# 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.

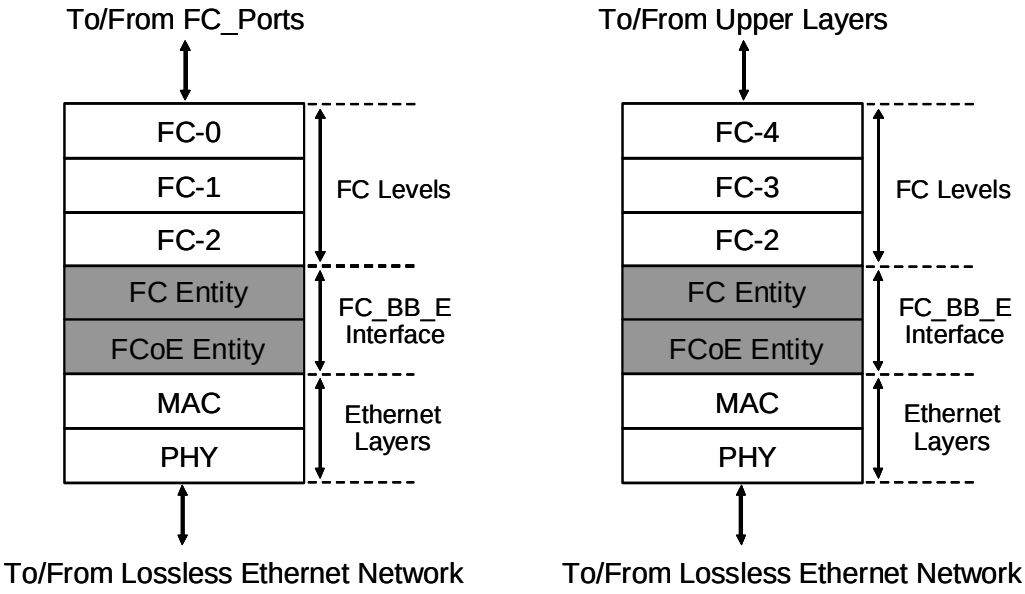


**Figure 24 – FC-BB_E protocol levels and layers**

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 (e.g., the Pause mechanism defined in 802.3-2005) 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-4) 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.

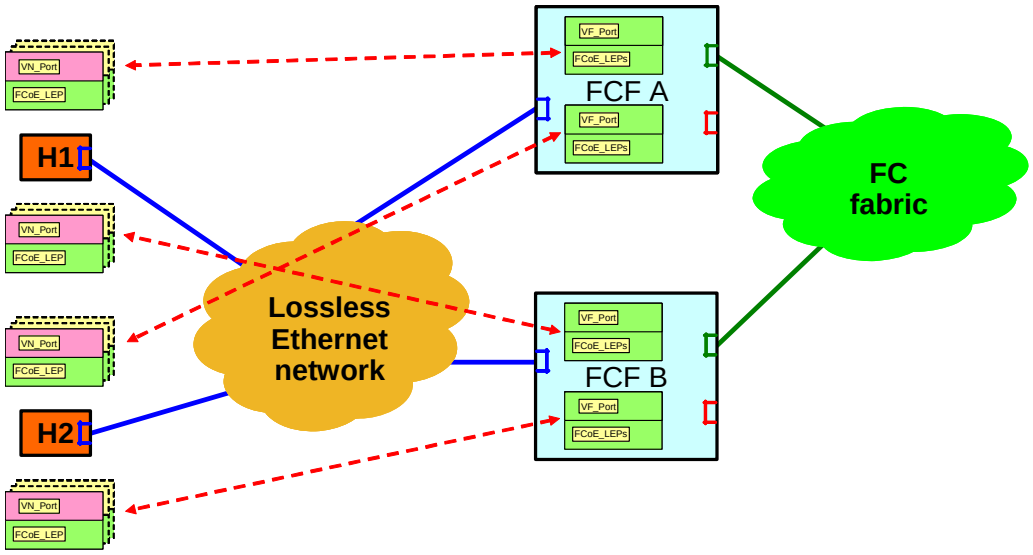


**Figure 25 – FCoE VN_Port to VF_Port Network Configuration**

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.

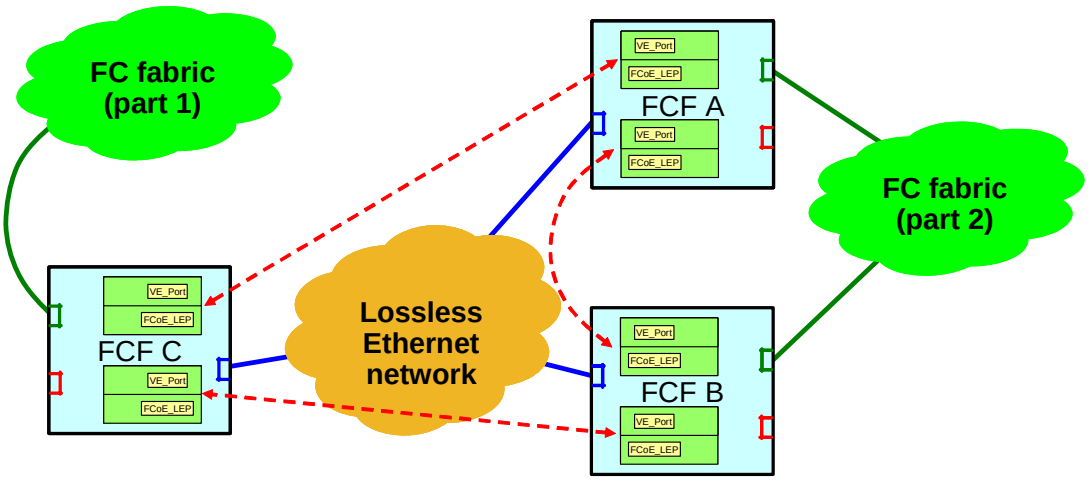


**Figure 26 – FCoE VE_Port to VE_Port Network Configuration**

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.

