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.

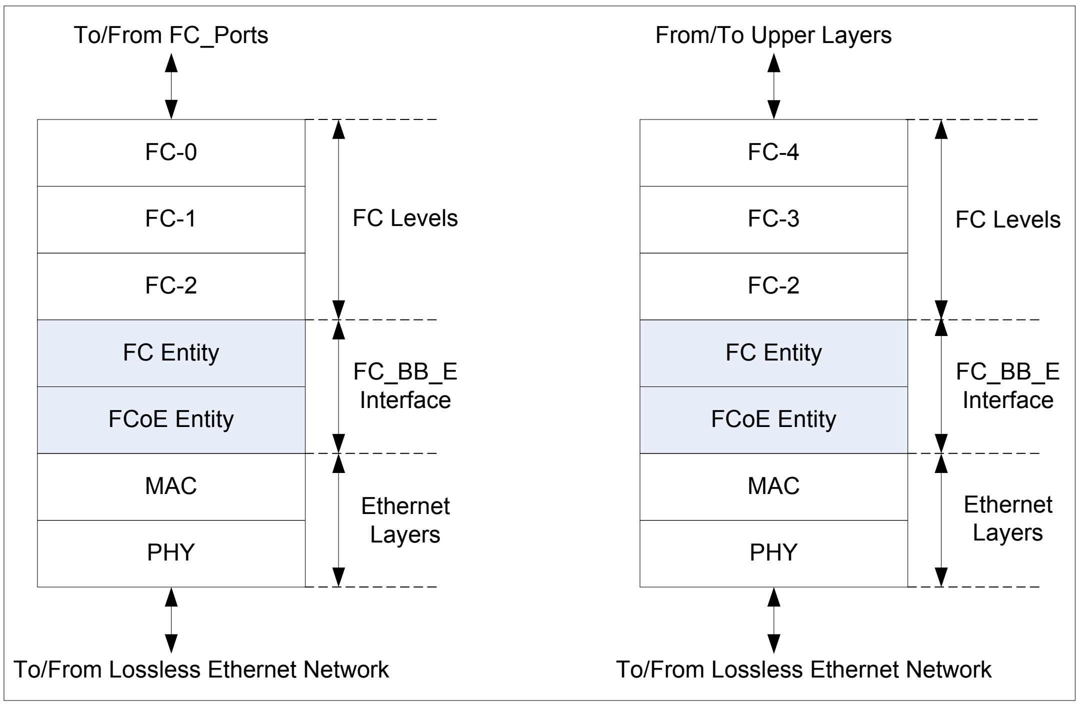

Figure 24 – FC-BB_E protocol levels and layers

FC-BB_E defines a direct mapping of Fibre Channel over Ethernet (FCoE). Although a generic Ethernet network may lose frames due to congestion, a proper implementation of appropriate Ethernet extensions (i.e, the Pause mechanism defined in 802.3-2005, Part 3) allows a full duplex Ethernet link to provide a lossless behavior similar to the one provided by the buffer to buffer credit mechanism in native Fibre Channel. The protocol mapping defined by FC-BB_E is referred to as Fibre Channel over Ethernet (FCoE) and requires the underlying Ethernet layer to be full duplex and lossless (i.e., to be composed only of full duplex links and to provide a lossless behavior when carrying Fibre Channel frames).

In native Fibre Channel, Fibre Channel Nodes (see FC-FS-3) and Switches (see FC-SW-5) communicate through FC_Ports. Fibre Channel links connect N_Ports to F_Ports and E_Ports to E_Ports.

In Fibre Channel over Ethernet, FCoE Nodes (ENodes) and FCoE Forwarders (FCFs) communicate through Ethernet ports supporting Lossless Ethernet MACs. FCoE Virtual Links replace the physical

Fibre Channel links by encapsulating FC frames in Ethernet frames. An FCoE Virtual Link is identified by the pair of MAC addresses of the two link end-points. FCoE supports VN_Port to VF_Port Virtual Links and VE_Port to VE_Port Virtual Links.

Figure 25 shows an FCoE VN_Port to VF_Port network configuration.

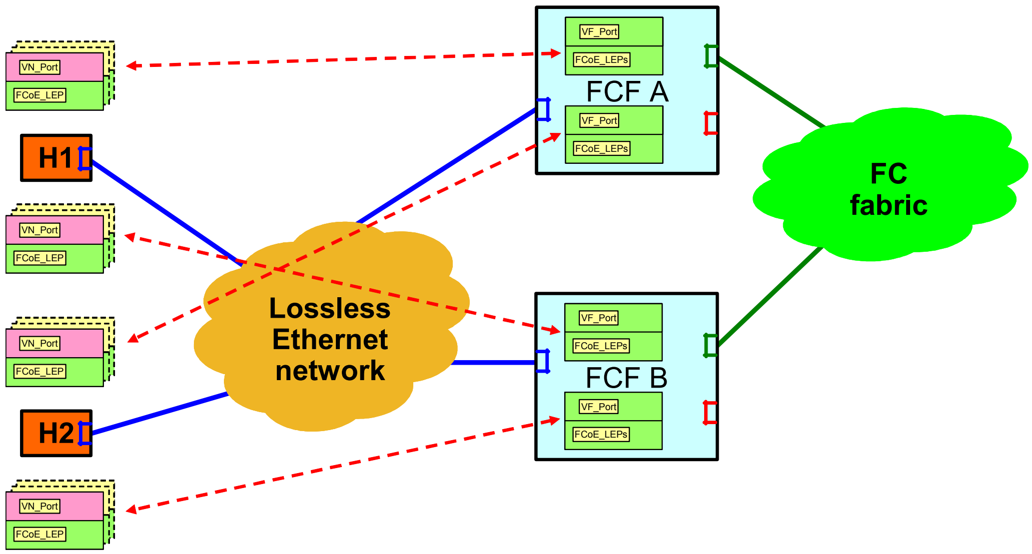

Figure 25 – FCoE VN_Port to VF_Port Network Configuration

Each of the two ENodes H1 and H2 depicted in figure 25 has a single physical Ethernet connection to the Lossless Ethernet network, as well as each of the two FCFs, FCF A and B. Each ENode may instantiate multiple VN_Ports, connected to VF_Ports instantiated by the FCFs through FCoE Virtual Links. The dotted lines in figure 25 depicts possible VN_Port to VF_Port Virtual Links. In this case, a multi-access Lossless Ethernet network is reduced by FCoE to a set of point-to-point VN_Port to VF_Port Virtual Links where the N_Port to F_Port Fibre Channel protocols are able to operate unchanged.

