

Configuration and Troubleshooting Guide

January 2015

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Contents

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

1. Introduction

1.1 Preface

This Fibre Channel over Ethernet (FCoE) troubleshooting guide is designed to help readers understand FCoE concepts and troubleshoot FCoE networks based on Cisco Nexus® and Cisco® MDS platforms.

This document will cover troubleshooting steps for FCoE on Cisco Nexus and MDS switches. The goal of this document is to assist our customers' networking teams in understanding the FCoE protocol; how it is configured on Cisco Nexus switches; and how to troubleshoot its various components. This document is not intended to be inclusive of native Fibre Channel and its associated functions, such as zoning, device aliases, domain IDs, and others. Those functions are managed by storage networking teams and, on Cisco Nexus switches, are similar to managing Fibre Channel in the MDS environment.

The following references identify detailed FCoE configuration documents on Cisco Nexus platforms. The configuration section of this document focuses mainly on Data Center Bridging (DCB) and its optimal configuration in order for FCoE to work properly.

1.2 References

- FCoE Initiation Protocol (FIP) White Paper
- Cisco Nexus 6000 Series FCoE Configuration Guide, Release 6.x
- Cisco Nexus 6000 FCoE Troubleshooting Guide
- Cisco Nexus 7000 Troubleshooting Guide - Troubleshooting FCoE
- Cisco Nexus 5500 to 7000 Multi-Hop FCoE Configuration Example
- FCoE Configuration Guide for Cisco Nexus 7000
- Cisco MDS FCoE Configuration Guide
- Ethanalyzer: Cisco NX-OS Software Built-In Packet Capture Utility
- Priority Flow Control: Build Reliable Layer 2 Infrastructure
- Information About Flow Control - Link Level and PFC
- End-End FCoE Design Guide

2. Overview of Fibre Channel over Ethernet (FCoE)

FCoE is technology that enables converged I/O, which encompasses data Ethernet traffic and Fibre Channel (FC) sharing the same Ethernet wire. There are several designs that can enable a hybrid of native Cisco MDS FC switches and Cisco Nexus FC and FCoE switches (for example, Cisco Nexus 5500 Series Switches). This guide, however, focuses on pure FCoE environments, where the environment consists of Cisco Nexus switches and FCoE host and storage end devices.

2.1 Enhancements to Ethernet

The T11 organization's FC-BB-5 standard defines FCoE, and also defines running FC over other media types. The IEEE 802.1 organization facilitates FCoE by defining enhancements to Ethernet. These enhancements fall under the DCB umbrella, specifically, three enabling standards for Ethernet to support FCoE:

1. Priority-based Flow Control (PFC)
2. Enhanced Transmission Selection (ETS)
3. Data Center Bridging Exchange (DCBX)

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Figure 1 identifies the T11 and IEEE standards.

Figure 1. T11 and IEEE 802.1 FCoE Standards

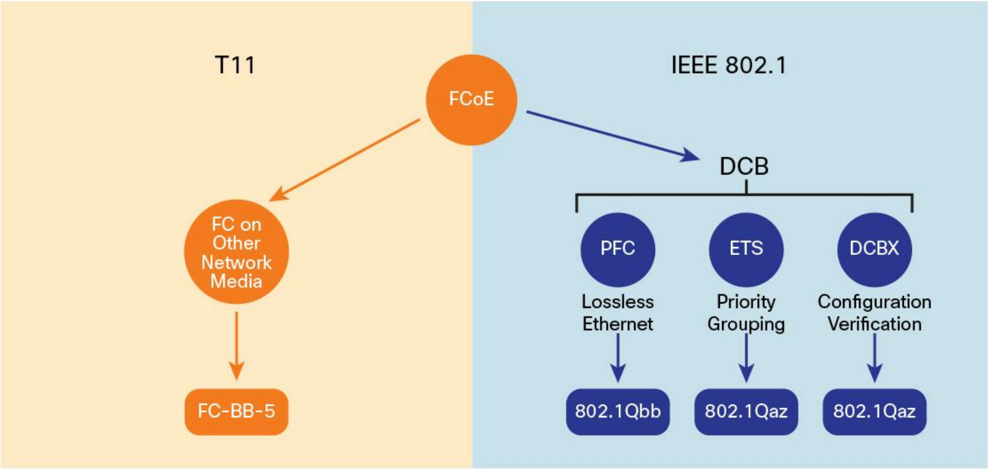

2.1.1 Priority-Based Flow Control (PFC)

PFC is defined in IEEE 802.1Qbb. Link sharing is critical to I/O consolidation. For link sharing to succeed, large bursts from one traffic type must not affect other traffic types, large queues of traffic from one traffic type must not starve other traffic types' resources, and optimization for one traffic type must not create large latency for small messages of other traffic types. The Ethernet pause mechanism can be used to control the effects of one traffic type over another.

PFC creates eight separate virtual links on the physical link and allows any of these links to be paused and restarted independently. This approach can enable the network to create a no-drop class of service for an individual virtual link that can coexist with other traffic types on the same interface. In native FC, the network is designed not to drop traffic (lossless). PFC can enable Ethernet to support FC by providing a lossless fabric.

Figure 2 shows the eight virtual data lanes on a single wire that make up PFC. One virtual lane of data (for example, FC) can be paused while the remaining lanes continue to transmit.

Figure 2. Priority Flow Control (PFC)

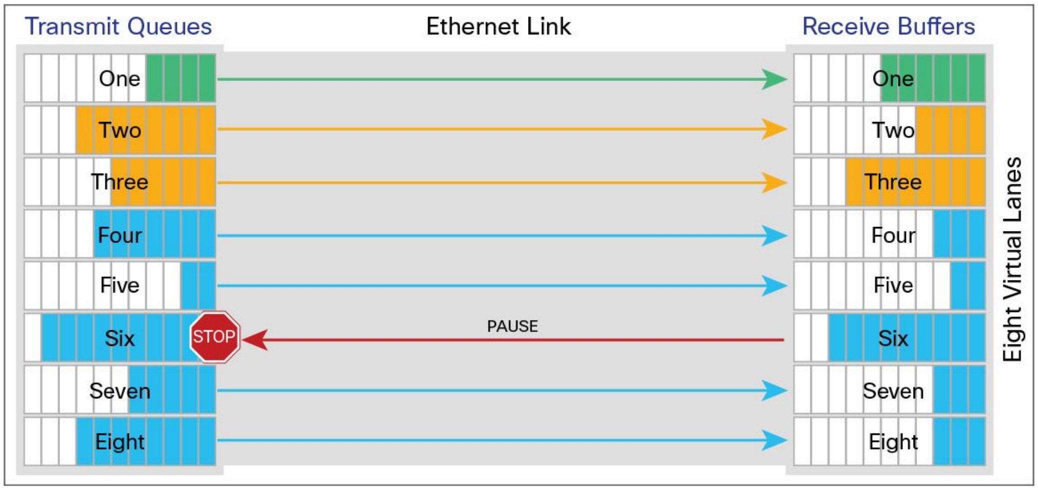

The Ethernet frames that are sent by the switch to the adapter may include the IEEE 802.1Q tag. This tag includes a field for the class of service (CoS) value used by the priority flow control (PFC). The IEEE 802.1Q tag also includes a virtual LAN (VLAN) field.

2.1.2 Enhanced Transmission Selection (ETS)

ETS is defined in IEEE 802.1Qaz. PFC can create eight distinct virtual link types on a physical link, and it can be advantageous to have different traffic classes defined within each virtual link. Traffic within the same PFC IEEE 802.1p class can be grouped together, yet treated differently within each group. ETS provides prioritized processing based on bandwidth allocation, low latency, or best effort, resulting in per-group traffic class allocation.

Extending the virtual link concept, the network interface controller (NIC) provides virtual interface queues: one for each traffic class. Each virtual interface queue is accountable for managing its allotted bandwidth for its traffic group, but has flexibility within the group to dynamically manage the traffic. For example, virtual link 3 for the IP class of traffic may have a high-priority designation and a best effort within that same class, with the virtual link 3 class-sharing a percentage of the overall link with other traffic classes. ETS allows differentiation among traffic of the same priority class, thus creating a priority group (Figure 3). The capability to apply differentiated treatment to different traffic within the same priority class is enabled by implementing ETS.

Figure 3. Enhanced Transmission Selection (ETS)

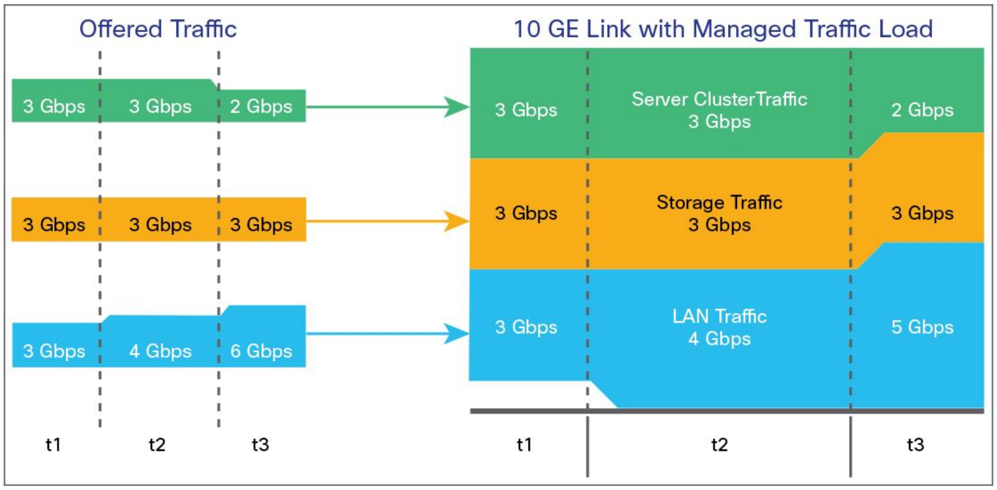

2.1.3 Data Center Bridging Exchange (DCBX)

DCBX is defined in IEEE 802.1Qaz. The DCBX Protocol is a discovery and capability exchange protocol that is used by IEEE DCBs to discover peers and exchange configuration information between DCB-compliant bridges (see Figure 4). The following parameters can be exchanged with DCBX:

- Priority groups in ETS
- PFC
- Congestion notification
- Applications
- Logical link-down
- Network interface virtualization

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

DCBX allows network devices to advertise their identities and capabilities over the network. It enables end devices to pick up proper configuration from the network and for switches to verify proper configuration.

Figure 4. Data Center Bridging Exchange (DCBX)

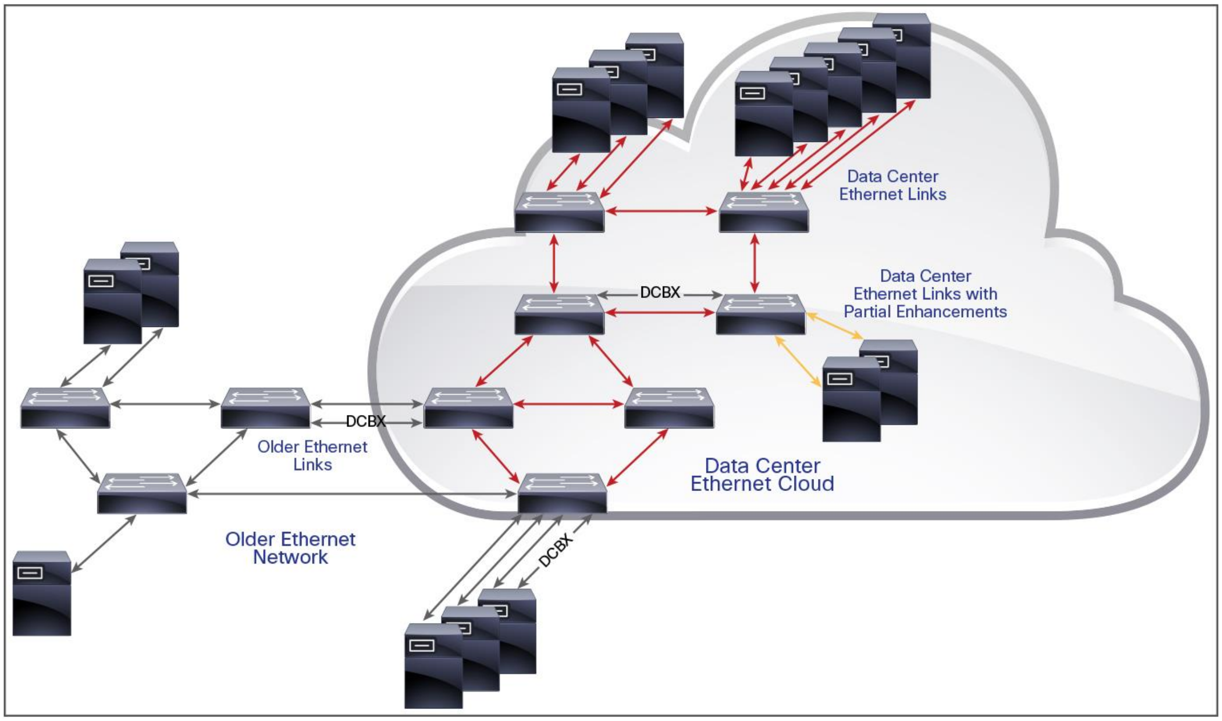

Data Center Ethernet Links
DCBX
Data Center Ethernet Links with Partial Enhancements
DCBX
Older Ethernet Links
Data Center Ethernet Cloud
Older Ethernet Network
DCBX

The DCBX protocol is an extension of the Link Layer Discovery Protocol (LLDP). DCBX endpoints 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 device 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 Cisco Nexus device are programmed to accept the configuration values sent by the switch, allowing the switch to distribute configuration values to all attached CNAs. This reduces the possibility of configuration errors and simplifies CNA administration.

2.2 FCoE Protocol

FCoE is two protocols, the FCoE data plane protocol and the FCoE Initialization Protocol (FIP) control plane protocol. Each protocol has different Ethertypes: the FCoE data plane is 8906 and FIP is type 8914. As well, the frame formats are different. Note that FIP, the control protocol, is based on the T11 FC standard and is not DCB (PFC, ETS, and DCBX are all strictly Ethernet standards). But for FIP to work properly, DCB must be configured correctly.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.2.1 FCoE Frame Format

FCoE is implemented by encapsulating an FC frame in an Ethernet packet with dedicated Ethertypes, 0x8906 and 0x8914. That 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.

An FC frame consists of 36 bytes of headers and up to 2112 bytes of data for a total maximum size of 2148 bytes. The encapsulated FC frame has all the standard headers, which allow it to be passed to the storage network without further modification. To accommodate the maximum FC frame in an FCoE frame, the class-fcoe is defined with a default maximum transmission unit (MTU) of 2158 bytes on Cisco Nexus 5000 and 6000 Series Switches while defined as 2112 bytes on Cisco Nexus 7000 Series Switches.

Figure 5 details the FC encapsulation into Ethernet frames.

Figure 5. FCoE Frame Size

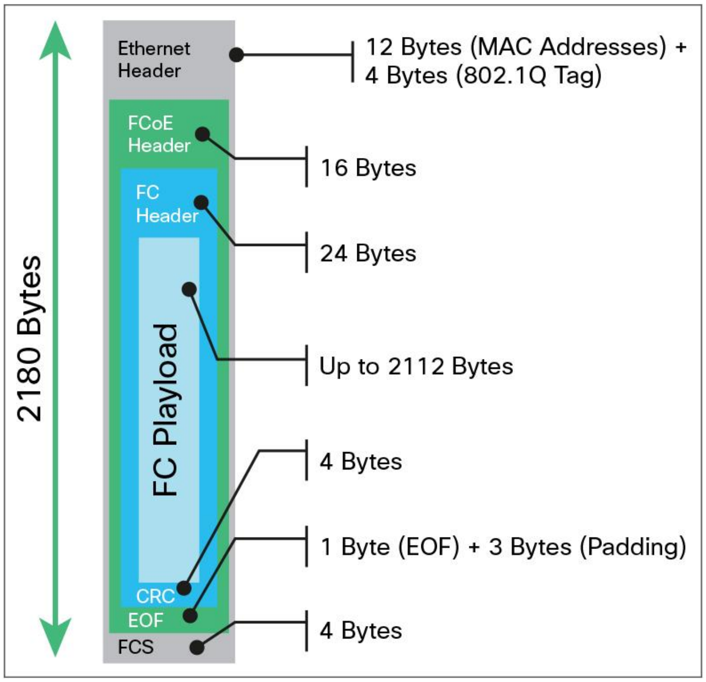

2.2.2 FCoE Virtual Interfaces

FCoE enables FC interfaces on Ethernet. FC interfaces are mapped to virtual interfaces in an Ethernet network. This virtualization essentially allows for management of an FCoE infrastructure in the same way as a native FC infrastructure. It is important to understand the virtual interfaces that make up an FCoE implementation and how they map to physical Ethernet interfaces. FCoE interfaces are defined as port types, identified in Table 1.

Table 1. FCoE Port Types

Port Type	FCoE Virtual Interface	Binds to	Example Interface
VF-port or virtual fabric port	vfc1	Ethernet interface	Ethernet101/1/1
VE-port or virtual expansion port (ISL)	vfc100	Ethernet interface or port-channel interface	Port-channel 100

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Note that in the case of port-channels, a virtual FCoE interface is created over a virtual port-channel interface (virtual mapped to virtual mapped to physical). It is the port-channel that is mapped to physical Ethernet ports.

Virtual FCoE interfaces are bound to physical interfaces. This binding is a one-to-one mapping. A physical Ethernet interface can only have one VFC bound to it (also true when bound to a port-channel).

Once defined, the configuration of virtual interfaces is similar to physical interfaces, that is, virtual interfaces need to have shutdown or no-shutdown commands issued to them. You can check status of virtual interfaces with the show command, for example, a "show interface vfc1" command. It should be evident that a VFC interface cannot be in the up state until the physical interface to which it is bound is in the up state. In order for virtual interfaces to come to the up state, proper configuration is required in the network, specifically pertaining to the DCB enhancements to Ethernet. The section, "Understanding the FIP Process" further in this document details the virtual interface instantiation.

FC communication between end devices (for example, a host and storage port) cannot occur until:

- Their associated physical port VFC interface is in the up state
- The FIP process has taken place, and
- The proper FC zoning is defined and active

2.3 Virtual Fibre Channel (VFC)

In a native FC storage area network (SAN), physical FC switches (such as the Cisco MDS 9513 Multilayer Director) and end devices (such as hosts with host bus adapters [HBAs]) are connected with fiber cable. The FC protocol runs natively on the SAN and both the switches and end devices communicate through it.

With FCoE, we are taking that FC SAN and overlaying it onto a physical Ethernet network, creating a VFC SAN running over Ethernet. In the previous section, enhancements to Ethernet needed to facilitate this were discussed. In addition to these enhancements, a new process is needed on Cisco Nexus switches to support FC. This process is known as FCoE Manager or fcoe_mgr. The fcoe_mgr process controls all the FCoE components and configuration. The fcoe_mgr process is started by installing the FCoE feature set. Virtual SANs (VSANs) are then associated to designated FCoE VLANs and bind VFC interfaces to physical Ethernet interfaces. FCoE configuration is discussed later in this document.

With FCoE, HBAs are replaced with CNAs. CNAs enable converged I/O by supporting both FC and classical Ethernet data traffic on the same Ethernet wire. CNAs, along with proper drivers on the host end device, support FCoE.

Figure 6 is an illustration of a VSAN over Ethernet.

