1 FCoE Functional Models

1.1 FC-BB_E Definitions

1.1.1 FC Entity: The interface between an FC Switching Element or an FC stack and the FCoE Entity. Each FC Entity contains a single instance of either a VE_Port, a VF_Port, or a VN_Port.

1.1.2 FCoE Controller: A functional entity, coupled with a Lossless Ethernet MAC, instantiating new VE_Ports, VF_Ports, and VN_Ports, and/or creating new FCoE_LEPs.

1.1.3 FCoE Entity: The interface between the FC Entity and a Lossless Ethernet MAC. Each FCoE Entity contains one or more FCoE_LEPs.

1.1.4 FCoE Forwarder (FCF): A Fibre Channel Switching Element (see FC-SW-4) with one or more Lossless Ethernet MACs, each coupled with an FCoE Controller, and optionally one or more Losseless Ethernet bridging elements and optionally an FC Fabric interface. An FCF forwards FCoE frames addressed to one of its FCF-MACs based on the D_ID of the encapsulated FC frames.

1.1.5 FCoE Link End-Point (FCoE_LEP): The data forwarding component of an FCoE Entity that handles FC frame encapsulation/decapsulation, and transmission/reception of encapsulated frames through a single Virtual Link.

1.1.6 FCoE Node (ENode): A Fibre Channel Node (see FC-FS-2) with one or more Lossless Ethernet MACs, each coupled with an FCoE Controller.

1.1.7 Lossless Ethernet Bridging Element: An Ethernet bridging function supporting the minimum required capabilities of Lossless Ethernet MACs.

1.1.8 Lossless Ethernet MAC: A full duplex Ethernet MAC supporting at least 2.5KB jumbo frames and implementing extensions to avoid Ethernet frame loss due to congestion (e.g., the Pause mechanism, see IEEE 802.3-2005).

1.1.9 Lossless Ethernet network: An Ethernet network composed only of full duplex links, Lossless Ethernet MACs, and Lossless Ethernet Bridging Elements.

1.1.10 VE_Port_Name: The Name_Identifier of a VE_Port.

1.1.11 VF_Port_Name: The Name_Identifier of a VF_Port.

1.1.12 Virtual E_Port (VE_Port): The data forwarding component of an FC Entity that emulates an E_Port and is dynamically instantiated on successful completion of an ELP Exchange. The term virtual indicates the use of a non Fibre Channel link connecting the VE_Ports.

1.1.13 Virtual F_Port (VF_Port): The data forwarding component of an FC Entity that emulates an F_Port and is dynamically instantiated on successful completion of an FLOGI Exchange. The term virtual indicates the use of a non Fibre Channel link connecting a VF_Port with a VN_Port.

1.1.14 Virtual Link: The logical link connecting two FCoE_LEPs over a Lossless Ethernet network and is identified by the pair of MAC addresses of the two link end-points.

1.1.15 Virtual N_Port (VN_Port): The data forwarding component of an FC Entity that emulates an N_Port and is dynamically instantiated on successful completion of an FLOGI or FDISC Exchange. The term virtual indicates the use of a non Fibre Channel link connecting a VN_Port to a VF_Port.

1.1.16 VN_Port_Name: The Name_Identifier of a VN_Port.

1.2 FC-BB_E Acronyms

ENode	FCoE Node
FCF	FCoE Forwarder
FCoE_LEP	FCoE Link End-Point
MAC	Media Access Control
VE_Port	Virtual E_Port
VF_Port	Virtual F_Port
VN_Port	Virtual N_Port

1.3 FC-BB_E Overview

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 IEEE 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-2) 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 1 shows an FCoE VN_Port to VF_Port network configuration.