Figure 26 shows an FCoE VE_Port to VE_Port network configuration.

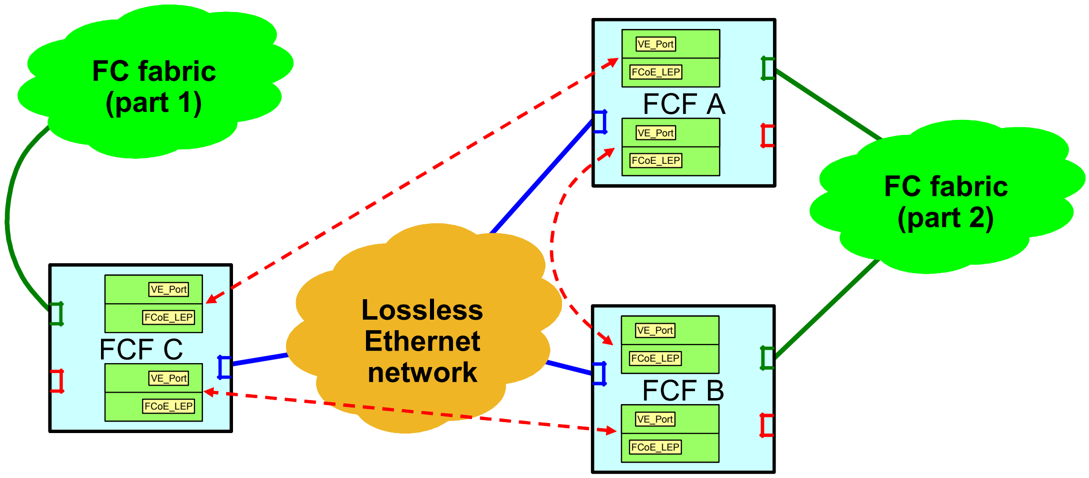

Figure 26 – FCoE VE_Port to VE_Port Network Configuration

Each of the three FCFs A, B, and C depicted in figure 26 has a single physical Ethernet connection to the Lossless Ethernet network. Each FCF may instantiate multiple VE_Ports, connected to other VE_Ports through FCoE Virtual Links. The dotted lines in figure 26 depicts possible VE_Port to VE_Port Virtual Links. In this case a multi-access Lossless Ethernet network is reduced by FCoE to a set of point-to-point VE_Port to VE_Port Virtual Links where the E_Port to E_Port Fibre Channel protocols are able to operate unchanged.

7.3 FC_BB_E VN_Port/ENode functional model

Figure 27 shows the functional model of an ENode, where the bracketed functional components are optional. An ENode is functionally composed of at least one Lossless Ethernet MAC (i.e., the ENode's MAC), coupled with an FCoE Controller function.

The FC-BB_E VN_Port/ENode functional model is specified in figure 27

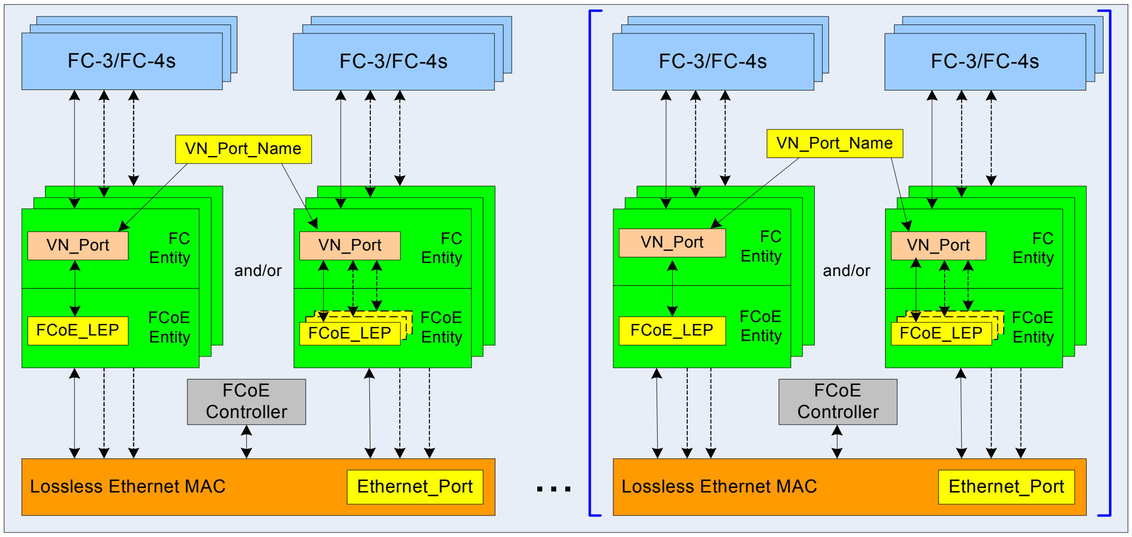

Figure 27 – FC-BB_E VN_Port/ENode functional model

An ENode's MAC shall support the instantiation of VN_Ports.

The FCoE Controller is the functional entity executing the FCoE Initialization Protocol (FIP) and instantiating VN_Ports and FCoE_LEPs as needed.

For an ENode's MAC, the FCoE Controller:

1) initiates the FIP Discovery protocol in order to discover VF_Port capable FCF-MACs connected to the same Lossless Ethernet network;
2) initiates FIP FLOGI Exchanges and instantiates a VN_Port/FCoE_LEP pair for any successful completion of a FIP FLOGI Exchange with a VF_Port capable FCF-MAC;
3) initiates FIP FDISC Exchanges and instantiates a VN_Port/FCoE_LEP pair for any successful completion of a FIP FDISC Exchange with a VF_Port capable FCF-MAC; and
4) de-instantiates a VN_Port/FCoE_LEP pair when that VN_Port is logged out.

VN_Ports instantiated by an ENode's MAC on successful completion of FIP NPIV FDISC Exchanges are all associated to the same VF_Port, instantiated by the VF_Port capable FCF-MAC on successful completion of a FIP FLOGI Exchange.