Figure 6. Virtual FC over Ethernet

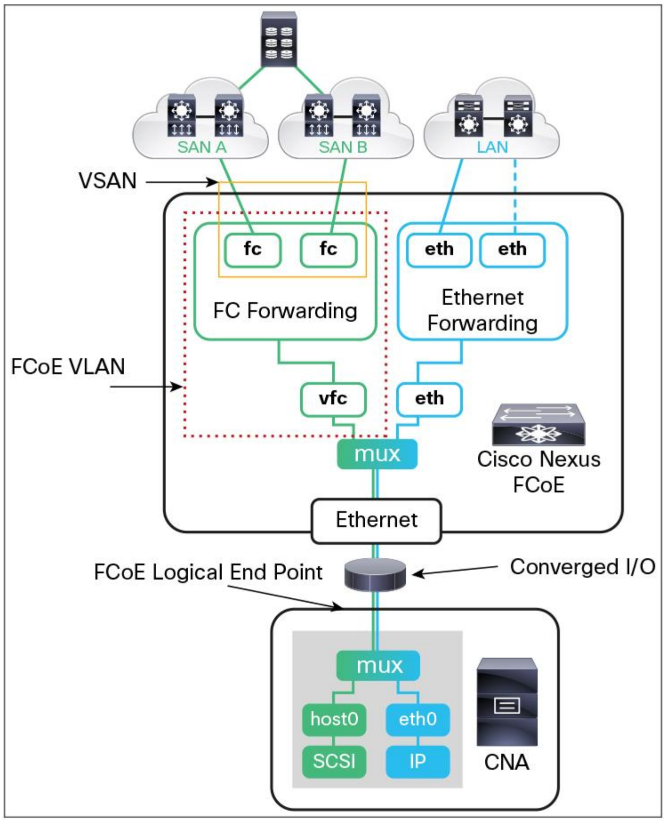

There are a number of command-line interface (CLI) commands that may be issued against fcoe_mgr to assist in troubleshooting. An example of the fcoe_mgr event commands for **Cisco Nexus 5000 and 6000 Series Switches** is shown here:

```
N6K# show platform software fcoe_mgr ?
  event-history  Show various event logs of FCOE_MGR
  info           Show internal data structure information
  mem-stats      Show memory allocation statistics of FCOE_MGR
```

The command for **Cisco Nexus 7000 Series Switches and Cisco MDS** is slightly different.

```
N7K-storage# show system internal fcoe_mgr event-history ?
  errors     Show error logs of FCOE_MGR
  interface  Enter interface info
  lock       Show internal locking event log
  module     Show module information
  msgs       Show various message logs of FCOE_MGR
```

These various fcoe_mgr show commands will be referenced throughout the remainder of this document.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.4 Capturing and Analyzing Ethernet Frames

There are tools available for capturing Ethernet frames on Cisco Nexus and MDS switches, including Ethanalyzer and switched port analyzer (SPAN). The resulting capture files with these tools may be analyzed with Wireshark. In addition to these tools, external in-line analyzers may be used to capture data.

2.4.1 Ethanalyzer

Ethanalyzer is useful for analyzing packets that are destined to the switch supervisor. Generally, this is control-plane-type traffic, such as Link Aggregation Control Protocol (LACP), spanning tree, Address Resolution Protocol (ARP), routing protocols, and others. Ethanalyzer will not capture data traffic that is processed directly by the ASIC. In the case of FCoE, Ethanalyzer may be used to analyze FIP traffic.

Different types of Ethernet interfaces (Table 2) are available for capture with Ethanalyzer.

Table 2. Ethanalyzer Interface Types

Interface	Switch	Ethanalyzer Interface	Protocol	Description
eth0 **eth1**	Nexus 5000, 6000 Nexus 7000, MDS	mgmt	Management interface	Management interface defined as “mgmt.” by Ethanalyzer
eth0	Nexus 7000, MDS	inband	All control packets	eth0 is inband interface that captures all control packets on N7K and MDS
eth3	Nexus 5000, 6000	inbound-low	Internet Group Management Protocol(IGMP), ARP, User Datagram Protocol (UDP)	Low-priority control packets destined to the switch CPU
eth4	Nexus 5000, 6000	inbound-high	LACP, DCBX, FCoE, Cisco Discovery Protocol	High-priority control packets destined to the switch CPU

Here is an example CLI to use to capture traffic with Ethanalyzer.

```
N6K# ethanalyzer local interface [inbound-hi|inbound-low|mgmt] (options)
```

Here is an example to capture 200 frames of FCoE control traffic (eth type 8914) and write the output to a file on bootflash.

```
N6K# ethanalyzer local interface inbound-hi display-filter "vlan.etype==0x8914"
limit-cap 200 write bootflash:etype8914.pcap
```

For additional information and examples of Ethanalyzer, refer to the white paper, Ethanalyzer: Cisco NX-OS Software Built-In Packet Capture Utility.

2.4.2 Switch Port Analyzer (SPAN)

The SPAN tool is used to mirror, or span, a source port to an unused port (destination port) on the switch where a capture utility collects all the packets traversing the source interface. The capture utility on the destination port can be Cisco SwitchProbe or an analysis device that has the ability to capture the data. The captured data can later be used with Wireshark or another protocol analyzer to analyze the traffic. With SPAN, you can choose traffic in the ingress direction (traffic entering the switch through the source port), egress direction (traffic exiting the switch through the source port), or both ingress and egress.

SPAN allows you to capture VFC interfaces on a source port. The destination port for the mirrored traffic will be an Ethernet port (FCoE).

SPAN does not capture pause frames in an FCoE network because pause frames sent from the virtual expansion (VE) port are generated and terminated by the outermost MAC layer. To capture pause frames, an in-line capture appliance is needed.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

For details on SPAN, reference the following guides:

- Cisco Nexus 7000 Series NX-OS Configuring SPAN
- Cisco Nexus 6000 Series NX-OS Configuring SPAN
- Cisco MDS 9000 Series NX-OS Configuring SPAN

2.4.3 CLI Debugging Commands

Available debugs depend on features enabled in Cisco NX-OS Software. There are many different options to select when turning on debugs.

Determine the destination of the output:

- Logfile - Data file in switch memory
- Capture to direct to screen through a console, Telnet, or SSH

You must have administrator privileges to run debugs. Debugs can only be run from the CLI. Normally, for troubleshooting with CLI debug, the Cisco Technical Assistance Center (TAC) will advise what data to capture.

2.5 Understanding the FIP Process

The FIP instantiates a VFC interface and allows an end device to perform a fabric login (FLOGI). Events are recorded for each phase of the configuration and eventual startup of the VFC interfaces. This section will detail the instantiation and associated events for each step of the process. Remember, VFC interfaces (including the VFC's VSAN) won't display and become operational unless the Ethernet interface is up and the DCB enhancements (PFC, ETS, and DCBX) are configured properly.

2.5.1 Detailed Steps and Status of FIP Virtual Link Instantiation

Cisco NX-OS supports the T11-compliant FIP on Cisco Nexus devices. FIP is used to perform device discovery, initialization, and link maintenance. FIP performs the following protocol steps:

- **FIP discovery** - When an FCoE device is connected to the fabric, it sends out 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 FLOGI, fabric discovery (FDISC), logout (LOGO), and exchange link parameters (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 the connection is still valid. This is referred to as the FCoE keepalive (FKA).

The fcoe_mgr process monitors and controls all FCoE traffic on the switch. Using show fcoe_mgr commands, triggered events during the FIP process can be analyzed. The Check fcoe_mgr Events for FIP Transitions section covered later in this document details the proper FIP events from fcoe_mgr.

The FIP process is summarized in Table 3. Note that the expected fcoe_mgr triggered event is identified in the last column.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Table 3. FIP Virtual Link Instantiation Summary

FIP Step	Action	Response	fcoe_mgr Event
VLAN discovery	End device (CNA) broadcasts a request for FCoE VLAN. The request occurs on the native VLAN.	Switch responds with FCoE VLAN	FCOE_MGR_VFC_EV_FIP_VLAN_DISCOVERY or FCOE_MGR_VFC_EV_BRING_UP
FCF discovery	CNA broadcasts a solicitation to find FCF to log into. Broadcasts go out on the FCoE VLAN.	Switch responds with Advertisement	FCOE_MGR_VFC_EV_FIP_SOLICITATION
FLOGI/DISC	CNA performs a FLOGI or with NPV FDISC.	Switch accepts FLOGI/FDISC	FCOE_MGR_PROTO_EV_FIP_FLOGI
FC commands	CNA begins normal FC data commands using ethertype 8906.	Switch forwards encapsulated FCoE frames	FCOE_MGR_PROTO_EV_FC2_DONE
Normal VFC state	The desired state of the VFC		FCOE_MGR_PROTO_ST_UP

You can capture the FIP process using Ethanalyzer on the switch and using Wireshark to analyze it. To capture EtherType 8914 packets, run the following command from the switch:

```
ethanalyzer local interface inbound-hi display-filter "vlan.etype==0x8914" limit-
cap 200 write bootflash:etype8914.pcap
```

The following steps detail the FIP process and include Wireshark output.

2.5.1.1 CNA Performs VLAN Discovery Request

The CNA initiates a FIP VLAN request, broadcasting to destination MAC 01:10:18:01:00:02. This is a well-known MAC address and is referred to as the ALL-FCF-MACs address, meaning FCoE-enabled switches will recognize and respond to it. Keep in mind that FIP uses Ether Type 8914. In this step, the CNA is requesting to know the FCoE VLAN. The protocol screen capture (Figure 7) for this communication is shown here.

Figure 7. CNA Request to Know the FCoE VLAN

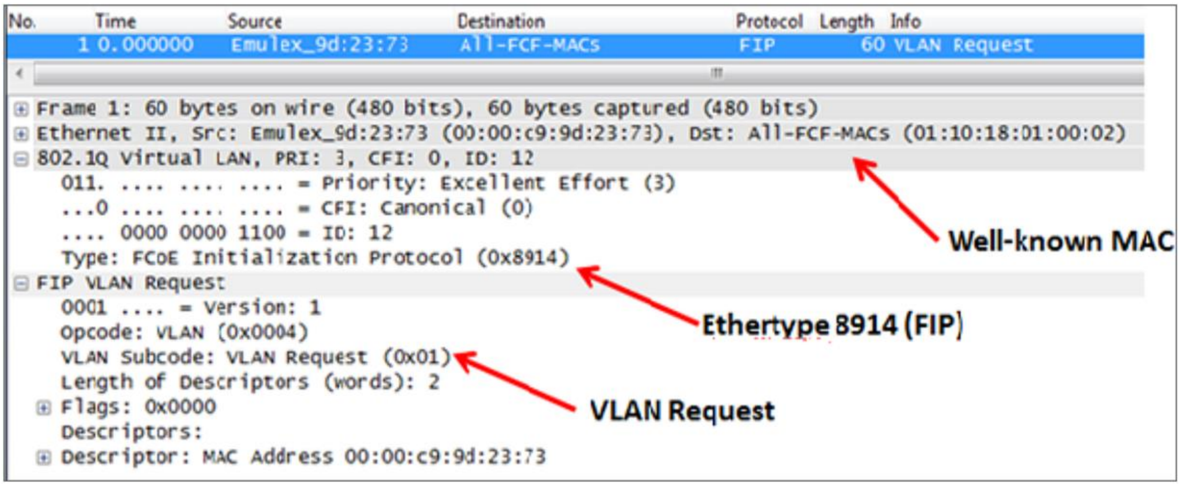

```
No.     Time        Source            Destination       Protocol  Length  Info
      1 0.000000    Emulex_9d:23:73   All-FCF-MACs      FIP           60  VLAN Request

⊞ Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
⊞ Ethernet II, Src: Emulex_9d:23:73 (00:00:c9:9d:23:73), Dst: All-FCF-MACs (01:10:18:01:00:02)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 12
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 1100 = ID: 12
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP VLAN Request
    0001 .... = Version: 1
    Opcode: VLAN (0x0004)
    VLAN Subcode: VLAN Request (0x01)
    Length of Descriptors (words): 2
  ⊞ Flags: 0x0000
    Descriptors:
  ⊞ Descriptor: MAC Address 00:00:c9:9d:23:73
```

In the capture, we can see the CNA broadcast to the ALL-FCF-MAC using FIP (eth type 8914) and requesting the FCoE VLAN. The VLAN Request from the host should be received on the native VLAN. The native VLAN cannot be an FCoE VLAN.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2.5.1.2 Switch Responds with VLAN

The screen capture in Figure 8 shows the switch response to the CNA request. The switch is responding with the FCoE VLAN of 100.

Figure 8. Switch Response to the CNA Request

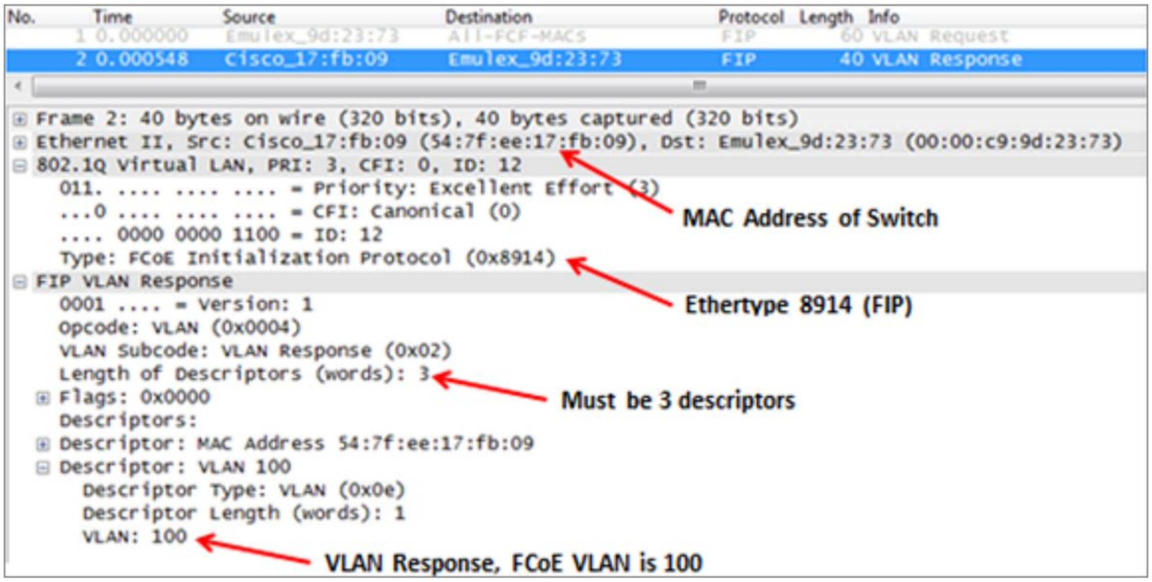

```
No.     Time        Source              Destination          Protocol  Length  Info
      1 0.000000    Emulex_9d:23:73     All-FCF-MACS         FIP           60  VLAN Request
      2 0.000548    Cisco_17:fb:09      Emulex_9d:23:73      FIP           40  VLAN Response

⊞ Frame 2: 40 bytes on wire (320 bits), 40 bytes captured (320 bits)
⊞ Ethernet II, Src: Cisco_17:fb:09 (54:7f:ee:17:fb:09), Dst: Emulex_9d:23:73 (00:00:c9:9d:23:73)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 12
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 1100 = ID: 12
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP VLAN Response
    0001 .... = Version: 1
    Opcode: VLAN (0x0004)
    VLAN Subcode: VLAN Response (0x02)
    Length of Descriptors (words): 3
  ⊞ Flags: 0x0000
    Descriptors:
  ⊞ Descriptor: MAC Address 54:7f:ee:17:fb:09
  ⊟ Descriptor: VLAN 100
      Descriptor Type: VLAN (0x0e)
      Descriptor Length (words): 1
      VLAN: 100
```

2.5.1.3 CNA Solicits FCF Discovery

The CNA next performs a FIP discovery by looking for a FCF switch to log into. In the screen capture shown in Figure 9, the CNA broadcasts again to the All-FCF-MACs address. This request, however, is transmitted on the FCoE VLAN that was learned from the previous request. In this discovery, the CNA provides information about itself, such as the maximum FCoE frame size it supports, it’s World Wide Name (WWN), and that it supports Fabric Provided MAC Address (FPMA).

Figure 9. Second CNA Broadcast to All-FCF-MACs Address

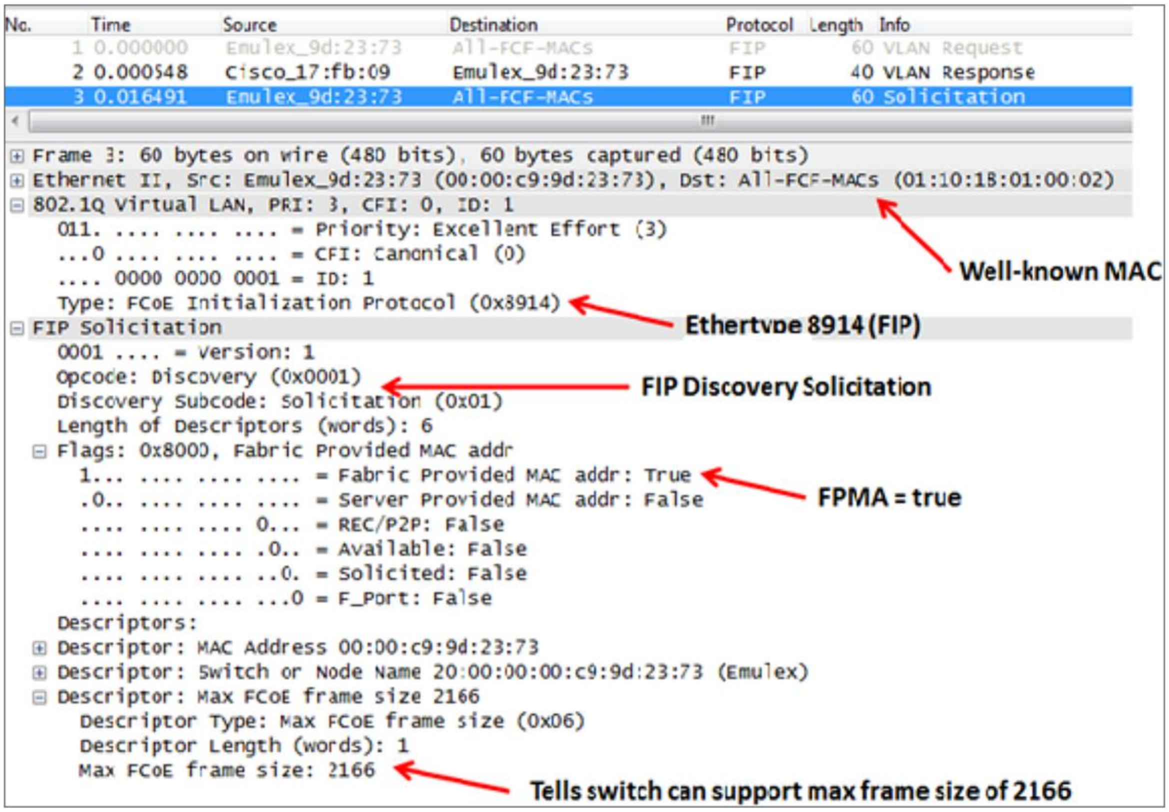

```
No.   Time        Source             Destination        Protocol  Length  Info
    1 0.000000    Emulex_9d:23:73    All-FCF-MACs       FIP           60  VLAN Request
    2 0.000548    Cisco_17:fb:09     Emulex_9d:23:73    FIP           40  VLAN Response
    3 0.016491    Emulex_9d:23:73    All-FCF-MACs       FIP           60  Solicitation

⊞ Frame 3: 60 bytes on wire (480 bits), 60 bytes captured (480 bits)
⊞ Ethernet II, Src: Emulex_9d:23:73 (00:00:c9:9d:23:73), Dst: All-FCF-MACs (01:10:18:01:00:02)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 1
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 0001 = ID: 1
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP Solicitation
    0001 .... = Version: 1
    Opcode: Discovery (0x0001)
    Discovery Subcode: Solicitation (0x01)
    Length of Descriptors (words): 6
  ⊟ Flags: 0x8000, Fabric Provided MAC addr
      1... .... .... .... = Fabric Provided MAC addr: True
      .0.. .... .... .... = Server Provided MAC addr: False
      .... .... .... 0... = REC/P2P: False
      .... .... .... .0.. = Available: False
      .... .... .... ..0. = Solicited: False
      .... .... .... ...0 = F_Port: False
    Descriptors:
  ⊞ Descriptor: MAC Address 00:00:c9:9d:23:73
  ⊞ Descriptor: Switch or Node Name 20:00:00:00:c9:9d:23:73 (Emulex)
  ⊟ Descriptor: Max FCoE frame size 2166
      Descriptor Type: Max FCoE frame size (0x06)
      Descriptor Length (words): 1
      Max FCoE frame size: 2166
```

Well-known MAC

Ethertvpe 8914 (FIP)

FIP Discovery Solicitation

FPMA = true

Tells switch can support max frame size of 2166

2.5.1.4 Switch Advertises Capabilities

The switch advertises its capabilities. The screen capture image in Figure 10 details the advertisement. It shows the virtual fabric ID (VSAN), the switch FC MAP ID (which, in this case, is the default), and the FC Keep Alive (FKA) period, which is 8000 ms (eight seconds). The total frame size in this advertisement equals the maximum FCoE frame size the CNA sent in its discovery. The switch pads the Advertisement frame to ensure it matches what the CNA expects and this will confirm the network path indeed supports full FC frame sizes (encapsulated in Ethernet).

