

Proxy Based Shortcut

Direct FCoE Data Transfer End to End

John L. Hufferd

10/7/2009

This paper describes how FCoE (Fibre Channel over Ethernet) ULP (Upper Layer Protocol) frames can be transferred from one FCoE network adapter to another (through one or more "Proxy" Capable Switches) without having to traverse through a Fiber Channel Forwarder (FCF) device. After the logical End-To-End connection is established via an FCF (Fibre Channel Forwarder), FCoE ULP messages originating at an FCoE network adapter can be diverted by one or more FSP (FCoE Shortcut Proxy) Switches and sent to an FCoE receiving adapter across "Ethernet" links and switches, but without having to pass through FCF devices.

Proxy Based Shortcut

(Direct FCoE Data Transfer End to End)

FCoE has been proposed for standardization by the T11 committee (called herein, the **FC-BB-5 Standard**). It defines a way for a **FCP** (Fibre Channel Protocol) capable switch to connect to a “**Lossless Ethernet**” (aka: **CEE** -- Converged Enhanced Ethernet, or **DCB** Data Center Bridging) network and permit **FCoE** (Fibre Channel over Ethernet) frames to travel from the Host adapter called a **CNA** (Converged Network Adapter) to a special switch (called an **FCF** – FC Forwarder) and be sent on to a target/peer device. This is accomplished by encapsulating the FCP packets into Ethernet Frames and sending them on a Lossless Ethernet Network. The FCP protocol is made up of both link and device management messages and **ULP** (Upper Layer Protocol) messages (e.g. SCSI Commands and Data, etc.).

This FCoE protocol also contains a sub-protocol called **FIP** (FCoE Initiation Protocol) which performs the discovery of FCoE devices and FCF ports in the Fabric. This FIP protocol has also been made extensible by having specific ways to include new FIP operations descriptor codes. And backward compatibility is also possible because FIP processing requires a receiver of unknown FIP descriptors to discard and ignore those descriptors in a FIP message (as long as they are in the "non critical" range). Also it continues the Standards' practice of ignoring any value set into a reserved field. (See the "New FIP Messages Compatibility" section below.)

Therefore, it is possible to make product extensions to the FCoE standard while keeping compatibility with current implementations. Likewise it is possible to create a backward compatible follow-on standard to support new functions without disruption to existing implementations.

There has also been a proposal called "**Adapter Based Shortcut**" which permits an FCoE environment to have central FCF processes (like that found in an FCF) to handle only the FLOGIs, PLOGIs, etc., and have "Shortcut" enabled CNA peers send FCoE ULP messages directly to each other in a bidirectional manner. That is, the goal of that proposal was to permit the ULP's commands, data, responses, etc. to flow, directly between the Source and the Target Peer CNAs (VN_Port to VN_Port), "Shortcutting" around the FCFs and leaving the FCFs to only perform connection processing, **Zoning** (isolation of port groups into administrative and data flow domains), etc.

This "Proxy Based Shortcut" proposal builds upon the "Adapter Based Shortcut" proposal and it will be assumed (herein) that the reader of this document is familiar with the "Adapter Based Shortcut" document. (We will sometimes refer (herein) to the "Adapter Based Shortcut" document as the "**ABS document**").

The "Adapter Based Shortcut" proposal described how NEW "Shortcut Enabled" CNAs would operate. However, there is a need to have a way that existing CNAs can participate in the shortcut process. To do that and not increase the path length, there needs to be a new type of Lossless Ethernet Switch that also contains the Shortcut capabilities and can act as a Proxy for the non Shortcut Enabled CNAs.

Proposal in a Nut Shell:

This proposal permits an FCoE environment to have central FCF processes (like that found in an FCF) to handle only the FLOGIs, PLOGIs, etc. Then a “Shortcut” enabled Ethernet Switch can act as a Proxy for non Shortcut enabled CNA peers by sending FCoE ULP messages directly to other such switches or Shortcut Enabled CNAs. This type of Switch is called (herein) an FCoE Shortcut Proxy (FSP). The FSP will send/receive the FCoE ULP messages in a bidirectional manner. That is, the goal of this proposal is to permit the ULP’s commands, data, responses, etc. (with the help of FSPs) to flow between the Source and the Target Peer CNAs (VN_Port to VN_Port); “Shortcutting” around the FCFs and leaving the FCFs to only perform connection processing, Zoning (isolation of port groups into administrative and data flow domains), etc.

This idea of a bidirectional “Shortcut” process will only work when the Source and the Target Peer CNAs are located within the same Lossless Ethernet Network Segment. If it is necessary for two or more Network Segments to be connected or if a real FC connection is needed then a fully functional FCF will be required to interconnect the fabrics and manage the ULP data flow.

In Fig. 1a below, the link L11 would be the Shortcut path. In Fig. 1b below, the Shortcut path is an internal connection within FSP Ethernet Switch A.