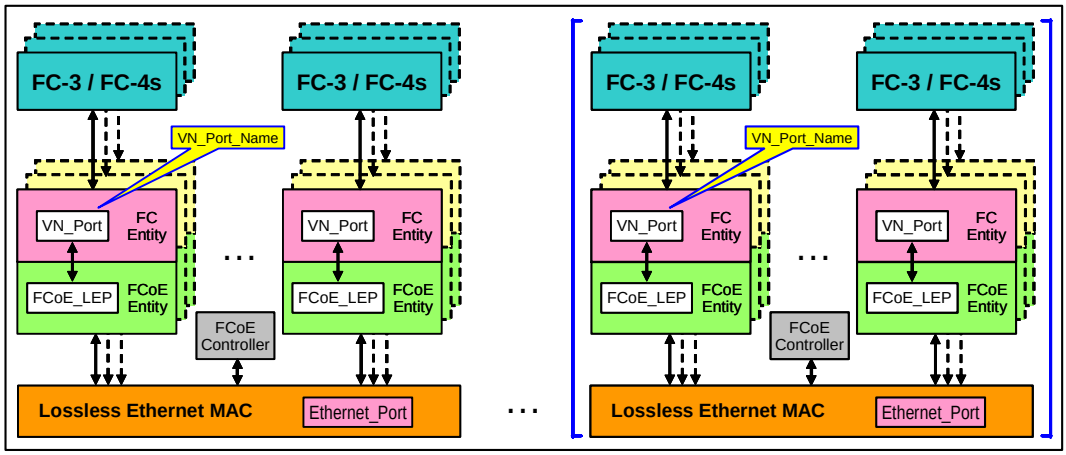


**Figure 27 – FC-BB_E VN_Port/ENode Functional Model**

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;
4) de-instantiates a VN_Port/FCoE_LEP pair when that VN_Port is logged out;
5) monitors the status of the instantiated VN_Ports and generates appropriate periodic FIP Keep Alive messages on behalf of them; and
6) monitors the status of the VF_Ports the instantiated VN_Ports are logged with by maintaining appropriate timers and by checking that periodic FIP Advertisements are received on time.

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-4) 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.

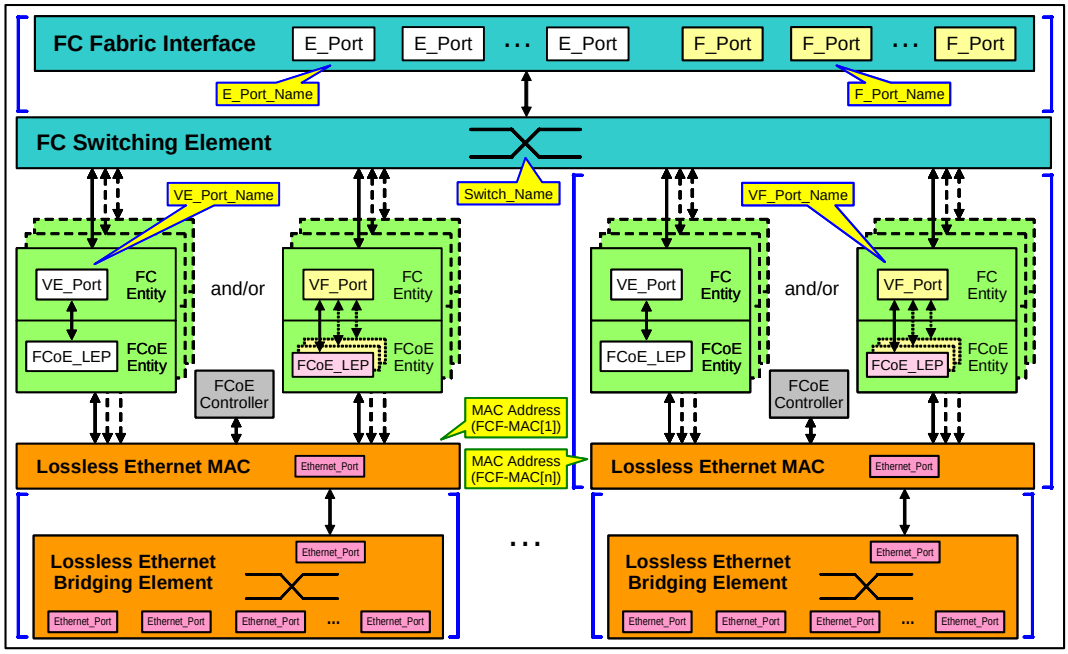


**Figure 28 – FC-BB_E VE_Port/VF_Port Functional Model**

NOTE 12 – Other combinations of Lossless Ethernet Bridging Elements with Lossless Ethernet MACs may be possible.

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;
2) instantiates a VE_Port/FCoE_LEP pair for any successful completion of a FIP ELP Exchange with a remote FCF-MAC;
3) de-instantiates a VE_Port/FCoE_LEP pair when appropriate;
4) monitors the status of the instantiated VE_Ports; and
5) monitors the status of remote VE_Ports by maintaining appropriate timers and by checking that periodic FIP Advertisements are received on time.

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;
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;
5) monitors the status of the instantiated VF_Ports; and
6) monitors the status of the logged in VN_Ports by maintaining appropriate timers and by checking that periodic FIP Keep Alive messages are received on time.

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-4, 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 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-4.

