7.12 Distributed FCF

7.12.1 Distributed FCF overview

A Distributed FCF is a set of FDFs associated with at least one Controlling FCF, that controls the operations of the set of FDFs. ~~Figure 45 shows an example of Distributed FCF composed of a Controlling FCF and two FDFs.~~An example of a Distributed FCF with one Controlling FCF and two FDFs is shown in figure 45.

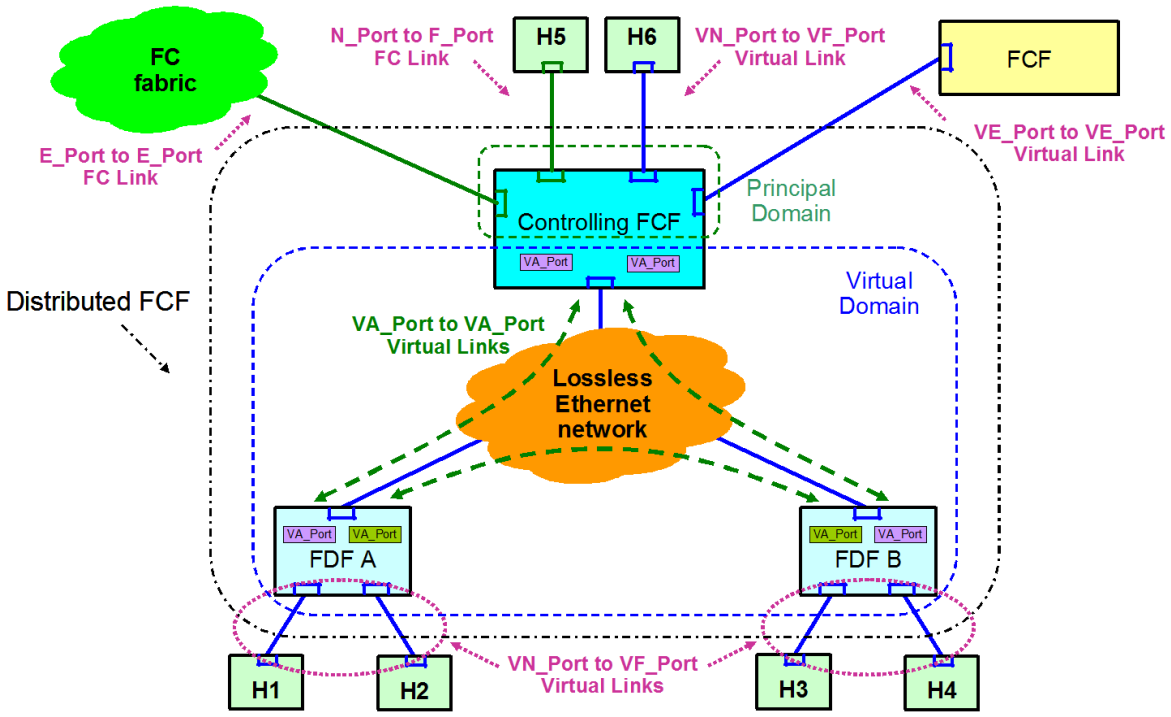

Figure 45 – ~~Example of~~ Distributed FCF ~~-~~example - one Controlling FCFand two FDFs

From an external point of view (i.e., outside the dotted and dashed black line in figure 45), a Distributed FCF behaves as an FCF. In particular, a Distributed FCF supports the instantiation of:

a) VN_Port to VF_Port Virtual Links with ENode MACs; and
b) VE_Port to VE_Port Virtual Links with FCF-MACs.

VN_Port to VF_Port Virtual Links are supported by both FDFs and Controlling FCFs, while VE_Port to VE_Port Virtual Links are supported only by Controlling FCFs. This means that it is possible to connect a Distributed FCF to another FCF or to a native FC fabric only through a Controlling FCF, not through an FDF.

From an internal point of view (i.e., inside the dotted and dashed black line in figure 45), VA_Port to VA_Port Virtual Links enable FCoE frame forwarding between the Controlling FCF and FDFs, as well as between the FDFs. VA_Port to VA_Port Virtual Links are also used to exchange control information between Controlling FCF and FDFs (see FC-SW-6).

The Controlling FCF uses one or more Virtual Domain_IDs to perform N_Port_ID allocations for N_Ports connected to the FDF Set of the Distributed FCF (i.e., a Virtual Domain_ID is used as the most significant byte in the N_Port_IDs allocated to N_Ports that are attached to the FDF Set). The

Controlling FCF uses also another Domain_ID, called the Principal Domain_ID, for its normal functions as an FCF. As a result, a Distributed FCF such as the one shown in figure 45 uses two or more Domain_IDs, one for the Principal Domain, and one or more for the Virtual Domain. To properly support the operations of a Virtual Domain, a Controlling FCF shall have at least one Switch_Name to associate with the Virtual Domain, in addition to its own Switch_Name.