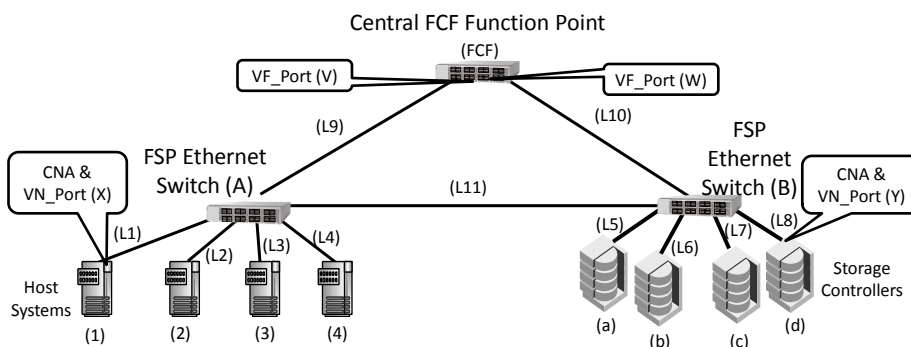


Fig. 1a – A FCoE Network Configuration with FSPs Example

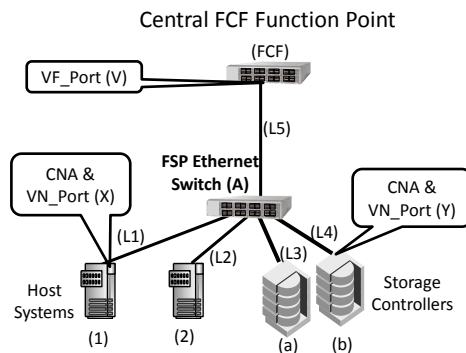


Fig. 1b – A FCoE Network Configuration with an FSP Example

Compatibility

It is a necessary requirement that the existing FCF Functions continue to operate without changes, even if shortcut capable CNAs or FSPs are present and operating. And it is also a necessary requirement that new Shortcut capable CNAs be able to operate with the new FSPs which are acting as proxies for existing FC-BB-5 Standards Base Peer CNAs that do not have this new Shortcut feature. These FSP should also work with non Shortcut Enabled CNAs that do not have the services of an FSP. In that case, the ULP messages will flow through the FCF as is the case today.

The FIP type messages that were defined in the "Adapter Based Shortcut" (ABS) document are the same ones used for this proposal. And there are no new compatibility issues with those FIP messages. (There will, however, be one new OPTIONAL FIP message that will be proposed in case updated FCFs want to be involved.)

The Dynamic and Static ACLs remain the same as in the "Adapter Based Shortcut" proposal.

The Proxy Concept

In this concept of "Shortcut", the VN_Port of one non Shortcut enabled CNA can send data directly to a VN_Port of another non Shortcut enabled CNA (or a Shortcut enabled CNA) by using a Proxy Ethernet Switch (an FSP) after the Logical Connection has been fully established via FC-BB-5 standards compliant FCFs.

The high level steps in this proposal are as follows:

1. The CNA peers will perform FLOGIs/FLOGI ACCs and PLOGIs/PLOGI ACCs as normal for an FCoE implementation; then the FSP will detect a PLOGI ACCs, being sent by a CNA/VN_Port toward an FCF, or being sent by an FCF toward a CNA/VN_Port. Information will then be extracted from those PLOGI ACCs messages as they flow through the FSP such that the MAC Address of the Logical Connection End point CNA/VN_Ports can be determined and saved.
2. The FSP should ensure that its internal Ethernet switch has the information/opportunity to get security processes/techniques in place (such as dynamic ACEs/ACLs) in order to protect the "Shortcut" Path.
3. The FSP will not send the "End-To-End_Connection_Established" FIP message, as it will only be sent by a Shortcut Enabled CNA and will not be sent by the FSP for any CNA/VN_Port for which it is providing Proxy services. But, the FSP may detect/alert on that FIP message being sent from a Shortcut enabled CNA/VN_Port and set up Switch ACEs/ACLs for the CNAs' Shortcuts.
4. The Shortcut enabled device peers (Shortcut Proxies or Shortcut enabled CNAs) will use the same Shortcut Path Establishment Protocol, specified in the ABS document, that was used between Shortcut enable CNAs. This Protocol, validates that a path exist between the Shortcut enabled device peers (which does not pass through an FCF), and that the path can carry jumbo Ethernet Frames of the required size.
5. ULP messages (e.g. commands, data, responses, etc.) will then flow on the Shortcut path directly between the Shortcut enabled device peers without traveling through an FCF.
6. The FIP Keep Alive functions and the FIP Advertisements messages will continue for the Virtual links between ENodes and the FCFs as is currently the case for any implementation consistent with the FC-BB-5 standard. In addition FIP Keep Alive messages will be sent periodically on the Shortcut path between the Shortcut enabled peers to ensure that their logical links' activity states are known. This will be done in a manner that is compatible with that defined in the ABS document.
7. End-To-End logical connections may be terminated via the normal FCoE FC-BB-5 Standardized processes (e.g. Fabric Logoff -- Fabric LOGO, or Clear Virtual Link FIP messages).
8. At logical connection termination, the dynamically built Ethernet Switch ACEs, if any, should be removed from the corresponding FSP Ethernet Switch ingress port's ACL (if any), in a manner suggested in the FC-BB-5 standard, plus any ACEs created for the Shortcut path as a result of step 2, or step 3 above, should be removed.

Zoning

The Zoning considerations are the same as those shown in the "ABS document".

Keep Alives

The non Shortcut enabled CNAs will need to ensure that the path to the FCF is kept alive, and the FSP needs to ensure that the Shortcut path to the remote VN_Port Peer is also kept alive. Therefore, Keep Alive FIP frames need to be sent to the FCF's VF_Ports as is normal in non Shortcut environments. That

is, sent from the CNA's FCoE Controller to the FCF ports, on behalf of the various instantiated VN_Ports. And the FCF's Advertisement also needs to be sent to the CNA's FCoE Controller as is normal with the current FCoE FC-BB-5 Specification.

