

7 FC-BB_E Structure and Concepts

7.1 Applicability

Clause 4 discussed the FC-BB_E reference model. This clause discusses the FC-BB_E functional models.

7.2 FC-BB_E overview

This clause discusses further aspects of FC-BB_E operation, including initialization, flow control, and procedures for the mapping of Fibre Channel frames over Ethernet.

Figure 24 illustrates the protocol levels and layers involved in FC-BB_E processes and devices.

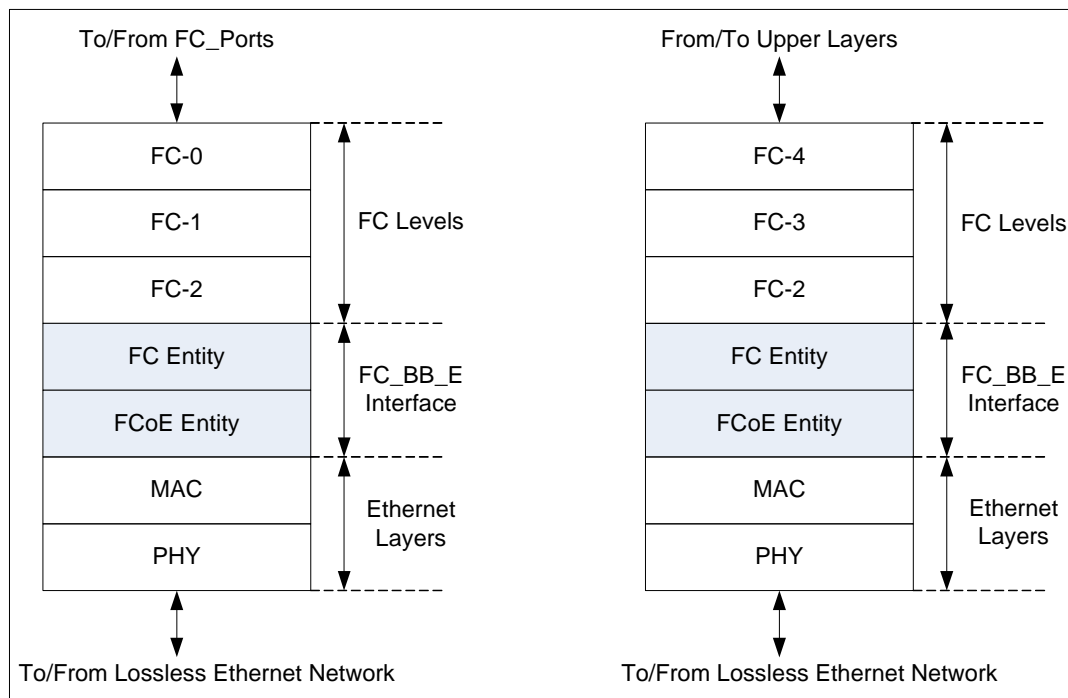


Figure 24 – FC-BB_E protocol levels and layers

7.3 FC_BB_E VN_Port/ENode functional model

The FC-BB_E VN_Port/ENode functional model is specified in figure 25.

Editor's Note: Add associated text and keys for figure 24 and figure 25, especially specifying required components and indicating optional functionality.