The FCoE_LEP is the functional entity performing the encapsulation of FC frames into FCoE frames in transmission and the decapsulation of FCoE frames into FC frames in reception. An FCoE_LEP operates according to the two parameters defining the Virtual Link: the MAC address of the local link end-point and the MAC address of the remote link end-point. When encapsulating FC frames into FCoE frames, the MAC address of the local link end-point shall be used as source address and the MAC address of the remote link end-point shall be used as destination address of the generated FCoE frame. When decapsulating FC frames from FCoE frames, the FCoE_LEP shall verify that the destination address of the received FCoE frame is equal to the MAC address of the local link end-point and should verify that the source address of the received FCoE frame is equal to the MAC address of the remote link end-point.

For an FCoE_LEP of an ENode's MAC, the MAC address of the local link end-point is the MAC address associated with its VN_Port and the remote link end-point address is the FCF-MAC address associtated with the remote VF_Port. The VN_Port may use an FPMA or an SPMA as MAC address.

A VN_Port is the data forwarding component of an FC Entity that emulates an N_Port and is dynamically instantiated on successful completion of a FIP FLOGI Exchange or a FIP NPIV FDISC Exchange. A VN_Port receives FC frames from the upper FC levels and sends them to its FCoE_LEP for encapsulation and transmission over the Lossless Ethernet network. In a similar way, a VN_Port sends FC frames received from its FCoE_LEP to the upper FC levels. A VN_Port may support one or more FC-4s. A VN_Port is uniquely identified by a VN_Port_Name Name_Identifier and is addressed by the address identifier the Fabric assigned to it. The VN_Port behavior shall be as specified in FC-LS and FC-FS-2, with the following exceptions:

a) a VN_Port does not perform buffer-to-buffer flow control; and
b) a VN_Port is instantiated on successful completion of a FIP FLOGI Exchange or a FIP NPIV FDISC Exchange, ignoring the buffer-to-buffer flow control parameters, rather than on completion of a native FLOGI or NPIV FDISC Exchange.

7.4 FC_BB_E VE_Port/VF_Port functional model

Figure 28 shows the functional model of an FCF, where the bracketed functional components are optional. An FCF is functionally composed by a Fibre Channel Switching Element (see FC-SW-5) with at least one Lossless Ethernet MAC (FCF-MAC). Each FCF-MAC shall be coupled with an FCoE Controller function. Each FCF-MAC may be coupled with a Losseless Ethernet bridging element. The Fibre Channel Switching Element may be coupled with a Fibre Channel Fabric interface, providing native E_Port and F_Port connectivity. An FCF forwards FCoE frames addressed to one of its FCF-MACs based on the D_ID of the encapsulated FC frames.

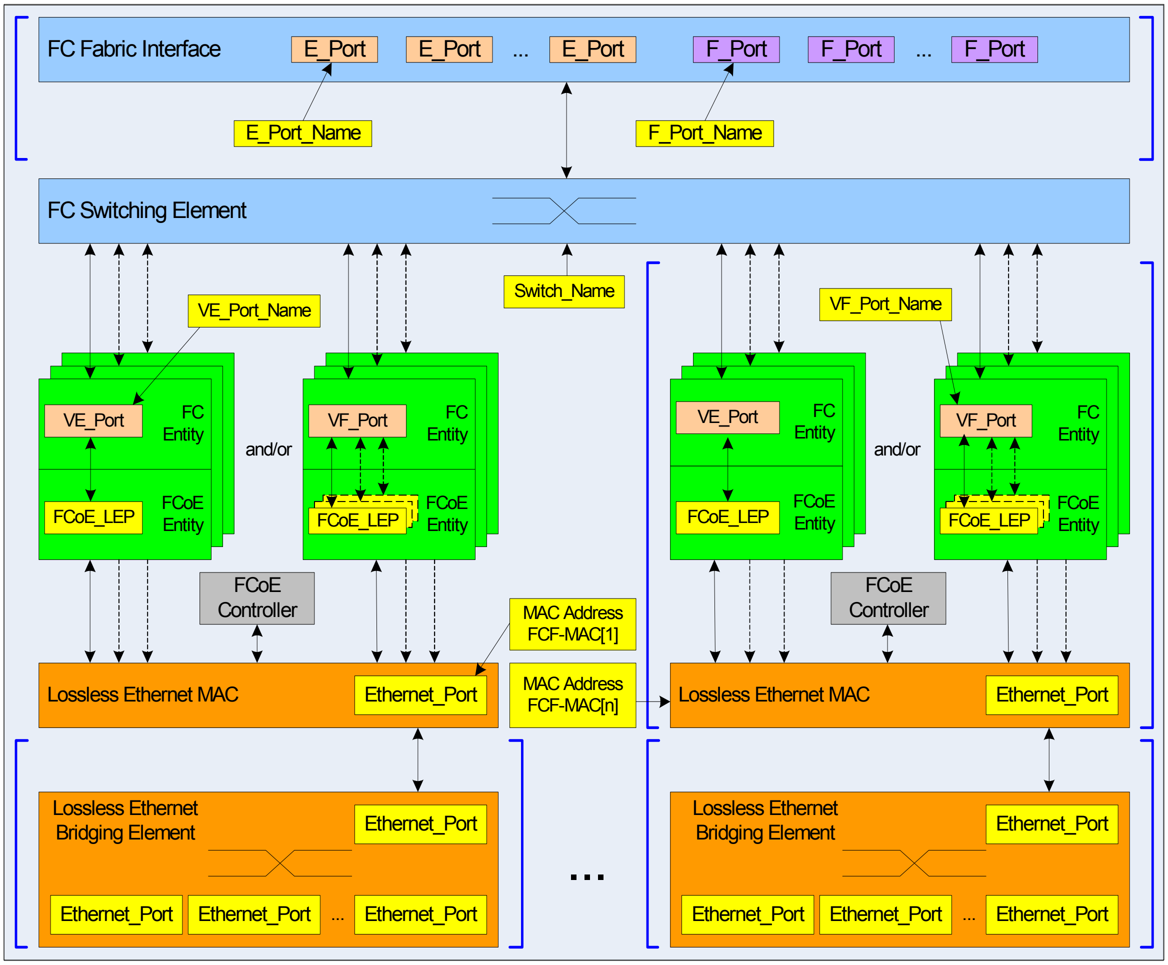

Figure 28 – FC-BB_E VE_Port/VF_Port functional model

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.

