

**\*\*Note: Information based on FC-BB-5 1.01 June 10, 2008 Version\*\***

## FCoE Tutorial

Fibre Channel over Ethernet (FCoE) is essentially a protocol for tunneling Fibre Channel frames over Ethernet frames. There are two parts to this process: 1. FCoE Initialization Protocol (FIP) and 2. Encapsulation. Through the encapsulation process the FC protocol remains unchanged; entire FC frames can be encapsulated into an Ethernet frame without any changes to the original frame. FCoE is different from previous technologies such as FCIP and iFCP in that there is no TCP involved. This confines FCoE to be solely in the realm of an enterprise data center. While limiting the range of this technology, it greatly simplifies it and focuses FCoE on its main benefit: to converge the data center on one fabric. Converging the data center to one fabric will allow enterprises to save on space and power, maintenance costs, and maintaining the knowledge required for multiple fabrics. FCoE is being defined in the FC-BB-5 standard. Prior to completion, the documents associated with this standard can be found within the document repository associated with the FC-BB-5 group.

FIP provides the following services: discovery of nodes and fabric login/logout. Specifically, FIP defines the encapsulation of FLOGI, FDISC, LOGO and ELP frames along with their reply frames. Of note, the encapsulation of FIP frames is not the same as the FCoE encapsulation covered later, and it is a violation of the standard if a FIP frame is encapsulated in an FCoE frame. Fibre Channel Forwarder (FCF) and FC-BB\_E Node (ENode) are the only devices in the FIP standard that are defined to use the FIP protocol. The behavior of each device is defined as it relates to the FIP protocol can be found in Table 1 below. In general, FCFs use FIP to discover the nodes on the fabric and other FCFs and ENodes use FIP to login with the fabric. Through FIP, the FCFs and ENodes on the fabric will establish virtual ports: VN\_Port, VF\_Port and VE\_Port. The virtual ports coincide with the standard Fibre Channel ports N\_Port, F\_Port and E\_Port.

Table 1:

Frame Type	Valid if Initiated From:	FIP Descriptors and order of inclusion:
Discovery Solicitation	ENode	MAC_Address, Name_Identifier, Max_Receive_Size
Discovery Solicitation	FCF	MAC_Address, FC-MAP, Name_Identifier, Max_Receive_Size
Discovery Advertisement	FCF	Priority, MAC_Address, FC-MAP, Name_Identifier, Fabric_Name
FLOGI Request	ENode	FLOGI, MAC_Address
FLOGI LS_ACC	FCF	FLOGI, MAC_Address
FLOGI LS_RJT	FCF	FLOGI
NPIV FDISC	ENode	FDISC_NPIV, MAC_Address
NPIV FDISC LS_ACC	FCF	FDISC_NPIV, MAC_Address
NPIV FDISC LS_RJT	FCF	FDISC_NPIV
Fabric LOGO	ENode or FCF	LOGO, MAC_Address
Fabric LOGO LS_ACC	ENode or FCF	LOGO, MAC_Address

Fabric LOGO LS_RJT	ENode or FCF	LOGO
ELP Request	FCF	ELP, MAC_Address
ELP SW_ACC	FCF	ELP, MAC_Address
ELP SW_RJT	FCF	ELP

FCoE provides two modes of addressing (FC-MAP): Server Provided MAC Addressing (SPMA) and Fabric Provided MAC Addressing (FPMA). If SPMA is being used the MAC address provided by the ENode in the FLOGI request will be used. If FPMA is being used the MAC address provided by the FCF in the FLOGI reply will be used. Valid FPMA addresses will have the most significant 24 bits based on the FCF and the least significant 24 bits uniquely identifying the ENode.

FIP is used to assign MAC addresses and de-assigning MAC addresses. FCoE also defines several well known MAC addresses used in FIP: all ENodes (ALL\_ENODE\_MACS) and all FCFs (ALL\_FCF\_MACS).