FDFs are not able to operate properly without a Controlling FCF, therefore the Controlling FCF is a single point of failure in a Distributed Switch configuration with only one Controlling Switch, as the one shown in figure 45. To avoid this issue, Distributed FCFs may support a redundant configuration consisting of two or more Controlling FCFs. with one Controlling FCF selected as a Primary Controlling FCF, and one Controlling FCF selected as a Secondary Controlling FCF~~a Primary one and a Secondary one~~. The Secondary Controlling FCF keeps its state synchronized with the Primary and is able to take its place in case of failure according to the Controlling Switch redundancy protocol (see FC-SW-6).

~~Figure 46 shows an example of Distributed FCF including a redundant pair of Controlling FCFs.~~An example of a Distributed FCF with two Controlling FCFs and two FDFs is shown in figure 46.

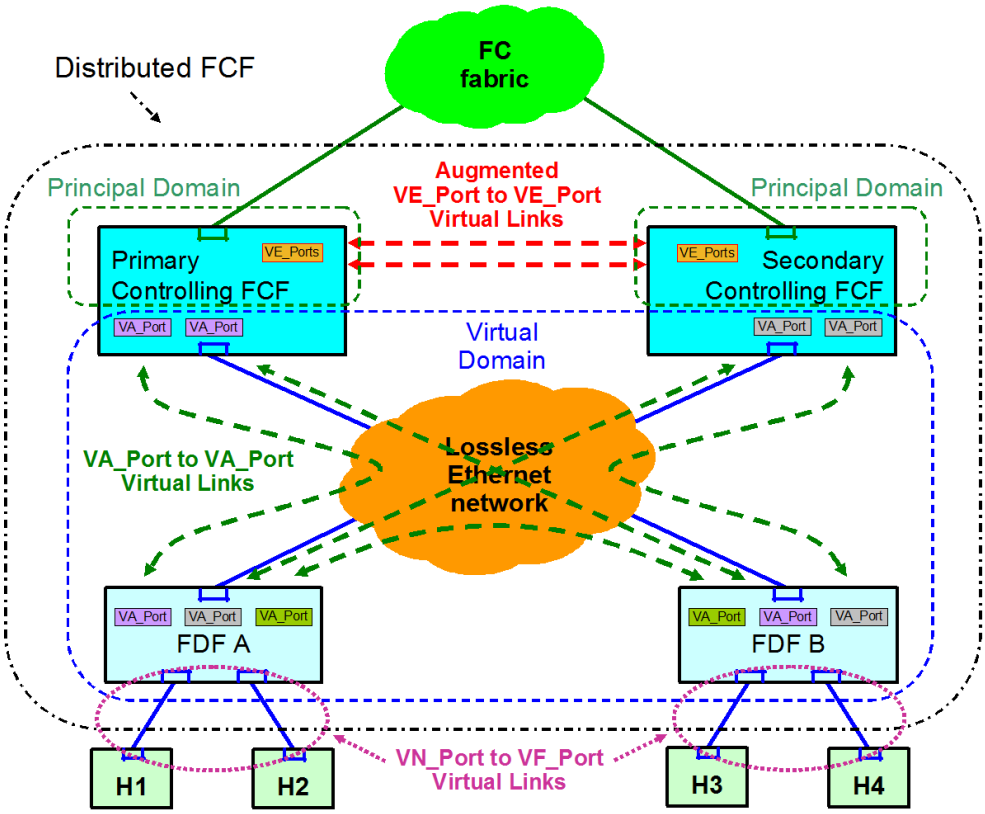

Figure 46 – ~~Example of~~ Redundant Distributed FCF example - two Controlling FCFs and two FDFs

~~The two~~ Controlling FCFs in a redundant Distributed FCF instantiate at least two ~~A~~augmented VE_Port to VE_Port Virtual Links between themselves, where the term 'augmented' indicates that Virtual Link is used also for the Controlling FCF redundancy protocol, in addition to normal VE_Port operation (see FC-SW-6).

The Controlling FCFs use one or more Virtual Domain_IDs to perform N_Port_ID allocations for N_Ports connected to the FDF Set of the Distributed FCF (i.e., a Virtual Domain_ID is used as the most significant byte in the N_Port_IDs allocated to N_Ports that are attached to the FDF Set). Using Virtual Domain_IDs to assign N_Port_IDs enables seamless operation in case of failures ~~of one of the two~~ redundant Controlling FCFs. Each Controlling FCF uses also another Domain_ID, called Principal Domain, for its normal functions as an FCF. As a result, a redundant Distributed FCF typically uses three or more Domain_IDs~~.~~: one for each Controlling FCF, and one for the Virtual Domain_ID. To properly support the operations of a Virtual Domain, a Controlling FCF shall have at least a Switch_Name to associate with the Virtual Domain, in addition to its own Switch_Name.