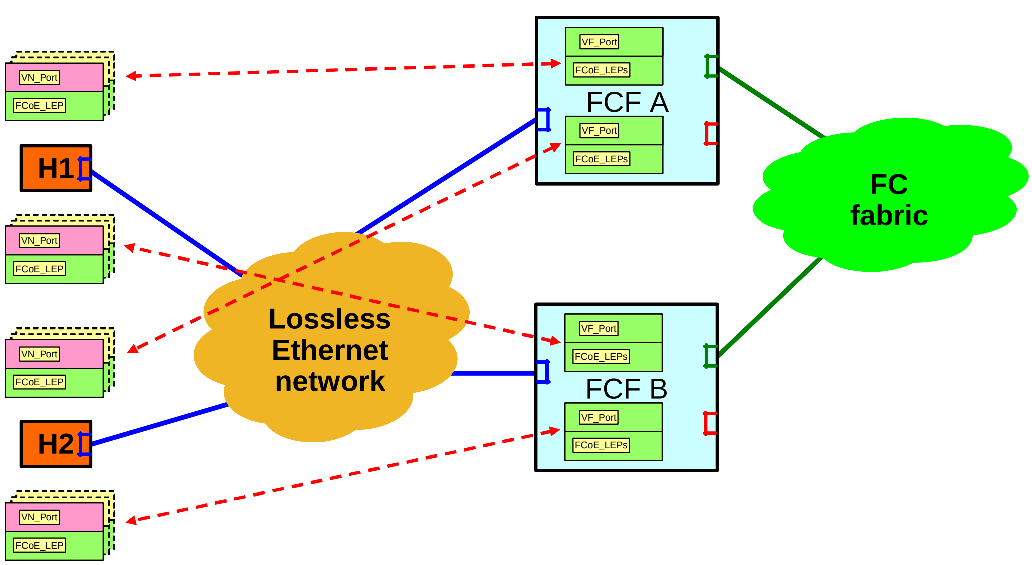

Figure 1 – FCoE VN_Port to VF_Port Network Configuration

Each of the two ENodes H1 and H2 depicted in figure 1 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 1 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 2 shows an FCoE VE_Port to VE_Port network configuration.

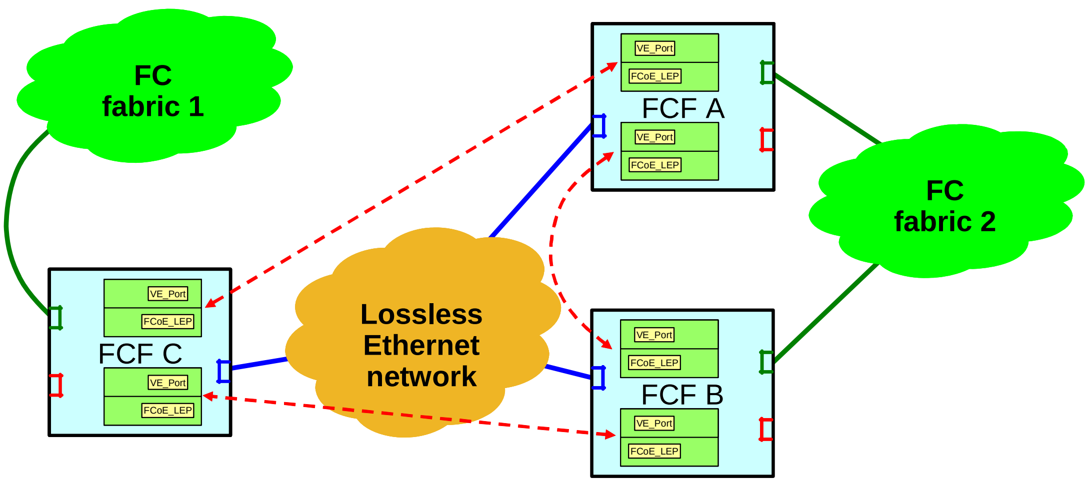

Figure 2 – FCoE VE_Port to VE_Port Network Configuration

Each of the three FCFs A, B, and C depicted in figure 2 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 2 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.

1.4 FCoE VN_Port/ENode Functional Model

Figure 3 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 (ENode's MAC), coupled with an FCoE Controller function.