## 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.

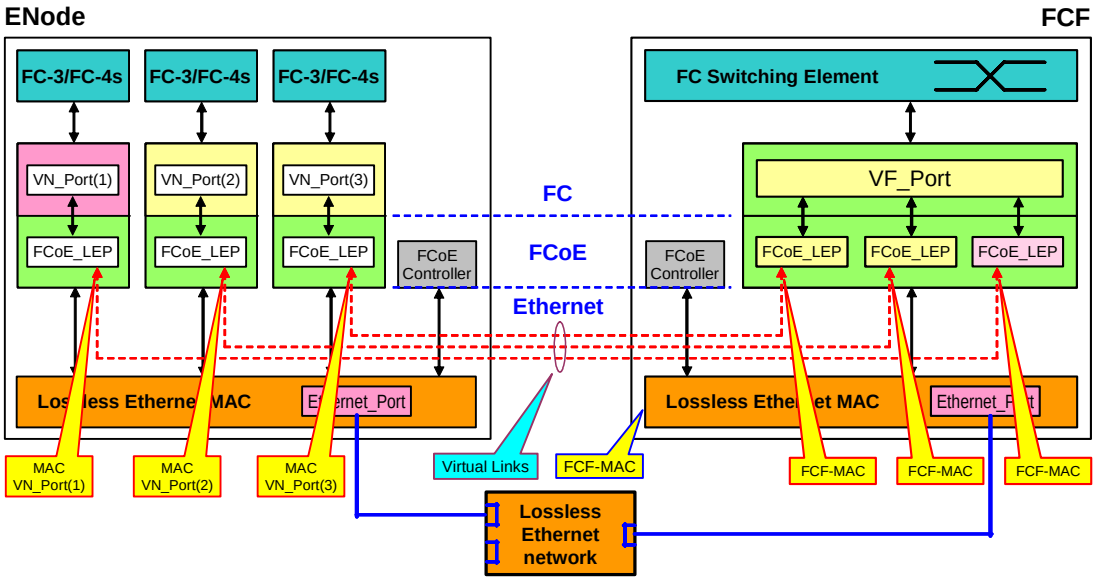


**Figure 29 – VN_Port to VF_Port Virtual Links**

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.

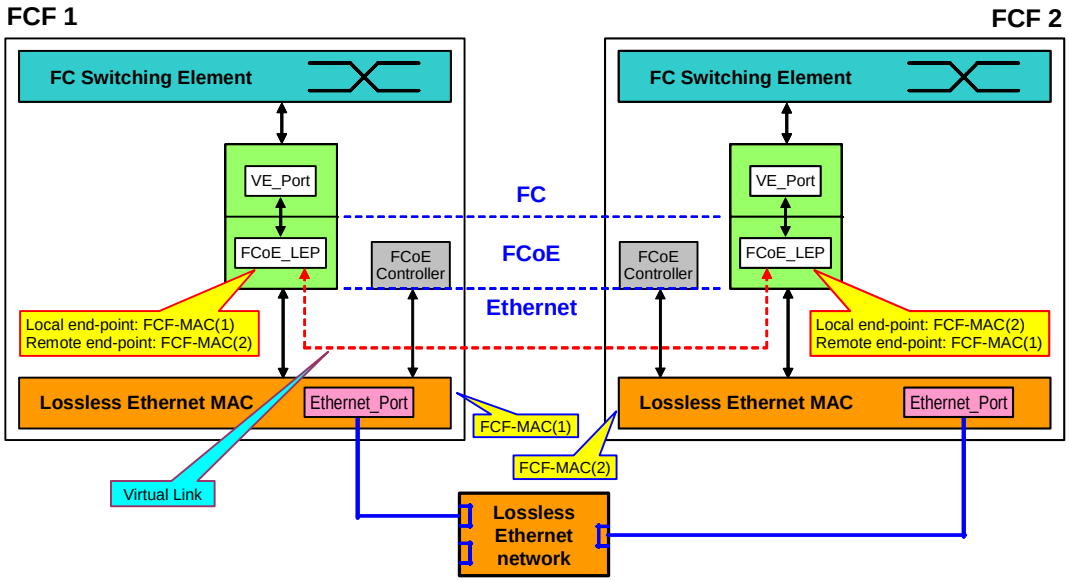


**Figure 30 – VE_Port to VE_Port Virtual Links**

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 FCoE frame format

The format of an FCoE frame is specified in table 22. The use of an 802.1Q tag header is optional.

**Table 22 – FCoE 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 | (IEEE 802.1Q Tag) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | Type | | | | | | | | | | | | | | | | Version | | | | Reserved | | | | | | | | | | | |
| 5 | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 6 | Reserved | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7 | Reserved | | | | | | | | | | | | | | | | | | | | | | | | SOF | | | | | | | |
| 8 | FC Frame (n words) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| n+8 | EOF | | | | | | | | Reserved | | | | | | | | | | | | | | | | | | | | | | | |
| n+9 | FCS | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

NOTE 13 – Additional IEEE 802.1 defined tags may be present in a FCoE frame.

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 45).

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 23.

**Table 23 – FCoE 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 24.

**Table 24 – FCoE EOF field**

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

## 7.7 FC-BB_E device initialization

### 7.7.1 FCoE Initialization Protocol (FIP) overview

The FCoE Initialization Protocol (FIP) is used to perform the function of FC-BB_E device discovery, initialization, and maintenance. To perform these functions, appropriate encapsulated FIP operations (see 7.7.5.2) are specified.

The FIP frame format (see 7.7.5.1) is different than the FCoE frame format (see 7.6) to enable the detection of discovery, initialization, and maintenance traffic from FCoE traffic.

FIP messages are used to perform the following protocols:

a) FIP Discovery (see 7.7.2);
b) FCoE Virtual Link instantiation (see 7.7.3); and
c) FCoE Virtual Link maintenance (see 7.7.4).

### 7.7.2 FIP Discovery protocol

#### 7.7.2.1 ENode/FCF discovery

VF_Port capable FCF-MACs periodically send multicast Advertisements to the All-ENode-MACs group address every FKA_ADV_PERIOD.

On receiving Advertisements, ENodes' FCoE Controllers create an entry per FCF-MAC in an internal FCF list, by default ordered on the basis of the value carried in the Priority descriptor of the Advertisements.

Each entry in the FCF list has some flags:

a) 'Max FCoE Size verified' (set to zero for entries created from multicast Advertisements, set to one when a unicast jumbo Advertisement is received); and
b) 'Available for Login' (set to one when the FCF is able to accept additional FIP FLOGI/FDISC requests, set to zero otherwise).