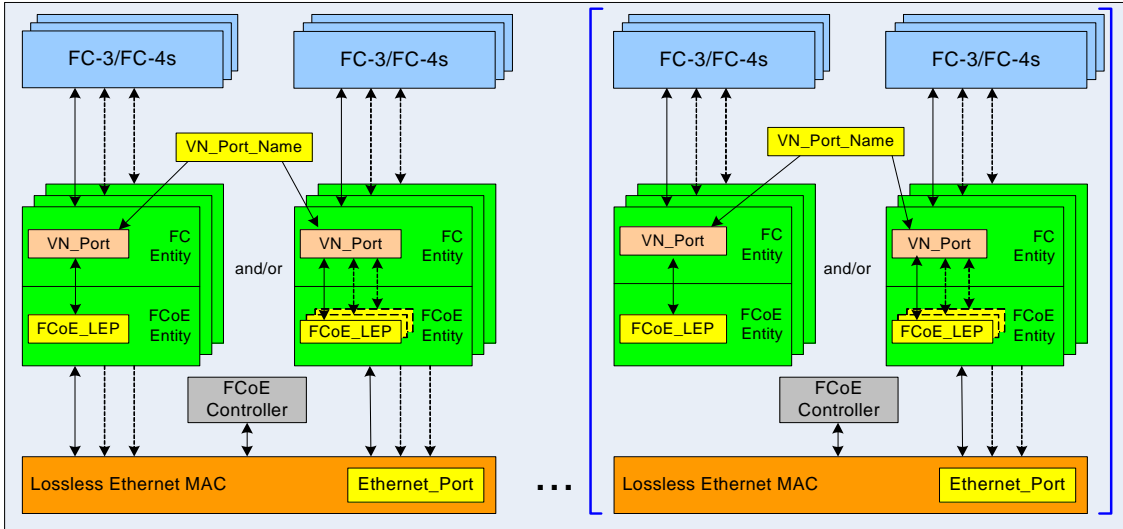


Figure 25 – FC-BB_E VN_Port/ENode functional model

7.4 FC_BB_E VE_Port/VF_Port functional model

The FC-BB_E VE_Port/VF_Port functional model is specified in figure 26.