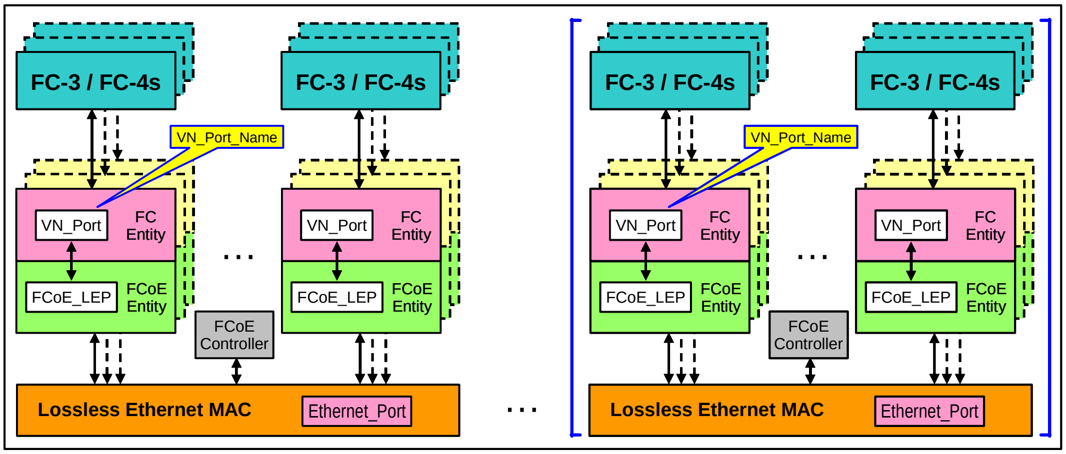

Figure 3 – FCoE VN_Port/ENode Functional Model

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 associated 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.

1.5 FCoE VE_Port/VF_Port Functional Model

Figure 4 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.

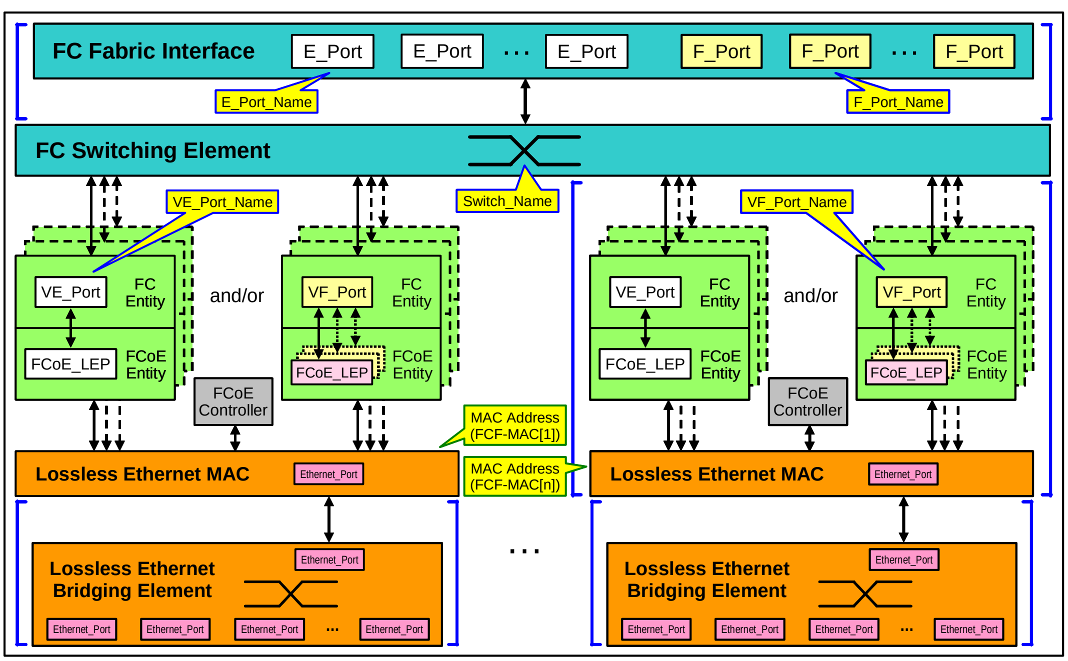

Figure 4 – FCoE VE_Port/VF_Port Functional Model

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-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-2, 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.

1.6 FCoE Virtual Links

Figure 5 shows how the models defined in 1.4 and 1.5 model VN_Port to VF_Port Virtual Links.