ENodes' FCoE Controllers select a subset of the Available FCF-MACs for Login (the 'FCF Login Set'), on the basis of a local policy (by default the one(s) with higher priority (i.e., lower priority value)). Each FCF-MAC of the 'FCF Login Set' shall be verified for Max FCoE Size support before performing the FIP FLOGI by sending a unicast Solicitation to a FCF-MAC and receiving a unicast Advertisement in response.

The periodic reception of multicast Advertisements allows ENode's FCoE Controllers to continuously verify the FCF-MACs reachability. The 'A' flag in received Advertisements carries the information that the sending FCF-MAC is available for FIP FLOGI/FDISC, information updated on the FCF list and on the 'FCF Login Set' on reception of Advertisements.

When an ENode's FCoE Controller becomes operational it discovers FCFs that it may perform FIP FLOGI with by transmitting a Solicitation (see 7.7.6.2) to the All-FCF-MACs group address. In response to a Solicitation from an ENode, an FCF transmits a solicited unicast Advertisement (see 7.7.6.3) to the ENode if it is configured to allow a FIP FLOGI from that ENode.

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

#### 7.7.2.2 FCF/FCF discovery

VE_Port capable FCF-MACs every FKA_ADV_PERIOD periodically send multicast Advertisements to the All-FCF-MACs group address.

On receiving Advertisements, FCoE Controllers of VE_Port capable FCF-MACs create an entry per FCF-MAC in an internal FCF list, by default ordered on the basis of the value carried in the Priority descriptor of the Advertisements.

Each entry in the FCF list has some flags:

a) 'Max FCoE Size verified' (set to zero for entries created from multicast Advertisements, set to one when a unicast jumbo Advertisement is received); and
b) 'Available for ELP' (set to one when the FCF is able to accept additional FIP ELP requests, set to zero otherwise).

Each FCF-MAC of the FCF list shall be verified for Max FCoE Size support before performing the FIP ELP by sending a unicast Solicitation to a FCF-MAC and receiving a unicast Advertisement in response.

The periodic reception of multicast Advertisements allows FCoE Controllers of VE_Port capable FCF-MACs to continuously verify the FCF-MACs reachability. The 'A' flag in received Advertisements carries the information that the sending FCF-MAC is available for FIP ELP, information updated on the FCF list on reception of Advertisements.

When a FCoE Controller of a VE_Port capable FCF-MAC becomes operational it discovers other VE_Port capable FCF-MACs by transmitting a Solicitation to the All-FCF-MACs group address. In response to a Solicitation from an FCF, a VE_Port capable FCF-MAC transmits a solicited unicast Advertisement to the FCF if it is configured to allow a Virtual Link with that FCF.

### 7.7.3 FCoE Virtual Link Instantiation protocol

#### 7.7.3.1 VN_Port to VF_Port Virtual Links

The FCoE Controller of an ENode's MAC instantiates VN_Port to VF_Ports Virtual Links by initiating the FIP Fabric login protocol. Fabric login (i.e., FLOGI, NPIV FDISC) is performed using the FIP frame format (see table 25) and the associated FIP descriptor type (see table 29), not using the FCoE frame format.

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

Fabric logout (i.e., LOGO) is performed by an ENode using the FIP frame format (see table 25) and the associated FIP descriptor type (see table 29), not using the FCoE frame format.

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

#### 7.7.3.2 VE_Port to VE_Port Virtual Links

The FCoE Controller of a VE_Port capable FCF-MAC instantiates VE_Port to VE_Ports Virtual Links by initiating the FIP ELP protocol. ELP is performed using the FIP frame format (see table 25) and the associated FIP descriptor type (see table 29), not using the FCoE frame format.

In addition to the ELP service, the FIP ELP provides a method to assign a MAC address for the VE_Port.

### 7.7.4 FCoE Virtual Link Maintenance protocol

#### 7.7.4.1 VN_Port to VF_Port Virtual Links

The FCoE Controllers of an ENode's MAC and of a VF_Port capable FCF-MAC continuously verify the state of a VN_Port to VF_Port Virtual Link by sending appropriate FIP messages and by checking received FIP messages.

The FCoE Controller of an ENode shall generate a FIP Keep Alive message with the ENode's MAC address as source MAC address toward each VF_Port capable FCF-MAC with which it has some VN_Ports logged in. This FIP Keep Alive message shall be generated every FKA_ADV_PERIOD. In addition, the FCoE Controller of an ENode shall generate for each VN_Port a FIP Keep Alive message with the VN_Port's MAC address as source MAC address toward the VF_Port capable FCF-MAC with which it the VN_Port is logged in. This FIP Keep Alive message shall be generated every 90 seconds.

The FCoE Controller of an ENode shall monitor the status of a VF_Port with which it has some VN_Ports logged in by checking the reception of FIP Advertisements from that VF_Port capable FCF-MAC, Advertisements that are generated every FKA_ADV_PERIOD. After three missing Advertisements all the VN_Port to VF_Port Virtual Links associated with that VF_Port shall be de-instantiated as well as the associated VN_Ports.

The FCoE Controller of a VF_Port capable FCF-MAC shall generate a FIP Advertisement message toward the All-ENode-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller of a VF_Port capable FCF-MAC shall monitor the status of an ENode with which it has some active Virtual Links by checking the reception of FIP Keep Alive messages from that ENode and its VN_Ports. VN_Port's related FIP Keep Alive messages (i.e., having the VN_Port MAC address as source address) are expected every 90 seconds and ENode's related FIP Keep Alive messages (i.e., having the ENode's MAC address as source address) are expected every FKA_ADV_PERIOD.

After three missing VN_Port's related FIP Keep Alive messages the associated VN_Port to VF_Port Virtual Link shall be de-instantiated. After three missing ENode's related FIP Keep Alive messages, all associated VN_Port to VF_Port Virtual Links shall be de-instantiated.

NOTE 14 – The use of a fast FKA_ADV_PERIOD period for the ENode's FIP Keep Alive messages allows to have fast response times against loss of connectivities on the Ethernet realm with a limited overhead. The use of VN_Ports' FIP Keep Alive messages generated at a slower pace allows to clear the state associated with each individual VN_Port when that VN_Port becomes not operational.