Figure 10. Advertisement

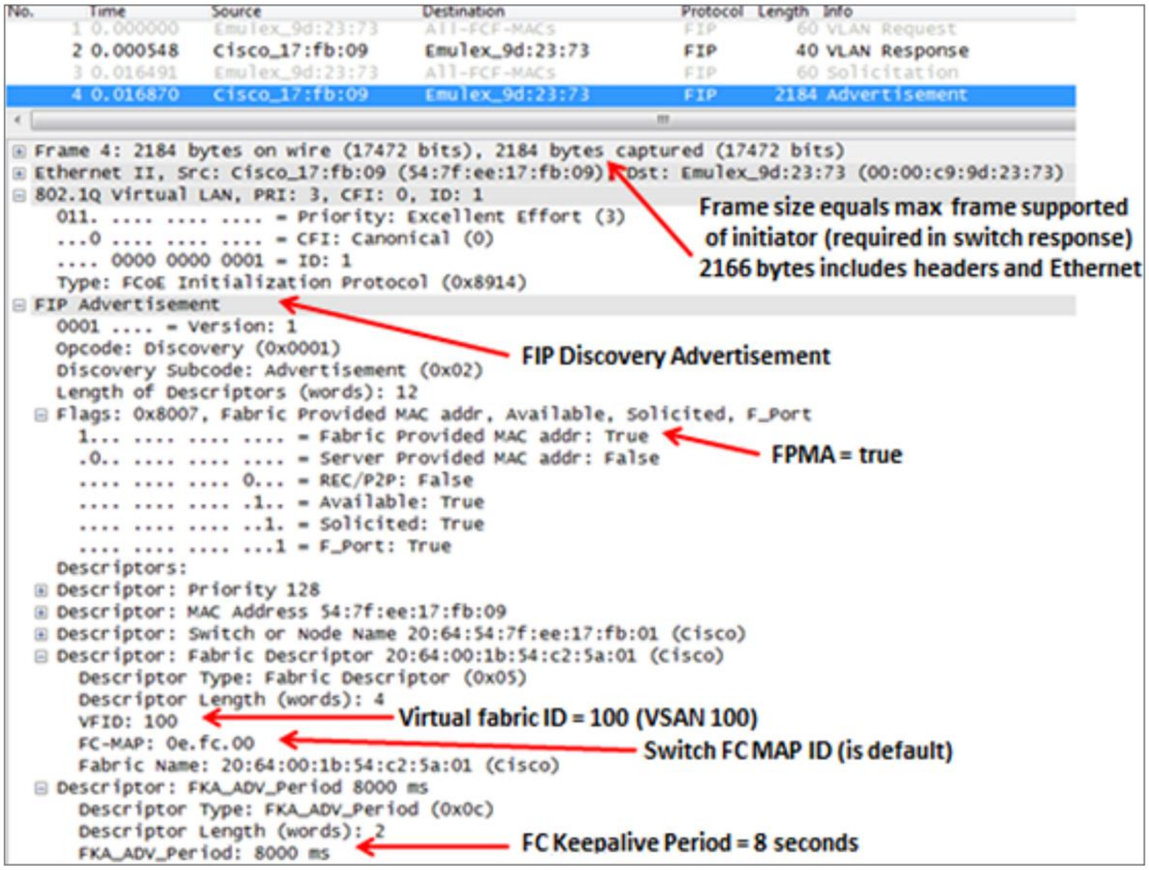

```
No.  Time       Source            Destination       Protocol Length Info
   1 0.000000   Emulex_9d:23:73   All-FCF-MACs      FIP          60 VLAN Request
   2 0.000548   Cisco_17:fb:09    Emulex_9d:23:73   FIP          40 VLAN Response
   3 0.016491   Emulex_9d:23:73   All-FCF-MACs      FIP          60 Solicitation
   4 0.016870   Cisco_17:fb:09    Emulex_9d:23:73   FIP        2184 Advertisement

⊞ Frame 4: 2184 bytes on wire (17472 bits), 2184 bytes captured (17472 bits)
⊞ Ethernet II, Src: Cisco_17:fb:09 (54:7f:ee:17:fb:09), Dst: Emulex_9d:23:73 (00:00:c9:9d:23:73)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 1
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 0001 = ID: 1
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP Advertisement
    0001 .... = Version: 1
    Opcode: Discovery (0x0001)
    Discovery Subcode: Advertisement (0x02)
    Length of Descriptors (words): 12
  ⊟ Flags: 0x8007, Fabric Provided MAC addr, Available, Solicited, F_Port
      1... .... .... .... = Fabric Provided MAC addr: True
      .0.. .... .... .... = Server Provided MAC addr: False
      .... .... .... 0... = REC/P2P: False
      .... .... .... .1.. = Available: True
      .... .... .... ..1. = Solicited: True
      .... .... .... ...1 = F_Port: True
    Descriptors:
  ⊞ Descriptor: Priority 128
  ⊞ Descriptor: MAC Address 54:7f:ee:17:fb:09
  ⊞ Descriptor: Switch or Node Name 20:64:54:7f:ee:17:fb:01 (Cisco)
  ⊟ Descriptor: Fabric Descriptor 20:64:00:1b:54:c2:5a:01 (Cisco)
      Descriptor Type: Fabric Descriptor (0x05)
      Descriptor Length (words): 4
      VFID: 100
      FC-MAP: 0e.fc.00
      Fabric Name: 20:64:00:1b:54:c2:5a:01 (Cisco)
  ⊟ Descriptor: FKA_ADV_Period 8000 ms
      Descriptor Type: FKA_ADV_Period (0x0c)
      Descriptor Length (words): 2
      FKA_ADV_Period: 8000 ms
```

Frame size equals max frame supported of initiator (required in switch response) 2166 bytes includes headers and Ethernet
FIP Discovery Advertisement
FPMA = true
Virtual fabric ID = 100 (VSAN 100)
Switch FC MAP ID (is default)
FC Keepalive Period = 8 seconds

2.5.1.5 CNA Initiates Fabric Login (FLOGI)

Now that the CNA has a valid FCF that will support fabric logins, it initiates the FLOGI. In the screen image shown in Figure 11, the EtherType is still 8914 (FIP). The rest of the frame contains standard FC FLOGI information.

Figure 11. CNA Initiates the FLOGI

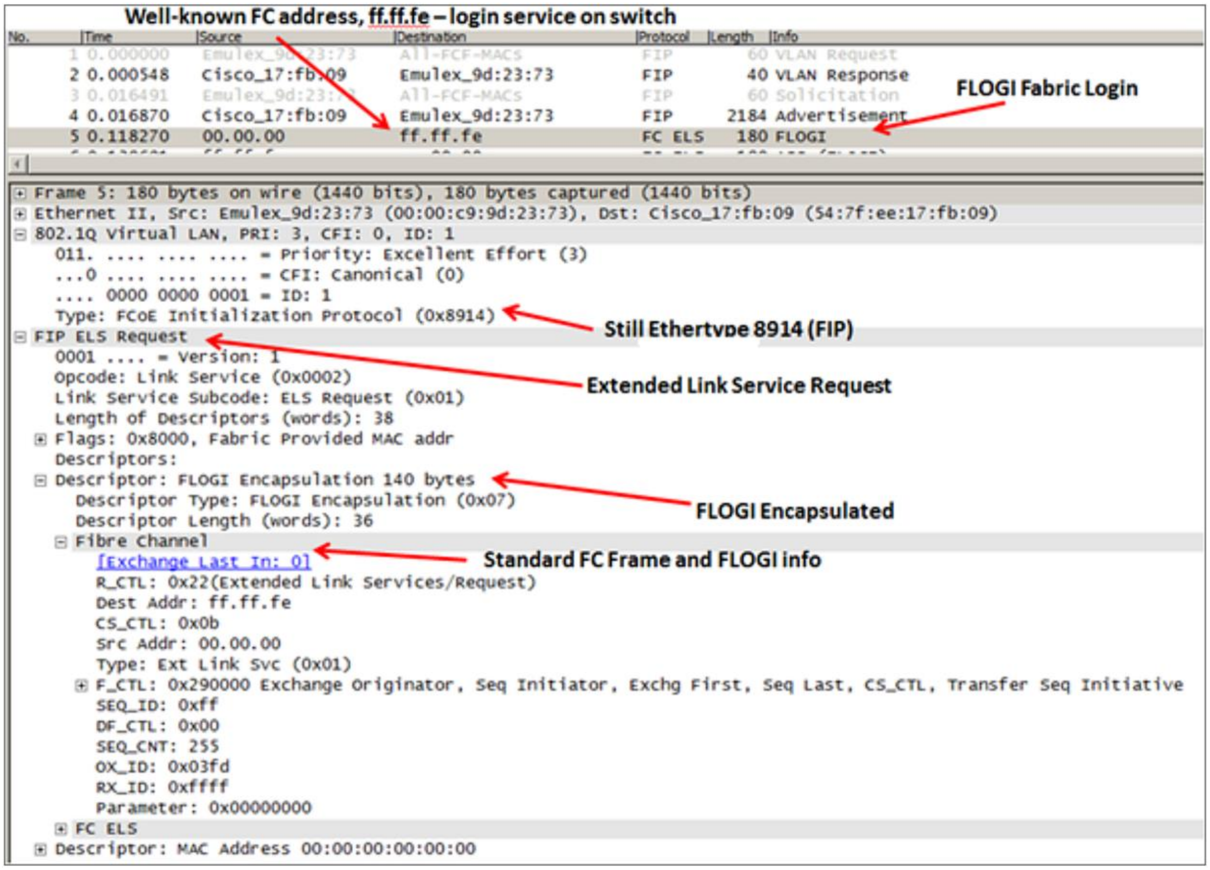

Well-known FC address, ff.ff.fe – login service on switch

FLOGI Fabric Login

Still Ethertype 8914 (FIP)

Extended Link Service Request

FLOGI Encapsulated

Standard FC Frame and FLOGI info

```
No.    Time       Source            Destination        Protocol  Length Info
     1 0.000000   Emulex_9d:23:73   All-FCF-MACs       FIP          60 VLAN Request
     2 0.000548   Cisco_17:fb:09    Emulex_9d:23:73    FIP          40 VLAN Response
     3 0.016491   Emulex_9d:23:73   All-FCF-MACs       FIP          60 Solicitation
     4 0.016870   Cisco_17:fb:09    Emulex_9d:23:73    FIP        2184 Advertisement
     5 0.118270   00.00.00          ff.ff.fe           FC ELS      180 FLOGI

⊞ Frame 5: 180 bytes on wire (1440 bits), 180 bytes captured (1440 bits)
⊞ Ethernet II, Src: Emulex_9d:23:73 (00:00:c9:9d:23:73), Dst: Cisco_17:fb:09 (54:7f:ee:17:fb:09)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 1
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 0001 = ID: 1
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP ELS Request
    0001 .... = Version: 1
    Opcode: Link Service (0x0002)
    Link Service Subcode: ELS Request (0x01)
    Length of Descriptors (words): 38
  ⊞ Flags: 0x8000, Fabric Provided MAC addr
    Descriptors:
  ⊟ Descriptor: FLOGI Encapsulation 140 bytes
      Descriptor Type: FLOGI Encapsulation (0x07)
      Descriptor Length (words): 36
    ⊟ Fibre Channel
        [Exchange Last In: 0]
        R_CTL: 0x22(Extended Link Services/Request)
        Dest Addr: ff.ff.fe
        CS_CTL: 0x0b
        Src Addr: 00.00.00
        Type: Ext Link Svc (0x01)
      ⊞ F_CTL: 0x290000 Exchange Originator, Seq Initiator, Exchg First, Seq Last, CS_CTL, Transfer Seq Initiative
        SEQ_ID: 0xff
        DF_CTL: 0x00
        SEQ_CNT: 255
        OX_ID: 0x03fd
        RX_ID: 0xffff
        Parameter: 0x00000000
    ⊞ FC ELS
  ⊞ Descriptor: MAC Address 00:00:00:00:00:00
```

2.5.1.6 Switch Accepts the CNA FLOGI

The switch accepts the CNA's FLOGI. This is the last step for the FIP VFC instantiation process and Ether Type 8914. All communication after this FLOGI Accept will be EtherType 8906, which is the FCoE's data plane. The screen capture in Figure 13 details the FLOGI accept. In the FC encapsulation, the switch provides the fabric MAC address (FPMA) of 0e:fc:00:aa:00:00. The first three bytes (0e: fc:00) are the switch's FC MAC seen in the switch advertisement capture in Figure 12. The last three bytes are the same as FCID (FC ID), which is aa:00:00.

Although not shown here, periodic FKA messages are EtherType 8914 and will occur every eight seconds between the end device and switch.

Figure 12. FLOGI Accept

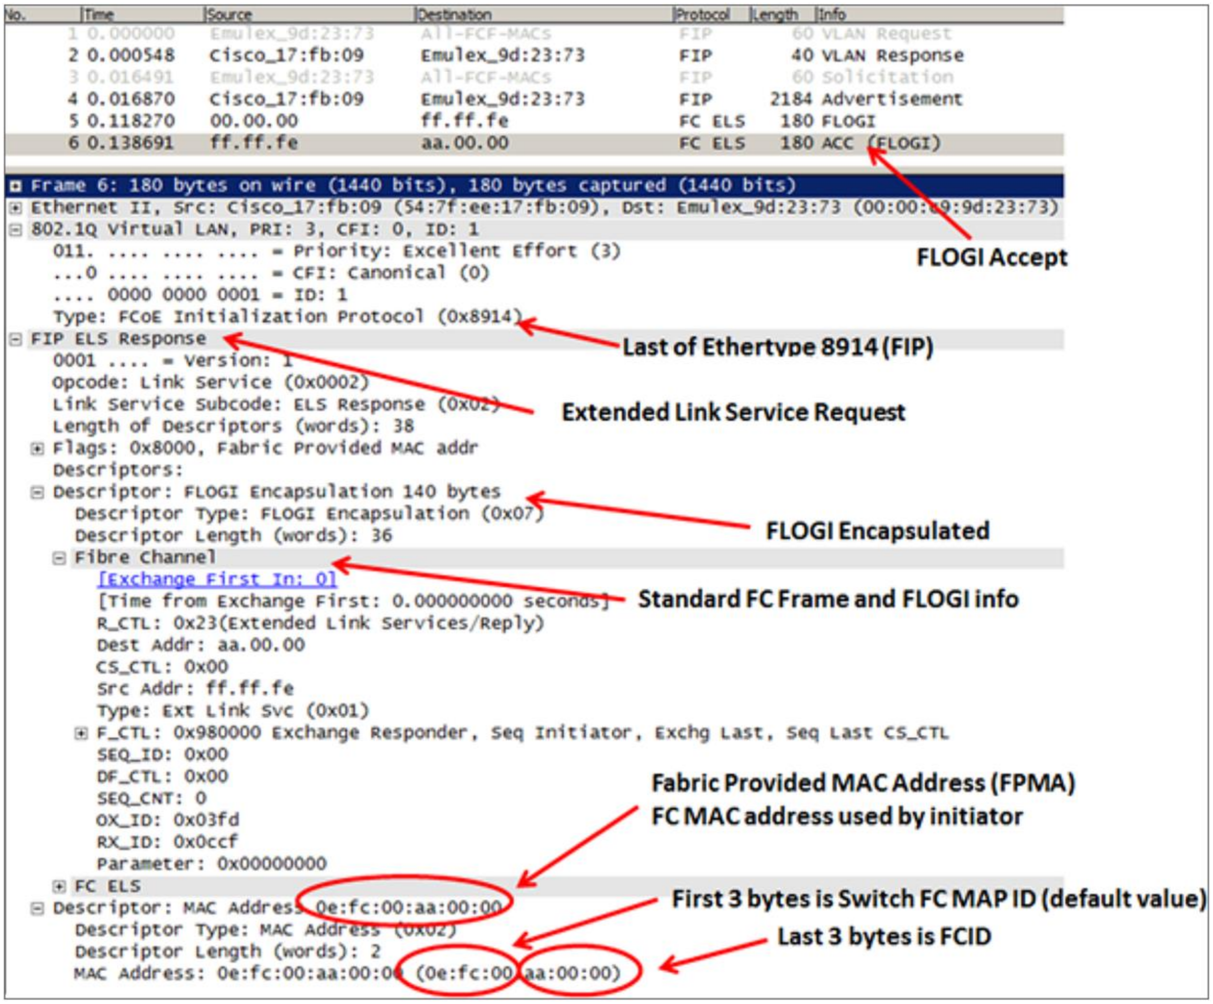

```
No.   Time       Source             Destination        Protocol  Length Info
    1 0.000000   Emulex_9d:23:73    All-FCF-MACs       FIP          60 VLAN Request
    2 0.000548   Cisco_17:fb:09     Emulex_9d:23:73    FIP          40 VLAN Response
    3 0.016491   Emulex_9d:23:73    All-FCF-MACs       FIP          60 Solicitation
    4 0.016870   Cisco_17:fb:09     Emulex_9d:23:73    FIP        2184 Advertisement
    5 0.118270   00.00.00           ff.ff.fe           FC ELS      180 FLOGI
    6 0.138691   ff.ff.fe           aa.00.00           FC ELS      180 ACC (FLOGI)

⊞ Frame 6: 180 bytes on wire (1440 bits), 180 bytes captured (1440 bits)
⊞ Ethernet II, Src: Cisco_17:fb:09 (54:7f:ee:17:fb:09), Dst: Emulex_9d:23:73 (00:00:c9:9d:23:73)
⊟ 802.1Q Virtual LAN, PRI: 3, CFI: 0, ID: 1
    011. .... .... .... = Priority: Excellent Effort (3)
    ...0 .... .... .... = CFI: Canonical (0)
    .... 0000 0000 0001 = ID: 1
    Type: FCoE Initialization Protocol (0x8914)
⊟ FIP ELS Response
    0001 .... = Version: 1
    Opcode: Link Service (0x0002)
    Link Service Subcode: ELS Response (0x02)
    Length of Descriptors (words): 38
  ⊞ Flags: 0x8000, Fabric Provided MAC addr
    Descriptors:
  ⊟ Descriptor: FLOGI Encapsulation 140 bytes
      Descriptor Type: FLOGI Encapsulation (0x07)
      Descriptor Length (words): 36
    ⊟ Fibre Channel
        [Exchange First In: 0]
        [Time from Exchange First: 0.000000000 seconds]
        R_CTL: 0x23(Extended Link Services/Reply)
        Dest Addr: aa.00.00
        CS_CTL: 0x00
        Src Addr: ff.ff.fe
        Type: Ext Link Svc (0x01)
      ⊞ F_CTL: 0x980000 Exchange Responder, Seq Initiator, Exchg Last, Seq Last CS_CTL
        SEQ_ID: 0x00
        DF_CTL: 0x00
        SEQ_CNT: 0
        OX_ID: 0x03fd
        RX_ID: 0x0ccf
        Parameter: 0x00000000
    ⊞ FC ELS
  ⊟ Descriptor: MAC Address 0e:fc:00:aa:00:00
      Descriptor Type: MAC Address (0x02)
      Descriptor Length (words): 2
      MAC Address: 0e:fc:00:aa:00:00 (0e:fc:00:aa:00:00)
```

FLOGI Accept

Last of Ethertype 8914 (FIP)

Extended Link Service Request

FLOGI Encapsulated

Standard FC Frame and FLOGI info

Fabric Provided MAC Address (FPMA)
FC MAC address used by initiator

First 3 bytes is Switch FC MAP ID (default value)

Last 3 bytes is FCID

2.5.1.7 FCID and Domain ID

In FC networks, an FCID is analogous to an IP address, while a WWN is analogous to a MAC address. The FCID is used for routing frames through a FC network. This concept is extended to FCoE environments as well. The FCID is made up of three bytes. In our example capture in Figure 13, aa:00:00, the first byte, aa, corresponds to the FC Domain ID (DID) for the VSAN on the switch. The DID is unique on each switch in a VSAN and is how frames are routed in a FC network. The next two bytes are assigned by the switch. FCIDs are unique to each end device and are only assigned by the switch during the FLOGI process.