An FCF-MAC shall support the instantiation of VE_Ports and/or VF_Ports. An FCF-MAC supporting the instantiation of VE_Ports is referred to as a VE_Port capable FCF-MAC. An FCF-MAC supporting the instantiation of VF_Ports is referred to as a VF_Port capable FCF-MAC. Support for both VE_Ports and VF_Ports on the same FCF-MAC is allowed but not recommended.

The FCoE Controller is the functional entity executing the FCoE Initialization Protocol (FIP) and instantiating VE_Ports, VF_Ports, and FCoE_LEPs as needed.

For a VE_Port capable FCF-MAC, the FCoE Controller:

1) discovers other VE_Port capable FCF-MACs connected to the same Lossless Ethernet network through the FIP Discovery protocol; and
2) instantiates a VE_Port/FCoE_LEP pair for any successful completion of a FIP ELP Exchange with a remote FCF-MAC.

For a VF_Port capable FCF-MAC, the FCoE Controller:

1) participates to the FIP Discovery protocol initiated by ENode's MACs;
2) instantiates a VF_Port and an FCoE_LEP for any successful completion of a FIP FLOGI Exchange initiated by an ENode's MAC;
3) instantiates an additional FCoE_LEP for any successful completion of a FIP NPIV FDISC Exchange initiated by an already logged in ENode's MAC; and
4) when a VN_Port is logged out, de-instantiates the FCoE_LEP associated to that VN_Port and the corresponding VF_Port if that FCoE_LEP was the only one associated with that VF_Port.

VN_Ports instantiated by an ENode's MAC on successful completion of FIP NPIV FDISC Exchanges are all associated to the same VF_Port, instantiated by the VF_Port capable FCF-MAC on successful completion of a FIP FLOGI Exchange.

The FCoE_LEP is the functional entity performing the encapsulation of FC frames into FCoE frames in transmission and the decapsulation of FCoE frames into FC frames in reception. An FCoE_LEP operates according to the two parameters defining the Virtual Link: the MAC address of the local link end-point and the MAC address of the remote link end-point. When encapsulating FC frames into FCoE frames, the MAC address of the local link end-point shall be used as source address and the MAC address of the remote link end-point shall be used as destination address of the generated FCoE frame. When decapsulating FC frames from FCoE frames, the FCoE_LEP shall verify that the destination address of the received FCoE frame is equal to the MAC address of the local link end-point and should verify that the source address of the received FCoE frame is equal to the MAC address of the remote link end-point.

For a VF_Port capable FCF-MAC, the MAC address of the local link end-point is the FCF-MAC address and the MAC address of the remote link end-point is the MAC address associated with the remote logged-in VN_Port. The remote VN_Port may use an FPMA or an SPMA as MAC address.

For a VE_Port capable FCF-MAC, the MAC address of the local link end-point is the FCF-MAC address and the MAC address of the remote link end-point is the MAC Address of the remote FCF-MAC with which a FIP ELP Exchange has been successfully completed.

A VE_Port is the data forwarding component of an FC Entity that emulates an E_Port and is dynamically instantiated on successful completion of a FIP ELP Exchange. A VE_Port receives FC frames from the FC Switching Element and sends them to its FCoE_LEP for encapsulation and transmission over the Lossless Ethernet network. In a similar way, a VE_Port sends FC frames received from its FCoE_LEP to the FC Switching element. A VE_Port is uniquely identified by a VE_Port_Name Name_Identifier and is addressed by the Fabric Controller address identifier (i.e., FFFFFDh). The VE_Port behavior shall be as specified in FC-SW-5, with the following exceptions:

a) a VE_Port does not perform buffer-to-buffer flow control; and
b) a VE_Port is instantiated on successful completion of a FIP ELP Exchange, ignoring the buffer-to-buffer flow control parameters, rather than on completion of a native ELP Exchange.

A VF_Port is the data forwarding component of an FC Entity that emulates an F_Port and is dynamically instantiated on successful completion of a FIP FLOGI Exchange. A VF_Port receives FC frames from the FC Switching Element and sends them to the proper FCoE_LEP for encapsulation and transmission over the Lossless Ethernet network. In a similar way, a VF_Port sends FC frames received from one of its FCoE_LEPs to the FC Switching element. A VF_Port is uniquely identified by a VF_Port_Name Name_Identifier and is addressed by the F_Port Controller address identifier (i.e., FFFFFEh). The VF_Port behavior shall be as specified in FC-LS-2 and FC-FS-3, with the following exceptions:

a) a VF_Port does not perform buffer-to-buffer flow control; and
b) a VF_Port is instantiated on successful completion of a FIP FLOGI Exchange, ignoring the buffer-to-buffer flow control parameters, rather than on completion of a native FLOGI Exchange.

The Fibre Channel Switching Element is the functional entity performing Fibre Channel switching among E_Ports, F_Ports, VE_Ports, and VF_Ports. A Fibre Channel Switching Element is uniquely identified by a Switch_Name Name_Identifier. The behavior of the Fibre Channel Switching Element shall be as specified in FC-SW-5.

7.5 FCoE Virtual Links

Figure 29 shows how the models defined in 7.3 and 7.4 model VN_Port to VF_Port Virtual Links.