Explicit VN_Port to VF_Port Virtual Link de-instantation is performed by generating a FIP Fabric logout by an ENode's MAC and by generating a FIP Clear Virtual Links message by a VF_Port capable FCF-MAC.

A FIP Clear Virtual Links message sent to an ENode's MAC carries the list of VN_Ports to be removed. The size of a FIP Clear Virtual Links message should not exceed the standard Ethernet size (i.e., 1500 bytes). If the list of VN_Ports to be removed does not fit, multiple frames may be generated instead than one.

#### 7.7.4.2 VE_Port to VE_Port Virtual Links

The FCoE Controllers of a VE_Port capable FCF-MAC continuously verify the state of a VE_Port to VE_Port Virtual Link by sending appropriate FIP Advertisements and by checking received FIP Advertisements.

The FCoE Controller of a VE_Port capable FCF-MAC shall generate a FIP Advertisement message toward the All-FCF-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller of a VE_Port capable FCF-MAC shall monitor the status of a VE_Port to VE_Port Virtual Link by checking the reception of FIP Advertisements. FIP Advertisements are expected every FKA_ADV_PERIOD. After three missing Advertisements the VE_Port to VE_Port Virtual Links associated with that FCF-MAC shall be de-instantiated as well as the associated VE_Ports.

Explicit VE_Port to VE_Port Virtual Link de-instantation is performed by generating a FIP Clear Virtual Links message by a VE_Port capable FCF-MAC. A FIP Clear Virtual Links message sent to a VE_Port capable FCF-MAC carries the destination FCF-MAC address.

### 7.7.5 FIP Messages

#### 7.7.5.1 FIP frame format

All FIP frames shall be formatted in accordance with 802.3-2005 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 25. The use of the 802.1Q tag header is optional.

**Table 25 – 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 | (IEEE 802.1Q Tag) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | Type = FIP | | | | | | | | | | | | | | | | Version | | | | Reserved | | | | | | | | | | | |
| 5 | Encapsulated FIP operation (n words) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| n+5 | FCS | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

NOTE 15 – Additional IEEE 802.1 defined tags may be present in a FIP frame.

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 FIP_TYPE (see table 45).

The Version field shall be set to 0001b.

The Encapsulated FIP operation field is specified in 7.7.5.2.

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

NOTE 16 – A FIP frame with no encapsulated FIP operation is not valid and should be ignored if received.

### 7.7.5.2 Encapsulated FIP operation

The format of the encapsulated FIP operation is specified in table 26.

**Table 26 – Encapsulated FIP operation 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 | FIP Operation Code | | | | | | | | | | | | | | | | Reserved | | | | | | | | FIP Subcode | | | | | | | |
| 1 | FIP Descriptor List Length | | | | | | | | | | | | | | | | FP | SP | Reserved | | | | | | | | | | | A | S | F |
| 2 | FIP Descriptor(s) (n words) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| n+m+2 | FIP_Pad (m words) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

The FIP Operation Code and FIP Subcode field values and operations are specified in table 27.

**Table 27 – FIP Operation Code and FIP Subcode field values and operations**

| FIP Operation Code | FIP Subcode | FIP Operation | Reference |
|---|---|---|---|
| 0001h | 01h | Discovery Solicitation | 7.7.6.2 |
| | 02h | Discovery Advertisement | 7.7.6.3 |
| 0002h | 01h | Virtual Link Instantiation Request | 7.7.6.4 |
| | 02h | Virtual Link Instantiation Reply | 7.7.6.4 |
| 0003h | 01h | FIP Keep Alive | 7.7.6.5 |
| | 02h | FIP Clear Virtual Links | 7.7.6.6 |
| FFF8h .. FFFEh | 00h .. FFh | Vendor Specific | 7.7.6.7 |
| 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. These bits shall be set as specified in table 28 and are reserved for all other FIP operations.

**Table 28 – FP bit and SP bit setting**

| Bit | FIP operation | Setting |
|---|---|---|
| FP | Discovery Solicitation<br>Discovery Advertisement | Set to 1 if originating device supports FPMA.<br>Set to 0 if originating device does not support FPMA. |
| | FLOGI Request[a]<br>FDISC_NPIV Request[a] | Set to 1 if FPMA is requested.<br>Set to 0 if FPMA is not supported. |
| | FLOGI LS_ACC<br>FDISC_NPIV LS_ACC | Set to 1 if FPMA granted.<br>Set to 0 if SPMA granted. |
| SP | Discovery Solicitation<br>Discovery Advertisement | Set to 1 if originating device supports SPMA.<br>Set to 0 if originating device does not support SPMA. |
| | FLOGI Request[a]<br>FDISC_NPIV Request[a] | Set to 1 if SPMA is requested.<br>Set to 0 if SPMA is not supported. |
| | FLOGI LS_ACC<br>FDISC_NPIV LS_ACC | Set to 1 if SPMA granted.<br>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 one in a Discovery Advertisement that is transmitted in response to a Discovery Solicitation. The S bit shall be set to zero 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 one in a Discovery Solicitation or Discovery Advertisement if the originating device is an FCF. The F bit shall be set to zero 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.

The ‘Available for Login’ (A) bit shall be set to one in a Discovery Advertisement if the originating FCF is available to accept Virtual Link instantiation requests. The A bit shall be set to zero in a Discovery Advertisement if the originating FCF is not available to accept Virtual Link instantiation requests. The A bit is reserved for all other FIP operations.

The FIP Descriptor(s) field contains one or more FIP descriptors (see 7.7.5.3).

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 7.7.6.3).

### 7.7.5.3 FIP descriptors

#### 7.7.5.3.1 FIP descriptor overview

The FIP descriptors are specified using a TLV format (i.e., Type, Length, Value). The length field value shall be specified as the number of 32-bit words in the FIP descriptor including the TLV format. The FIP descriptor type space is split in two ranges, critical and non-critical, as follows:

a) 0 .. 127: Critical. An FCoE Controller receiving a FIP message carrying an unknown critical descriptor shall ignore the entire message; and

b) 128 .. 255: Non-critical. An FCoE Controller receiving a FIP message carrying an unknown non-critical descriptor shall ignore the unknown descriptor and process the message.