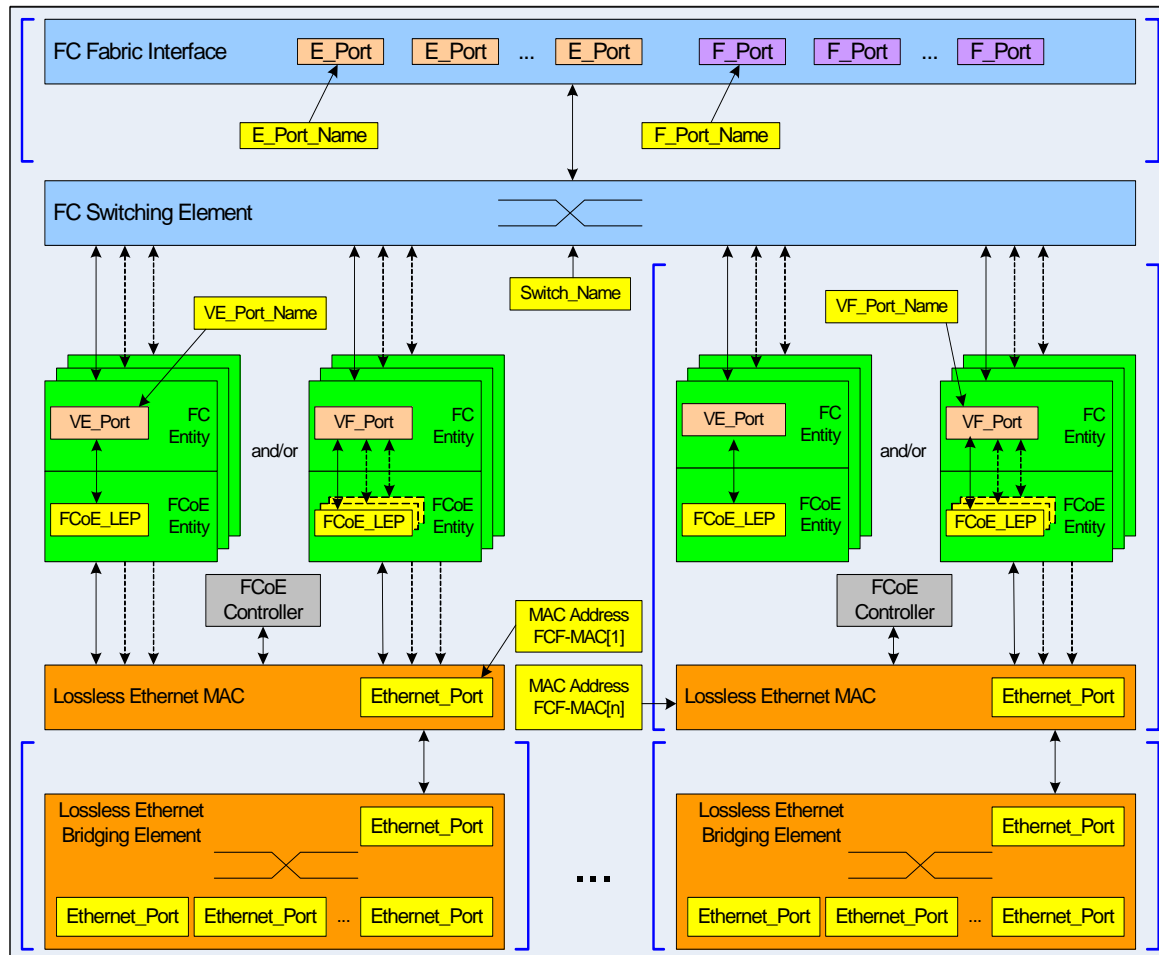


Figure 26 – FC-BB_E VE_Port/VF_Port functional model

Editor’s Note: Is each Lossless Ethernet MAC actually required to have a separate Lossless Ethernet Bridging Element as figure 25 suggests?

Editor’s Note: The Lossless Ethernet Bridging Element should not be part of the VE_Port/VF_Port functional model.

7.5 FC-BB_E device initialization

7.5.1 FCoE initialization protocol (FIP) overview

The FCoE initialization protocol (FIP) is used to perform the function of FC-BB_E device discovery, initialization, and exchange of FC-BB_E device information and parameters. To perform these functions, encapsulated FIP operations (see 7.5.4.1) are specified for discovery and specific Link Service Requests and Replies.

The FIP frame format (see 7.5.4) is different than the FC-BB_E frame format (see 7.6) to enable the detection of FC-BB_E discovery and initialization traffic from normal FC-BB_E traffic.

NOTE 12 – FIP allows for the snooping of frames (e.g., FLOGI with Request S_ID=0, FDISC with Request S_ID=0, and LOGO with D_ID or S_ID = F_Port Controller) by intermediate switches to allow for the generation and modification of Access Control Lists.

7.5.2 FIP discovery overview

When an ENode is added to a FC-BB_E fabric it discovers FCFs that it may perform login with by transmitting a Discovery Solicitation (see x.x.x) to the ALL_FCF_MACS multicast group address. In response to a Discovery Solicitation from an ENode, an FCF transmits a unicast Discovery Advertisement (see x.x.x) to the ENode if it is configured to allow a Fabric login from that ENode.

Editor's Note: Should we allow a unicast Discovery Solicitation from an ENode and/or FCF?

When an FCF is added to a FC-BB_E fabric it discovers other FCFs by transmitting a Discovery Solicitation to the ALL_FCF_MACS multicast group address. In response to a Discovery Solicitation from an FCF, an FCF transmits a solicited unicast Discovery Advertisement to the FCF if it is configured to allow an ISL with that FCF.

In addition, an FCF periodically transmits an unsolicited Discovery Advertisement to inform ENodes and other FCFs that the FCF exists on the FC-BB_E fabric.

Editor's Note: Are ENodes allowed to periodically transmit Discovery Solicitations?

Editor's Note: Insert example discovery figure here.

7.5.3 Link Service Request and Reply overview

Fabric login (i.e., FLOGI, NPIV FDISC), Fabric logout, and ELP shall be performed using the FIP frame format (see table 22) and the associated FIP descriptor type (see table 26) (i.e., Fabric login, Fabric logout, and ELP shall not be performed using the normal FC-BB_E frame format).

The FIP Operation Code and FIP Subcode field values and operations are specified in table 24.

Table 24 – FIP Operation Code and FIP Subcode field values and operations

FIP Operation Code	FIP Subcode	FIP Operation	Reference
00000001h	01h	Discovery Solicitation	7.5.4.2
	02h	Discovery Advertisement	7.5.4.3
00000002h	01h	Link Service Request	7.5.4.4
	02h	Link Service Reply	7.5.4.4
All others	All others	Reserved	

The FIP Descriptor List Length field shall be set to the length in words of all FIP descriptor(s) that follow.