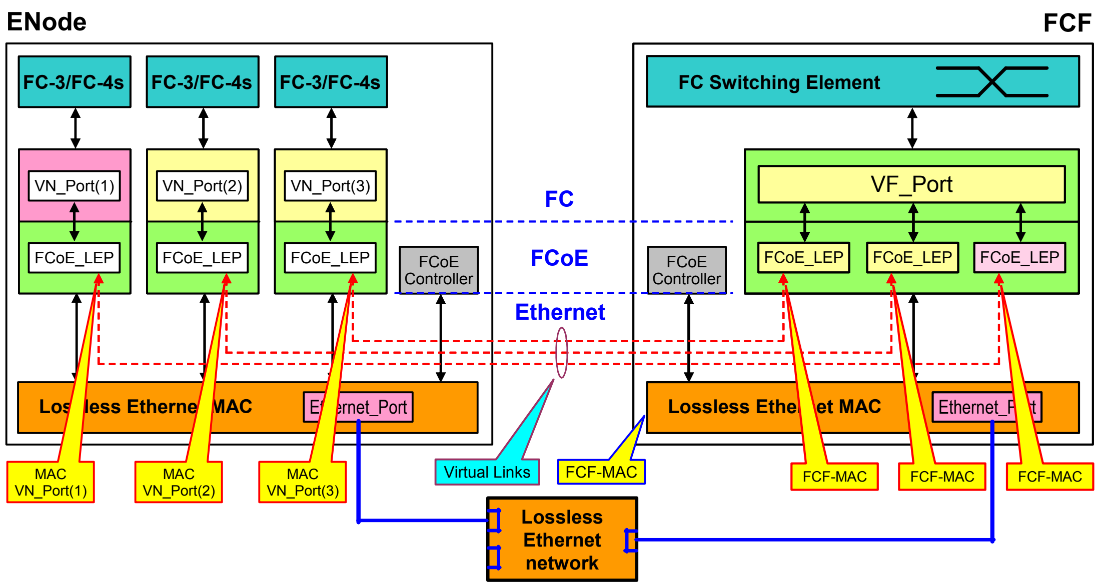

Figure 29 – VN_Port to VF_Port Virtual Links

On successful completion of a FIP FLOGI Exchange, the FCoE Controller of an ENode's MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(1) in figure 29) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates a VF_Port/FCoE_LEP pair.

On successful completion of a FIP FDISC Exchange, the FCoE Controller of an ENode's MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(2) in figure 29) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates an additional FCoE_LEP to the instantiated VF_Port.

On successful completion of an additional FIP FDISC Exchange, the FCoE Controller of an ENode's MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(3) in figure 29) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates an additional FCoE_LEP to the instantiated VF_Port.

Figure 29 shows the Virtual Links end-points, that are the MAC addresses used by the VN_Ports (i.e., MAC VN_Port(1), MAC VN_Port(2), and MAC VN_Port(3)), and the FCF-MAC address.

Figure 30 shows how the model defined in 7.4 model VE_Port to VE_Port Virtual Links.

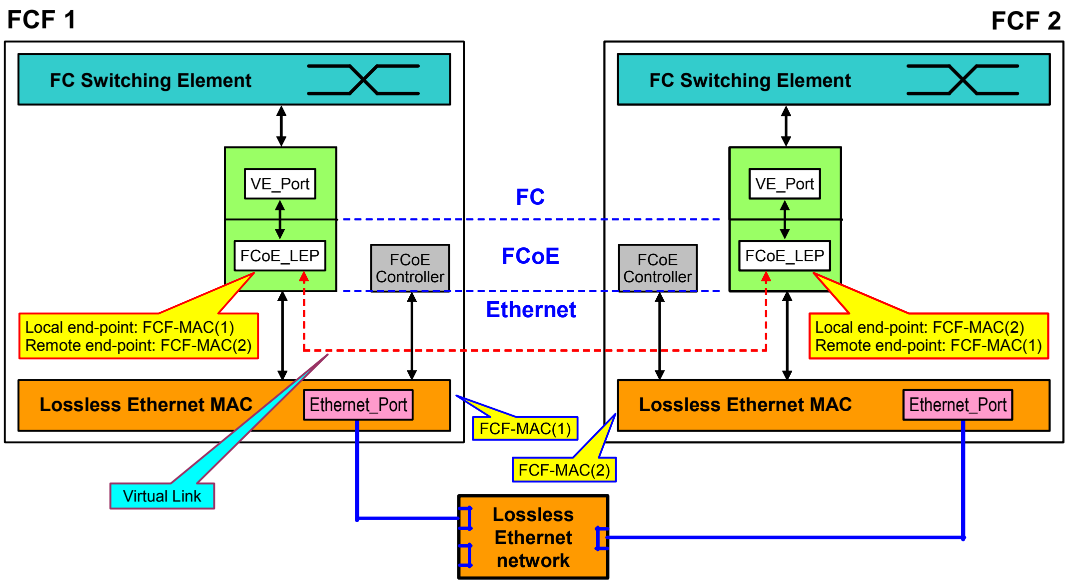

Figure 30 – VE_Port to VE_Port Virtual Links

On successful completion of a FIP ELP Exchange, the FCoE Controllers of the two involved VE_Port capable FCF-MACs instantiate a VE_Port/FCoE_LEP pair. Figure 30 shows the Virtual Links end-points, that are the MAC addresses of the two involved VE_Port capable FCF-MACs (i.e., FCF-MAC(1) and FCF-MAC(2)).

7.6 FC-BB_E device initialization

7.6.1 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 7.6.4.1) are specified for discovery and specific Link Service Requests and Replies.

The FIP frame format (see 7.6.4) is different than the FC-BB_E frame format (see 7.7) 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.

7.6.2 FIP discovery overview

When an ENode is added to a FC-BB_E fabric it discovers FCFs that it may perform Fabric login with by transmitting a Discovery Solicitation (see 7.6.5.2) 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 7.6.5.3) to the ENode if it is configured to allow a Fabric login from that ENode.

The ENode MAC address is used for all FIP operations and shall remain valid.

Editor's Note: Should we allow a unicast Discovery Solicitation froman 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 the ALL_ENODE_MACS and ALL_FCF_MACS multicast group addresses 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.

7.6.3 Link Service Request and Reply overview

Fabric login (i.e., FLOGI, NPIV FDISC), Fabric logout, and ELP are 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 are not performed using the normal FC-BB_E frame format).

In addtition to the Fabric login service, the FIP Fabric login provides a method to assign a MAC address for the VN_Port.

In addition to the Fabric logout service, the FIP Fabric logout provides a method to deassign a MAC address for the VN_Port.

Editor's Note: Add a statement on ELP if it remains in FIP.

7.6.4 FIP frame format

All FIP frames shall be formatted in accordance with 802.3-2005, Part 3 and the MAC Client Data field within the 802.3 frame shall contain the encapsulated FIP operation.

The format of a FIP frame is specified in table 22. The use of an 802.1Q tag header is optional and is not shown in table 22.

Table 22 – FIP frame format

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	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	Type																Version				Reserved											
4	Encapsulated FIP operation (n words)																															
n+4	FCS																															

Editor's Note: Add the optional use of 802.1Q tag header to the FIP frame format.

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 Type field shall be set to 8914h (i.e., FCoE initialization protocol).

The Version field shall be set to 0001b.

The Encapsulated FIP operation field is specified in 7.6.4.1.

The FCS field shall be set as specified in 802.3-2005, Part 3.

Editor's Note: Is it valid to send a FIP frame with no encapsulated FIP operation?

7.6.4.1 Encapsulated FIP operation

7.6.4.1.1 Encapsulated FIP operation overview

The format of the encapsulated FIP operation is specified in table 23.