But in addition, Keep Alive FIP frames also need to be sent/received on the Shortcut path between FSPs and their peer Shortcut enabled devices (FSPs or Shortcut Enabled CNAs) .

When an FSP is connected to a Shortcut enabled CNA the CNA's ENode FCoE Controller will be able to detect a link failure in the shortcut paths as well as the paths to the FCFs. When a link failure is detected in the Shortcut path by a Shortcut enabled CNA, it will be able to de-instantiate the corresponding VN_Port, and the FCF will be sent a Logout (LOGO) to terminate the connection.

However, when the link failure is detected by an FSP, it can only shut down the flow of messages coming from the local VN_Port that is headed to the related VF_Port. Then the FSP will rely upon the Keep Alive failure detection capabilities of the FCF to detect the problem and shut down the connection.

Details of the Proxy Concept

The following is the more detailed description of the processes for this Shortcut enabled CNA Concept:

1. At the end of the FC-BB-5 standardized FCoE process of Fabric and Port Logins, a logical End-To-End connection is established, through the FCF, from one CNA VN_Port to another CNA VN_Port.

That is, the CNA VN_Port peers will perform FLOGIs/FLOGI ACCs and PLOGIs/PLOGI ACCs as normal for an FCoE FC-BB-5 implementation (but including the FLOGI Shortcut Capable informational Flag -- the "C" Flag -- detailed in the ABS document); thus establishing a CNA to CNA logical connection through one or more FCFs.

The FSP will be able to notice the "C" Flag on the FLOGI FIP Message and then setup to inspect/detect any PLOGI ACCs not coming from or headed to that port. (The Shortcut Capable CNA that sent that FLOGI FIP message, with the "C" Flag, will extract the PLOGI information on its own and issue the "End-To-End_Connection_Established FIP message without any assistance by the FSP.)

The FSP will be able to notice/detect the absence of the "C" flag on the FLOGI FIP Message headed toward the FCF and will add the "C" Flag as the Frame is sent on to the FCF.

Then when the PLOGI ACC FCoE frame is seen/detected at the FSP port which is connected to a non Shortcut enabled CNA, the PLOGI ACC FCoE frame will either be coming from a CNA headed toward an FCF or is coming from an FCF headed to a CNA.

When a FSP detects that a PLOGI ACC is being received (from the FCF or from the CNA), the FSP will save and build the appropriate MAC Addresses of the VN_Port Peers. The source of

the information is the DA or SA that contains an FC-MAP, and the S_ID or D_ID which are located in the FC header of the arriving FCoE PLOGI ACC packet. This is done as follows:

- a. If the Destination (DA) MAC Address of the arriving FCoE PLOGI ACC frame contains the FC-MAP then the frame is arriving from an FCF and the Destination (DA) MAC Address needs to be saved as the FPMA of the local CNA/VN_Port peer, and the sending (remote) CNA/VN_Port peer's FPMA needs to be built (and saved) by concatenating the FC-MAP and the S_ID. This also means that the (Local) CNA/VN_Port peer at the Destination (DA) MAC Address is what is called (herein) the "ACTIVE" CNA/VN_Port peer, and the (Remote) CNA/VN_Port peer (whose FPMA was built by the FSP) is what is called (herein) the "PASSIVE" CNA/VN_Port peer. (The "ACTIVE" CNA/VN_Port peer sent the PLOGI Request, and the "PASSIVE" CNA/VN_Port peer sent the PLOGI ACC.)
- b. If the Source (SA) MAC Address of the arriving FCoE PLOGI ACC frame contains the FC-MAP then the frame is being sent to an FCF and the Source (SA) MAC Address needs to be saved as the FPMA of the local peer's CNA/VN_Port and the Destination (remote) CNA/VN_Port peer's FPMA needs to be built (and saved) by concatenating the FC-MAP and the D_ID. This also means that the (Local) CNA/VN_Port peer at the Source (SA) MAC Address is what is called (herein) the "PASSIVE" CNA/VN_Port peer, and the (Remote) CNA/VN_Port peer (whose FPMA was built and saved within the FSP) is what is called (herein), the "ACTIVE" CNA/VN_Port peer. (The "ACTIVE" CNA/VN_Port peer sent the PLOGI Request, and the "PASSIVE" CNA/VN_Port peer sent the PLOGI ACC.)

The Remote peer's CNA/VN_Port FPMA can then be used as a Destination (DA) MAC Address for FCoE ULP and FIP frames that the local FSP will send on the Shortcut path to the Remote CNA/VN_Port Peer.

It is also possible that by implementation approach, by Standardized Specification, by a new Flag in the FLOGI ACC FIP message, or by instructions from an Administrator -- the Proxy will anticipate and use the optional "End-to End_Connection_Established Advertisement FIP message instead of the PLOGI ACC to obtain the information about the VN_Port peers' MAC Addresses (refer to the "New Optional FIP Message" section below).

2. Sometime after the PLOGI ACC -- (which establishes the logical End-To-End connection) and before any ULP messages have been sent, the FSP should ensure that its internal Ethernet switch has the information/opportunity to get security processes/techniques in place (such as dynamic ACEs/ACLs) in order to obtain adequate protection for the "Shortcut" Path.

3. The Logical Link's (FCF) VF_Port will not be sent any "End-To-End_Connection_Established" FIP message originated from any FSP since those messages will only be sent by a Shortcut Enabled CNA and will not be sent by the FSP for any CNA/VN_Port for which it is providing Proxy services. Therefore, any new FCF functions that attempt to be Shortcut aware will need to be able to operate correctly without this FIP message.