It is a best practice to statically assign DIDs to each VSAN on each switch in the fabric. In order to statically assign a DID, the VSAN must be restarted. This is a disruptive event and should normally only occur during the initial VSAN configuration on that switch. To verify the DID settings for each VSAN on a switch, run the following command:

```
N6K# show running-config vsan
snip . . .
vsan database
  vsan 2930
```

```
fcdomain domain 22 static vsan 1
fcdomain domain 22 static vsan 2930
```

Next, verify the configured DID is the actual running DID with this command:

```
N6K# show fcdomain domain-list vsan 2930
VSAN 2930
Number of domains: 8
Domain ID              WWN
---------    -----------------------
 0x18(24)    2b:72:00:2a:6a:4e:de:41 [Principal]
  0x01(1)    2b:72:54:7f:ee:ea:f9:01
  0x02(2)    2b:72:54:7f:ee:ec:79:01
 0x0d(13)    2b:72:00:2a:6a:64:dc:01
 0x15(21)    2b:72:00:2a:6a:5b:52:81
 0x0e(14)    2b:72:54:7f:ee:eb:cf:01
 0x16(22)    2b:72:00:2a:6a:66:ad:81 [Local]
 0x17(23)    2b:72:00:2a:6a:66:a9:81
```

This is the switch we are on (Local). The DID matches running-config. You can see all the other switches (domains) in the VSAN as well.

To determine switches by switch WWN or VSAN DID, run the following command:

```
N6K# show fcs ie vsan 2930
IE List for VSAN: 2930
-------------------------------------------------------------------------------
IE-WWN                   IE     Mgmt-Id  Mgmt-Addr (Switch-name)
-------------------------------------------------------------------------------
2b:72:00:2a:6a:4e:de:41  S(Rem) 0xfffc18 172.29.1.20 (kgmtnc20gsadcr49)
2b:72:00:2a:6a:5b:52:81  S(Rem) 0xfffc15 172.29.1.20 (kgmtnc20gsadcr52)
2b:72:00:2a:6a:64:dc:01  S(Rem) 0xfffc0d 130.6.56.36 (yce293d013)
2b:72:00:2a:6a:66:a9:81  S(Rem) 0xfffc17 172.29.1.20 (kgmtnc20gsadcr50)
2b:72:00:2a:6a:66:ad:81  S(Loc) 0xfffc16 172.29.1.20 (kgmtnc20gsadcr51)
2b:72:54:7f:ee:ea:f9:01  S(Adj) 0xfffc01 130.6.56.32 (yce293d001)
2b:72:54:7f:ee:eb:cf:01  S(Rem) 0xfffc0e 130.6.56.37 (yce293d014)
2b:72:54:7f:ee:ec:79:01  S(Adj) 0xfffc02 130.6.56.33 (yce293d002)
[Total 8 IEs in Fabric]
```

Note the local "Loc" switch. The last byte in the Mgmt-ID is 16, which is the hex DID. This is the same data seen in the previous command output.

To see the FCID for all devices that are logged into the fabric with their corresponding port WWN, look at the Fibre Channel Name Service (FCNS) database. You can also use this command to see what switch a WWN is logged into:

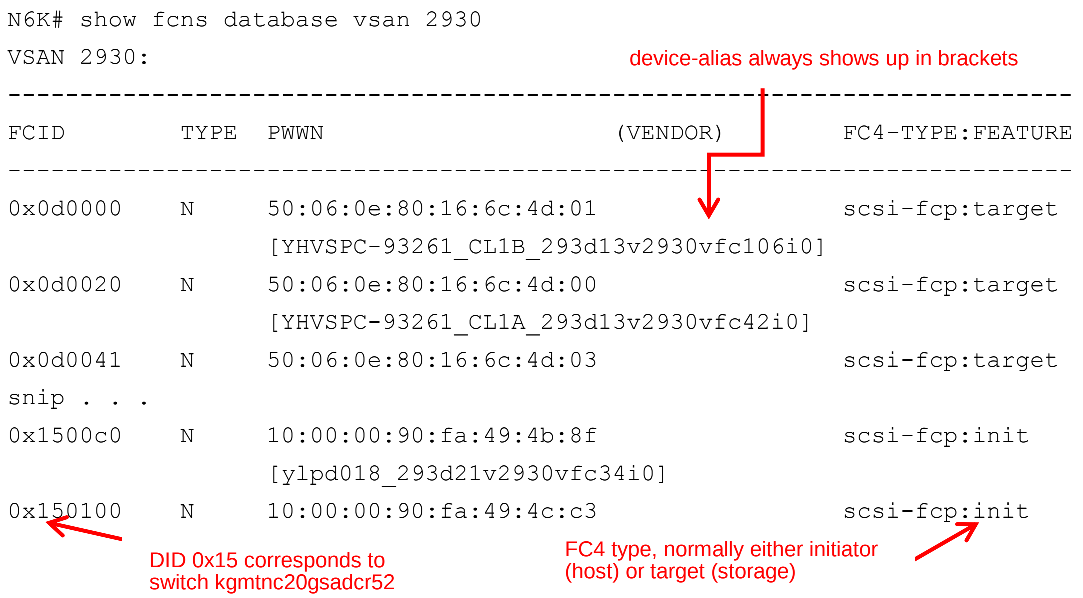

```
N6K# show fcns database vsan 2930
VSAN 2930:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x0d0000    N     50:06:0e:80:16:6c:4d:01                 scsi-fcp:target
                  [YHVSPC-93261_CL1B_293d13v2930vfc106i0]
0x0d0020    N     50:06:0e:80:16:6c:4d:00                 scsi-fcp:target
                  [YHVSPC-93261_CL1A_293d13v2930vfc42i0]
0x0d0041    N     50:06:0e:80:16:6c:4d:03                 scsi-fcp:target
snip . . .
0x1500c0    N     10:00:00:90:fa:49:4b:8f                 scsi-fcp:init
                  [ylpd018_293d21v2930vfc34i0]
0x150100    N     10:00:00:90:fa:49:4c:c3                 scsi-fcp:init
```

device-alias always shows up in brackets

DID 0x15 corresponds to switch kgmtnc20gsadcr52

FC4 type, normally either initiator (host) or target (storage)

3. FCoE Configuration on Cisco Nexus Switches

Section 2 reviewed the enhancements to Ethernet that make FCoE possible. On Cisco Nexus switches, the various enhancements must be configured properly. This configuration section focuses mainly on DCB settings and VFC creation. There are configuration guides for Cisco MDS and Cisco Nexus 6000 and 7000 Series Switches that provide detailed FCoE configurations. Those guides are identified in the References section of this document.

3.1.1 Configuring PFC and ETS

On Cisco Nexus and MDS switches, system-defined class maps for class-fcoe and class-default are enabled. These two classes cannot be deleted. The class-fcoe is defined as no-drop (pause-enabled) and maximum transmission unit (MTU) configuration of 2158 bytes on the Cisco Nexus 6000. On the Cisco Nexus 7000 and MDS the MTU is set to 2112. This MTU helps to ensure the Ethernet frame will encapsulate the largest FC frame and associated FCoE headers. All other traffic falls under class-default and may be dropped. The MTU for class-default is set to 1500 bytes, but it is recommended to be changed to 9216 bytes. Here is the policy with the "show policy-map system:"

```
N6K# show policy-map system type network-qos

  Type network-qos policy-maps
  ===============================

  policy-map type network-qos fcoe-default-nq-policy
    class type network-qos class-fcoe
      match qos-group 1

      pause no-drop
      mtu 2158
    class type network-qos class-default
      match qos-group 0

      mtu 1500
```

Note that class-fcoe is assigned to qos-group 1 while class-default is assigned to qos-group 0.

The following output is for the Cisco Nexus 7000:

```
N7K-storage# show policy-map system type network-qos

  Type network-qos policy-maps
  ============================
  policy-map type network-qos default-nq-7e-policy
    class type network-qos c-nq-7e-drop
      match cos 0-2,4-7
      congestion-control tail-drop
      mtu 1500
    class type network-qos c-nq-7e-ndrop-fcoe
      match cos 3
      match protocol fcoe
      pause
      mtu 2112
```

Here is the output for Cisco MDS switches:

```
MDS9513-A# show policy-map type network-qos default-nq-7e-policy


Type network-qos policy-maps
  ============================
  policy-map type network-qos default-nq-7e-policy template 7e
    class type network-qos c-nq-7e-drop
      congestion-control tail-drop
      mtu 1500
    class type network-qos c-nq-7e-ndrop-fcoe
      pause
      mtu 2112
```

By default, PFC is enabled on all interfaces with a setting of “auto.” To check the PFC status on interfaces, run the “show interface priority-flow-control” command. (The FCoE Troubleshooting section of this document will identify the various show commands and comment on pertinent output.) PFC is the enabling feature that allows a receiver to issue a Pause frame to a transmitter, thus allowing for lossless (no-drop) traffic.

When class-fcoe is not included in the Quality of Service (QoS) policies, VFC interfaces do not come up and increased drops occur. Specifically, the VFC VSAN will stay in an initializing state.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

You need to create a policy map to specify the policies for any user-defined class. In the policy map, you can configure the QoS parameters for each class. You can use the same policy map to modify the configuration of the default classes.

Default bandwidth allocations for ETS are also configured through maps. Of the eight possible classes of service, FCoE is assigned to CoS 3. The default bandwidth percentage assigned to CoS 3 is 50 percent. This allocates 50 percent of the link bandwidth to FCoE traffic during periods of congestion. During periods of non-congestion, other classes can use this bandwidth if FCoE is not.

Although not normally required, the following example shows how to change the allocated bandwidth in different classes:

```
dcn-j-nx5k-1(config)# policy-map type queuing class-fcoe
dcn-j-nx5k-1(config-pmap-que)# class type queuing class-fcoe
dcn-j-nx5k-1(config-pmap-c-que)# bandwidth percent 60
dcn-j-nx5k-1(config-pmap-c-que)# class type queuing class-default
dcn-j-nx5k-1(config-pmap-c-que)# bandwidth percent 40
```

On the Cisco Nexus 7000, you will need to apply network-qos type "default-nq-7e-policy" under system qos.

To see the queueing policy-map, enter the “show policy-map system” command:

```
n7k# show policy-map system type queuing

  Service-policy (queuing) input:   default-in-policy
    policy statistics status:   disabled

    Class-map (queuing):   class-fcoe (match-any)
      Match: qos-group 1
      bandwidth percent 50

    Class-map (queuing):   class-default (match-any)
      Match: qos-group 0
      bandwidth percent 50

  Service-policy (queuing) output:   default-out-policy
    policy statistics status:   disabled

    Class-map (queuing):   class-fcoe (match-any)
      Match: qos-group 1
      bandwidth percent 50

    Class-map (queuing):   class-default (match-any)
      Match: qos-group 0
      bandwidth percent 50
```

In the previous example, qos-group 0 and 1 are both assigned 50 percent of the I/O bandwidth.

3.1.2 Configuring DCBX (LLDP)

The Data Center Bridging Exchange Protocol (DCBXP) is an extension of Link Layer Discovery Protocol (LLDP). It is used to announce, exchange, and negotiate node parameters between peers. DCBXP parameters are packaged into a specific DCBXP Type Length Value (TLV). This TLV is designed to provide an acknowledgement to the received LLDP packet.

DCBXP is enabled by default when you enable LLDP. When LLDP is enabled, DCBXP can be enabled or disabled using the [no] lldp tlv-select dcbxp command. DCBXP is disabled on ports where LLDP transmit or receive is disabled.

Enable LLDP on each FCoE switch by issuing the feature lldp command. On the Cisco Nexus 7000, LLDP is enabled when the FCoE feature-set is installed (in the storage VDC). You cannot disable LLDP while the FCoE feature is installed.

3.1.2.1 DCBX Feature Negotiation

The switch and CNA exchange capability information and configuration values. Cisco Nexus devices support the following capabilities:

- FCoE - If the CNA supports FCoE capability, the switch sends the IEEE 802.1p CoS value to be used with FCoE packets
- PFC - If the adapter supports PFC, the switch sends the IEEE 802.1p CoS values to be enabled with PFC
- Priority group TLV
- Ethernet logical link up and down signal
- FCoE logical link up and down signal for pre-FIP CNAs

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

3.2 Configuring FCoE Interfaces

Ensure the Ethernet configuration pertaining to system and network QoS (PFC and ETS) and LLDP (DCBX) are properly configured and enabled.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

3.2.1 Create FCoE VLAN

The FCoE VLAN will be used for FCoE data and control plane traffic. First create the VSAN and then map it to the FCoE VLAN with the following commands:

```
N6K# configure terminal
N6K(config)# vsan database
N6K#(config-vsan-db) vsan 101
N6K#(config-vsan-db) exit
N6K(config)# vlan 101
N6K(config-vlan)# fcoe vsan 101
N6K(config-vlan)# end

N6K# show vlan fcoe
VLAN VSAN Status
-------- -------- --------
101  101  Operational
```

3.2.2 Create a VFC for the Host (Initiator)

A VFC port is bound to a specific Ethernet port. First configure the Ethernet interface to which the VFC will be bound.

```
N6K# configure terminal
N6K(config)# interface Ethernet101/1/1
N6K(config-if)# description server101
N6K(config-if)# switchport mode trunk
N6K(config-if)# switchport trunk allowed vlan [data vlan and 101]
N6K(config-if)# spanning-tree port type edge trunk
N6K(config-if)# no shut
```

Now create the VFC and bind to Ethernet interface with the following commands:

```
N6K# configure terminal
N6K(config)# interface vfc 1
N6K(config)# vsan database
N6K(config-vsan-db)# vsan 101 interface vfc1
N6K(config)# interface vfc 1
N6K(config-if)# switchport trunk allowed vsan 101
N6K(config-if)# bind interface Ethernet101/1/1
N6K(config-if)# no shut
```

Assuming the CNA on the host is configured properly, connectivity is good, and all switch configuration is proper, the physical and virtual interfaces should come up. An easy check can be done with the following command:

```
N6K# show interface Ethernet101/1/1 fcoe
  Ethernet101/1/1 is FCoE UP
  vfc1 is Up
  FCID is 0x490100
  PWWN is 21:00:00:1b:32:0a:e7:b8
  MAC addr is 00:c0:dd:0e:5f:76
```

3.2.3 Create VFC for Storage (Target)

```
N7K# configure terminal
N7K(config)# interface Ethernet1/1
N7K(config-if)# description HDS_array01_01
N7K(config-if)# switchport trunk allowed vlan [data vlan and 101]
N7K(config-if)# switchport mode trunk
N7K(config-if)# spanning-tree port type edge trunk
N7K(config-if)# no shut
```

Now create the VFC and bind to the Ethernet interface with the following commands:

```
N7K# configure terminal
N7K(config)# interface vfc 1/1
N7K(config)# vsan database
N7K(config-vsan-db)# vsan 101 interface vfc1
N7K(config)# interface vfc 1/1
N7K(config-if)# switchport trunk allowed vsan 101
N7K(config-if)# bind interface Ethernet1/1
N7K(config-if)# no shut
```

Assuming the array port (CNA) is configured properly, connectivity is good, and all switch configurations are proper, the physical and virtual interfaces should come up. An easy check can be done with the following command:

```
N7K# show interface Ethernet1/1 fcoe
  Ethernet1/1 is FCoE UP
  vfc1/1 is Up
  FCID is 0x100100
  PWWN is 50:00:00:1b:32:0a:10:20
  MAC addr is 00:c0:dd:0d:1a:b2
```

An output similar to the above example indicates the physical and VFC interfaces are up, and the FIP process completed successfully with the end device having performed a fabric login (FLOGI). The command above provides summary output to the more detailed, individual commands that follow. These commands and resulting output are similar on both the Cisco Nexus 6000 and 7000.

To check the status of the physical Ethernet interface, use the following command:

Back to Troubleshooting

```
N6K# show interface ethernet 101/1/1
Ethernet101/1/1 is up
  Hardware: 1000/10000 Ethernet, address: c8f9.f920.c102 (bia c8f9.f920.c102)
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA
  Port mode is trunk
  full-duplex, 10 Gb/s, media type is 10G
  Beacon is turned off
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
Input flow-control is off, output flow-control is on
Rate mode is dedicated
Switchport monitor is off
EtherType is 0x8100
Last link flapped 2d06h
Last clearing of "show interface" counters never
19 interface resets
30 seconds input rate 96 bits/sec, 0 packets/sec
30 seconds output rate 160 bits/sec, 0 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 96 bps, 0 pps; output rate 208 bps, 0 pps
RX
  51472 unicast packets  10888 multicast packets  200 broadcast packets
  62560 input packets  6843556 bytes
  0 jumbo packets  0 storm suppression bytes
  0 runts  0 giants  0 CRC  0 no buffer
  0 input error  0 short frame  0 overrun   0 underrun  0 ignored
  0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
  0 input with dribble  0 input discard
  0 Rx pause
TX
  193 unicast packets  52308 multicast packets  0 broadcast packets
  53400 output packets  9581244 bytes
  914 jumbo packets
  0 output error  0 collision  0 deferred  0 late collision
  0 lost carrier  0 no carrier  0 babble 0 output discard
  0 Tx pause
```

Always verify no CRC, discards, or errors have occurred.

Rx and Tx Pause is valid for Nexus 7000 PFC status. Nexus 6000 PFC pause counters are available with “show queuing” output.