Table 23 – Encapsulated FIP operation format

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="16">FIP Operation Code</td><td colspan="8">Reserved</td><td colspan="8">FIP Subcode</td></tr>
<tr><td>1</td><td colspan="16">FIP Descriptor List Length</td><td>FP</td><td>SP</td><td colspan="12">Reserved</td><td>S</td><td>F</td></tr>
<tr><td>2</td><td colspan="32">FIP Descriptor(s) (n words)</td></tr>
<tr><td>n+m+2</td><td colspan="32">FIP_Pad (m words)</td></tr>
</table>

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	7.6.5.2
	02h	Discovery Advertisement	7.6.5.3
00000002h	01h	Link Service Request	7.6.5.4
	02h	Link Service Reply	7.6.5.4
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 and is reserved for all other FIP operations.

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 requested. 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 requested. 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 is reserved for 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 is reserved for all other FIP operations.

Editor's Note: Should the FP, SP, S, and F bits be specified as reserved or ignored for all other FIP operations?

The FIP Descriptor(s) field contains one or more FIP descriptors (see 7.6.4.1.2).

The FIP_Pad field is used in solicited Discovery Advertisements to extend the frame length to indicate the maximum frame length supported by the originator (see 7.6.5.3).

7.6.4.1.2 FIP descriptors

7.6.4.1.2.1 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 including the TLV format. The FIP descriptor types are specified in table 26.

Table 26 – FIP descriptor types

Value	FIP Descriptor	Reference
1	Priority	7.6.4.1.2.2
2	MAC address	7.6.4.1.2.3
3	FC-MAP	7.6.4.1.2.4
4	Name_Identifier	7.6.4.1.2.5
5	Fabric_Name	7.6.4.1.2.6
6	Max Receive Size	7.6.4.1.2.7
7	FLOGI[a]	7.6.4.1.2.8
8	FDISC_NPIV[a]	7.6.4.1.2.9
9	LOGO[a]	7.6.4.1.2.10
10	ELP[a]	7.6.4.1.2.11
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.

7.6.4.1.2.2 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).

The FIP Priority descriptor format is specified in table 27.

Table 27 – FIP Priority descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type (1)								Length (1)								Reserved								Priority							

7.6.4.1.2.3 FIP MAC address descriptor

The FIP MAC address descriptor is used in Discovery Solicitations, Discovery Advertisements, and specific Link Service Requests and Replies.

The FIP MAC address descriptor format is specified in table 28.

Table 28 – FIP MAC address descriptor format

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (2)</td><td colspan="8">Length (2)</td><td colspan="8">MAC address[0]</td><td colspan="8">MAC address[1]</td></tr>
<tr><td>1</td><td colspan="8">MAC address[2]</td><td colspan="8">MAC address[3]</td><td colspan="8">MAC address[4]</td><td colspan="8">MAC address[5]</td></tr>
</table>

7.6.4.1.2.4 FIP FC-MAP descriptor

The FIP FC-MAP descriptor is used in Discovery Solicitations and Discovery Advertisements originated by an FCF.

The FIP FC-MAP descriptor format is specified in table 28.

Table 29 – FIP FC-MAP descriptor format

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (2)</td><td colspan="8">Length (2)</td><td colspan="16">Reserved</td></tr>
<tr><td>1</td><td colspan="8">Reserved</td><td colspan="8">FC-MAP[0]</td><td colspan="8">FC-MAP[1]</td><td colspan="8">FC-MAP[2]</td></tr>
</table>

7.6.4.1.2.5 FIP Name_Identifier descriptor

The FIP Name_Identifier descriptor is used in Discovery Solicitations originated by an ENode and Discovery Solicitations and Discovery Advertisements originated from an FCF.

The FIP Name_Identifier descriptor format is specified in table 30.

Table 30 – FIP Name_Identifier descriptor format

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (4)</td><td colspan="8">Length (3)</td><td colspan="16">Reserved</td></tr>
<tr><td>1</td><td colspan="8">MSB</td><td colspan="16" rowspan="2">Name_Identifier</td><td colspan="8"></td></tr>
<tr><td>2</td><td colspan="8"></td><td colspan="8">LSB</td></tr>
</table>

7.6.4.1.2.6 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

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (5)</td><td colspan="8">Length (3)</td><td colspan="16">Reserved</td></tr>
<tr><td>1</td><td colspan="8">MSB</td><td colspan="24" rowspan="2">Fabric_Name</td></tr>
<tr><td>2</td></tr>
<tr><td></td><td colspan="24"></td><td colspan="8">LSB</td></tr>
</table>

7.6.4.1.2.7 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

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (6)</td><td colspan="8">Length (1)</td><td colspan="16">Max_Receive_Size</td></tr>
</table>

7.6.4.1.2.8 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

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">Type (7)</td><td colspan="8">Length</td><td colspan="16">Reserved</td></tr>
<tr><td>1</td><td colspan="8">MSB</td><td colspan="24" rowspan="2">FLOGI Request or FLOGI
LS_ACC/LS_RJT</td></tr>
<tr><td>n</td></tr>
<tr><td></td><td colspan="24"></td><td colspan="8">LSB</td></tr>
</table>

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.

7.6.4.1.2.9 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

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type (8)								Length								Reserved															
1	MSB								NPIV FDISC Request or NPIV FDISC LS_ACC/LS_RJT																							
n																									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.

7.6.4.1.2.10 FIP LOGO descriptor

The FIP LOGO descriptor is used for Fabric logout requests and replies.

The FIP LOGO descriptor format is specified in table 33

Table 35 – FIP LOGO descriptor format

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	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																									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.

7.6.4.1.2.11 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

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	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																									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.

7.6.5 FIP operations

7.6.5.1 FIP operations overview

Each FIP operation shall contain one or more FIP descriptors with the payload and order as specified in table 37.