~~The two r~~Redundant Controlling FCFs instantiate VA_Port to VA_Port Virtual Links to enable the forwarding of FCoE frames and the communication of control information between Controlling FCFs and FDFs. In a redundant configuration, FDFs instantiate VA_Port to VA_Port Virtual Links to each of the Controlling FCFs and between themselves, if they are directly reachable through the Ethernet topology.

Editor's Note: Add text and example figure for n Controlling Switches. Also need the source for figure 46 etc.

A Distributed FCF may have a cascaded FDF configuration when FDFs with at least two VA_Port capable FDF-MACs and independent Lossless Ethernet Bridging Elements are used (see 7.12.3).

Figure 47 shows an example of such a configuration.

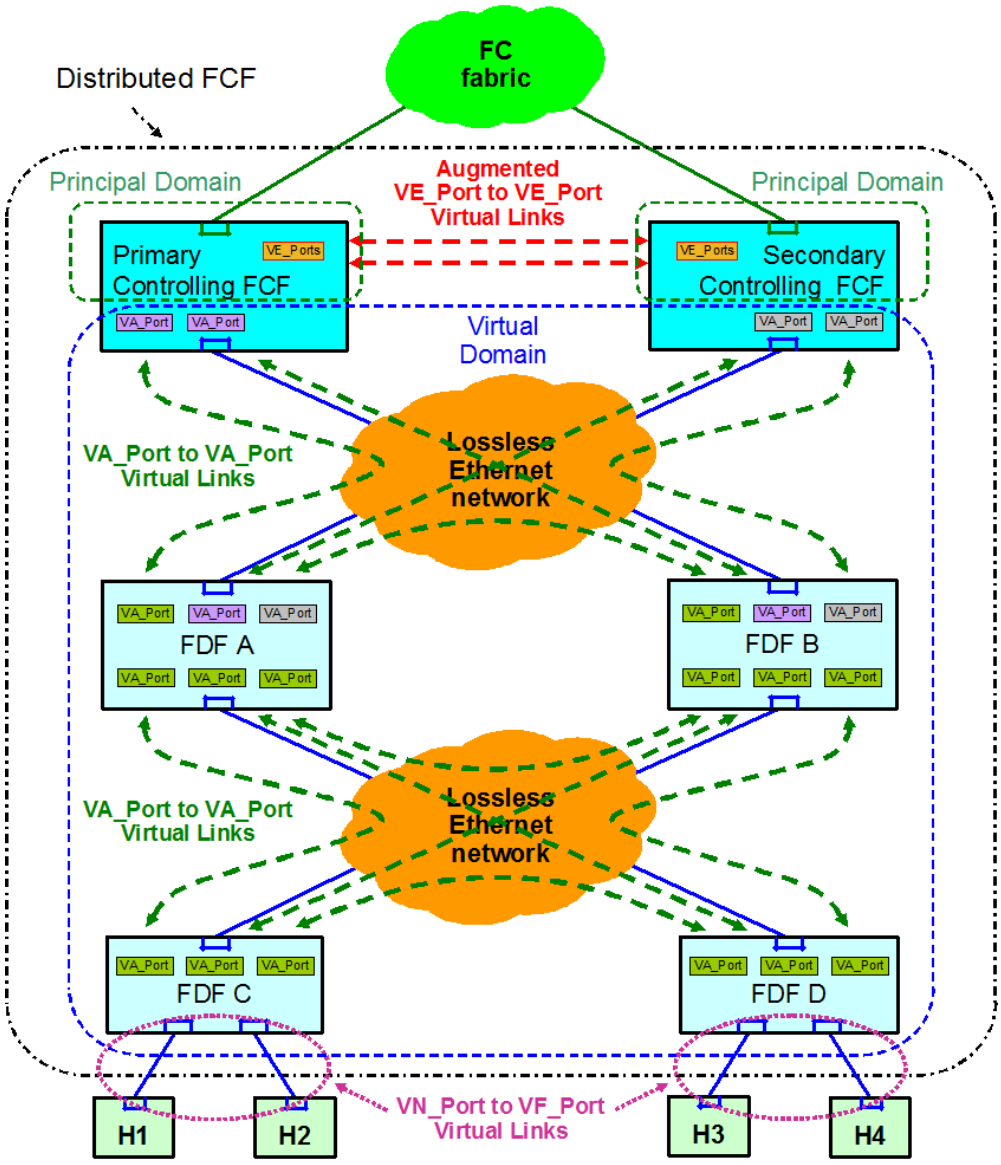

Figure 47 – ~~Example of~~ Distributed FCF with cascaded FDFs example

A Controlling FCF is uniquely identified by its Switch_Name Name_Identifier, as an FCF. An FDF is uniquely identified by its Switch_Name Name_Identifier. A Distributed FCF is defined by an administrative configuration on the Controlling FCFs, listing:

a) the Switch_Names of the ~~two~~ Controlling FCFs that are part of the~~act as the Primary/Secondary pair for that~~ Distributed FCF (i.e., the Controlling FCF Set); and
b) the Switch_Names of the FDFs that are part of that Distributed FCF (i.e., the FDF Set).