The FIP descriptor types are specified in table 29.

**Table 29 – FIP descriptor types**

| Criticality | Type | FIP Descriptor | Reference |
|---|---|---|---|
| Critical | 0 | Reserved | |
| | 1 | Priority | 7.7.5.3.2 |
| | 2 | MAC address | 7.7.5.3.3 |
| | 3 | FC-MAP | 7.7.5.3.4 |
| | 4 | Name_Identifier | 7.7.5.3.5 |
| | 5 | Fabric_Name | 7.7.5.3.6 |
| | 6 | Max FCoE Size | 7.7.5.3.7 |
| | 7 | FLOGI[a] | 7.7.5.3.8 |
| | 8 | FDISC_NPIV[a] | 7.7.5.3.9 |
| | 9 | LOGO[a] | 7.7.5.3.10 |
| | 10 | ELP[a] | 7.7.5.3.11 |
| | 11 | Vx_Port Identification | 7.7.5.3.12 |
| | 12 | FKA_ADV_Period | 7.7.5.3.13 |
| | 13 | Vendor_ID | 7.7.5.3.14 |
| | 14 .. 127 | Reserved | |
| Non-critical | 128 .. 240 | Reserved | |
| | 241 .. 254 | Vendor Specific | 7.7.5.3.15 |
| | 255 | Reserved | |

a The FC CRC, SOF, and EOF shall not be included in the FIP descriptor.

### 7.7.5.3.2 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 45). 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 30.

**Table 30 – 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 | | | | | | | |

#### 7.7.5.3.3 FIP MAC address descriptor

The FIP MAC address descriptor is used in Discovery Solicitations, Discovery Advertisements, FIP Keep Alive messages, FIP Clear Virtual Links messages, and specific Virtual Link Instantiation Requests and Replies.

The FIP MAC address descriptor format is specified in table 31.

**Table 31 – FIP MAC address 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 (2) | | | | | | | | Length (2) | | | | | | | | MAC address[0] | | | | | | | | MAC address[1] | | | | | | | |
| 1 | MAC address[2] | | | | | | | | MAC address[3] | | | | | | | | MAC address[4] | | | | | | | | MAC address[5] | | | | | | | |

#### 7.7.5.3.4 FIP FC-MAP descriptor

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

The FIP FC-MAP descriptor format is specified in table 31.

**Table 32 – FIP FC-MAP 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 (3) | | | | | | | | Length (2) | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | Reserved | | | | | | | | FC-MAP[0] | | | | | | | | FC-MAP[1] | | | | | | | | FC-MAP[2] | | | | | | | |

#### 7.7.5.3.5 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 by an FCF.

The FIP Name_Identifier descriptor format is specified in table 33.

**Table 33 – FIP Name_Identifier 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 (4) | | | | | | | | Length (3) | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | MSB | | | | | | | | Name_Identifier | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | LSB | |

#### 7.7.5.3.6 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 34.

**Table 34 – FIP Fabric_Name 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 (5) | | | | | | | | Length (3) | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | MSB | | | | | | | | Fabric_Name | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | LSB | |

#### 7.7.5.3.7 FIP Max FCoE Size descriptor

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

The FIP Max FCoE Size descriptor format is specified in table 35.

**Table 35 – FIP Max FCoE Size 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 (6) | | | | | | | | Length (1) | | | | | | | | Max_FCoE_Size | | | | | | | | | | | | | | | |

#### 7.7.5.3.8 FIP FLOGI descriptor

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

The FIP FLOGI descriptor format is specified in table 36.

**Table 36 – FIP FLOGI 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 (7) | | | | | | | | Length | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | MSB | | | | | | | | FLOGI Request or FLOGI LS_ACC/LS_RJT | | | | | | | | | | | | | | | | | | | | | | | |
| n | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | 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.

#### 7.7.5.3.9 FIP FDISC_NPIV descriptor

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

The FIP FDISC_NPIV descriptor format is specified in table 37.

**Table 37 – FIP FDISC_NPIV 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 (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.

#### 7.7.5.3.10 FIP LOGO descriptor

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

The FIP LOGO descriptor format is specified in table 38.

**Table 38 – FIP LOGO 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 (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.

#### 7.7.5.3.11 FIP ELP descriptor

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

The FIP ELP descriptor format is specified in table 39.

**Table 39 – FIP ELP 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 (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.

#### 7.7.5.3.12 FIP Vx_Port Identification descriptor

The FIP Vx_Port Identification descriptor is used in FIP Clear Virtual Links messages.

The FIP Vx_Port Identification descriptor format is specified in table 40.

**Table 40 – FIP Vx_Port Identification 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 (11) | | | | | | | | Length (5) | | | | | | | | MAC address[0] | | | | | | | | MAC address[1] | | | | | | | |
| 1 | MAC address[2] | | | | | | | | MAC address[3] | | | | | | | | MAC address[4] | | | | | | | | MAC address[5] | | | | | | | |
| 2 | Reserved | | | | | | | | Port_ID | | | | | | | | | | | | | | | | | | | | | | | |
| 3 | MSB | | | | | | | | Port_Name | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | | | | | | | | | | | | | | | | | | | | | | | | | LSB | | | | | | | |

#### 7.7.5.3.13 FIP FKA_ADV_Period descriptor

The FIP FKA_ADV_Period descriptor is used in Discovery Advertisements.

The FIP FKA_ADV_Period descriptor format is specified in table 41.

**Table 41 – FIP FKA_ADV_Period 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 (12) | | | | | | | | Length (2) | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | FKA_ADV_PERIOD | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

#### 7.7.5.3.14 FIP Vendor_ID descriptor

The FIP Vendor_ID descriptor is used in Vendor Specific FIP messages.

The FIP Vendor_ID descriptor format is specified in table 42.