Table 37 – FIP operation payload and order

FIP Operation	Originator	Payload and order	Reference
Discovery Solicitation	ENode	1) MAC address 2) Name_Identifier 3) Max Receive Size	
Discovery Solicitation	FCF	1) MAC address 2) FC-MAP 3) Name_Identifier 4) Max Receive Size	
Discovery Advertisement	FCF	1) Priority 2) MAC address 3) FC-MAP 4) Name_Identifier 5) Fabric_Name	
FLOGI Request	ENode	1) FLOGI 2) MAC address	
FLOGI LS_ACC	FCF	1) FLOGI 2) MAC address	
FLOGI LS_RJT	FCF	1) FLOGI	
NPIV FDISC	ENode	1) FDISC_NPIV 2) MAC address	
NPIV FDISC LS_ACC	FCF	1) FDISC_NPIV 2) MAC address	
NPIV FDISC LS_RJT	FCF	1) FDISC_NPIV	
Fabric LOGO	ENode or FCF	1) LOGO 2) MAC address	
Fabric LOGO LS_ACC	ENode or FCF	1) LOGO 2) MAC address	
Fabric LOGO LS_RJT	ENode or FCF	1) LOGO	
ELP Request	FCF	1) ELP 2) MAC address	
ELP SW_ACC	FCF	1) ELP 2) MAC address	
ELP SW_RJT	FCCF	1) ELP	

Unless otherwise specified, the FIP descriptors may be received in any order. Additional FIP descriptors may be present. If additional FIP descriptors are present, then they shall be ignored. This is to provide flexibility for future standardization.

Editor's Note: Not sure why the above paragraph is needed since TLV format is used, and is confusing for the novice reader.

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.

If a non-FCF device receives a Discovery Solicitation, then the Discovery Solicitation shall be discarded.

7.6.5.2 FIP Discovery Solicitation

7.6.5.2.1 ENode FIP Discovery Solicitation

ENodes may transmit Discovery Solicitations to FCFs to request the FCF to reply with a Discovery Advertisement. The Discovery Solicitation shall be addressed to a specific FCF (i.e., unicast) or to the ALL_FCF_MACS multicast group address.

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent Discovery Advertisements from an FCF.

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

The Max_Receive_Size field in the Max Receive Size descriptor 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.

7.6.5.2.2 FCF FIP Discovery Solicitation

FCFs may transmit Discovery Solicitations to FCFs to request the FCF to reply with a Discovery Advertisement. The Discovery Solicitation shall be addressed to a specific FCF (i.e., unicast) or to the ALL_FCF_MACS multicast group address.

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent Discovery Advertisements from an FCF.

For FCFs that support FPMA, the FC-MAP field in the FC-MAP descriptor shall be set to the FC-MAP value the FCF is using. If the FC-MAP value is not administratively configured, then the FC_MAP value shall be set to DEFAULT_FC-MAP (see table 41).

For FCFs that only support SPMA, the FC-MAP field in the FC-MAP descriptor is reserved.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF.

The Max_Receive_Size field in the Max Receive Size descriptor shall be set to the maximum 802.3 frame size that the FCF 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.

After receiving a valid Discovery Solicitation originated by an FCF (i.e., the F bit is set to one), an FCF shall perform the following verification steps:

a) the Name_Identifier field value in the Discovery Solicitation is different than the Switch_Name of the recipient FCF; and
b) the FC-MAP field value in the Discovery Solicitation is zero or is the same as the FC-MAP value of the recipient FCF.

If any verification step is false, then the Discovery Solicitation shall be discarded.

NOTE 13 – It is possible for an FCF to receive a Discovery Solicitation that it originated because Discovery Solicitations sent to the ALL_FCF_MACS multicast group address may be forwarded to other ports on the same FCF by intermediate ethernet bridges.

7.6.5.3 FIP Discovery Advertisements

An FCF shall transmit a solicited Discovery Advertisement in response to a received Discovery Solicitation (see 7.6.5.2). A solicited Discovery Advertisement shall be sent to the MAC address field value specified in the MAC address descriptor in the received Discovery Solicitation. The solicited Discovery Advertisement shall be sent within D_A_TOV (see table 41) seconds upon reception of the Discovery Solicitation.

An FCF may receive Discovery Solicitations from the same FC-BB_E device on multiple FCF-MACs. In this case, a seperate solicited Discovery Advertisement shall be transmitted on each of the FCF-MACs that received the Discovery Solicitation. The FC-BB_E device that transmitted the Discovery Solicitation may determine that it received multiple solicited Discovery Advertisements from the same FCF since the value of the Name_Identifer field in the Name_Identifier descriptor will be the same in each of the solicited Discovery Advertisements.

In addition, a VF_Port capable FCF-MAC and a VE_Port capable FCF-MAC shall transmit an unsolicited Discovery Advertisement to the ALL_ENODE_MACS (i.e., for VF_Port capable FCF-MAC) or ALL_FCF_MACS (i.e., for VE_Ports capable FCF-MAC) multicast address group every D_A_TOV * 10 seconds +/- 50%. FCFs should randomize the time that Discovery Advertisements are sent within this window to avoid large bursts of multicast traffic within the network.

The Priority field in the Priority descriptor shall be set to the value the FCF is using. If the priority value is not administratively configured, then the priority value shall be set to DEFAULT_FIP_PRIORITY (see table 41).

The MAC address field in the MAC address descriptor shall be set to the FCF-MAC address.

For FCFs that support FPMA, the FC-MAP field in the FC-MAP descriptor shall be set to the FC-MAP value the FCF is using. If the FC-MAP value is not administratively configured, then the FC_MAP value shall be set to DEFAULT_FC-MAP (see table 41).

For FCFs that only support SPMA, the FC-MAP field in the FC-MAP descriptor is reserved.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF.

The Fabric_Name field in the Fabric_Name descriptor shall be set to the Fabric_Name for the FCF.

For a solicited Discovery Advertisement, the FIP_Pad field shall be set to the length required to create an 802.3 frame with an overall length that matches the Max_Receive_Size field value in the Max Receive Size descriptor in the received Discovery Solicitation. The FIP_Pad field values shall be set to reserved. For an unsolicited Discovery Advertisement, the FIP_Pad field shall be of zero length (i.e. not present).

Unless otherwise administratively configured, FCFs and ENodes shall not form VN_Ports, VF_Ports or VE_Ports, as appropriate, with peer FC-BB_E devices unless the Discovery Advertisement contains the highest priority received from all received Discovery Solicitations within at least 2 * D_A_TOV seconds after transmitting the Discovery Solicitation. If the same highest Discovery Advertisement priority is received from multiple peer FC-BB_E devices, then such Ports may be established with any set of FC-BB_E devices from which these Discovery Advertisements were received. Discovery Advertisements received after this period may, but are not required to, be considered.

Editor's Note: Not quite clear what this paragraph is trying to say, and may also be better to move it to a Vx_Port/virtual link creation subclause.