7.12.2 Controlling FCF functional model

Figure 48 shows the functional model of a Controlling FCF, where the bracketed functional components are optional. A Controlling FCF is an FCF that supports the instantiation of VA_Ports

over its Lossless Ethernet MACs (FCF-MACs), in addition to VE_Ports and VF_Ports. As for any FCF, each FCF-MAC of a Controlling FCF may be coupled with a Lossless Ethernet bridging element. A Controlling FCF may have a Fibre Channel Fabric interface, providing native E_Port, A_Port, and F_Port connectivity. Supporting the operation of a Controlling Switch (see FC-SW-6), a Controlling FCF is functionally modeled as having two FC Switching Elements, one for the Principal Domain and one for the Virtual Domain, connected by an internal VE_Port to VE_Port link. The Switching Element associated with the Principal Domain supports the instantiation of VF_Ports and VE_Ports, the Switching Element associated with the Virtual Domain supports the instantiation of VA_Ports.

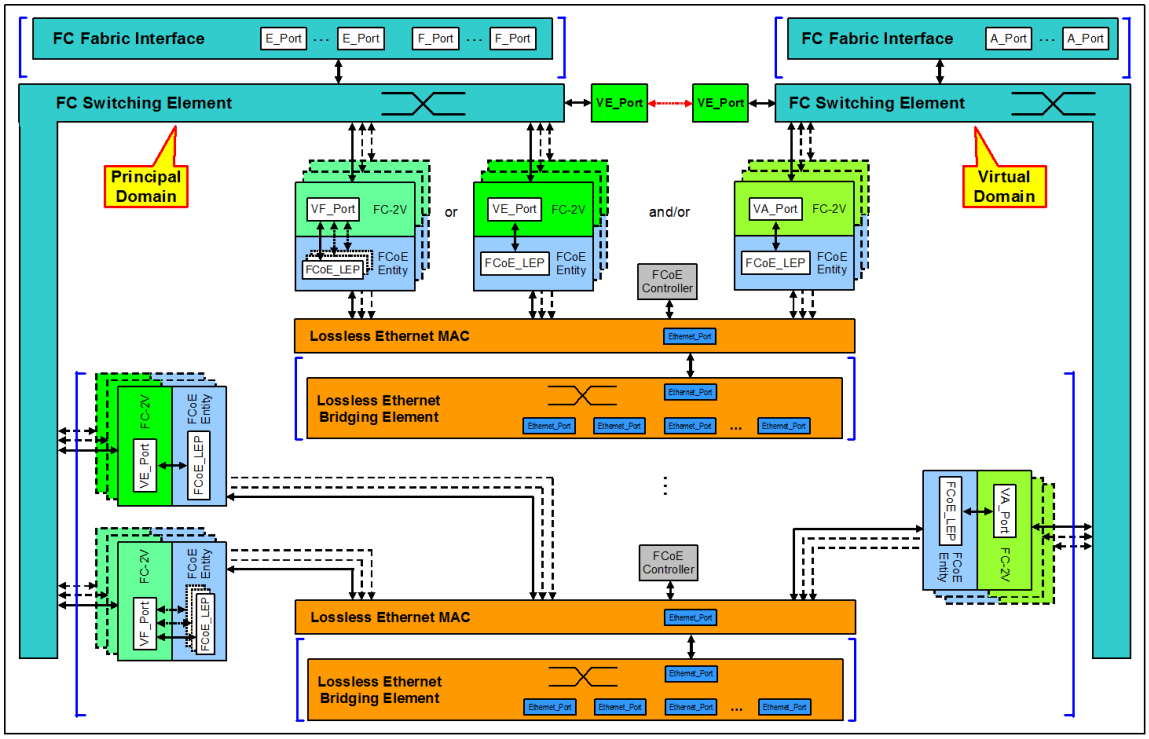

Figure 48 – Controlling FCF ~~F~~functional ~~M~~model

NOTE 13 – Other combinations of Lossless Ethernet bridging elements and Lossless Ethernet MACs connections are allowed.