To check the status of the VFC interface run the following command:

```
N6K# show int vfc1
vfc1 is trunking
    Bound interface is Ethernet101/1/1
    Hardware is Ethernet
    Port WWN is 20:00:00:2a:6a:35:a5:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 101
    Trunk vsans (admin allowed and active) (101)
    Trunk vsans (up)                       (101)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
        1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
        1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
          1137 frames input, 199416 bytes
            0 discards, 0 errors
          142 frames output, 15312 bytes
            0 discards, 0 errors
        last clearing of "show interface" counters Tue Jun 17 21:47:00 2014

        Interface last changed at Thu Jun 19 19:26:46 2014
```

To check the status of the FC FLOGI, enter the following command:

```
N6K# show flogi database interface vfc1
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
vfc1             101   0xd10000  20:00:74:26:ac:17:2a:b1 10:00:74:26:ac:17:2a:b1

Total number of flogi = 1.
```

3.2.4 Create Zoning

Create zones and add port World Wide Name (PWWN) members to it with the following commands:

```
N6K(config)# zone name zlpd018_293d21v2930vfc1 vsan 2930
N6K(config-zone)# member pwwn 10:00:00:90:fa:49:12:43
N6K(config-zone)# member pwwn 50:06:0e:80:16:6c:43:21
```

Add zone to zoneset and activate the zoneset with the following command:

```
N6K(config)# zoneset name zce293v2930 vsan 2930
N6K(config-zoneset)# member zlpd018_293d21v2930vfc1
N6K(config-zoneset)# zoneset activate name zce293v2930 vsan 2930
N6K(config)# zone commit vsan 2930
```

Verify the status of the newly activated zone:

```
N6K# show zoneset active vsan 2930
snip. . .
zoneset name yce293v2930 vsan 2930
  zone name zlpd018_293d21v2930vfc1 vsan 2930
  * fcid 0x1500c0 [pwwn 10:00:00:90:fa:49:12:43]
  * fcid 0x0d0141 [pwwn 50:06:0e:80:16:6c:43:21]
```

Successful FLOGI will be identified with*.
Note that you can see the FCID.

3.3 Configuration Troubleshooting

The previous section, Overview of FCoE, discussed the various Ethernet enhancements that are required for FCoE. This Configuration Troubleshooting section outlines the discrete steps for configuring VFC interfaces, and will again review commands to ensure these DCB enhancements are configured properly.

3.3.1 Verify Ethernet Interface Status and FCoE Support

FCoE's VFC interface will not come up if the Ethernet interface it is bound to is not up and operational. Check the interface is in the "up" state with the command, "show interface Ethernet 101/1/1 fcoe". Note the keyword, fcoe, at the end of the show statement. The following example shows the status of the VFC interface when the physical interface is up and all other configurations are correct. We know this is true since we can see the VFC is up with an assigned FCID, PWWN, and the MAC address.

```
N6K# show interface e101/1/1 fcoe
Ethernet101/1/1 is FCoE UP
    vfc1 is Up
        FCID is 0xd10000
        PWWN is 20:00:74:26:ac:17:2a:b1
        MAC addr is 74:26:ac:17:2a:b1
```

The output above, though short, is comprehensive since it indicates that the end device has actually performed a fabric flogi (FLOGI). The output reveals that all configuration of PFC, ETS, and DCBX are correct. If any of those features is not configured properly, the VFC will not have a successful FLOGI. You should also confirm that no discards are seen on the interface. Incrementing discards may indicate a PFC configuration problem, that is, that QoS or CoS is not configured correctly.

The following output from the same command indicates that there is a configuration or other problem preventing the VFC instantiation and FIP process from successfully coming up:

```
N6K# show interface e101/1/1 fcoe
Ethernet101/1/1 is FCoE UP
    vfc1 is Up
```

The previous output shows that the VFC is up but there is no successful FLOGI information. As a result, you know that the Ethernet interface is up and the VFC is bound correctly. There could be various reasons why the VFC does not have a successful FLOGI. Some of the causes for this will be identified in the next subsections.

3.3.2 Verification of DCB Ethernet Enhancements

Here are steps to follow in order to verify DCB functions are configured properly.

3.3.2.1 PFC and ETS

PFC allows for per-priority flow control. FCoE requires the no-drop policy be set for its CoS. ETS allows QoS assignment on a CoS. The FCoE CoS is a system-defined QoS value. The default ETS QoS assignment for FCoE is 50 percent and 50 percent for the default CoS. The FCoE QoS requires an MTU setting of 2158 (2112 on the Cisco Nexus 7000). If QoS is not set correctly, the VFC instantiation will not come up (note that the VFC may be trunking but its VSAN will be in initializing state). There will also be no FLOGI. This will be discussed further in section 4.3. To see what the QoS policy is set to, run the following command on Cisco Nexus 6000 switches (Cisco Nexus 7000 commands follow the Nexus 6000 commands).

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
N6K # show policy-map

  Type qos policy-maps
  ====================

  policy-map type qos default-in-policy
    class type qos class-fcoe
      set qos-group 1
    class type qos class-default
      set qos-group 0
```

FCoE class set to qos-group 1
Default class set to qos-group 0

```
  Type queuing policy-maps
  ========================

  policy-map type queuing default-in-policy
    class type queuing class-fcoe
      bandwidth percent 50
    class type queuing class-default
      bandwidth percent 50
  policy-map type queuing default-out-policy
    class type queuing class-fcoe
      bandwidth percent 50
    class type queuing class-default
      bandwidth percent 50
```

FCoE bandwidth set to 50 for in-policy; as well as for the out-policy

```
  Type network-qos policy-maps
  ===============================
  policy-map type network-qos default-nq-policy
    class type network-qos class-fcoe

      pause no-drop
      mtu 2158
    class type network-qos class-default
      mtu 1500
```

FCoE class is set to no-drop; Pause is enabled for PFC; and MTU is set to 2158.

Use the "show queuing interface" command to verify PFC and ETS settings on the FCoE-designated Ethernet interface (Cisco Nexus 6000).

```
N6K# show queuing interface e1/1
Ethernet1/1 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0       WRR             50
        1       WRR             50
```

ETS
Group 0 is default set to 50 percent
Group 1 is fcoe-class at 50 percent

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
  RX Queuing
    qos-group 0
    q-size: 243200, HW MTU: 1600 (1500 configured)
    drop-type: drop, xon: 0, xoff: 243200
    Statistics:
        Pkts received over the port              : 0
        Ucast pkts sent to the cross-bar         : 0
        Mcast pkts sent to the cross-bar         : 0
        Ucast pkts received from the cross-bar   : 0
        Pkts sent to the port                    : 0
        Pkts discarded on ingress                : 0
        Per-priority-pause status                : Rx (Inactive), Tx (Inactive)

    qos-group 1
    q-size: 76800, HW MTU: 2240 (2158 configured)
    drop-type: no-drop, xon: 20480, xoff: 38400
    Statistics:
        Pkts received over the port              : 0
        Ucast pkts sent to the cross-bar         : 0
        Mcast pkts sent to the cross-bar         : 0
        Ucast pkts received from the cross-bar   : 0
        Pkts sent to the port                    : 0
        Pkts discarded on ingress                : 0
        Per-priority-pause status                : Rx (Inactive), Tx (Inactive)

  Total Multicast crossbar statistics:
    Mcast pkts received from the cross-bar       : 0
```

Group 0 COS is set to drop policy

FCoE Group must be set to 2158

FCoE MTU must be set to 2158

xon/xoff verifies qos-group 1 can transmit Pause frame; Group 1 COS set to no-drop

Cisco Nexus 7000 show policy map output:

```
N7K# show policy-map

  Type queuing policy-maps
  ========================

  policy-map type queuing default-in-policy
    class type queuing in-q1
      queue-limit percent 50
      bandwidth percent 80
    class type queuing in-q-default
      queue-limit percent 50
      bandwidth percent 20
  policy-map type queuing default-out-policy
    class type queuing out-pq1
      priority level 1
      queue-limit percent 16
    class type queuing out-q2
      queue-limit percent 1
    class type queuing out-q3
      queue-limit percent 1
    class type queuing out-q-default
      queue-limit percent 82
      bandwidth remaining percent 25
  policy-map type queuing default-4q-7e-in-policy
    class type queuing c-4q-7e-drop-in
      service-policy type queuing default-4q-7e-drop-in-policy
      queue-limit percent 70
    class type queuing c-4q-7e-ndrop-in
      service-policy type queuing default-4q-7e-ndrop-in-policy
      queue-limit percent 30
  policy-map type queuing default-4q-7e-out-policy
    class type queuing c-4q-7e-drop-out
      service-policy type queuing default-4q-7e-drop-out-policy
      bandwidth remaining percent 80
    class type queuing c-4q-7e-ndrop-out
      service-policy type queuing default-4q-7e-ndrop-out-policy
      bandwidth remaining percent 20
  policy-map type queuing default-4q-7e-drop-in-policy
    class type queuing 4q4t-7e-in-q1
      queue-limit percent 10
      bandwidth percent 25
    class type queuing 4q4t-7e-in-q-default
      queue-limit percent 45
      bandwidth percent 25
    class type queuing 4q4t-7e-in-q3
      queue-limit percent 45
      bandwidth percent 25
  policy-map type queuing default-4q-7e-drop-out-policy
    class type queuing 1p3q1t-7e-out-pq1
      priority level 1
    class type queuing 1p3q1t-7e-out-q3
      bandwidth remaining percent 50
    class type queuing 1p3q1t-7e-out-q-default
      bandwidth remaining percent 50
  policy-map type queuing default-4q-7e-ndrop-in-policy
```

```
    class type queuing 4q4t-7e-in-q4
      queue-limit percent 100
      bandwidth percent 25
  policy-map type queuing default-4q-7e-ndrop-out-policy
    class type queuing 1p3q1t-7e-out-q2
      bandwidth remaining percent 100

  Type network-qos policy-maps
  ============================
snip . . .

  policy-map type network-qos default-nq-7e-policy template 7e
    class type network-qos c-nq-7e-drop
      congestion-control tail-drop
      mtu 1500
    class type network-qos c-nq-7e-ndrop-fcoe
      pause
      mtu 2112                                 FCoE MTU set to 2112, pause enabled

snip . . .
```

Cisco Nexus 7000 show queuing interface output:

```
show queuing interface ethernet 1/1
slot  1
=======

Egress Queuing for Ethernet1/1 [System]
-------------------------------------------
Template: 8Q7E
----------------------------------
Group Bandwidth% PrioLevel Shape%
----------------------------------                   FCoE group1 at 50%
    0         50         -      -
    1         50         -      -
-----------------------------------------------------------------------------
     Queue                    Group   Bandwidth% PrioLevel Shape%   CoSMap
-----------------------------------------------------------------------------
7e-4q8q-out-q4                     0         16          -      -          4
7e-4q8q-out-q2                     0         16          -      -          7
7e-4q8q-out-q6                     0         16          -      -          2
7e-4q8q-out-q7                     0         16          -      -          1
7e-4q8q-out-q1                     0          -        High     -          5
7e-4q8q-out-q5                     1        100          -      -          3
7e-4q8q-out-q3                     0         16          -      -          6
7e-4q8q-out-q-default              0         16          -      -          0
```

```
Ingress Queuing for Ethernet1/1 [System]
-------------------------------------------
Trust: Trusted
-------------
Group Qlimit%
-------------
    0      70
    1      30
DSCP to Ingress Queue : Enabled
--------------------------------------------------------------------------------
      Queue                  Group Qlimit% IVL      CoSMap             DSCPMap
--------------------------------------------------------------------------------
7e-4q8q-in-q-default              0      45   0         0-1                0-15
7e-4q8q-in-q1                     0      10   5         5-7               40-63
7e-4q8q-in-q4                     1     100   3           3                   -
7e-4q8q-in-q3                     0      45   2         2,4               16-39
```

FCoE CoS 3 in group1

3.3.2.2 DCBX

DCBX is the protocol that allows network devices to exchange configuration information. DCBX uses the LLDP protocol. To verify that LLDP is enabled, check the feature with the following command:

```
N6K# show feature | include lldp
lldp                  1         enabled
```

Verify that the FCoE-designated Ethernet interface is configured properly (these are default and correct values in the example following) using the following command:

```
N6K# show run interface eth101/1/1 all | include "lldp|priority-flow"
  priority-flow-control mode auto
  lldp transmit
  lldp receive
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

4. FCoE Troubleshooting

The following chart is a guide to assist in identifying an FCoE problem.

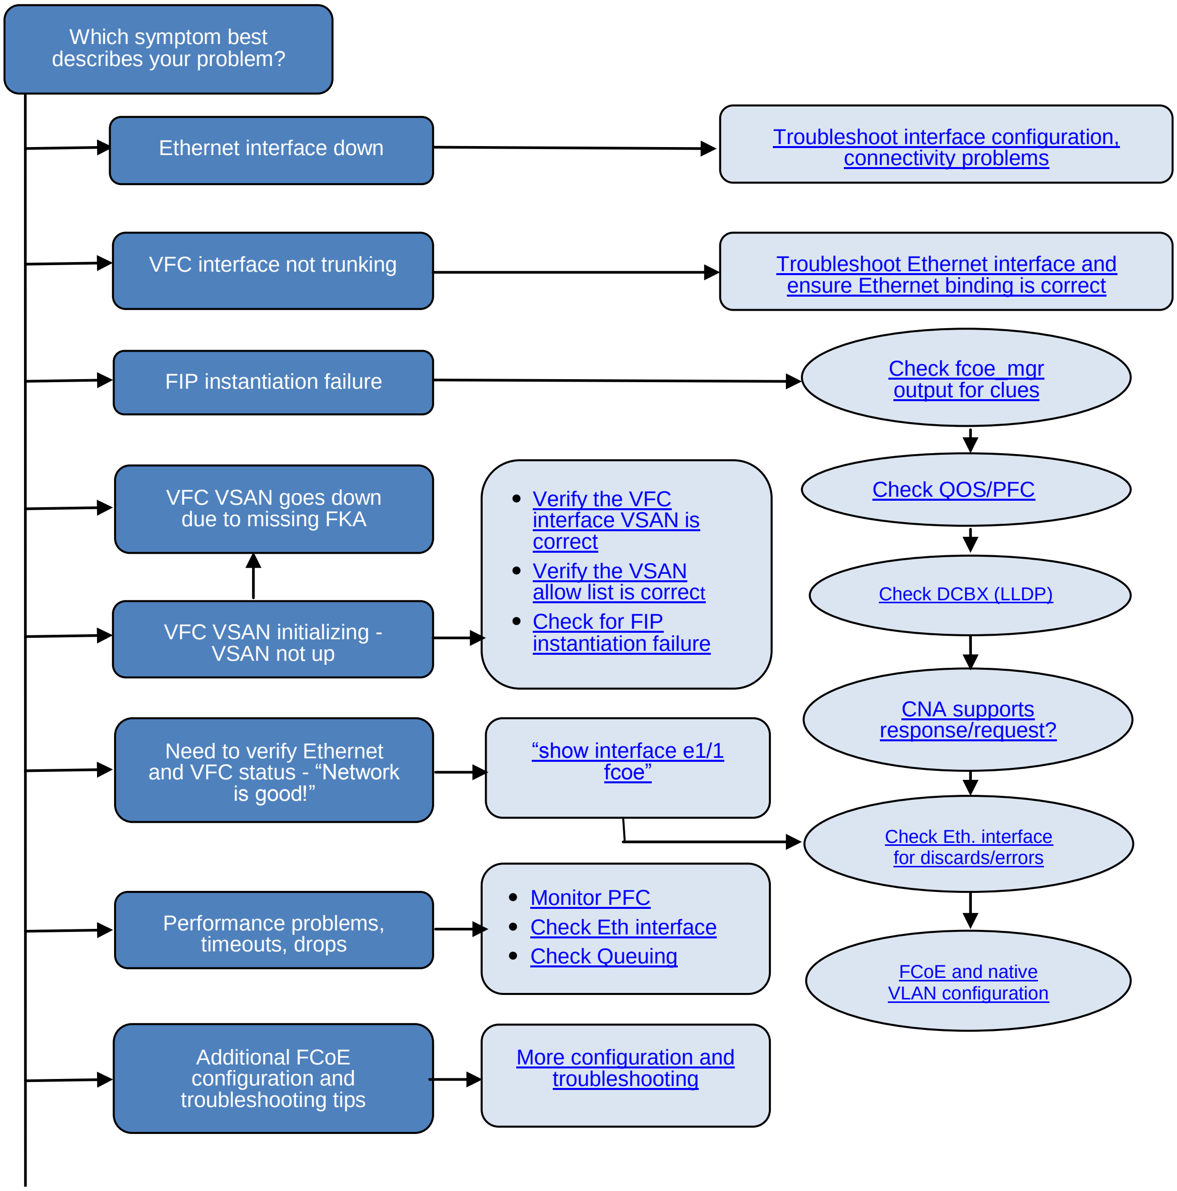
Which symptom best describes your problem?
Ethernet interface down
Troubleshoot interface configuration, connectivity problems
VFC interface not trunking
Troubleshoot Ethernet interface and ensure Ethernet binding is correct
FIP instantiation failure
Check fcoe_mgr output for clues
Check QOS/PFC
Check DCBX (LLDP)
CNA supports response/request?
Check Eth. interface for discards/errors
FCoE and native VLAN configuration
VFC VSAN goes down due to missing FKA
VFC VSAN initializing - VSAN not up
Verify the VFC interface VSAN is correct
Verify the VSAN allow list is correct
Check for FIP instantiation failure
Need to verify Ethernet and VFC status - "Network is good!"
"show interface e1/1 fcoe"
Performance problems, timeouts, drops
Monitor PFC
Check Eth interface
Check Queuing
Additional FCoE configuration and troubleshooting tips
More configuration and troubleshooting

Appendix A to this document provides a guide for recommended steps in how you approach troubleshooting an FCoE problem.

Common troubleshooting steps are listed in Section 4.1. These are referenced in the previous flowchart.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

4.1.1 Understanding the Topology

One of the most useful troubleshooting tools is a topology diagram which details your connectivity from the source to target. The detail should include the physical interfaces (Ethernet interfaces), logical interfaces (VFCs and port-channels), WWN and MAC addresses, etc. Once you have the topology diagram, you can simply go from point to point and check for expected status, always looking for abnormal signs such as errors, discards, pause frames, etc. The topology provided by Cisco Data Center Network Manager (DCNM) may assist in an overall connectivity view.

If you're investigating a host problem, you may need to verify connectivity, starting at the host and going through all the connectivity and switch points to the storage array interfaces. You can look at the active zone set to determine what ports are zoned together.

4.1.2 VFC Interface Not Trunking

The VFC interface will not show as trunking until the Ethernet interface is up and the VFC is bound to it. The VFC interface should be assigned to its VSAN and in a no-shutdown state. Check Ethernet interface status and verify the VFC configuration.

Back to Troubleshooting

```
N6K# show run interface vfc1
interface vfc1
  bind interface Ethernet1/1
  switchport trunk allowed vsan 101
  no shutdown
```

4.1.3 VFC VSAN Is in Initializing State

If you find a VFC with its VSAN in the initializing state, check the conditions in the following example. It is also an example of 'show interface vfc1' output. Although the VFC is showing up because "vfc1 is trunking," the specified VSAN is not up. As a result, FCoE traffic will not traverse the interface. In the following example, VSAN 101 is the configured VSAN for the VFC.

Back to Troubleshooting

```
N6K# show interface vfc1
vfc1 is trunking
    Bound interface is Ethernet101/1/1
    Hardware is Ethernet
    Port WWN is 20:00:00:2a:6a:35:a5:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 101
    Trunk vsans (admin allowed and active) (101)
    Trunk vsans (up)                       ()
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (101)
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
```

You want VSAN to be up, and **not** initializing as in the example here.

```
  12 frames input, 1308 bytes
    0 discards, 0 errors
  13 frames output, 1400 bytes
    0 discards, 0 errors
last clearing of "show interface" counters Tue Jun 17 21:47:00 2014

Interface last changed at Tue Jun 17 21:47:22 2014
```

You should also investigate any discards or errors in the output above, as well as check the bound Ethernet interface for discards and errors.

The switch syslog will also indicate that the trunked VSAN is down, as shown in the following output. Note that this does not indicate the VFC interface is down, but rather, it is the VFC status. Again, this is an indication the VSAN is in an initializing state.