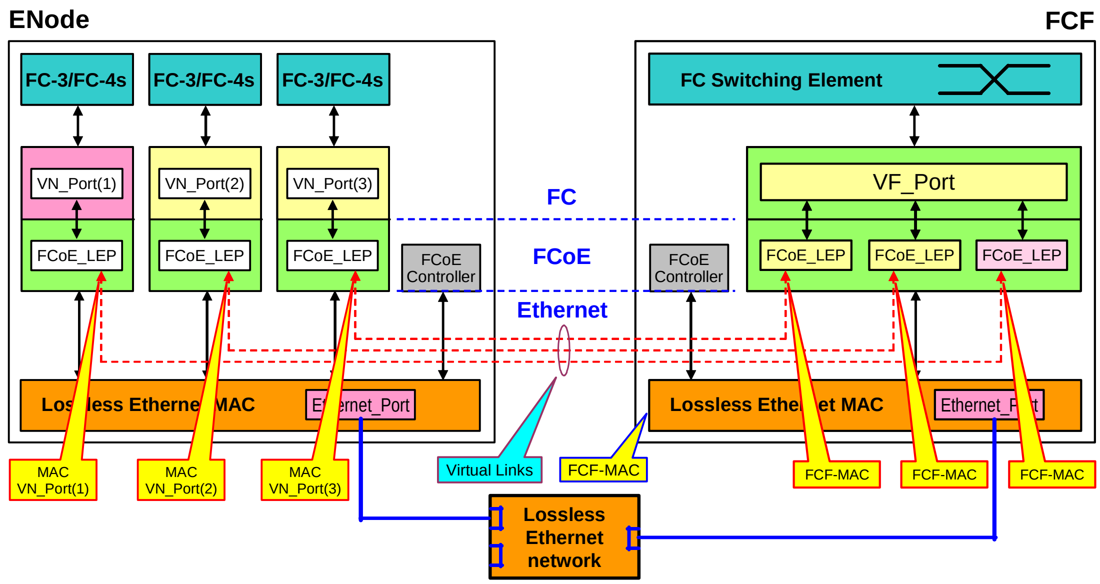

Figure 5 – VN_Port to VF_Port Virtual Links

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 5) 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 5) 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 5) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates an additional FCoE_LEP to the instantiated VF_Port.

Figure 5 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 6 shows how the model defined in 1.5 model VE_Port to VE_Port Virtual Links.

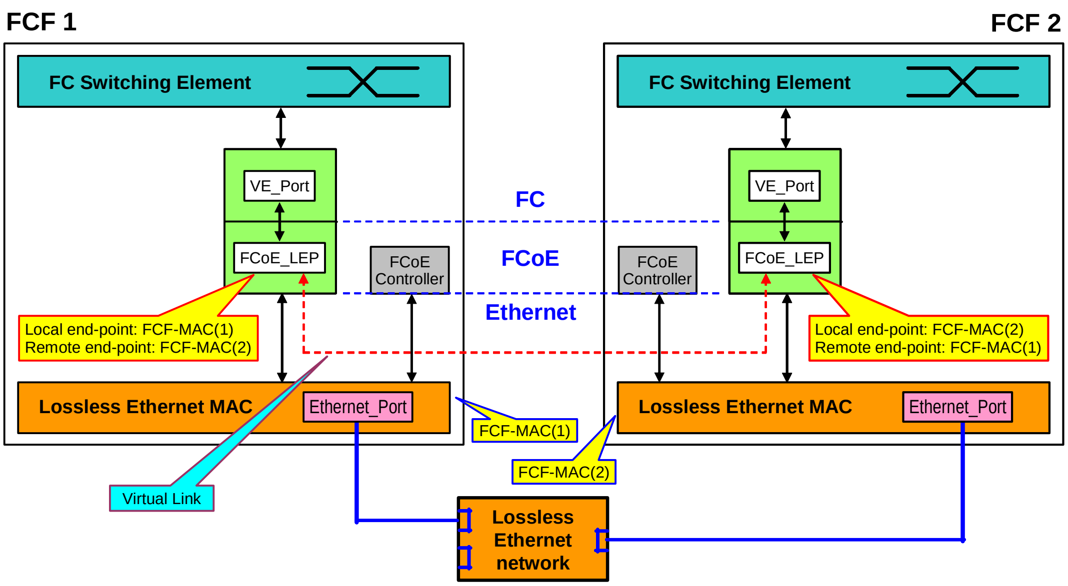

Figure 6 – VE_Port to VE_Port Virtual Links

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 6 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)).