VA_Ports and VE_Ports may coexist over the same FCF-MAC. VF_Ports coexist with neither VA_Ports nor VE_Ports over the same FCF-MAC. An FCF-MAC supporting the instantiation of VA_Ports and VE_Ports is referred to as a VA_Port/VE_port capable FCF-MAC.

The FCoE Controller is the functional entity that performs the FCoE Initialization Protocol (FIP) and instantiates or de-instantiates VA_Ports, VE_Ports, or VF_Ports, as needed.

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

a) optionally performs the FIP VLAN discovery protocol to discover FCoE VLANs;
b) discovers other VA_Port/VE_Port capable FCF-MACs, VE_Port capable FCF-MACs, and VA_Port capable FDF-MACs connected to the same Lossless Ethernet network using the FIP discovery protocol;

c) instantiates a VE_Port/FCoE_LEP pair on successful completion of each FIP ELP Exchange with a remote VA_Port/VE_Port capable FCF-MAC or a remote VE_Port capable FCF-MAC;
d) de-instantiates a VE_Port/FCoE_LEP pair on receiving a FIP Clear Virtual Link request;
e) instantiates a VA_Port/FCoE_LEP pair on successful completion of each FIP ELP Exchange with a remote VA_Port capable FDF-MAC;
f) de-instantiates a VA_Port/FCoE_LEP pair on receiving a FIP Clear Virtual Link request;
g) monitors the status of the instantiated VE_Port/FCoE_LEP and VA_Port/FCoE_LEP pairs;
h) initiates FIP Clear Virtual Link requests as needed to terminate Virtual Links to other VE_Ports or VA_Ports;
i) transmits periodic FIP Discovery Advertisements to the All-FCF-MACs address every FKA_ADV_PERIOD; and
j) monitors the status of remote VE_Ports and VA_Ports by maintaining timers and verifying that periodic FIP Discovery Advertisements are received within every FKA_ADV_PERIOD.

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 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 shall verify that the source address of the received FCoE frame is equal to the MAC address of the remote link end-point. If either check fails the FCoE frame shall be discarded.

For a VA_Port/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 or FDF-MAC with which a FIP ELP Exchange has been successfully completed.

A VA_Port is an instance of the FC-2V sublevel of Fibre Channel that is dynamically instantiated together with its FCoE_LEP on successful completion of a FIP ELP Exchange, according to the rules specified in 7.12.5.2. A VA_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 VA_Port sends FC frames received from its FCoE_LEP to the FC Switching element. A VA_Port is uniquely identified by an A_Port_Name Name_Identifier and is addressed by the A_Port Controller address identifier (i.e., FFFFF9h).

7.12.3 FDF functional model

An FDF is a simplified FCoE switching entity that forwards FC frames among VA_Ports and VF_Ports through a FCDF Switching Element (see FC-SW-6). Figure 49 shows the functional model of an FDF, where the bracketed functional components are optional. An FDF is functionally composed of an FCDF Switching Element with at least one Lossless Ethernet MAC (FDF-MAC). Each FDF-MAC shall be coupled with an FCoE Controller function. Each FDF-MAC may be coupled with a Lossless Ethernet bridging element. An FDF supports the instantiation of VA_Ports or VF_Ports over its FDF-MACs. The FCDF Switching Element may be coupled with a Fibre Channel

Fabric interface, providing native A_Port and F_Port connectivity. An FDF forwards FCoE frames addressed to one of its FDF-MACs based on the D_ID of the encapsulated FC frames.