```
2014 Jun 18 20:33:04 N6K %PORT-5-IF_TRUNK_DOWN: %$VSAN 101%$ Interface vfc1, vsan
101 is down (waiting for flogi)
```

4.1.3.1 Check QoS and PFC

If the FCoE class in network-qos settings is not configured, the interface will not come out of an initializing state. Since this is a system-qos value, none of the VFC interfaces will come up since they all rely on the same valid QoS setting.

Back to Troubleshooting

First, check the QoS settings. The following example is from a Cisco Nexus 6000. The different platforms may have slightly different outputs:

```
N6K# show running-config ipqos
system qos
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type qos input fcoe-default-in-policy
  service-policy type network-qos fcoe-default-nq-policy
```

This is the output from a Cisco Nexus 7000:

```
N7K# show running-config ipqos
system qos
  service-policy type network-qos default-nq-7e-policy
```

This is the output from Cisco MDS:

```
MDS9513-A# show running-config eth-qos all

!Command: show running-config eth-qos all
!Time: Mon Dec 15 12:41:03 2014

version 6.2(7)
system qos
  service-policy type network-qos default-nq-7e-policy
```

You can also verify PFC at the interface level with this command (remember that PFC is Ethernet, so look at the Ethernet interface):

```
N6K# show interface ethernet 101/1/1 priority-flow-control
============================================================
Port                Mode Oper(VL bmap)  RxPPP      TxPPP
============================================================

Ethernet101/1/1     Auto On   (8)         0          0
```

If PFC is correct, VL bmap should be 8

Also check detailed queuing information on the interface with this command:

```
N6K# show queuing interface ethernet 101/1/1
if_slot 33, ifidx 0x1f640000
Ethernet101/1/1 queuing information:
  Input buffer allocation:
  Qos-group: 1
  frh: 3
  drop-type: no-drop
  cos: 3
  xon       xoff      buffer-size
  ---------+---------+-----------
  8960      14080     24320

  Qos-group: 0
  frh: 8
  drop-type: drop
  cos: 0 1 2 4 5 6
  xon       xoff      buffer-size
  ---------+---------+-----------
  0         117760    126720

  Queueing:
  queue    qos-group    cos            priority  bandwidth mtu
  --------+------------+--------------+---------+---------+----
  2        0            0 1 2 4 5 6     WRR        50       1600
  3        1            3               WRR        50       2240

  Queue limit: 66560 bytes

  Queue Statistics:
  queue  rx               tx
  ------+---------------+---------------
  2      200              1656
  3      18724            12026
```

Make sure there is group 1 and no-drop

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

Priority-flow-control enabled: yes
Flow-control status:
cos     qos-group   rx pause  tx pause  masked rx pause
-------+-----------+---------+---------+---------------
0              0    xon       xon       xon
1              0    xon       xon       xon
2              0    xon       xon       xon
3              1    xon       xon       xon
4              0    xon       xon       xon
5              0    xon       xon       xon
6              0    xon       xon       xon
7            n/a    xon       xon       xon
```

4.1.3.2 Check DCBX (LLDP)

DCBX is enabled on LLDP-enabled interfaces. LLDP is enabled globally when the feature is enabled. It is possible to disable LLDP (transmit or receive or both on a per interface basis). If either transmit or receive is disabled, DCBX is automatically disabled. Check LLDP/DCBX status on the interface with this command:

[Back to Troubleshooting]

```
N6K# show lldp interface e101/1/1
Interface Information:
  Enable (tx/rx/dcbx): Y/Y/Y    Port Mac address: c8:f9:f9:20:c1:02

Peer's LLDP TLVs:
Type Length Value
---- ------ -----
001  007    047426ac 172aac
002  007    037426ac 172ab0
003  002    0078
127  055    001b2102 020a0000 00000001 0000000e 080a0000 c0008906 001b2108
            06060000 c000ff08 04110000 c000ffff ffff0000 00000000 000008
127  005    00014201 01
000  000
```

Note LLDP and DCBX are enabled

Note TLV info received from peer (good!)

Following is an example of a further check of LLDP/DCBX status.

```
N6K# show lldp dcbx interface e101/1/1

Local DCBXP Control information:
Operation version: 00  Max version: 00  Seq no: 1  Ack no: 1
Type/
Subtype     Version     En/Will/Adv Config
003/000      000         Y/N/Y       0808
```

If there is no type 003 on the switch, then a PFC configuration problem exists.

```
004/000     000        Y/N/Y      8906001b21 08
002/000     000        Y/N/Y      0001000032 32000000 00000002

Peer's DCBXP Control information:
Operation version: 00  Max version: 00  Seq no: 1  Ack no: 1
Type/      Max/Oper
Subtype    Version    En/Will/Err Config
004/000     000/000    Y/Y/N      8906001b21 08
003/000     000/000    Y/Y/N      ff08
002/000     000/000    Y/Y/N      ffffffff00 00000000 00000008
```

Subtype 004
should match.

Check that DCBX packets are incrementing. This command has a lot of detailed information. To look at just the packet status, use the begin command:

```
N6K# show system internal dcbx info interface Ethernet 101/1/1 | begin "DCBX pkt"

DCBX pkt stats:

    Total frames out: 9041
    Total Entries aged: 38
    Total frames in: 9002
    DCBX frames in: 8989
    Total frames received in error: 0
    Total frames discarded: 0
    Total TLVs unrecognized: 0
```

Run this command several times and verify the frames are incrementing.

4.1.3.3 CNA Settings Must Support FIP Request/Response

The CNA settings of the server need to be set correctly in order to support VFC instantiation. Since there are multiple CNA vendors with various models, it is not possible to list all the permutations of settings. The key is that CNA is configured to support FIP and FIP Request/Response. The server system administrator has tools to query the CNA for its current configuration.

Note: FIP issues, abnormal pause frames, and discards may be due to problems with the CNA or the CNA driver. The server administrator should check that the driver or firmware for the CNA is the tested and verified version.

You can verify CNA is communicating with the switch properly by looking at fcoe_mgr information and LLDP. Use the following command to do so:

Back to Troubleshooting

```
N6K# show platform software fcoe_mgr info interface vfc1
vfc1(0x8461fa4), if_index: 0x1e000000, VFC RID vfc1
  FSM current state: FCOE_MGR_VFC_ST_PHY_UP
  PSS Runtime Config:-
      Type: 3
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
      Bound IF: Eth101/1/1
      FCF Priority: 128 (Global)
      Disable FKA: 0
  PSS Runtime Data:-
      IOD: 0x00000000, WWN: 20:00:00:2a:6a:35:a5:3f
      Created at: Tue Jun 17 21:46:53 2014

      FC Admin State: up
      Oper State: up, Reason: down
      Eth IF Index: Eth101/1/1
      Port Vsan: 101
      Port Mode: F port
      Config Vsan: 101
      Oper Vsan: 101
      Solicits on vsan: 101
      Isolated Vsan:
      FIP Capable ? : TRUE
      UP using DCBX ? : FALSE
      Peer MAC : 74:26:ac:17:2a:b1
  PSS VN Port data:-
     FC ID 0xD10000 -
     vfc  index 503316480 vfc name vfc1
     vsan id 101
     enode_mac 74:26:ac:17:2a:b1
     vfc wwn 20:00:74:26:ac:17:2a:b1
snip . . .
```

If there is no VSAN “Solicits on vsan” then FIP is not configured properly on the CNA. The value should show the FCoE VSAN you configured for the VFC.

FIP Capable = TRUE doesn’t necessarily mean FIP is configured properly on CNA. FIP VLAN Request/Response is required. Older drivers require FCoE VLAN to be manually set rather than using Request/ Response.

With successful FLOGI, you can identify the WWN as well as the MAC address. With this CNA model, enode_mac and vfc_wwn are same, except for the leading two bytes in WWN.

4.1.4 CNA Not Receiving a VLAN Response from Switch During FIP

The switch may be sending out a VLAN response, but the response is not received by the CNA. In this case, the VFC will be initializing. This could happen due to various reasons:

FCOE_Troubleshooting_Section

- A bound interface native VLAN ID should be a non-FCoE VLAN. Check native VLAN configuration of the parent Ethernet interface where the VFC is bound. The VLAN request from the host should be received on the native VLAN. The native VLAN cannot be an FCoE VLAN.
- Packet drops are occurring on the network. See the Packet Drop section in 5.1.1.2.

Back to Troubleshooting

4.1.5 CNA Not Sending FIP Keepalives (FKA) as Specified

Cisco Nexus and MDS switches use a default setting of eight seconds for end devices to send FKA frames to switch. A CNA that is not sending FKA frames during the specified switch FKA period is a problem which should be addressed. When a switch does not receive an FKA frame within about 2.5 times the configured setting (about 20 seconds when FKA period set to eight seconds), it will bring down the VFC-trunked VSAN, thus showing that the VFC VSAN is initializing. You may disable FKA at the VFC interface as a workaround while investigating the cause, but this should be done only if absolutely necessary.

When a CNA does not send FKAs in the specified period (default of eight seconds), a driver or firmware version problem may be the cause. The server administrator should verify that the version is the tested and approved version. Although FKA can be disabled at the VFC interface, it is not recommended. Since a large number of missing FKAs may be affect service (VFC flapping), FKA can be disabled while the issue is being investigated at the server.

To check if missing FKAs are causing the trunked VSAN to go down on a VFC, use the following fcoe_mgr command:

Back to Troubleshooting

N6K# show platform software fcoe_mgr event-history errors

Find the most recent entry for the VFC.

```
1) Event:E_DEBUG, length:93, at 203962 usecs after Wed Jun 18 20:33:04 2014
        [102] fcoe_mgr_vfc_ac_eval(4946): DEBUG:shut:Sending event to delete protos
    of vfc1 due to 62


    2) Event:E_DEBUG, length:119, at 197505 usecs after Wed Jun 18 20:33:04 2014
        [102] fcoe_mgr_fc2_msg_handler(5706): proto if_index 1e000000 p_proto (nil)
    and oxid 8805fc2 usrhandle 0[0] iuhdr type:1


    3) Event:E_DEBUG, length:91, at 197053 usecs after Wed Jun 18 20:33:04 2014
        [102] fcoe_mgr_proto_ac_eval(1847): >Bringing down PROTO 1e000000 due to
    truly missing fka
```

The protocol is brought down due to a missing FKA.

To determine the interface index (if_index), run the following command:

```
N6K# show port internal info interface vfc1

vfc1 - if_index: 0x1E000000, phy_port_index: 0x1000
     local_index: 0xffff
  Admin Config - state(up), mode(F), speed(auto), trunk(on)
    beacon(off), snmp trap(on), tem(false)
    description()
snip . . .
```

Check that FKAs are incrementing and check the timestamp of the last FKA event:

```
N6K(config)# show platform software fcoe_mgr info interface vfc1 | begin PROTOS
next 13

PROTOS Info:
vfc1(0x846918c), if_index: 0x1e000000, Proto RID 101, 74:26:ac:17:2a:b1
  FSM current state: FCOE_MGR_PROTO_ST_UP
  PSS Runtime Data:-
      Eth IF Index: Eth101/1/1
      Port Mode: Unknown(0)
      FKA check enabled ? : TRUE
      Recv Multicast solicitation from peer? : FALSE
      Recv Unicast advertisement from peer? : FALSE
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
      Advertisement period from peer? : 0 ms
      Proto number of  devices : 0
  FIP FKA event count : 1679
  FIP FKA last event time stamp : Thu Jun 19 23:21:06 2014
```

Should be incrementing every 8 seconds

Last FKA time stamp

Here is the output of the same command on the Cisco MDS and Nexus 7000:

```
MDS9513-A# show system internal fcoe_mgr info interface vfc610 | begin PROTOS
next 13
PROTOS Info:
vfc610(0x103ff594), if_index: 0x1e000261, VEProto RID vfc610, vsan 10
  FSM current state: FCOE_MGR_VE_PROTO_ST_UP
  PSS Runtime Data:-
      Eth IF Index: Eth6/1
      Port Mode: Unknown(0)
      FKA check enabled ? : TRUE
      Recv Multicast solicitation from peer? : TRUE
      Recv Unicast advertisement from peer? : TRUE
      Advertisement period from peer? : 8000 ms
  FIP FKA event count : 54065
  FIP FKA last event time stamp : Mon Dec 15 12:47:48 2014
```

4.1.6 Check fcoe_mgr Events for FIP Transitions

In addition to the command in 4.1.7 for the FKA status, the fcoe_mgr information output will provide the various state transitions for the VFC. This will help in determining if the VFC is attempting to login with the FIP process and where it may be having issues. The following command is for Cisco Nexus 5000 and 6000 platforms. The Cisco Nexus 7000 command directly follows the output. The following output is a normal, successful fcoe_mgr state transition to the desired FCOE_MGR_VFC_ST_PHY_UP state.

Back to Troubleshooting

```
N6K# show platform software fcoe_mgr info interface vfc1

398) FSM:<vfc1> Transition at 624238 usecs after Thu Jun 26 15:23:10 2014
    Previous state: [FCOE_MGR_VFC_ST_PHY_UP]
    Triggered event: [FCOE_MGR_VFC_EV_BRING_UP_EVAL]
    Next state: [FSM_ST_NO_CHANGE]

399) FSM:<vfc1> Transition at 624270 usecs after Thu Jun 26 15:23:10 2014
    Previous state: [FCOE_MGR_VFC_ST_PHY_UP]
    Triggered event: [FCOE_MGR_VFC_EV_BRING_UP] or
                                            [FCOE_MGR_VFC_EV_FIP_VLAN_DISCOVERY]
    Next state: [FSM_ST_NO_CHANGE]

400) FSM:<vfc1> Transition at 148566 usecs after Thu Jun 26 15:23:16 2014
    Previous state: [FCOE_MGR_VFC_ST_PHY_UP]
    Triggered event: [FCOE_MGR_VFC_EV_FIP_SOLICITATION]
    Next state: [FSM_ST_NO_CHANGE]

401) FSM:<101, 74:26:ac:17:2a:b1> Transition at 150810 usecs after Thu Jun 26
15:23:20 2014
    Previous state: [FCOE_MGR_PROTO_ST_INIT]
    Triggered event: [FCOE_MGR_PROTO_EV_FIP_FLOGI]
    Next state: [FCOE_MGR_PROTO_ST_BRUP_WAIT]

402) FSM:<101, 74:26:ac:17:2a:b1> Transition at 150879 usecs after Thu Jun 26
15:23:20 2014
    Previous state: [FCOE_MGR_PROTO_ST_BRUP_WAIT]
    Triggered event: [FCOE_MGR_PROTO_EV_FIP_FLOGI]
    Next state: [FCOE_MGR_PROTO_ST_FC2_SEND]

403) FSM:<101, 74:26:ac:17:2a:b1> Transition at 158336 usecs after Thu Jun 26
15:23:20 2014
    Previous state: [FCOE_MGR_PROTO_ST_FC2_SEND]
    Triggered event: [FCOE_MGR_PROTO_EV_FC2_DONE]
    Next state: [FCOE_MGR_PROTO_ST_UP]

    Curr state: [FCOE_MGR_PROTO_ST_UP]
```

In the previous output, the current state (Curr state) is the desired state for the VFC. The same output may provide a clue on where the CNA is having issues during the initialization process, such as a configuration error that does not allow FIP to complete. You may have to use Ethanalyzer and capture traffic to investigate where FIP is failing. Refer to the Detailed Steps and Status of FIP Virtual Link Instantiation section in this document for more information.

The command on the Cisco Nexus 7000 and MDS is slightly different, but the output is the same. Use the following command:

```
N7K-storage# show system internal fcoe_mgr info interface vfc101
```

4.2 Best Status Command for FCoE Interface Status

This command has been reviewed in previous sections, and is a simple command. It can be run on a questionable interface that allows for a quick check that FCoE is configured properly between the CNA and switch. This output indicates the VFC interface is completely in the up state, showing the end device has performed a FLOGI (that it, it successfully went through the FIP process). An output similar to the following example normally means any storage issue is probably not network-related:

Back to Troubleshooting

```
N6K# show interface ethernet 101/1/1 fcoe
Ethernet101/1/1 is FCoE UP
    vfc1 is Up
        FCID is 0xd10000
        PWWN is 20:00:74:26:ac:17:2a:b1
        MAC addr is 74:26:ac:17:2a:b1
```

To be certain there are no other underlying network issues, check the output of show interface Ethernet 101/1/1 to ensure there are no discards, pause frames (Nexus 7000), or cyclic redundancy check (CRC) errors occurring. On the Cisco Nexus 6000, verify pause frame counters do not look abnormally high with the show interface priority-flow-control command (outlined in Section 4.3). This command is also valid on the Cisco Nexus 7000 and MDS.

4.3 Monitoring Priority Flow Control (PFC)

Pause frames issued through PFC may be an indication of performance issues at the end device. To check the status of PFC at a glance on all interfaces on either a Cisco Nexus 5000 or 6000 switch, use the following command:

Back to Troubleshooting

```
N6K# show interface priority-flow-control
============================================================
Port               Mode Oper(VL bmap)  RxPPP      TxPPP
============================================================

Ethernet1/1        Auto On   (8)       0          0
Ethernet1/2        Auto On   (8)       0          0
Ethernet1/3        Auto Off            0          4
Ethernet101/1/1    Auto On   (8)       10         0
```

The interfaces above that have a VL bmap of eight are FCoE-enabled interfaces. A large count of RxPPP may indicate a host that is busy or experiencing congestion with FCoE traffic and data traffic. To clear the counters for the previous output, issue the command clear qos statistics. This is a global command and will clear the counters for all interfaces.

It is important to understand that pause frames are normal with FCoE. There is a limited amount of buffering on each switch interface and once reached, pause frames will be transmitted by the device receiving data. There will always be two pause frames sent for a PFC pause event. The first pause tells the device to pause with quanta of 66535. The second pause, with quanta of 0, tells the device to restart transmission immediately.

Interfaces that have a large number of TxPPP may indicate they are experiencing congestion and are asking the transmitter (end device, or in the case of an ISL port, the adjacent switch) to stop sending data. Normally, on Cisco Nexus switches, each interface can handle the full line rate throughput it receives from its connected device. This may not always be the case on oversubscribed line cards and, in that case, suspect the port-group for the line card.

More often, a switch sending numerous pause frames would be seen on ISL ports (VE ports) to its neighbor switch, indicating either the port-channel may be reaching its maximum throughput or that a slow-drain activity is occurring. Slow-drain is a native FC phenomenon related to depletion of buffer credits on MDS switches due to devices that are not able to process traffic as fast as the switch is presenting it. Slow-drain also occurs with FCoE and is related to PFC pause frames. Slow-drain troubleshooting is presented in separate documentation.

On the Cisco Nexus 7000 and MDS, PFC pause frame counters are shown in the output of “show interfaces eth x/x.” Counters can be cleared at the individual interface level with “clear counter interface eth x/x.” The following example shows the output for pause frames on the Cisco Nexus 7000:

Back to Troubleshooting

```
N7K# show interface ethernet 1/5
Ethernet1/5 is up
```

```
admin state is up, Dedicated Interface

  snip . . .

  RX
    787684853 unicast packets  1123408 multicast packets  0 broadcast packets
    788764877 input packets  1107347999422 bytes
    504923610 jumbo packets  0 storm suppression packets
    0 runts  0 giants  0 CRC/FCS  0 no buffer
    0 input error  0 short frame  0 overrun   0 underrun  0 ignored
    0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
    0 input with dribble  0 input discard
    4 Rx pause  ←  Pause frames received from device
  TX
    1049563562 unicast packets  1123385 multicast packets  0 broadcast packets
    1050644471 output packets  1667351027759 bytes
    774888865 jumbo packets
    0 output error  0 collision  0 deferred  0 late collision
    0 lost carrier  0 no carrier  0 babble  0 output discard
    912 Tx pause  ←  Pause frames sent to device
```

To see the Rx and Tx pause counters for all Ethernet interfaces on the Cisco Nexus 7000 and MDS, run the following command:

```
MDS# show interface | include pause|Ether | no-more
snip . . .
Ethernet1/9 is up
  Hardware: 1000/10000 Ethernet, address: 0c68.0329.4808 (bia 0c68.0329.4808)
  EtherType is 0x8100
    1350 Rx pause
    62 Tx pause
Ethernet1/10 is up
  Hardware: 1000/10000 Ethernet, address: 0c68.0329.4809 (bia 0c68.0329.4809)
  EtherType is 0x8100
    950 Rx pause
    24 Tx pause
Ethernet1/11 is down (Administratively down)
  Hardware: 1000/10000 Ethernet, address: 0c68.0329.480a (bia 0c68.0329.480a)
  EtherType is 0x8100
    0 Rx pause
    0 Tx pause

snip . . .
```

These counters are cleared with the “clear counter interface eth x/x” command.