On the other hand the FSP will need to be able to inspect/detect on any "End-To-End_Connection_Established" FIP message, passing through the FSP Switch to the FCF, that may originate with a Shortcut enabled CNA/VN_Port; and the FSP should be able to put in place additional dynamic ACEs/ACLs for the FSP Port which connects to that Shortcut enabled CNA/VN_Port -- as is suggested by the FC-BB-5 standard. (The assumption is that the Administrator had permitted the FSP's to previously perform inspection/detection of the FLOGI ACC FIP message sent from the FCF's VF_Port, to this Shortcut Enabled CNA/VN_Port, and build dynamic ACEs/ACLs in a manner suggested by the FC-BB-5 standard.)

It is also possible that a Shortcut aware FCF might optionally send the "End-To-End_Connection_Established" Advertisement FIP message (refer to the "New Optional FIP Message" section below) to the Shortcut enabled and non Shortcut enabled CNA/VN_Ports; in which case the FSP should inspect/detect on that message and use information, therein, to build the additional dynamic ACEs/ACLs to protect that CNA/VN_Port's Shortcut path instead of extracting the information from a PLOGI ACC.

It should be understood that it is not required that FSPs have the ability to dynamically build or modify their ACEs/ACLs. In fact various Administrators may not want their FSPs to do that, or they may want to statically build the ACEs/ACLs themselves. (Refer to the "Static ACLs" section below.)

4. Following after the above process steps a bidirectional Shortcut path, between the CNA/VN_Port peers, will be established by having FSPs exchanging the "Shortcut_Path_Verification" FIP message and the "Shortcut_Path_Verification" Response FIP message with remote FSPs (or with Shortcut enabled CNA/VN_Ports). These Shortcut_Path_Verification (and Response) FIP messages (see the "New FIP Messages" section below) sent from the FSPs will be formatted and used in the same manner as when they are used just between Shortcut enabled CNAs (refer to the Appendix within the ABS document and the Appendix within this document).

Referring to the "New FIP Messages" section below, and the "New FIP Messages" section within the ABS document, the "Shortcut_Path_Verification" FIP message and the "Shortcut_Path_Verification" Response FIP message are padded out to be the same size as the largest sized FCoE frames (intended to ensure that the shortcut bidirectional path not only exists but that it also has the appropriate jumbo frame capability since if they do not,

the FCoE frames will be dropped by the switches). These FIP messages will be sent to the Destination (DA) MAC Address derived in step 1 above for the Remote CNA/VN_Port.

The details of the bidirectional Shortcut Path establishment process are shown in the Appendix of this document. The process shown, however, is that which might occur between peer FSPs. To see how things might work with a Shortcut enabled CNA and an FSP; one needs to reference both the Appendix in this document, but also the Appendix in the ABS document.

(Note: when an FSP is involved, its local (non Shortcut enabled) CNA will not see any of the Shortcut Path Establishment processes or message Flows. That is, the arriving "Shortcut_Path_Verification" FIP message or the "Shortcut_Path_Verification" Response FIP message will be discarded after being processed by the FSP.)

5. With the establishment of the bidirectional Shortcut Path, all FCoE frames containing ULP Commands/Data/Response messages will be diverted by the local FSPs so that they travel via the Shortcut Path between the local CNA/VN_Port and the remote peer CNA/VN_Ports. That is, the ULP messages will travel on the Shortcut Path in FCoE Frames that have an Ethernet Destination (DA) MAC Address which was derived in step 1 above for the remote CNA/VN_Port. The ULP messages will not (normally) be sent, by the Shortcut enabled peers, through the FCF.

(Note: There are some scenarios, when the Shortcut Path establishment protocol may have very rare **FCS** (Frame Check Sequence) errors and then it is possible that some FCoE ULP messages may travel in one direction via the FCF and the other direction via the shortcut path. This should not cause any integrity errors but may slow down the total response time to a small amount.)

This diversion of the ULP messages onto the Shortcut path by an FSP is accomplished by the FSP replacing the Destination (DA) MAC address, which is the FCF's MAC address, with the MAC Address of the Remote CNA/VN_Port peer (derived in Step 1 above). Likewise, for FCoE frames arriving from that Remote CNA/VN_Port the Source (SA) MAC Address will be replaced with the related FCF's MAC address.

To ensure that they are dealing with a FCoE ULP frame, the FSP must inspect/detect on the Frame Type being FCoE (Ether Type = 8906h), and then inspecting/detecting that the encapsulated FC Frame is a ULP Message.

(Note: unless the imbedded switch within the FSP can adjust to dynamic path changes by moving the detection and diversion capability, the installation should prevent issues of dynamic changes in the path by ensuring the FSP is the switch directly connected to the CNA or that the path to the FSP from the CNA be forced not to change (probably through a Administrator involved process).

6. The ENode's FCoE Controller and the FCF should continue to send Keep Alive messages and monitor the status of the original Virtual link between their ENode/VN_Port and the corresponding FCF VF_Port which had been used to log in. The sending and monitoring is necessary to ensure that the original virtual links remain operational and their path through intervening Ethernet Switches is kept refreshed. (This is important since the traffic will be mostly occurring on the Shortcut path, yet the original Virtual links needs to continue, or the entire End-to-End logical connection will be brought down.) The frequency periods and format of these Keep-Alive FIP Messages and the Advertisement (used as Keep-Alive) FIP messages will continue as specified by the FC-BB-5 FCoE standard. The FSP will not affect the normal FC-BB-5 Keep-Alive process in any way.