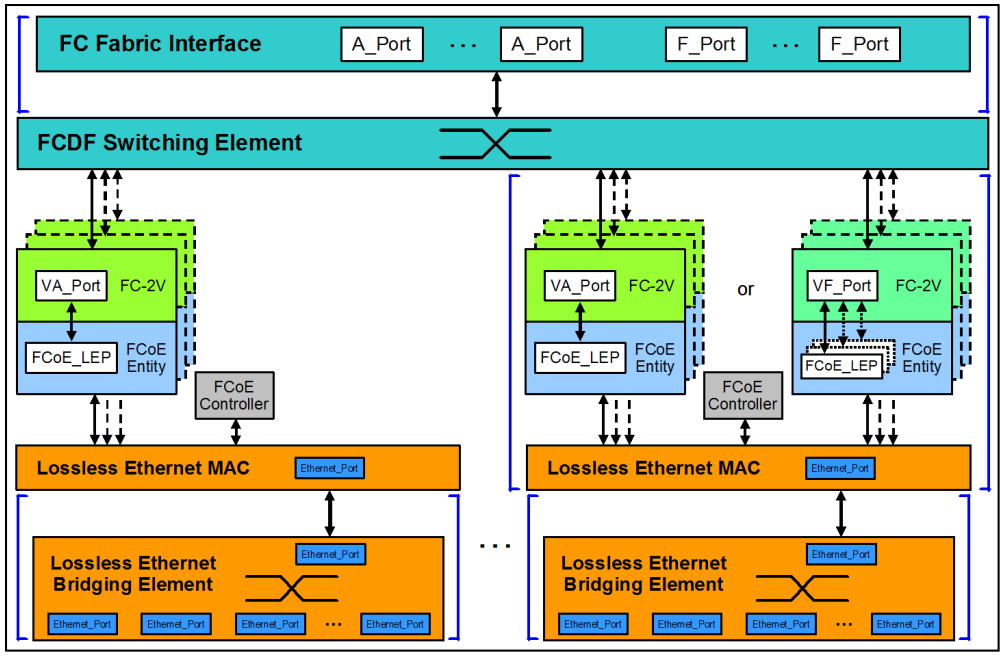

Figure 49 – FDF functional model

NOTE 14 – Other combinations of Lossless Ethernet bridging elements and Lossless Ethernet MACs connections are allowed.

An FDF does not support VE_Ports. VA_Ports and VF_Ports do not coexist over the same FDF-MAC. An FDF-MAC supporting the instantiation of VA_Ports is referred to as a VA_Port capable FDF-MAC. An FDF-MAC supporting the instantiation of VF_Ports is referred to as a VF_Port capable FDF-MAC. A VF_Port capable FDF-MAC behaves as a VF_Port capable FCF-MAC. An FDF shall support at least one VA_Port capable FDF-MAC and may support one or more VF_Port capable FDF-MAC.

The FCoE Controller is the functional entity that performs the FCoE Initialization Protocol (FIP) and instantiates or de-instantiates VA_Ports or VF_Ports, as needed.

For a VA_Port capable FDF-MAC, the FCoE Controller:

a) optionally performs the FIP VLAN discovery protocol to discover FCoE VLANs;
b) discovers other VA_Port capable FDF-MACs and VA_Port/VE_Port capable FCF-MACs connected to the same Lossless Ethernet network using the FIP discovery protocol;
c) instantiates a VA_Port/FCoE_LEP pair on successful completion of each FIP ELP Exchange with a remote VA_Port capable FDF-MAC or VA_Port/VE_Port capable FCF-MAC;
d) de-instantiates a VA_Port/FCoE_LEP pair on receiving a FIP Clear Virtual Link request;
e) monitors the status of the instantiated VA_Port/FCoE_LEP pairs;
f) initiates FIP Clear Virtual Link requests as needed to terminate Virtual Links to other VA_Ports;
g) transmits periodic FIP Discovery Advertisements to the All-FCF-MACs address every FKA_ADV_PERIOD; and
h) monitors the status of remote VA_Ports by maintaining timers and verifying that periodic FIP Discovery Advertisements are received within every FKA_ADV_PERIOD.

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 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 shall verify that the source address of the received FCoE frame is equal to the MAC address of the remote link end-point. If either check fails the FCoE frame shall be discarded.

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

A VA_Port is an instance of the FC-2V sublevel of Fibre Channel that is dynamically instantiated together with its FCoE_LEP on successful completion of a FIP ELP Exchange, according to the rules specified in 7.12.5.2. A VA_Port receives FC frames from the FCDF Switching Element and sends them to its FCoE_LEP for encapsulation and transmission over the Lossless Ethernet network. In a similar way, a VA_Port sends FC frames received from its FCoE_LEP to the FCDF Switching element. A VA_Port is uniquely identified by an A_Port_Name Name_Identifier and is addressed by the A_Port Controller address identifier (i.e., FFFFF9h).

7.12.4 VA_Port to VA_Port Virtual Links

Figure 50 shows how the functional models defined in 7.12.2 and 7.12.3 model a VA_Port to VA_Port Virtual Link between two FDFs.