You can run the “show interface priority-flow-control” command on the Cisco Nexus 7000 and MDS, but it will not show the PFC counters, as seen on the Cisco Nexus 6000. What the output will tell you is that PFC is on or off for an interface. Here is an example of the output when the command is run:

```
N7K# show interface priority-flow-control
=====================================
Interface        Admin   Oper
=====================================


port-channel100 Auto  Auto
port-channel101 Auto  Off
port-channel200 Auto  Off
port-channel203 Auto  Auto
port-channel300 Auto  Auto
Ethernet1/1     Auto  Off
Ethernet1/2     Auto  On
Ethernet1/3     Auto  On
snip . . .
```

5. Additional FCoE Configuration and Troubleshooting

This section details initial configuration and verification steps on the switches to help ensure FCoE is configured properly. It also extends troubleshooting steps to more detailed analysis of interfaces that may not be in their correct state.

Back to Troubleshooting

5.1.1 Cisco Nexus 7000 Feature Set and License

Ensure the FCoE feature set is enabled on the Cisco Nexus 7000 before creating a storage VDC with the following command:

```
N7K# show feature-set
Feature Set Name      ID        State
--------------------  --------  --------
fcoe                   1          installed
```

Verify the license installed on the FCoE module(s) is used in storage VDC with this command:

```
N7K# show license fcoe
-------------------------------------
Module-Number    Package-Name
-------------------------------------
      8          FCOE-N7K-F132XP
```

5.1.2 Cisco Nexus 5000 and 6000 License

Verify the FC_FEATURES license is installed with this command:

```
N6K# show license usage
Feature                      Ins  Lic   Status Expiry Date Comments
                                 Count
-------------------------------------------------------------------------
FCOE_NPV_PKG                  No    -   Unused             -
FM_SERVER_PKG                 No    -   Unused             -
ENTERPRISE_PKG                No    -   Unused             -
FC_FEATURES_PKG               Yes   -   In use             -
VMFEX_FEATURE_PKG             No    -   Unused             -
ENHANCED_LAYER2_PKG           Yes   -   Unused Never       -
LAN_BASE_SERVICES_PKG         Yes   -   In use Never       -
LAN_ENTERPRISE_SERVICES_PKG   Yes   -   Unused Never       -
-------------------------------------------------------------------------
```

5.1.3 Cisco Nexus 7000 Ethernet Interface Allocation to Storage VDC

Allocate interfaces to the storage VDC and check that they are configured properly.

```
N7K(config-vdc)# allocate interface ethernet 8/1-2

N7K(config-vdc)# show vdc storage membership

vdc_id: 2 vdc_name: storage interfaces:
        Ethernet8/1           Ethernet8/2
```

5.1.4 Cisco Nexus 7000 FCoE VLAN Allocation to the Storage VDC

In order to configure VSANs in the storage VDC, FCoE VLANs must be allocated from the system/default VDC. Use this command:

```
N7K(config-vdc)# allocate fcoe-vlan-range 101

N7K(config-vdc)# show vdc fcoe-vlan-range
Storage VDC: 2
Reserved Vlans: 101
```

5.1.5 Cisco Nexus VLAN and VSAN Association

On the Cisco Nexus 5000 and 6000, MDS, and Nexus 7000 platforms you must associate the VSAN to an FCoE-designated VLAN. This is the VLAN that will be used for the FIP control plane and FCoE data plane. On the Cisco Nexus 7000, all VSAN and FC-related commands are configured in the storage VDC. Create the VSAN first, then associate with the VLAN.

```
! Create VSAN
MDS(config)# vsan database
MDS(config-vsan-db)# vsan 101 name "FCoE VSAN 101"

! Create VLAN, associate to VSAN
Nexus(config)# vlan 101
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
Nexus(config-vlan)# fcoe vsan 101

! Check status
N7K-storage# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active
101  VLAN0101                      active    Po101, Eth8/1, Eth8/2

VLAN Type  Vlan-mode
---- ----- ----------
1    enet  CE
101  enet  CE

N6K# show vsan
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:down

vsan 101 information
         name:FCoE VSAN 101  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

! Check FCoE VLAN/VSAN association
N7K-storage# show vlan fcoe

Original VLAN ID        Translated VSAN ID      Association State
----------------        ------------------      -----------------

      101                       101              Operational
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

If, after checking the above status and it is still not operational, check the global locked state of the VLAN with this command:

```
MDS# show system internal assoc bitmaps

1. FCoE VLANs:
101

2. FCoE VSANs:
101

3. Created VLANs:
1,101

4. Active VLANs:
1,101

5. No Shut VLANs:
1,101

6. Created VSANs:
1,101

7. Active VSANs:
1,101

8. Global Locked VLANs:

9. Global Locked VSANs:
```

No locked VLANs or VSANs

5.1.6 Verifying FEX FCoE Configuration

Verify FEX is configured properly to support FCoE with this command:

```
N6K# show running-config fex

feature fex

fex 101
  pinning max-links 1
  description "FEX0101"
  fcoe

fex 102
  pinning max-links 1
```

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
  description "FEX0102"

interface Ethernet1/3
  fex associate 101

interface Ethernet1/4
  fex associate 102
```

Note that when using vPC with dual-attached FEXs, only one of the FEXs can be FCoE-enabled from the respective switch’s perspective (in the previous output, FEX 101 is the designated FCoE switch).

Check the status of the FEX with this command:

```
N6K# show fex detail
FEX: 101 Description: FEX0101   state: Online
  FEX version: 6.0(2)N2(3) [Switch version: 6.0(2)N2(3)]
  FEX Interim version: 6.0(2)N2(3)
  Switch Interim version: 6.0(2)N2(3)
  Extender Serial: SSI160309DX
  Extender Model: N2K-C2232PP-10GE,  Part No: 73-12533-05
  Card Id: 82, Mac Addr: c8:f9:f9:20:c1:02, Num Macs: 64
  Module Sw Gen: 12594  [Switch Sw Gen: 21]
  post level: complete
 pinning-mode: static    Max-links: 1
  Fabric port for control traffic: Eth1/3
  FCoE Admin: true
  FCoE Oper: true
  FCoE FEX AA Configured: false
  Fabric interface state:
    Eth1/3 - Interface Up. State: Active
  Fex Port          State  Fabric Port
       Eth101/1/1     Up      Eth1/3
       Eth101/1/2   Down        None
snip . . .
```

FCoE enabled

In the case of a second-level vPC (enhanced vPC), ensure the physical interface of vPC is bound to the VFC and not the port-channel.

5.1.7 Verifying Ethernet Interface Status

If, after verifying physical connectivity and that the interface has been put in the admin up state (no shut), check the following items to determine the issue:

Back to Troubleshooting

```
! Verify configuration of interface
interface Ethernet101/1/1
  switchport mode trunk
  switchport trunk native vlan 10
```

Switchport mode trunk and
native VLAN configured

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
    switchport trunk allowed vlan 10,101
    spanning-tree port type edge
    no shutdown
```

Allowed list is correct;
It includes FCoE VLAN;
STP port type is edge

```
! Verify Ethernet interface show output.  Verify VLAN allowed list:
N6K# show interface ethernet 101/1/1 switchport
Name: Ethernet101/1/1
  Switchport: Enabled
  Switchport Monitor: Not enabled
  Operational Mode: trunk
  Access Mode VLAN: 1 (default)
  Trunking Native Mode VLAN: 10 (native)
  Trunking VLANs Allowed: 10,101
  Voice VLAN: none
  Extended Trust State : not trusted [COS = 0]
  Administrative private-vlan primary host-association: none
  Administrative private-vlan secondary host-association: none
  Administrative private-vlan primary mapping: none
  Administrative private-vlan secondary mapping: none
  Administrative private-vlan trunk native VLAN: none
  Administrative private-vlan trunk encapsulation: dot1q
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk private VLANs: none(0 none)
  Operational private-vlan: none
  Unknown unicast blocked: disabled
  Unknown multicast blocked: disabled
```

If the previous output is not in the expected state, look at the following internal event history to determine if any failure occurred. The normal up state is: ETH_PORT_FSM_ST_TRUNK_UP

```
N6K# show system internal ethpm event-history interface ethernet 101/1/1 |
include "Curr state"
    Curr state: [ETH_PORT_FSM_ST_TRUNK_UP]
```

Some of the states that may be reported other than the normal up state in the previous output are listed in the immediately following output.

FSM might be in one of the following states:

ETH_PORT_FSM_ST_NOT_INIT

ETH_PORT_FSM_ST_DOWN

ETH_PORT_FSM_ST_INIT_EVAL

ETH_PORT_FSM_ST_SPAN_EVAL

ETH_PORT_FSM_ST_WAIT_PRE_CFG

ETH_PORT_FSM_ST_LINK_INIT

ETH_PORT_FSM_ST_WAIT_BRINGUP

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

ETH_PORT_FSM_ST_WAIT_LOGICAL_UP

ETH_PORT_FSM_ST_L2_UP

ETH_PORT_FSM_ST_L3_UP

ETH_PORT_FSM_ST_PROTOCOL_DOWN

ETH_PORT_FSM_ST_SPAN_DEST_UP

ETH_PORT_FSM_ST_WAIT_PROTOCOL_DOWN

ETH_PORT_FSM_ST_WAIT_PHYSICAL_DOWN

ETH_PORT_FSM_ST_WAIT_APPLY_CONFIG

ETH_PORT_FSM_ST_WAIT_LOGICAL_DOWN

ETH_PORT_FSM_ST_WAIT_LOGICAL_CHANGE_TRUNK

ETH_PORT_FSM_ST_NOT_UP

ETH_PORT_FSM_ST_BUNDLE_MEMBER_UP

ETH_PORT_FSM_ST_WAIT_BUNDLE_PRE_CFG

ETH_PORT_FSM_ST_WAIT_BUNDLE_LOGICAL_UP

ETH_PORT_FSM_ST_WAIT_BUNDLE_LOGICAL_DOWN

ETH_PORT_FSM_ST_WAIT_BUNDLE_MEMBER_DOWN

ETH_PORT_FSM_ST_ERROR_DISABLED_LEVEL_1

ETH_PORT_FSM_ST_ERROR_DISABLED_LEVEL_2

ETH_PORT_FSM_ST_AUTH_FAIL

ETH_PORT_FSM_ST_WAIT_LOGICAL_DOWN_RNF 30

ETH_PORT_FSM_ST_WAIT_PROTOCOL_DOWN_RNF 31

There may be possibility that a sequence error occurred, which would be logged to the syslog similar to the following output on the console:

2014 Jun 22 15:01:37 DCE-1 %$ VDC-1 %$ %ETHPORT-2-IF_SEQ_ERROR: Error ("sequence timeout") while communicating with component MTS_SAP_ETH_PORT for opcode MTS_OPC_ETHPM_PORT_BRINGUP (RID_PORT: Ethernet1/16)

2014 Jun 22 06:39:11 dist-B %$ VDC-1 %$ %ETHPORT-2- IF_DOWN_ERROR_DISABLED: Interface Ethernet1/3 is down (Error disabled. Reason:sequence timeout)

2014 Jun 22 06:39:51 dist-B %$ VDC-1 %$ %ETHPORT-2-IF_SEQ_ERROR: Error ("sequence timeout") communicating with MTS_SAP_ETH_PORT_CHANNEL_MGR for opcode MTS_OPC_ETHPM_PORT_CLEANUP

If the Ethernet interface is still down, collect the “show tech-support ethpm” and contact the Cisco TAC.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

5.1.8 Verifying Ethernet Port-Channel Status

First, verify physical connections on the port-channel member interfaces. Check the cable between the port-channel members. If the cable is not connected “show interface ethernetx/y” will show as:

```
show int ethernet 1/1
Ethernet1/1 is down (Link not connected)
```

If Ethernet ports between the Cisco Nexus switch are up, but Ethernet port-channels are not up, check the running configuration on the core and edge switches with this command:

```
show run interface port-channel 101

interface port-channel101
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 101
```

If the running configuration looks correct, check the trunking status of the allowed VLANs by executing this command:

```
N7K-storage# show interface port-channel 101 switchport
Name: port-channel101
  Switchport: Enabled
  Switchport Monitor: Not enabled
  Operational Mode: trunk
  Access Mode VLAN: 1 (default)
  Trunking Native Mode VLAN: 1 (default)
  Trunking VLANs Allowed: 101
  FabricPath Topology List Allowed: 0
  Administrative private-vlan primary host-association: none
  Administrative private-vlan secondary host-association: none
  Administrative private-vlan primary mapping: none
  Administrative private-vlan secondary mapping: none
  Administrative private-vlan trunk native VLAN: none
  Administrative private-vlan trunk encapsulation: dot1q
  Administrative private-vlan trunk normal VLANs: none
  Administrative private-vlan trunk private VLANs: none
  Operational private-vlan: none
```

Next, check the VLAN membership of the port-channel interface with this command:

```
N7K-storage# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active
101  VLAN0101                         active    Po101, Eth8/1, Eth8/2
```

Verify no sequence timeout errors exist in the syslog.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

2014 Jun 22 15:01:37 DCE-1 %$ VDC-1 %$ %ETHPORT-2-IF_SEQ_ERROR: Error ("sequence timeout") while communicating with component MTS_SAP_ETH_PORT_CHANNEL_MGR for opcode MTS_OPC_ETHPM_PORT_BRINGUP (RID_PORT: Ethernet1/1)

2014 Jun 22 06:39:11 dist-B %$ VDC-1 %$ %ETHPORT-2-IF_DOWN_ERROR_DISABLED: Interface Ethernet1/3 is down (Error disabled. Reason:sequence timeout)

2014 Jun 22 06:39:51 dist-B %$ VDC-1 %$ %ETHPORT-2-IF_SEQ_ERROR: Error ("sequence timeout") communicating with MTS_SAP_ETH_PORT_CHANNEL_MGR for opcode MTS_OPC_ETHPM_PORT_CLEANUP

Similar to when we checked individual interfaces above, use the following command to check the current state of the port-channel interface. It should be in the state, ETH_PORT_FSM_ST_TRUNK_UP.

```
N7K-storage# show system internal ethpm event-history interface port-channel 101
| inc “Curr state”


Curr state: [ETH_PORT_FSM_ST_TRUNK_UP]
```

Refer to Section 4.6.6 to see other possible states of the interface.

5.1.9 Verifying DCBX/LLDP for VFC Port-Channels

If DCBX negotiation fails between switches for port-channels, most likely there is a QoS PFC configuration incompatibility between the switches. There may be a syslog event:

```
2014 Jun 18 10:17:18 N7K-storage %IPQOSMGR-2-QOSMGR_DCBXP_PFC_CMP_FAIL_MSG:
Ethernet8/2 - qos config 'Priority-flow-control' not compatible with the peer
```

Check the DCBX status on the interfaces making up the port-channel with this command:

```
N7K-storage# show system internal dcbx info interface ethernet 8/1


Interface info for if_index: 0x1a380000(Eth8/1)
tx_enabled: TRUE
rx_enabled: TRUE
dcbx_enabled: TRUE
DCX Protocol: CEE
```

All interfaces should be set to TRUE.

Verify the VFC port-channel interface (FCoE) comes up with this command:

```
N7K-storage# show interface vfc101
vfc101 is trunking
    Bound interface is port-channel101
    Hardware is Ethernet
    Port WWN is 20:64:00:26:98:08:fb:3f
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 20 Gbps
    Trunk vsans (admin allowed and active)  (101)
    Trunk vsans (up)                        (101)
```

VSAN is up.

```
Trunk vsans (isolated)                ()
Trunk vsans (initializing)            ()
256419 fcoe in packets
21936804 fcoe in octets
256418 fcoe out packets
28800216 fcoe out octets
Interface last changed at Sun Jun 11 08:04:18 2014
```

If the VFC and Ethernet configuration look correct, verify LLDP neighbors with this command:

```
N7K-storage# show lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID            Local Intf      Hold-time  Capability  Port ID
N6K                  Eth8/1          120        BR          Eth1/1
N6K                  Eth8/2          120        BR          Eth1/2
Total entries displayed: 2
```

The Cisco Nexus 7000’s neighbor is the Nexus 6000.

Verify LLDP DCBX exchange information for the local and peer interface reflect the same with this command:

```
N7K-storage# show lldp dcbx interface ethernet 8/1

Local DCBXP Control information:
Operation version: 00  Max version: 00  Seq no: 20  Ack no: 34
Type/
Subtype     Version     En/Will/Adv Config
003/000      000         Y/N/Y      0808
004/000      000         Y/N/Y      8906001b21 08
002/000      000         Y/N/Y      1123200019 19191900 00000004

Peer's DCBXP Control information:
Operation version: 00  Max version: 00  Seq no: 34  Ack no: 20
Type/       Max/Oper
Subtype     Version     En/Will/Err Config
003/000      000/000     Y/N/N      0808
004/000      000/000     Y/N/N      8906001b21 08
002/000      000/000     Y/N/Y      0001000032 32000000 00000002
```

In the previous output, class subtype 3 refers to FCoE APP. En refers to enable. Will refers to willing. For FCoE, En should be Y after a successful DCBX exchange.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Verify fcfwd mpmap shows the MAC address for the respective VFCs (Cisco Nexus 7000 and MDS command) with this command:

```
N7K-storage# show system internal fcfwd mpmap vfcs

                           FCoE VFC Information

--------------------------------------------------------------------------
     ID     | if-index |S|M|T|          Members          |      MAC(s)
------------+----------+-+-+-+---------------------------+-----------------
vfc101      |0x1e000064|D|E|-|0x16000064|eth-pc 101 (U)*|00:26:98:08:fb:00
            |          | | | |          |               |00:2a:6a:35:a5:18
------------+----------+-+-+-+---------------------------+-----------------
```

In the previous output, the MAC address should reflect both the MAC addresses, that is, the peer and the local interfaces' MAC addresses.

Verify port-channel VFC is in the up state. This command is valid on the Cisco Nexus 7000 and MDS. The Cisco Nexus 6000 is slightly different and can be seen in the following output:

```
N7K-storage# show system internal fcoe_mgr info interface vfc101
vfc101(0x83b6c54), if_index: 0x1e000064, VFC RID vfc101
  FSM current state: FCOE_MGR_VFC_ST_PHY_UP
  PSS Runtime Config:-
      Type: 3
      Bound IF: Po101
      FCF Priority: 128 (Global)
  PSS Runtime Data:-
      IOD: 0x00000000, WWN: 20:64:00:26:98:08:fb:3f
      Created at: Sun Jan 11 08:03:52 2009

      FC Admin State: up
      Oper State: up, Reason: down
      Eth IF Index: Po101
      Port Vsan: 1
      Port Mode: E port
      Config Vsan: 101
      Oper Vsan: 101
      Solicits on vsan: 101
      Isolated Vsan:
      FIP Capable ? : TRUE
      UP using DCBX ? : FALSE
  PSS VN Port data:-
```

This is the desired state.

ISLs are E ports.

Use this command to run on the Cisco Nexus 6000:

```
N6K# show platform software fcoe_mgr info interface vfc101
vfc101(0x845b264), if_index: 0x1e000064, VFC RID vfc101
  FSM current state: FCOE_MGR_VFC_ST_PHY_UP
  PSS Runtime Config:-
      Type: 3
      Bound IF: Po101
      FCF Priority: 128 (Global)
      Disable FKA: 0
  PSS Runtime Data:-
      IOD: 0x00000000, WWN: 20:64:00:2a:6a:35:a5:3f
      Created at: Mon Jan 16 16:30:20 2012

      FC Admin State: up
      Oper State: up, Reason: down
      Eth IF Index: Po101
      Port Vsan: 1
      Port Mode: E port
      Config Vsan: 101
      Oper Vsan: 101
      Solicits on vsan: 101
      Isolated Vsan:
      FIP Capable ? : TRUE
      UP using DCBX ? : FALSE
      Peer MAC : 00:26:98:08:fb:00
  PSS VN Port data:-
```

To run this command on the Cisco Nexus 7000, use the following:

```
N7K# show system internal fcoe_mgr info interface vfc101
```

The other possible FSM states that may warrant further investigation include:

No transition

FSM_ST_NO_CHANGE

FSM_ST_AN

FCOE_MGR_VE_PROTO_ST_INIT

FCOE_MGR_VE_PROTO_ST_BRUP_WAIT

FCOE_MGR_VE_PROTO_ST_BRDOWN_WAIT

FCOE_MGR_VE_PROTO_ST_DOWN

FCOE_MGR_VE_PROTO_ST_DELETE_WAIT

If VFC is up and FCF discovery has gone through, the FCoE database will show the VFC interfaces in the FCoE database. Here is the output:

```
N6K# show fcoe database
```

Normal end-device VFC logins show up here.