The Fabric Provided (FP) bit and Server Provided (SP) bit setting is dependent on the FIP operation and the bits shall be set as specified in table 25.

Table 25 – FP bit and SP bit setting

Bit	FIP operation	Setting
FP	Discovery Solicitation Discovery Advertisement	Set to 1 if originating device supports FPMA. Set to 0 if originating device does not support FPMA.
	FLOGI Request ^a FDISC_NPIV Request ^a	Set to 1 if FPMA is request. Set to 0 if FPMA is not supported.
	FLOGI LS_ACC FDISC_NPIV LS_ACC	Set to 1 if FPMA granted. Set to 0 if SPMA granted.
SP	Discovery Solicitation Discovery Advertisement	Set to 1 if originating device supports SPMA. Set to 0 if originating device does not support SPMA.
	FLOGI Request ^a FDISC_NPIV Request ^a	Set to 1 if SPMA is request. Set to 0 if SPMA is not supported.
	FLOGI LS_ACC FDISC_NPIV LS_ACC	Set to 1 if SPMA granted. Set to 0 if FPMA granted.
a Both the FP bit and SP bit may be set to 1 in a FLOGI Request or FDISC_NPIV Request, but at least one of the bits shall be set to 1.		

The Solicited (S) bit shall be set to 1 in a Discovery Advertisement that is transmitted in response to a Discovery Solicitation. The S bit shall be set to 0 in a Discovery Advertisement that is not transmitted in response to a Discovery Solicitation. The S bit shall be set to 0 in all other FIP operations.

The FCF (F) bit shall be set to 1 in a Discovery Solicitation or Discovery Advertisement if the originating device is an FCF. The F bit shall be set to 0 in a Discovery Solicitation or Discovery Advertisement if the originating device is not an FCF. The F bit shall be set to 0 in all other FIP operations.

Editor’s Note: Should the S and F bit be specified as reserved or ignored for all other FIP operations?

The FIP Descriptor(s) field contains one or more FIP descriptors (see 7.5.4.1.2).

Two multicast group addresses (see table 41) are used for Discovery Solicitations and Discovery Advertisements. The ALL_ENODE_MACS multicast group address is used to address all FC-BB_E ENode devices and the ALL_FCF_MACS multicast group address is used to address all FC-BB_E FCF devices.

7.5.4.1.2 FIP descriptors

7.5.4.1.2.1 FIP descriptor overview

The FIP descriptors are specified using a TLV format. The length field value shall be specified as the number of words in the FIP descriptor. The FIP descriptor types are specified in table 26.

Table 26 – FIP descriptor types

Value	FIP Descriptor	Reference
1	Priority	7.5.4.1.2.2
2	MAC address	7.5.4.1.2.3
3	FC-MAP	7.5.4.1.2.4
4	Name_Identifier	7.5.4.1.2.5
5	Fabric_Name	7.5.4.1.2.6
6	Max Receive Size	7.5.4.1.2.7
7	FLOGI	7.5.4.1.2.8
8	FDISC_NPIV	7.5.4.1.2.9
9	LOGO	7.5.4.1.2.10
10	ELP	7.5.4.1.2.11
All others	Reserved	
a The FC CRC, SOF, and EOF shall not be included in the FIP descriptor.		

Editor’s Note: Should descriptor type 4 be Name_Identifier or Node_Name/Switch_Name?

Editor’s Note: The ELP descriptor type is questionable for FIP.

7.5.4.1.2.2 FIP Priority descriptor

The FIP Priority descriptor is used in Discovery Advertisements originated by an FCF to indicate a priority to an ENode when multiple Discovery Advertisements are received. The default value for the Priority field is DEFAULT_FIP_PRIORITY (see table 41). The highest priority value is 0 and the lowest priority value is 255 (i.e., lower numerical values indicate higher priorities).

7.5.4.1.2.6 FIP Fabric_Name descriptor

The FIP Fabric_Name descriptor is used in Discovery Advertisements originated by an FCF.