FIP Keep Alive messages are also to be sent and monitored on the Shortcut path from the Local FSP port to the Remote FSP port peer (or with an actual Shortcut enabled CNA/VN_Port) in a manner that is consistent with the FC-BB-5 standard but at an interval consistent with the monitoring periods specified in the FC-BB-5 standard for the ENode FCoE Controller called the "FKA_ADV_PERIOD" (which has a default value of 8 seconds).

These Shortcut path Keep Alive FIP messages should be laid out as specified in "New FIP Message" section below and in the "Shortcut Path Keep Alive FIP Message" section within the ABS document. And these Shortcut path Keep Alive FIP messages should have the Ethernet Destination (DA) MAC Address set to the MAC Address of the Remote VN_Port peer.

Note: The "Shortcut path Keep Alive" FIP message will be discarded if it ever reaches a non Shortcut enabled CNA since the Source MAC Address (SA) of the "Shortcut Path Keep Alive" Ethernet messages frames, when they arrive at the CNA, will have the MAC Address of their Remote CNA/VN_Port peer and not the MAC Address of the related FCF VF_Port as a FC-BB-5 compliant CNA requires.

The FSP ports can exchange these Keep Alive messages without concern about the content of the Keep Alive message itself. In other words, the FSP should not worry about obtaining the Sending ENode's MAC Address, the Sending VN_Port N_Port_ID, or the Sending VN_Port_Name. The FSP should set those fields to zeros in the "Shortcut_Path Keep Alive" FIP message.

7. The logical End-To-End Connection will be terminated via normal FC_BB-5 standard FCoE processes i.e. Fabric Logout (LOGO, Clear Virtual Link, etc.) sent from the CNA to FCF (LOGO) or from the FCF to the CNA (Clear Virtual Link). At that point the FSP should detect these LOGO and Clear Virtual Link FIP Message and remove any state that the FSP is carrying that relates the Shortcut path, with the Virtual Link or Logical Connection.

It should be noted that the FSP will not issue any termination messages themselves.

8. At logical connection termination, the dynamically built Ethernet switch ACEs, if any, should be removed from the FSP related CNA ingress port's ACL (if any), in a manner suggested in the FC-BB-5 standard, plus any ACEs created for the FSP Shortcut process as a result of step 2 above, should also be removed. (Note: There may not be any dynamic ACLs or ACEs to remove if the Administrator created static ACLs or permitted none at all.)

New FIP Messages

The FSP will use the same 4 new FIP messages that were described in the ABS document. They are:

1. The "**End-To-End_Connection_Established**" FIP Message (refer to Fig. 3 in the ABS document)
2. The "**Shortcut_Path_Verification**" FIP message (refer to Fig. 4 in the ABS document)
3. The "**Shortcut_Path_Verification_Response**" FIP message (refer to Fig. 5 in the ABS document)
4. The "**Shortcut_Path_Keep_Alive**" FIP message (also refer to Fig. 4 in the ABS document)

Likewise, the layout of the slightly modified "FLOGI" FIP Message is also be as described in the ABS document (refer to Fig. 2 in that document).

For FIP messages 2 - 4, above, it is only required to provide a valid "Sending VN_Port MAC Address" value. The creations of those messages are then quite easy for an FSP. Therefore, FSPs should set the other value fields in those descriptors to zeros when sent; these value fields are named the Sending ENode's MAC Address, the Sending VN_Port N_Port_ID, and the Sending VN_Port_Name. These "other" fields have no value in these "Shortcut_Path_Verification"/"Shortcut_Path_Verification_Response" FIP Messages nor in the "Shortcut_Path_Keep_Alive" FIP message, and are quite difficult for an FSP to obtain. The inclusion of the Sending VN_Port MAC Address, however, may be of value, perhaps for reporting/monitoring or debugging, and should be included.

New Optional FIP Message

The following is an example of the New FIP Message that may optionally be sent by an FCF; it is called (herein):

- The "**End-To-End_Connection_Established Advertisement FIP Message**" (Refer to Fig. 10 below)

The following is a description of this message:

The "End-To-End_Connection_Established Advertisement" FIP Message (Refer to Fig. 10 below) will be formatted similar to the FC-BB-5 Standard FIP Advertisement message. The FIP Advertisement message was chosen as a base because, except for an indication of the link being alive and indicating the FCFs availability, they have no other significant impactful function and can almost be sent at any time. Even though the FIP Advertisement format is used, there will be an additional descriptor, a new flag, and a variation on usage.