7.6.5.4 FIP Link Service Requests and Replies

7.6.5.4.1 FIP Link Service Requests and Replies overview

FIP Link Service Requests and Replies shall be transmitted:

a) using the FIP frame format (see 7.6.4);
b) as a single-frame Sequence; and
c) contain no Login Extension Data (see FC-FS-3), if applicable.

7.6.5.4.2 Fabric login (FLOGI or NPIV FDISC)

When an ENode transmits a FIP FLOGI Request or NPIV FDISC Request it shall indicate the addressing mode it supports (i.e., FPMA, SPMA, or both).

For a Fabric login operation using FLOGI, the FLOGI descriptor shall contain a FLOGI Request, FLOGI LS_ACC, or FLOGI LS_RJT payload.

For a Fabric login operation using NPIV FDSIC, the FDISC_NPIV descriptor shall contain a NPIV FDISC Request, NPIV FDISC LS_ACC, or NPIV FDISC LS_RJT payload.

If an ENode only supports SPMA or supports both SPMA and FPMA, the MAC address field in the MAC address descriptor shall be set to the proposed MAC address to use for subsequent FC-BB_E frames.

If an ENode only supports FPMA, the MAC address field in the MAC address descriptor shall be set to the proposed MAC address to use for subsequent FC-BB_E frames, or to all zeroes to indicate no MAC address is proposed.

If the ENode only supports SPMA, the MAC address specified in the FIP FLOGI Request or NPIV FDISC Request shall be returned in the FIP FLOGI Reply and shall be used as the VN_Port MAC address for all subsequent FC-BB_E frames.

If the ENode only supports 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. The assigned MAC address shall be a properly formed FPMA MAC address. In this case, the assigned MAC address should be the MAC address proposed in the FIP FLOGI Request or NPIV FDISC Request, if the proposed MAC address is a properly formed FPMA MAC address.

If the ENode supports both FPMA and SPMA, the assigned MAC address shall be either the MAC address specified in the FIP FLOGI Request or NPIV FDISC Request, or a properly formed FPMA MAC address assigned by the FCF.

A properly formed FPMA MAC address is one in which the 24 most significant bits equal the Fabric's FC-MAP value and the least significant 24 bits are unique within the Fabric for all FPMA addresses assigned with the same FC-MAP.

If both the FCF and ENode support both SPMA and FPMA, the FCF may assign an address of either form.

FCFs shall reject FIP FLOGI Requests and NPIV FDISC Requests from ENodes that support only SPMA and propose a MAC address that is not a unicast address. In addition, FCFs shall reject FIP FLOGI Requests and NPIV FDISC Requests for an addressing mode (i.e., SPMA or FPMA) not supported by the FCF. Finally, if the FCF supports both FPMA and SPMA and the ENode supports only SPMA, the FCF shall reject FIP FLOGI and NPIV FDISC requests that contain a proposed MAC address in which the 24 most significant bits match the FC-MAP in use by the FCF.

The 802.3 frame source address in a FIP NPIV_FDISC Request shall be the same as the 802.3 frame source address of the prior FIP FLOGI Request associated with the FIP NPIV FDISC. A successful FIP FLOGI operation creates a VF_Port. Subsequent FIP NPIV FDISCs with the same 802.3 frame source address as the FIP FLOGI associate additional VN_Ports to the single VF_Port.

7.6.5.4.3 Fabric logout

The LOGO descriptor shall contain a Fabric LOGO Request, Fabric LOGO LS_ACC, or Fabric LOGO LS_RJT payload.

For a Fabric LOGO Request and Fabric LOGO LS_ACC, the MAC address field in the MAC address descriptor shall be set to the MAC address assigned to the VN_Port that is being logged out.

7.6.5.4.4 Exchange Link Parameters

The ELP descriptor shall contain an ELP Request, ELP SW_ACC, or ELP SW_RJT payload.

For an ELP Request and ELP SW_ACC, the MAC address field in the MAC address descriptor shall be set to the FCF-MAC address

7.7 FC-BB_E frame format

7.7.1 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

<table>
<tr><th>Bit
Word</th><th>31</th><th>30</th><th>29</th><th>28</th><th>27</th><th>26</th><th>25</th><th>24</th><th>23</th><th>22</th><th>21</th><th>20</th><th>19</th><th>18</th><th>17</th><th>16</th><th>15</th><th>14</th><th>13</th><th>12</th><th>11</th><th>10</th><th>9</th><th>8</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>1</th><th>0</th></tr>
<tr><td>0</td><td colspan="8">DA[0]</td><td colspan="8">DA[1]</td><td colspan="8">DA[2]</td><td colspan="8">DA[3]</td></tr>
<tr><td>1</td><td colspan="8">DA[4]</td><td colspan="8">DA[5]</td><td colspan="8">SA[0]</td><td colspan="8">SA[1]</td></tr>
<tr><td>2</td><td colspan="8">SA[2]</td><td colspan="8">SA[3]</td><td colspan="8">SA[4]</td><td colspan="8">SA[5]</td></tr>
<tr><td>3</td><td colspan="16">Type</td><td colspan="4">Version</td><td colspan="12">Reserved</td></tr>
<tr><td>4</td><td colspan="32">Reserved</td></tr>
<tr><td>5</td><td colspan="32">Reserved</td></tr>
<tr><td>6</td><td colspan="24">Reserved</td><td colspan="8">SOF</td></tr>
<tr><td>7</td><td colspan="32">FC Frame (n words)</td></tr>
<tr><td>n+7</td><td colspan="8">EOF</td><td colspan="24">Reserved</td></tr>
<tr><td>n+8</td><td colspan="32">FCS</td></tr>
</table>

Editor's Note: Add an optional 802.1Q tag header to the above table.

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 Type field shall be set to FCoE_TYPE (see table 41).

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 shall contain:

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

7.8 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	Reference
FIP_TYPE	8914h	The value specified in the 802.3 Type field for a FIP frame.	7.6.4
FCoE_TYPE	8906h	The value specified in the 802.3 Type field for an FC-BB_E frame.	7.7.1
FC-BB_E_FRAME_VER	0000b	The value specified in the Version field for an FC-BB_E frame.	7.7.1
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.	7.6.4.1.2.2
DEFAULT_FC-MAP	0EFC00h	The default value for the FC-MAP field in a FIP FC-MAP descriptor.	7.6.4.1.2.4
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.