```
--------------------------------------------------------------------------

INTERFACE       FCID            PORT NAME               MAC ADDRESS

--------------------------------------------------------------------------

vfc1            0xd10000        20:00:74:26:ac:17:2a:b1 74:26:ac:17:2a:b1

Total number of flogi count from FCoE devices = 1.

VE Ports:
```

VE ports are FCoE ISLs.

```
--------------------------------------------------------------------------

INTERFACE       MAC ADDRESS             VSAN

--------------------------------------------------------------------------

vfc101          00:26:98:08:fb:00       101
```

5.1.10 Verify the VFC Interface and FLOGI Status

If the Ethernet interface is up, but VFC is not, execute the “show port internal event-history interface vfc1” command and check for the FSM state. The following status shows a normal interface status:

```
N6K# show port internal event-history interface vfc 1

>>>>FSM: <vfc1> has 200 logged transitions<<<<<

snip . . .

200) FSM:<vfc1> Transition at 623770 usecs after Thu Jun 26 15:23:18 2014
    Previous state: [PI_FSM_ST_TXPORT_INIT_TRUNKING_ENABLED]
    Triggered event: [PI_FSM_EV_PACER_TIMER_EXPIRED]
    Next state: [FSM_ST_NO_CHANGE]

    Curr state: [PI_FSM_ST_TXPORT_INIT_TRUNKING_ENABLED]
```

Desired state

If the state is not correct, execute “show port internal event-history errors” and observe any error associated with VFC port.

```
N6K# show port internal event-history errors

snip . . .
47) Event:E_DEBUG, length:245, at 835116 usecs after Mon Jan 16 16:30:20 2012
    [102] pm_error_disable_port: parent_function pi_fsm_ac_port_init_resp_rcvd,
Ifindex (vfc101)0x1e000064,  Err disabled VLAN L2 down on Eth
interface(0x42070010) event 0x138 reason (pre_i
nit_from_fcoe_mgr_failed),  cfg_wait_str: cfg wait for none
```

Interface is shut down

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Verify the Ethernet interface belongs to the correct VLAN(s), namely FCoE VLAN, with this command:

```
N6K# show vlan id 101

VLAN Name                             Status    Ports
---- -------------------------------- --------- --------------------------
101  VLAN0101                         active    Po101, Eth1/1, Eth1/2
                                                Eth101/1/1

VLAN Type  Vlan-mode
---- ----- ----------
101  enet  CE
```

Confirm the VFC belongs to the correct VSAN with this command:

```
N6K# show vsan 101 membership
vsan 101 interfaces:
    vfc1
```

5.1.11 Verify Queuing on the Ethernet Interface

Back to Troubleshooting

Interface queuing details can be found using the “show queuing interface” command. QoS group 1 is the no-drop CoS 3 group used for FCoE. From the following Cisco Nexus 6000 output, it is evident that the buffer size for the interface and the buffer limit before pause frames are transmitted. “xoff” below the 14080 column identifies the buffer space of 14080 bytes available for holding received frames. If that limit is reached, the switch will send a pause frame to the end device. It will not accept frames (it will no longer issue pause frames) until the xon value of 8960 is reached (8960 bytes of free buffer space). The following output is a FEX interface:

```
N6K# show queuing interface ethernet 101/1/1
if_slot 33, ifidx 0x1f640000
Ethernet101/1/1 queuing information:
  Input buffer allocation:
  Qos-group: 1
  frh: 3
  drop-type: no-drop
  cos: 3
  xon       xoff      buffer-size
  ---------+---------+-----------
  8960      14080     24320

  Qos-group: 0
  frh: 8
  drop-type: drop
  cos: 0 1 2 4 5 6
  xon       xoff      buffer-size
  ---------+---------+-----------
  0         117760    126720
```

QoS group 1 is no-drop and CoS is 3.

If more than 14080 bytes are in queue, send Pause frame -- xoff = transmit off

When 8960 bytes are available in queue, we resume (stop sending Pause frame) -- xon = transmit on

```
Queueing:
queue    qos-group    cos            priority  bandwidth mtu
--------+------------+--------------+---------+---------+----
2        0            0 1 2 4 5 6     WRR        50       1600
3        1            3               WRR        50       2240
```

CoS 3 (FCoE)
MTU is 2240

```
Queue limit: 66560 bytes

Queue Statistics:
queue  rx              tx
------+---------------+---------------
2      200             3113
3      19805           100065
```

There should be no drops.

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

Priority-flow-control enabled: yes
Flow-control status:
cos     qos-group   rx pause  tx pause  masked rx pause
-------+-----------+---------+---------+---------------
0             0     xon       xon       xon
1             0     xon       xon       xon
2             0     xon       xon       xon
3             1     xon       xon       xon
4             0     xon       xon       xon
5             0     xon       xon       xon
6             0     xon       xon       xon
7           n/a     xon       xon       xon
```

CoS 3 belongs to group 1

To show the difference between a FEX port and a physical Cisco Nexus 6000 interface, see a portion of the "show queuing interface" in the following output. This is an uplink port to the FEX. It has an xoff buffer size of 88320 compared to 14080 that was seen on the FEX port.

```
N6K# show queuing interface Ethernet 3/1/1
Ethernet1/1 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0        WRR            15
        1        WRR            50
        2        WRR             0
```

QoS group 1 (FCoE) with 50 percent of bandwidth

```
        3       WRR              2
        4       WRR             15
        5       priority        15

snip . . .

    qos-group 1
    q-size: 165120, HW MTU: 2158 (2158 configured)
    drop-type: no-drop, xon: 62720, xoff: 88320
    Statistics:
        Pkts received over the port         : 226937448
        Ucast pkts sent to the cross-bar    : 226970619
        Mcast pkts sent to the cross-bar    : 0
        Ucast pkts received from the cross-bar      : 831906773
        Pkts sent to the port               : 831906449
        Pkts discarded on ingress           : 0
        Per-priority-pause status           : Rx (Active), Tx (Inactive)

snip . . .
```

Note the buffer size of more than the 14080 seen on the FEX interface.

Confirm the VFC interface is trunking and that VSAN is up. The VSAN must be in the up state for FCoE traffic to traverse the VFC.

```
N6K# show interface vfc1
vfc1 is trunking
    Bound interface is Ethernet101/1/1
    Hardware is Ethernet
    Port WWN is 20:00:00:2a:6a:35:a5:3f
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 101
    Trunk vsans (admin allowed and active)  (101)
    Trunk vsans (up)                        (101)
    Trunk vsans (isolated)                  ()
    Trunk vsans (initializing)              ()
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1154 frames input, 201320 bytes
        0 discards, 0 errors
      160 frames output, 17332 bytes
        0 discards, 0 errors
    last clearing of "show interface" counters Tue Jun 17 21:47:00 2014

    Interface last changed at Thu Jun 26 15:23:10 2014
```

VSAN up

Always verify no discards or errors

Verify VFC has performed a FLOGI with the following command:

```
N6K# show flogi database interface vfc1
--------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------
vfc1             101   0xd10000  20:00:74:26:ac:17:2a:b1 10:00:74:26:ac:17:2a:b1
```

If VFC is up, but “show flogi database” doesn’t show an entry for it, check the FSM state using “show flogi internal event-history interface vfc”.

```
N6K# show flogi internal event-history interface vfc1

>>>>FSM: <[101]20:00:74:26:ac:17:2a:b1> has 15 logged transitions<<<<<

snip . . .

15) FSM:<[101]20:00:74:26:ac:17:2a:b1> Transition at 158180 usecs after Thu Jun
26 15:23:20 2014
    Previous state: [FLOGI_ST_PERFORM_CONFIG]
    Triggered event: [FLOGI_EV_CONFIG_DONE_COMPLETE]
    Next state: [FLOGI_ST_FLOGI_DONE]

    Curr state: [FLOGI_ST_FLOGI_DONE]
```

Desired state

If the FLOGI is in any of the following states, check the FLOGI event-history error in the previous output to investigate.

FLOGI_ST_FLOGI_RECEIVED

FLOGI_ST_GET_FCID

FLOGI_ST_PERFORM_CONFIG

FLOGI_ST_FLOGI_DONE

FLOGI_ST_CLEANUP

FLOGI_ST_DESTROY_FSM

FLOGI_ST_PERFORM_FCFWD_CONFIG

FLOGI_ST_FETCH_PRECFG_INFO

FLOGI_ST_QUERY_PORT_NUMBER

FLOGI_ST_CHECK_SECURITY_NEGOTIATION

FLOGI_ST_FCSP_READY

FLOGI_ST_FCSP_HANDSHAKE_RCVD_FLOGI_CFG

FLOGI_ST_WAIT_FCSP_DONE

FLOGI_ST_CHECK_PORT_LOCK

FLOGI_ST_DPVM_CHECK

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Confirm FCoE I/O packets are incrementing. You can also verify counters on the physical Ethernet interface using this command:

```
N6K# show interface vfc1 counters

vfc1
    1154 fcoe in packets
    201320 fcoe in octets
    160 fcoe out packets
    17332 fcoe out octets
```

5.1.12 Verify Interface Status with fcoe_mgr

The fcoe_mgr info global output command shows detailed status of the fcoe_mgr module running in software as well as the state transitions of all the VFC interfaces (including port-channels). On switches with a large number of VFC interfaces, you will need to parse the output with the “begin” utility within the show command. The fcoe_mgr commands are slightly different between the Cisco Nexus 5000 and 6000, MDS, and Cisco Nexus 7000 platforms, however, the output is very similar. The first command is what is issued on the Cisco Nexus 5000 and 6000, and the second is for Cisco Nexus 7000 and MDS with its subsequent output.

```
N6K# show platform software fcoe_mgr info global

N7K-storage# show system internal fcoe_mgr info global
FCOE-Mgr module
Low Priority Pending queue: len(0), max len(1) [Thu Jun 26 17:48:19 2014]
High Priority Pending queue: len(0), max len(2) [Thu Jun 26 17:48:19 2014]
Log Buffer:
dequeued timer msg: rid (0x6), event_id (0)

FCOE-Mgr database
================
------ Global Config Data ------
================
    FCMAP: 0xefc00
    FCF Priority: 128
    FKA Adv Period: 8
    VE Loopback : disabled
------VLAN Info------
================
Info for VLAN 101
fcoe_enabled 1
vsan_id: 101
------FCF Info------
================
FCF
    FCF MAC Addr: 00:00:00:00:00:00
    FCF Num Pinned by NPM: 0
    FCF Num Pinned by FIP: 0
```

VLAN enabled = 1

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

```
    FCF Priority Offset : 0
          List of Active VSANS
    FCF Vsan: 101
    FCF Switch WWN: 20:65:00:26:98:08:fb:01
    FCF Fabric WWN: 20:65:00:26:98:08:fb:01
vfc101(0x83b6c54), if_index: 0x1e000064, VFC RID vfc101
```

In this global information output there will be a section for every VFC interface on the switch. Here we are starting at vfc101.

```
  FSM current state: FCOE_MGR_VFC_ST_PHY_UP
  PSS Runtime Config:-
      Type: 3
      Bound IF: Po101
      FCF Priority: 128 (Global)
  PSS Runtime Data:-
```

The normal state for fcoe_mgr for vfc101

```
>>>>FSM: <vfc101> has 15 logged transitions<<<<<

1) FSM:<vfc101> Transition at 154738 usecs after Sun Jan 11 08:03:52 2009
    Previous state: [FCOE_MGR_VFC_ST_INIT]
    Triggered event: [FCOE_MGR_VFC_EV_CREATE]
    Next state: [FCOE_MGR_VFC_ST_CREATE_WAIT]

2) FSM:<vfc101> Transition at 156412 usecs after Sun Jan 11 08:03:52 2009
    Previous state: [FCOE_MGR_VFC_ST_CREATE_WAIT]
    Triggered event: [FCOE_MGR_VFC_EV_CREATE_SUCC]
    Next state: [FCOE_MGR_VFC_ST_CREATED]

snip . . .

    Curr state: [FCOE_MGR_VFC_ST_PHY_UP]

PROTOS Info:
vfc101(0x83b9ce4), if_index: 0x1e000064, VEProto RID vfc101, vsan 101
  FSM current state: FCOE_MGR_VE_PROTO_ST_UP
  PSS Runtime Data:-
      Eth IF Index: Po101
      Port Mode: Unknown(0)
      FKA check enabled ? : TRUE
      Recv Multicast solicitation from peer? : TRUE
      Recv Unicast advertisement from peer? : TRUE
      Advertisement period from peer? : 8000 ms
  FIP FKA event count : 90106
  FIP FKA last event time stamp : Thu Jun 26 17:57:23 2014
```

At the end of the transition states, and before the next VFC/port-channel interface are details about the FKA.

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Appendix A: Recommended Steps in Troubleshooting

Table 4 is a recommended approach on the steps to take when you encounter an FCoE problem. Details on the suggested commands and actions to look for may be found on the referenced page number. Each step has sub-steps that may be needed when the parent step is not in the expected state. For example, use step 3 to verify the status of the Ethernet interface, “show interface Ethernet [int]”. If not in the expected up state, try using the sub-steps, for instance, steps 3a, 3b, etc., to diagnose the problem. This table of steps is meant to act a starting point and a guide for possible next steps. They are not mandatory and are open to flexibility when needed.

Table 4. Recommended Steps in Troubleshooting

Step	Protocol/Process	Platform	Show Command	Page
1	FCoE overall status	Nexus 6000, Nexus 7000, MDS	show interface [eth interface] fcoe	26, 40
2	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show interface [vfc interface]	26, 34
2a	FCoE interface configuration	Nexus 6000, Nexus 7000, MDS	show run interface [vfc interface]	34
2b	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show fcoe database	51
2c	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show vsan 101 membership	52
2d	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show port internal event-history errors	52
2e	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show system internal fcfwd mpmap vfcs	50
2e	FCoE interface status	Nexus 6000, Nexus 7000, MDS	show port internal event-history interface [vfc int]	52
2f	fcoe_mgr	Nexus 6000	show platform software fcoe_mgr ?	14, 56
3	Ethernet interface status	Nexus 6000, Nexus 7000, MDS	show interface [eth interface]	27
3a	Ethernet interface status	Nexus 7000, MDS	show system internal ethpm event-history interface ethernet [eth interface] | include "Curr state"	47
3b	Ethernet interface configuration (PC)	Nexus 6000, Nexus 7000, MDS	show run interface port-channel [po interface]	48
3c	Ethernet interface status (PC)	Nexus 6000, Nexus 7000, MDS	show interface port-channel [po interface] switchport	48
3d	Ethernet interface status (PC)	Nexus 7000, MDS	show system internal ethpm event-history interface port-channel [po interface] | inc “Curr state”	49
4	FIP status	Nexus 6000	show platform software fcoe_mgr info interface [vfc]	37, 51, 56
4	FIP status	Nexus 7000, MDS	show system internal fcoe_mgr info interface [vfc]	40
4a	FIP status	Nexus 6000	show platform software fcoe_mgr event-history errors	39
4a	FIP status	Nexus 7000, MDS	show system internal fcoe_mgr event-history errors	39
4b	FIP/FKA status	Nexus 6000	show platform software fcoe_mgr info interface [vfc] | begin PROTOS next 13	39
4b	FIP/FKA status	Nexus 7000, MDS	show system internal fcoe_mgr info interface [vfc] | begin PROTOS next 13	39
4c	fcoe_mgr	Nexus 7000, MDS	show system internal fcoe_mgr ?	14
5	DCB PFC and ETS	Nexus 6000, Nexus 7000, MDS	show policy-map	23, 24, 29
5a	DCB PFC	Nexus 6000, Nexus 7000, MDS	show interface priority-flow-control	41
5a	DCB PFC	Nexus 6000, Nexus 7000	show running-config ipqos	35
5a	DCB FC	MDS	Show running-config eth-qos all	35

© 2015 Cisco and/or its affiliates. All rights reserved. This document is Cisco Public Information.

Step	Protocol/Process	Platform	Show Command	Page
5b	DCB PFC	Nexus 6000, Nexus 7000, MDS	show interface [eth interface] priority-flow-control	35, 41
5c	DCB PFC	Nexus 6000, Nexus 7000, MDS	show policy-map system type network-qos	23
6	DCB ETS	Nexus 6000, Nexus 7000,	show queuing interface [eth interface]	30, 35
6a	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show lldp neighbors	49
6b	DCB ETS	Nexus 6000, Nexus 7000, MDS	show policy-map system type queuing	24
7	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show lldp interface [eth interface]	36
7a	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show run interface eth101/1/1 all | include "lldp|priority-flow"	32
7b	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show lldp dcbx interface [eth interface]	37
7c	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show system internal dcbx info interface [eth interface] | begin "DCBX pkt"	37
7d	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show system internal dcbx info interface Ethernet [eth]	49
7e	DCB DCBX	Nexus 6000, Nexus 7000, MDS	show feature | include lldp	32
8	FCoE FLOGI	Nexus 6000, Nexus 7000, MDS	show flogi database interface [vfc interface]	28, 55
8a	FCoE FLOGI	Nexus 6000, Nexus 7000, MDS	show flogi internal event-history interface [vfc int]	50
9	FCoE configuration	Nexus 6000, Nexus 7000, MDS	show vlan fcoe	25, 44
9a	FCoE configuration	Nexus 6000, Nexus 7000,	show running-config fex	45
9b	FCoE configuration	Nexus 6000, Nexus 7000	show fex detail	46
9c	FCoE configuration	Nexus 6000	show feature | include fcoe	41
9d	FCoE configuration	Nexus 7000, MDS	show feature-set	42
9e	FCoE configuration	Nexus 7000	show vdc membership	44
9f	FCoE configuration	Nexus 7000	show vdc fcoe-vlan-range	44
9g	FCoE configuration	Nexus 7000, MDS	show system internal assoc bitmaps	45
9h	FCoE configuration	Nexus 6000, Nexus 7000, MDS	show license usage	43
9i	FCoE configuration	Nexus 7000	show license fcoe	43
11	FC zone status	Nexus 6000, Nexus 7000, MDS	show zoneset active vsan [vsan]	28

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C07-733622-00 01/15