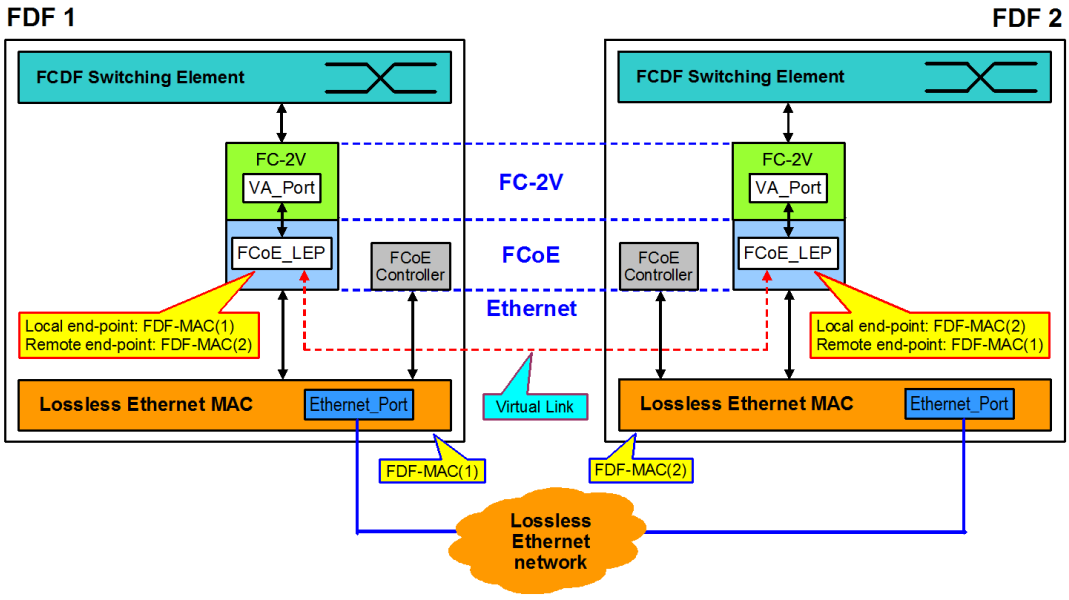

Figure 50 – VA_Port to VA_Port Virtual Link ~~E~~example

On successful completion of a FIP ELP Exchange, the FCoE Controllers of the two involved VA_Port capable FDF-MACs instantiate a VA_Port/FCoE_LEP pair. Figure 50 shows the Virtual Link end-points, that are the MAC addresses of the two involved VA_Port capable FDF-MACs (i.e., FDF-MAC(1) and FDF-MAC(2)).

7.12.5 FIP support

7.12.5.1 FIP discovery

Controlling FCFs and FDFs behave as FCFs from a point of view of FIP discovery, therefore they discover each other through the FIP discovery protocol, as specified for FCF to FCF discovery (see 7.9.3.3). Two bits are specified in the encapsulated FIP operation (see 7.9.7.2) to indicate if the originator of a FIP frame is a Controlling FCF or an FDF.

An FDF is part of a Distributed FCF internal topology if the initialization exchanges with the Primary controlling FCF are completed. See FC-SW-6.

If an FDF is not part of a Distributed FCF internal topology:

a) all VA_Port capable FDF-MACs on that FDF shall transmit Discovery Advertisements with the Fabric_Name field of the Fabric Descriptor set to zero; and
b) all VF_Port capable FDF-MACs on that FDF shall not transmit Discovery Advertisements.

If an FDF is part of a Distributed FCF internal topology, all VA_Port capable and VF_Port capable FDF-MACs on that FDF shall have the Fabric_Name received from the Primary Controlling FCF in the Fabric_Name field of the Fabric Descriptor in all transmitted Discovery Advertisements.

7.12.5.2 FCoE Virtual Link instantiation

Controlling FCFs and FDFs establish FCoE Virtual Links between themselves on successful completion of FIP ELP Exchanges. Two bits in the ELP payload indicate if the originator of the FIP ELP Request or SW_ACC is a Controlling FCF or an FDF.

Bit 13 of the Flags field of the ELP payload is the Controlling FCF/Switch bit. This bit set to one indicates that the originator of the FIP ELP Request or SW_ACC is a VA_Port/VE_Port capable FCF-MAC. This bit set to zero indicates that the originator of the FIP ELP Request or SW_ACC is not a VA_Port/VE_Port capable FCF-MAC.

Bit 12 of the Flags field of the ELP payload is the FDF/FCDF bit. This bit set to one indicates that the originator of the FIP ELP Request or SW_ACC is a VA_Port capable FDF-MAC. This bit set to zero indicates that the originator of the FIP ELP Request or SW_ACC is not a VA_Port capable FDF-MAC.

A received FIP ELP Request or SW_ACC having both these bits set to one is invalid and shall be ignored.

Editor's Note: Receiving a FIP ELP Request or SW_ACC with both bits set is an error. What is the behavior if ignored?

A VA_Port/VE_Port capable FCF-MAC instantiates VE_Port to VE_Port Virtual Links with other FCF-MACs through a FIP ELP Exchange. VE_Port to VE_Port Virtual Links are instantiated when the FDF/FCDF bit is set to zero in both FIP ELP Request and FIP ELP SW_ACC.

A VA_Port/VE_Port capable FCF-MAC shall instantiate VE_Port to VE_Port Virtual Links with other VA_Port/VE_Port capable FCF-MACs and VE_Port capable FCF-MACs discovered by FIP discovery on the Lossless Ethernet network. VE_Port to VE_Port Virtual Links with VA_Port/VE_Port capable FCF-MACs belonging to the other Controlling FCF of the Primary/Secondary pair defining the Distributed FCF are used for the redundancy protocol of the Distributed FCF (see FC-SW-6). VE_Port to VE_Port Virtual Links are established according to the normal ELP rules (see FC-SW-6).