When an ENode comes onto the fabric it uses the following process: 1. FIP Discovery, 2. FLOGI, 3. FDISC, 4. Logout/De-instantiate. During FIP discovery, an ENode sends a Discovery Solicitation (DS) frame to all FCFs in the fabric. When an FCF receives a DS it responds with a Discovery Advertisement (DA) to the ENode that sent the DS if it is configured to allow that ENode on the fabric.

When an FCF is added to the fabric, it sends a DS to all FCFs in the fabric. When an FCF receives a DS from an FCF it sends a DA if it is configured to allow an Inter Switch Link (ISL) with the originating FCF. Additionally, FCFs periodically transmit DAs to all ENodes and FCFs as a heartbeat.

Discovery Solicitations may be sent to a single FCF (unicast) or all FCFs (multicast). The MAC address field shall be set to the the senders MAC address so that all response DA can be sent back. The max receive size field shall be set to the largest 802.3 frame size than can be received. An ENode should discard all DS it receives. An FCF receiving a valid DS shall verify that the node name is not the FCFs and that the FC-MAP value is zero or the same as the FCF supports, if either is not true the FCF will discard the DS.

Discovery Advertisement shall only originate from an FCF. The DA shall be addressed to the MAC that coincides with the MAC reported in the DS. An FCF that receives a valid DS must send out a valid DA before D\_A\_TOV (default 5sec). An FCF might receive multiple DS from the same BB\_E device with different MACs, in this case the FCF must send a valid DA to each MAC on the device requesting a DA. An FCF shall send out periodic DAs to the all ENode and all FCF MAC addresses every D\_A\_TOV \* 10sec +- 50% (the +-50% is so that the FCF can avoid high traffic periods). The priority of DAs shall be that of the priority of FCFs unless not configured it will use DEFAULT\_FIP\_PRIORITY (128). For solicited DAs, the FCF shall fill the pad field until the DA is the same size as the Max\_Receive\_Size in the DS. For unsolicited DAs, the pad field will be empty.

FCoE encapsulation defines how all frames, other than the FIP frames, will be transmitted on the Ethernet network. The frame format can be seen below in Table 2. The DA fields contain the Destination MAC address. The SA fields contain the Source MAC address. The Type field shall be set to FCoE\_TYPE. The version field is 0000b. The FC-BB-5 standard has defined the values for the following SOF and EOF fields: SOFf, SOFi2, SOFn2, SOFi3, SOFn3, EOFn, EOFt, EOFrt, EOFni, EOFrti, EOFa. The FC Frame contains the headers, data, and CRC of an FC-FS-3 compliant frame.

Table 2: FC-BB\_E Frame Format

0	DA[0]	DA[1]	DA[2]	DA[3]
1	DA[4]	DA[5]	SA[0]	SA[1]
2	SA[2]	SA[3]	SA[4]	SA[5]
3	Type		Version	Reserved
4	Reserved			
5	Reserved			
6	Reserved			SOF
7	FC Frame (n words)			
N+7	EOF	Reserved		
N+8	FCS			

Table 3: Important Values

ALL_FCOE_MACS	01:10:18:01:00:00
ALL_ENODE_MACS	01:10:18:01:00:01
ALL_FCF_MACS	01:10:18:01:00:02
D_A_TOV	5sec
FIP TYPE field	8914h
FCoE_TYPE field	8906h
DEFAULT_FIP_PRIORITY	128

To test an FCoE system all that is required is to connect the initiator to a target through an FCoE forwarder. Once connected, scan the drives on the initiator using whatever means is provided by the operating system of your choice. Because FCoE does not affect the FC stack and above, any standard SCSI tools can be used to drive traffic. Typically, the UNH-IOL uses Iometer to generate SCSI reads and writes or we just mount the drives and read or write files to the network drive.

Because FCoE works over standard network interface controllers (NIC) along with specialized FCoE hardware, you can do more extensive testing utilizing a simple off the shelf NIC. In programming terms, you would need access to a “raw socket”, this is a connection below the standard TCP socket that lets you build an Ethernet frame from the ground up. Through a raw socket you can implement conformance tests with either a partial or complete implementation of an FCoE stack. This will allow a tester to walk any device through the initialization process and inject errors along the way.