FIP Operation Code = 001h		Reserved	SubCode = 02h
Descriptor List Length = 14		E	Flags E A
Type = 1	Length = 1	Reserved	Priority
Type = 2	Length = 2	FCF's MAC Address	
Type = 4	Length = 3	Reserved	
Switch Name			
Type = 5	Length = 4	VF_ID	
Reserved	FC_MAP		
Fabric_Name			
Type = 12	Length = 2	Reserved	
FKA_ADV_Period			
Type = 128	Length = 2	Shortcut Peer's VN_Port's MAC Address	

Fig. 10 The Operation Section of an End-To-End_Connection_Established Advertisement FIP Message

The End-To-End_Connection_Established Advertisement FIP message is an FC-BB-5 Standard FIP Advertisement Frame with:

- a. A new flag that specifies this as an End-To-End_Connection_Established Advertisement FIP message (refer to the “E” Flag in Fig. 10). This means that any intervening FSP switch will be able to detect on this message frame as something it needs to use to build dynamic ACEs. (It should ignore any other Advertisement FIP frames that do not have this flag set.)
- b. An existing Advertisement FIP message flag, called the “S” (Solicited) flag will (in this example) not be set (since it was not solicited) but it will be a Unicast message sent directly to the subject VN_Port instead of the normal Multicast unsolicited message that is usually sent to ENodes to which it advertises its existence. (Note: it should be possible to set the “S” Flag without affecting the processes in any significant way but that is not an example included here.)
- c. One (new) “non critical” descriptor with descriptor Type = 128 (refer to Fig. 10). This additional descriptor will have the same form as the Type = 2 FIP Descriptor (MAC Address), but will carry the MAC Address of the Remote VN_Port peer. This MAC Addresses will be used by the FSP switch to build the Source and Destination MAC Addresses (SA & DA) values in ACEs which will “Permit” FCoE and FIP messages to be sent from the attached VN_Port to its Remote Peer VN_Port without being blocked by the switch.

As an alternative layout (refer to Fig. 11 below) it is possible for this FIP message to include both of the peers’ MAC addresses in a new descriptor (Type = 129) for this FIP Message, so that the Message could be sent to the ENode FCoE Controller’s MAC Address instead of the VN_Port MAC Address (this maybe more compatible for some CNA implementations).

FIP Operation Code = 001h		Reserved	SubCode = 02h	
Descriptor List Length = 14		E	Flags	
Type = 1	Length = 1	Reserved	E	A
Type = 2	Length = 2	Reserved		
FCF's MAC Address				
Type = 4	Length = 3	Reserved		
Switch Name				
Type = 5	Length = 4	VF_ID		
Reserved	FC_MAP			
Fabric_Name				
Type = 12	Length = 2	Reserved		
FKA_ADV_Period				
Type = 129	Length = 4	Reserved		
VN_Port's MAC Address				
Reserved				
Shortcut Peer's VN_Port's MAC Address				

Fig. 11 The Operation Section of an End-To-End_Connection_Established Advertisement Alternate FIP Message

Note: One of these formats of "End-To-End_Connection_Established Advertisement" FIP message will be sent but only to the VN_Ports that declared their Shortcut Capability (with the "C" flag in the FLOGI Request FIP message) or for which the FSP declared their Shortcut capability when it added the "C" flag to their FLOGI Request FIP message. Therefore no compatibility issue should occur with the non shortcut enabled CNAs, however, because this is a correctly formed Advertisement Message, and (in these example layouts) because a "non critical" descriptor code of 128 (or 129) is used, this descriptor (if not understood) will be ignored by the ENode/VN_Port when the message is received. And, because the "E" flag bit is currently reserved, the "E" flag can also be ignored (not checked) by the ENode/VN_Port. Therefore, this End-To-End_Connection_Established Advertisement FIP message will not be seen as an error and will not cause an error message to be issued, if passed through by the FSP, but it will be treated as a normal FIP Advertisement message.

Dynamic ACLs

The dynamic ACLs/ACEs establishment is logically the same as is shown in the ABS document; however, since it is the FSP that is actually creating and sending the Shortcut Path Verification FIP Messages and the Shortcut path Keep Alive messages, the internal vendor specific implementation of permitting or denying this path for FIP and FCoE ULP messages is probably different for each vendor.

There are also new considerations for the handling of the Optional FIP message known as the "End-to-End_Connection_Established Advertisement" FIP message. If this Advertisement FIP message is used,

the inspection of the PLOGI ACC FCoE messages can be avoided since the same needed information can be extracted from the "End-to-End_Connection_Established Advertisement" FIP message as can be extracted from the PLOCI ACC FCoE message. However, it requires that the FCF become a partner in this process, and not just an uninvolved device.

This means that the same dynamic ACEs can logically be built as was described in the ABS document with the exception of the vendor specific consideration mentioned above.

Static ACLs

The static ACLs that can or should be built for an FSP Ethernet Switch are the same as should be built for normal Ethernet Switches on behalf of a Shortcut enabled CNA. Of course, the method that is used to establish the ACEs for FSP initiated Shortcut Path FIP messages will probably be vendor specific, and may not require any external Administrative action at all.

Logical FCFs

Note: If an installation has only FCoE Shortcut enabled CNAs and/or FSPs (that provide proxy services for all its non Shortcut enabled CNAs) and a single Lossless Ethernet Segment, it would be possible to run the FCF processes in a Logical FCF function operating on a Host system. At that point, the only Real Network would be the Lossless Ethernet Fabric, and no ULP data will be flowing through the Logical FCF.