The FIP Fabric_Name descriptor format is specified in table 31.

Table 31 – FIP Fabric_Name descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (5)					Length (3)					Reserved																					
1	MSB																															
2	Fabric_Name																LSB															

7.5.4.1.2.7 FIP Max Receive Size descriptor

The FIP Max Receive Size descriptor is used in Discovery Solicitations originated by an ENode or FCF.

The FIP Max Receive Size descriptor format is specified in table 32.

Table 32 – FIP Max Receive Size descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (6)					Length (1)					Max_Receive_Size																					

7.5.4.1.2.8 FIP FLOGI descriptor

The FIP FLOGI descriptor is used in Fabric login requests and replies.

The FIP FLOGI descriptor format is specified in table 33

Table 33 – FIP FLOGI descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (7)					Length					Reserved																					
1	MSB																															
n	FLOGI Request or FLOGI LS_ACC/LS_RJT																LSB															

The Length field value shall be set to 36 for a FLOGI Request and FLOGI LS_ACC, or to 9 for a FLOGI LS_RJT.

7.5.4.1.2.9 FIP FDISC_NPIV descriptor

The FIP FDISC_NPIV descriptor is used in NPIV based Fabric login requests and replies.

The FIP FDISC_NPIV descriptor format is specified in table 33

Table 34 – FIP FDISC_NPIV descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (8)								Length								Reserved															
1	MSB								NPIV FDISC Request or NPIV FDISC																							
n									LS_ACC/LS_RJT																LSB							

The Length field value shall be set to 36 for an FDISC Request and FDISC LS_ACC, or to 9 for an FDISC LS_RJT.

7.5.4.1.2.10 FIP LOGO descriptor

The FIP LOGO descriptor is used NPIV based Fabric login requests and replies.

The FIP LOGO descriptor format is specified in table 33

Table 35 – FIP LOGO descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (9)								Length								Reserved															
1	MSB								LOGO Request or LOGO LS_ACC/LS_RJT																							
n									LOGO Request or LOGO LS_ACC/LS_RJT																LSB							

The Length field value shall be set to 11 for a LOGO Request, 10 for a LOGO LS_ACC, or to 9 for a LOGO LS_RJT.

7.5.4.1.2.11 FIP ELP descriptor

The FIP ELP descriptor is used in Exchange Link Parameter requests and replies.

The FIP ELP descriptor format is specified in table 33

Table 36 – FIP ELP descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type (10)								Length								Reserved															
1	MSB								ELP Request or ELP SW_ACC/SW_RJT																							
n									ELP Request or ELP SW_ACC/SW_RJT																LSB							

The Length field value shall be set to 33 for an ELP Request and ELP SW_ACC, or to 9 for an ELP SW_RJT.

7.5.4.2 FIP Discovery Solicitation

7.5.4.2.1 ENode FIP Discovery Solicitation

The Discovery Solicitation contains the MAC address to use for subsequent FCF Discovery Advertisements, the addressing mode capability of the ENode (i.e., FPMA and/or SPMA), the Node_Name of the ENode and the maximum receive size of the ENode.

The ENode Discovery Solicitation format is specified in table..

Table 37 – ENode FIP Discovery Solicitation format

Word	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0
0	FIP Operation Code (0001h)								Reserved								FIP Subcode (01h)							
1	FIP Descriptor List Length (6)								F	S	Reserved										S	F		
2	Type (2)				Length (2)				MAC address[0]				MAC address[1]											
3	MAC address[2]				MAC address[3]				MAC address[4]				MAC address[5]											
4	Type (4)				Length (3)				Reserved															
5	MSB																							
6	Name_Identifier												LSB											
7	Type (6)				Length (1)				Max_Receive_Size															
n	Padding																							

The FP bit shall be set to 1 if the ENode supports FPMA. The FP bit shall be set to 0 if the ENode does not support FPMA.

The SP bit shall be set to 1 if the ENode supports SPMA. The FP bit shall be set to 0 if the ENode does not support SPMA.