A VA_Port/VE_Port capable FCF-MAC instantiates VA_Port to VA_Port Virtual Links with VA_Port capable FDF-MACs through a FIP ELP Exchange. These VA_Port to VA_Port Virtual Links are instantiated when the FDF/FCDF bit is set to one in the FIP ELP SW_ACC. A VA_Port/VE_Port capable FCF-MAC shall reject a received FIP ELP Request with the FDF/FCDF bit set to one with Reason Code 'Protocol Error' and Reason Code Explanation 'Invalid Request'.

When operational (i.e., when the Controlling FCF is in state P2 or S2 of the Controlling Switch redundancy protocol, see FC-SW-6), a VA_Port/VE_Port capable FCF-MAC shall instantiate VA_Port to VA_Port Virtual Links with VA_Port capable FDF-MACs belonging to FDFs that are part of the Distributed FCF's FDF Set and discovered by FIP discovery on the Lossless Ethernet network. In this case, the FIP ELP Exchange shall be initiated by the VA_Port/VE_Port capable FCF-MAC and the VA_Port capable FDF-MAC shall process it irrespective of the value of the Switch_Name field in the ELP payload (i.e., acceptance or rejection shall be based on the other ELP parameters, not on the involved Switch_Names). A VA_Port/VE_Port capable FCF-MAC shall not establish VA_Port to VA_Port Virtual Links with VA_Port capable FDF-MACs belonging to FDFs that are not part of the Distributed FCF's FDF Set. A VA_Port capable FDF-MAC shall not initiate a FIP ELP Exchange with a VA_Port/VE_Port capable FCF-MAC.

Editor's Note: Reword the text in red above once the Controlling Switch redundancy protocol is finalized.

A VA_Port capable FDF-MAC instantiates VA_Port to VA_Port Virtual Links with other FDF-MACs through a FIP ELP Exchange. These VA_Port to VA_Port Virtual Links are instantiated when the FDF/FCDF bit is set to one in both FIP ELP Request and FIP ELP SW_ACC.

A VA_Port capable FDF-MAC shall initiate a FIP ELP Exchange with a discovered VA_Port capable FDF-MAC only if:

a) it has already at least a VA_Port to VA_Port Virtual Link with the Primary Controlling FCF or another FDF;
b) it has received the Distributed FCF's FDF Set through the DFMD SW_ILS (see FC-SW-6) from the Primary Controlling FCF; and
c) the discovered FDF-MAC belongs to an FDF in the Distributed FCF's FDF Set.

A VA_Port capable FDF-MAC shall not initiate a FIP ELP Exchange with a discovered VA_Port capable FDF-MAC if the discovered FDF-MAC does not belong to an FDF in the FDF Set. A VA_Port capable FDF-MAC that has not initiated a FIP ELP Exchange shall reply to a received FIP ELP Request, irrespective of the value of the Switch_Name field in the ELP payload. A VA_Port capable FDF-MAC that has initiated a FIP ELP Exchange (i.e., sent a FIP ELP Request) shall reply to a received FIP ELP Request according to the normal ELP rules (i.e., acceptance or rejection includes considering the involved Switch_Names).

NOTE 15 – These rules enable an ordered establishment of VA_Port to VA_Port Virtual Links from the Controlling FCF(s) to the peripheral FDFs in a Distributed FCF with cascaded FDFs.

An FDF does not establish VE_Port to VE_Port Virtual Links, therefore an FDF-MAC shall reject a received FIP ELP Request with both Controlling FCF/Switch bit and FDF/FCDF bit set to zero (i.e., a FIP ELP Request coming from an FCF that is not a Controlling FCF) with Reason Code 'Protocol Error' and Reason Code Explanation 'Invalid Request'. An FDF-MAC shall also reject a received FIP ELP Request coming from a Controlling FCF other than the Controlling FCFs that define its Distributed FCF, with Reason Code 'Logical Error' and Reason Code Explanation 'Not Authorized'.

7.12.5.3 FCoE Virtual Link maintenance

VA_Port to VA_Port Virtual Link maintenance is performed as for VE_Port to VE_Port Virtual Links.

7.12.6 Distributed FCF operations

A Distributed FCF operates as a Distributed Switch (see FC-SW-6), with the only difference that FIP discovery is performed before instantiating Virtual Links. This enables Controlling FCFs and FDFs to instantiate VA_Port to VA_Port Virtual Links only with the discovered FDFs that are directly reachable and part of the FDF Set (i.e., VA_Port to VA_Port Virtual Links are not established with FDFs that are not part of the FDF Set, see 7.12.5.2).