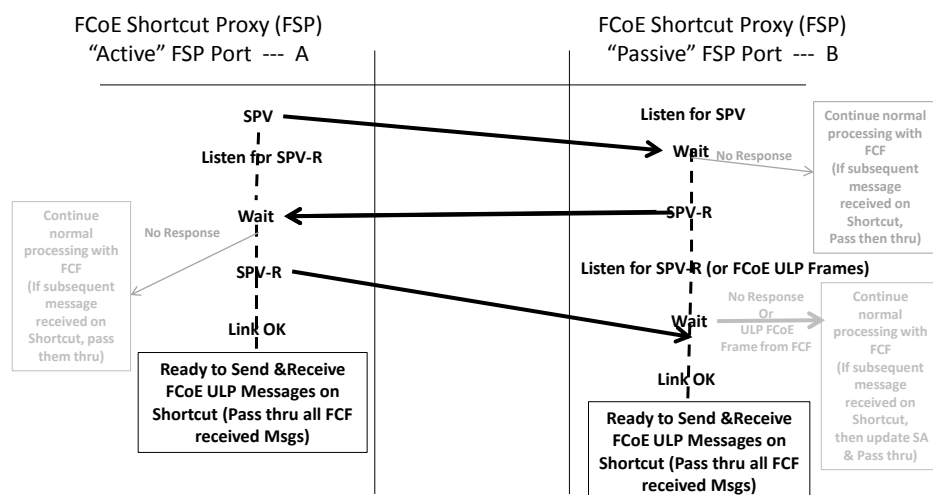
Logical FCFs (or very small FCFs) and Shortcut enabled CNAs and/or FSPs should permit FCoE to go "down market" to the midrange (and below) installations. This should be directly competitive to iSCSI in the data center of those installations.

Note: With only logical FCFs it is possible that an implementation (possibility with input from the Administrator) may not want the FCF to ever be used for switching FCoE ULP messages. This could be the case when there is no Real FCF in the configuration and the Pseudo (Logical) FCF has not enough Bandwidth capability. Therefore, a small FCF or a Software version of an FCF may just reject the arrival of FCoE ULP messages and issue a "Clear Virtual Link" FIP message if they ever do arrive.

Appendix

Detail example of the possible process for establishing the Shortcut Path

There can be many approaches to establishing the “Shortcut” Path; the following is such an approach. (Refer to Fig. 6 - Flow Ladder for a Shortcut Path Establishment Protocol.)



Note: the FSP Ports will ignore any reception of a duplicate SPV or SPV-R

SPV=Shortcut_Path_Verification SPV-R=Shortcut_Path_Verification Response

Note: If FCoE ULP Msg arrives on both paths (FCF & SC) pass it through w/o any SA update

Fig. 6 – FSP Flow Ladder for a Shortcut Path Establishment Protocol

When reading the following description also refers to Figs 6, 7, 8, 9a and 9b:

- The term ACTIVE will make reference to that side (of a logical connection) from which the PLOGI request was initiated.
- The term PASSIVE will make reference to that side (of a logical connection) from which the PLOGI ACC was issued.

During the Shortcut path establishment (as described below) when the “Shortcut_Path_Verification” FIP message or the “Shortcut_Path_Verification Response” FIP messages are sent, they may optionally be doubled (sent twice, one after the other—possibly with a small pause in between) in case the first message of the set has an intermittent FCS (Frame Check Sequence) error. An FCS error during transmission will cause the frame to be discarded. Therefore, if it was just an intermittent error, the second message (the duplicate) should get through. This means that whether or not the sender doubles the message, during reception -- if more than one (a duplicate) of the “Shortcut_Path_

Verification FIP message or the “Shortcut_Path_Verification Response FIP messages are received, the second message (the duplicate) should be ignored and discarded.

Note: All of the following occurs after the normal FCoE End-To-End Logical Connection has been established by the end CNAs (after the PLOGI ACC is sent/received) and after the End-To-End_Connection_Established FIP message has been seen by the FSP or sent by the CNA to their corresponding FCF VF_Ports (as specified in step 2 & 3 of the "Details of the Proxy Concept" section above) or in, some cases, after the Optional "End-To-End_Connection_Established Advertisement" FIP message has been received.

The following description will be depicting two peer FSP which are performing proxy services on behalf of non Shortcut enabled CNAs. Since the protocol is suppose to be compatible with the Shortcut Path Establishment performed by Shortcut enabled CNAs, it is also possible that either the ACTIVE side or the PASSIVE side could be a Shortcut enabled CNA instead of an FSP. However, the details of any non FSP will be found in the Appendix of the ABS document.

Also in the following description when we say the ACTIVE or PASSIVE FSP port will "listen" for a Shortcut_Path_Verification FIP message, a Shortcut_Path_Verification Response FIP message, and/or an FCoE ULP Frame to arrive on the Shortcut path it should be read to mean: "listen" for the arrival of one of those frames which are headed for the subject FSP port, which is connected to a non Shortcut enabled CNA and the message has an FPMA in both the Ethernet Frame SA and DA fields and those SA & DA MAC Address are the MAC Address of the CNA Port peers.

In order to ensure that no "Race" conditions exists while the Shortcut path is being established, any FCoE ULP messages that arrive from the local side's related CNA VN_Port, should be queued until the appropriate path is established.

The PASSIVE FSP port (Refer to Fig. 6) -- will begin to listen for a Shortcut_Path_Verification FIP message from the ACTIVE FSP port peer. As shown in Fig. 6, the PASSIVE FSP port will then Wait for the possible arrival of the Shortcut_Path_Verification FIP message from the ACTIVE FSP Port peer.

Note: Any non FCoE ULP messages being sent to or received (by either the ACTIVE or PASSIVE FSP port) from the appropriate FCF VF_Port, will be passed on through without any modification.

Then (continuing to refer to Fig. 6) when the ACTIVE FSP port peer is ready, the ACTIVE FSP port will send to the PASSIVE FSP port peer the “Shortcut_Path_Verification” FIP message and then listen for the “Shortcut_Path_Verification” Response FIP message from the PASSIVE FSP port peer. As shown in Fig. 6, the ACTIVE FSP port will then Wait for the possible arrival of the Shortcut_Path_Verification Response FIP message from the PASSIVE FSP port peer.