**Table 42 – FIP Vendor_ID 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 (13) | | | | | | | | Length (3) | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | MSB | | | | | | | | Vendor_ID | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | | | | | | | | | | | | | | | | | | | | | | | | | LSB | | | | | | | |

#### 7.7.5.3.15 FIP Vendor Specific descriptors

The FIP Vendor Specific descriptors are identified by a type value in the range 241 .. 254. FIP Vendor Specific descriptors are non-critical and may be used in any FIP message.

The FIP Vendor Specific descriptor format is specified in table 43.

**Table 43 – FIP Vendor Specific 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 | | | | | | | | Length | | | | | | | | Reserved | | | | | | | | | | | | | | | |
| 1 | MSB | | | | | | | | Vendor_ID | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | LSB | |
| 3 | MSB | | | | | | | | Vendor Specific Information | | | | | | | | | | | | | | | | | | | | | | | |
| n | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | LSB | |

An FC-BB_E device is never required to generate a Vendor Specific descriptor for normal operation.

### 7.7.6 FIP operations

#### 7.7.6.1 FIP operations overview

Each FIP operation shall contain one or more FIP descriptors as specified in table 44. Table 44 specifies the FIP descriptors required in each FIP operation and the order in which they should be encapsulated by a transmitting FCoE Controller. Additional descriptors (e.g., Vendor Specific descriptors) may be present. A receiving FCoE Controller shall process unknown descriptors according to the criticality of the descriptor (see 7.7.5.3.1). Unless otherwise specified (e.g., for a FIP FLOGI Request), a receiving FCoE Controller shall be able to process the FIP descriptors in any order.

NOTE 17 – This is to provide flexibility for future protocol extensions.

**Table 44 – FIP operation payload and order**

| FIP Operation | Originator | Payload and order | Reference |
|---|---|---|---|
| Discovery Solicitation | ENode | 1) MAC address<br>2) Name_Identifier<br>3) Max FCoE Size | |
| Discovery Solicitation | FCF | 1) MAC address<br>2) FC-MAP<br>3) Name_Identifier<br>4) Max FCoE Size | |
| Discovery Advertisement | FCF | 1) Priority<br>2) MAC address<br>3) FC-MAP<br>4) Name_Identifier<br>5) Fabric_Name<br>6) FKA_ADV_Period | |
| FIP FLOGI Request[a] | ENode | 1) FLOGI<br>2) MAC address | |
| FIP FLOGI LS_ACC[a] | FCF | 1) FLOGI<br>2) MAC address | |

a Strict ordering of descriptors is required.

**Table 44 – FIP operation payload and order**

| FIP Operation | Originator | Payload and order | Reference |
|---|---|---|---|
| FIP FLOGI LS_RJT[a] | FCF | 1) FLOGI | |
| FIP NPIV FDISC Request[a] | ENode | 1) FDISC_NPIV<br>2) MAC address | |
| FIP NPIV FDISC LS_ACC[a] | FCF | 1) FDISC_NPIV<br>2) MAC address | |
| FIP NPIV FDISC LS_RJT[a] | FCF | 1) FDISC_NPIV | |
| FIP Fabric LOGO Request[a] | ENode | 1) LOGO<br>2) MAC address | |
| FIP Fabric LOGO LS_ACC[a] | FCF | 1) LOGO<br>2) MAC address | |
| FIP Fabric LOGO LS_RJT[a] | FCF | 1) LOGO | |
| FIP ELP Request[a] | FCF | 1) ELP<br>2) MAC address | |
| FIP ELP SW_ACC[a] | FCF | 1) ELP<br>2) MAC address | |
| FIP ELP SW_RJT[a] | FCF | 1) ELP | |
| FIP Keep Alive | ENode | 1) MAC address<br>2) Vx_Port Identification | |
| FIP Clear Virtual Links | FCF | 1) MAC address<br>2) Name_Identifier<br>3) Vx_Port Identification | |
| Vendor Specific | ENode or FCF | 1) Vendor_ID<br>2) List of Descriptors | |

a Strict ordering of descriptors is required.

Two group addresses (see table 45) are used for Discovery Solicitations and Discovery Advertisements. The All-ENode-MACs group address is used to address all ENode devices and the All-FCF-MACs group address is used to address all FCF devices.

### 7.7.6.2 FIP Discovery Solicitation

#### 7.7.6.2.1 ENode FIP Discovery Solicitation

ENode's MACs may transmit Discovery Solicitations to VF_Port capable FCF-MACs to request the FCF-MAC to reply with a solicited Discovery Advertisement. The Discovery Solicitation shall be addressed to a specific FCF-MAC (i.e., unicast) or to the All-FCF-MACs group address.

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent solicited Discovery Advertisements from VF_Port capable FCF-MACs.

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

The Max_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE payload size that the ENode is able to receive. The Max_FCoE_Size value shall be specified as the number of octets starting from the Version field to the FCS, exclusive (see table 22).

If an ENode receives a Discovery Solicitation, then the Discovery Solicitation shall be discarded.

#### 7.7.6.2.2 FCF FIP Discovery Solicitation

VE_Port capable FCF-MACs may transmit Discovery Solicitations to VE_Port capable FCF-MACs to request the FCF-MACs to reply with a solicited Discovery Advertisement. The Discovery Solicitation shall be addressed to a specific FCF-MAC (i.e., unicast) or to the All-FCF-MACs group address.

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent solicited Discovery Advertisements from VE_Port capable FCF-MACs.

For FCF-MACs that support FPMA, the FC-MAP field in the FC-MAP descriptor shall be set to the FC-MAP value the FCF-MAC 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 45).

For FCF-MACs 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_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE payload size that the VE_Port capable FCF-MAC is able to receive. The Max_FCoE_Size value shall be specified as the number of octets starting from the Version field to the FCS, exclusive (see table 22).

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

#### 7.7.6.3 FIP Discovery Advertisements

VF_Port capable FCF-MACs shall transmit periodically unsolicited Discovery Advertisements to the All-ENode-MACs group address. The period shall be FKA_ADV_PERIOD, randomized with a random delay uniformly distributed between 0 and 100 ms to avoid large bursts of multicast traffic within the network.