The S bit shall be set to 0.

The F bit shall be set to 0.

The MAC address field shall be set to the MAC address that the ENode uses for subsequent Discovery Advertisements.

The Name_Identifier field shall be set to the Node_Name of the ENode.

The Max_Receive_Size field shall be set to the maximum 802.3 frame size that the ENode is able to receive. The Max_Receive_Size value shall be specified as the number of octets starting from the first octet of the destination MAC address to the last octet of the FCS, inclusive.

7.5.4.2.2 FCF FIP Discovery Solicitation

7.5.4.3 FIP Discovery Advertisement

7.5.4.4 FIP Link Service Request and Replies

7.5.4.4.1 Fabric login

When an ENode transmits a FIP FLOGI Request it shall indicate the addressing mode it supports (i.e., FPMA, SPMA, or both). If an ENode only supports SPMA it shall specify a MAC address.

TBD

The FIP FLOGI Reply shall provide the address identifier for the VN_Port and:

- a) for SPMA, the ENode MAC address used during FIP operation shall be returned in the FIP FLOGI Reply frame and shall be used as the VN_Port MAC address for all subsequent FC-BB_E frames; or
- b) for FPMA, the MAC address specified in the FIP FLOGI Reply frame shall be used as the VN_Port MAC address for all subsequent FC-BB_E frames.

Each VN_Port MAC address used during discovery shall remain valid and may be used for subsequent FIP operations.

7.5.4.4.2 Fabric logout

TBD

7.5.4.4.3 Exchange Link Parameters

TBD

7.6 FC-BB_E frame format

7.6.1 FC-BB_E frame format overview

The format of an FC-BB_E frame is specified in table 38. The use of an 802.1Q tag header is optional and is not shown in table 38.

Table 38 – FC-BB_E frame format

Word	Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
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		Ethertype								Version		Reserved																			
4		Reserved																													
5		Reserved																													
6		Reserved																				SOF									
7		FC Frame (n words)																													
n+7		EOF				Reserved																									
n+8		FCS																													

The DA[0-5] fields shall be set to the destination MAC address.

The SA[0-5] fields shall be set to the source MAC address.

The Ethertype field shall be set to 8906h (i.e., FCoE - Fibre Channel over Ethernet).

The Version field shall be set to 0000b.

The SOF field specifies the SOF Ordered Set that is associated with encapsulated frame. The value of the SOF field shall be compliant with FC-FS-3 and the SOF field shall be set as specified in table 39.

Table 39 – FC-BB_E SOF field

Value	Description
28h	SOFf
2Dh	SOFi2
35h	SOFn2
2Eh	SOFi3
36h	SOFn3

The FC Frame field contains:

- a) FC Extended_Header(s) (see FC-FS-3), if any;
- b) the FC Frame_Header (see FC-FS-3);
- c) the FC Data_Field (see FC-FS-3); and

d) the FC CRC (see FC-FS-3).

The EOF field specifies the EOF Ordered Set that is associated with encapsulated frame. The value of the EOF field shall be compliant with FC-FS-3 and the EOF field shall be set as specified in table 40.

Table 40 – FC-BB_E EOF field

Value	Description
41h	EOFn
42h	EOFt
44h	EOFrt
49h	EOFni
4Fh	EOFrti
50h	EOFa

7.7 Timers and constants

FC-BB_E timers and constants are specified in table 41.

Table 41 – FC-BB_E timers and constants

Timer/Constant	Value	Description
FIP_TYPE	8914h	The value specified in the 802.3 Type field for a FIP frame.
ALL_FCF_MACS	TBD	The multicast group address for all FCFs.
ALL_ENODE_MACS	TBD	The multicast group address for all ENodes.
DEFAULT_FIP_PRIORITY	128	The default value specified in the FIP Priority descriptor.
D_A_TOV	5	The default value that specifies the number of seconds between discovery advertisements transmitted by FCFs. This value may be administratively configured to any value between 1 and 60, inclusive.

Editor's Note: Need to specify the ALL_FCF_MACS and ALL_ENODE_MACS multicast group address values.