VE_Port capable FCF-MACs shall transmit periodically unsolicited Discovery Advertisements to the All-FCF-MACs group address. The period shall be FKA_ADV_PERIOD, randomized with a random delay uniformly distributed between 0 and 100 ms to avoid large bursts of multicast traffic within the network.

An FCF-MAC shall transmit a solicited Discovery Advertisement in response to a received Discovery Solicitation (see 7.7.6.2). 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 45) seconds upon reception of the Discovery Solicitation. Discovery Advertisements sent in response to a multicast Discovery Solicitation should be delayed by a random time uniformly distributed between 0 and 100 ms to avoid large bursts of multicast traffic within the network.

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.

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 45).

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 45).

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 payload length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the received Discovery Solicitation. The FIP_Pad field values shall be set to reserved. For an unsolicited Discovery Advertisements, the FIP_Pad field shall be of zero length (i.e. not present).

#### 7.7.6.4 FIP Virtual Link Instantiation Requests and Replies

##### 7.7.6.4.1 FIP Virtual Link Instantiation Requests and Replies overview

FIP Virtual Link Instantiation Requests and Replies shall be transmitted:

a) using the FIP frame format (see 7.7.5.1);
b) as a single-frame Sequence; and
c) without Login Extension Data (see FC-FS-3), if applicable.

##### 7.7.6.4.2 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 FDISC, 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 FCoE 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 FCoE 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 FCoE 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 FCoE 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 equal the N_Port_ID assigned to the VN_Port by the FCF. This guarantees that FPMAs are unique within the Fabric.

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.

#### 7.7.6.4.3 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.

#### 7.7.6.4.4 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.

### 7.7.6.5 FIP Keep Alive

The FCoE Controller of ENode's MACs logged in with one or more VF_Port capable FCF-MACs shall generate a FIP Keep Alive message every FKA_ADV_PERIOD. These FIP Keep Alive messages shall have the ENode's MAC address as source address and shall carry only the MAC address descriptor. The MAC address field in the MAC address descriptor shall be set to the ENode's MAC address. The FKA_ADV_PERIOD timer is taken from the Discovery Advertisements.

The FCoE Controller of ENode's MACs logged in with one or more VF_Port capable FCF-MACs shall also generate a FIP Keep Alive message on behalf of each logged in VN_Port every 90 seconds. These FIP Keep Alive messages shall have the VN_Port's MAC address as source address and shall carry both the MAC address descriptor and the Vx_Port Identification descriptor. The MAC address field in the MAC address descriptor shall be set to the ENode's MAC address. In the Vx_Port Identification descriptor the MAC address field shall be set to the VN_Port's MAC address, the Port_ID field shall be set to the VN_Port's N_Port_ID, and the Port_Name field shall be set to the VN_Port's VN_Port_Name.

### 7.7.6.6 FIP Clear Virtual Links

#### 7.7.6.6.1 FIP Clear Virtual Links toward an ENode

The FCoE Controller of a VF_Port capable FCF-MAC may turn down one or more VN_Port to VF_Port Virtual Links by sending a FIP Clear Virtual Links to an ENode's MAC. This FIP Clear Virtual Links message shall contain one MAC address descriptor, one Name_Identifier descriptor, and a list of Vx_Port Identification descriptors, one per each VN_Port the Virtual Link with it is requested to be turned down. The MAC address field in the MAC address descriptor shall be set to the FCF-MAC address. The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF. Per each Vx_Port Identification descriptor, the MAC address field shall be set to the VN_Port's MAC address, the Port_ID field shall be set to the VN_Port's N_Port_ID, and the Port_Name field shall be set to the VN_Port's VN_Port_Name.

#### 7.7.6.6.2 FIP Clear Virtual Links toward an FCF

The FCoE Controller of a VE_Port capable FCF-MAC may turn down a VE_Port to VE_Port Virtual Links by sending a FIP Clear Virtual Links to a VE_Port capable FCF-MAC. This FIP Clear Virtual Links message contains one MAC address descriptor, one Name_Identifier descriptor, and one Vx_Port Identification descriptor. The MAC address field in the MAC address descriptor shall be set to the FCF-MAC address. The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF. In the Vx_Port Identification descriptor, the MAC address field shall be set to the remote FCF-MAC address, the Port_ID field shall be set to FFFFFDh, and the Port_Name field shall be set to the remote VE_Port_Name.

### 7.7.6.7 FIP Vendor Specific messages

FIP Vendor Specific messages may be generated by both ENodes and FCFs. A FIP Vendor Specific message shall include a Vendor_ID descriptor as first descriptor, followed by a list of additional

descriptors. An unknown received Vendor Specific message shall be ignored. An FC-BB_E device is never required to generate a Vendor Specific message for normal operation.

### 7.8 Timers and constants

FC-BB_E timers and constants are specified in table 45.

**Table 45 – 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. | 7.7.5.1 |
| FCoE_TYPE | 8906h | The value specified in the 802.3 Type field for an FCoE frame. | 7.6 |
| All-FCoE-MACs | 01-10-18-01-00-00 | The group address for all FCoE device MACs. | |
| All-ENode-MACs | 01-10-18-01-00-01 | The group address for all ENode MACs. | |
| All-FCF-MACs | 01-10-18-01-00-02 | The group address for all FCF MACs. | |
| DEFAULT_FIP_PRIORITY | 128 | The default value specified in the FIP Priority descriptor. | 7.7.5.3.2 |
| DEFAULT_FC-MAP | 0EFC00h | The default value for the FC-MAP field in a FIP FC-MAP descriptor. | 7.7.5.3.4 |
| D_A_TOV | 4 | The default value in seconds of the maximum time between a discovery Solicitation and the responding solicited Discovery Advertisements. This value may be administratively configured to any value between 1 and 60, inclusive. | |
| FKA_ADV_PERIOD | 8000 | The default value in milliseconds for the periodicity of Discovery Advertisements and FIP Keep Alive messages. | |