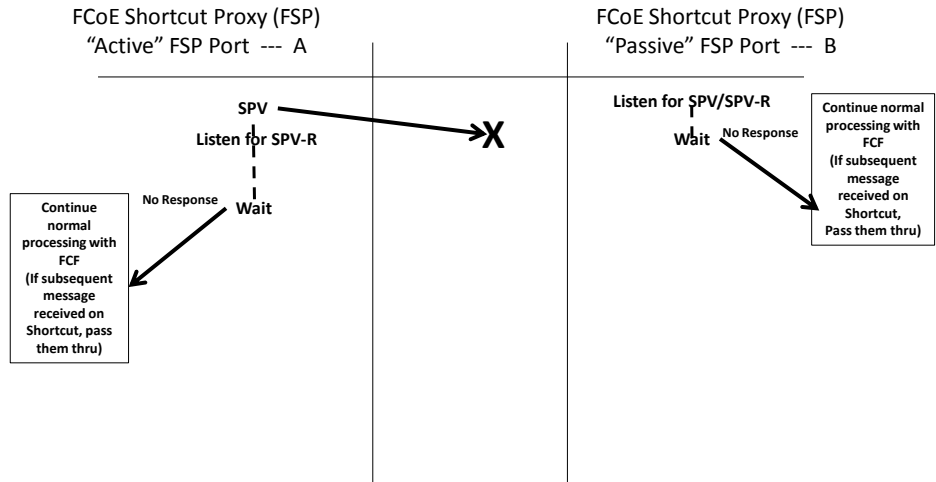
Continuing to refer to Fig. 6 -- When the PASSIVE FSP port receives a “Shortcut_Path_Verification” FIP message from the ACTIVE FSP (ignoring more than one), it will respond by Sending the “Shortcut_Path_Verification” Response FIP message to the ACTIVE FSP port peer, and then set up to listen for a “Shortcut_Path_Verification” Response FIP message or the first FCoE (ULP) frame, on the

Shortcut path; then it will Wait for a one of these responses from the ACTIVE FSP port Peer. (Note: the PASSIVE FSP port will also be watching all messages from the related FCF VF_Port to see if it receives an FCoE ULP messages from it before receiving one on the shortcut path.)

Referring to Fig 6, after the ACTIVE FSP port receives the “Shortcut_Path_Verification” Response FIP message (ignoring more than one) from the PASSIVE FSP port peer, it should send its own “Shortcut_Path_Verification” Response FIP message to the PASSIVE FSP port peer. At this point the ACTIVE FSP port should consider the Shortcut path Verified & Enabled. The ACTIVE FSP port should send subsequent (or Queued) FCoE ULP messages directly via the Shortcut path to the PASSIVE FSP port. That is, place the Remote VN_Port peer’s MAC Address into the Ethernet DA field of FCoE ULP frames replacing the MAC address of the FCF’s VF_Port and then sending the message on to the Remote CNA/VN_Port specified in the Ethernet DA MAC Address which was just inserted. The ACTIVE FSP port should also be prepared to receive FCoE frames directly from the PASSIVE FSP port. That means, when the ACTIVE FSP receives an FCoE ULP message from the FSP PASSIVE port's peer, it should replace the Ethernet SA field of the arriving FCoE ULP frame with the MAC address of the related ACTIVE FCF's VF_Port. And then the FSP should send that FCoE ULP message on to the Local CNA/VN_Port specified by the Ethernet DA MAC Address. Any FCoE or FIP messages arriving from the FCF headed to that local CNA/VN_Port should be passed through without any modification. This means that even though the ACTIVE FSP port has determined that it should use the Shortcut path to send its FCoE ULP messages, its Remote PASSIVE Peer, may have determined that it needs to send FCoE ULP messages via the FCF (perhaps because of a Ethernet FCS error, that caused a “Shortcut_Path_Verification” Response FIP message to be discarded).

Referring to Fig 6, after the PASSIVE FSP port receives the “Shortcut_Path_Verification” Response FIP message (ignoring more than one) or an FCoE ULP frame from the ACTIVE FSP port on the Shortcut path, the PASSIVE FSP port should consider the Shortcut path Verified & Enabled. The PASSIVE FSP port should then send further ULP messages in FCoE frames directly to the ACTIVE CNA/VN_Port. That is, place the ACTIVE CNA/VN_Port’s MAC Address into the Ethernet DA field of FCoE frames replacing the MAC address of the corresponding FCF’s VF_Port. And then the PASSIVE FSP should send the ULP message on to the Remote CNA/VN_Port specified in the Ethernet DA MAC Address which was just inserted. The PASSIVE FSP port should also be prepared to receive FCoE frames directly from the ACTIVE FSP port. That means that when an FCoE ULP message is received from the ACTIVE FSP port, it should replace the Ethernet SA field of the arriving FCoE ULP frame with the MAC address of the related FCF's VF-Port. And then the PASSIVE FSP should send the ULP message on to the Local CNA/VN_Port specified by the Ethernet DA MAC Address. Any FCoE or FIP message arriving from the FCF headed for that local CNA/VN_Port should be passed through without any modification.

The following examines the Error paths (Wait Time-outs).



Note: These are the same actions that are taken by an FSP Port when the other side is a non Shortcut CNA, and has no FSP

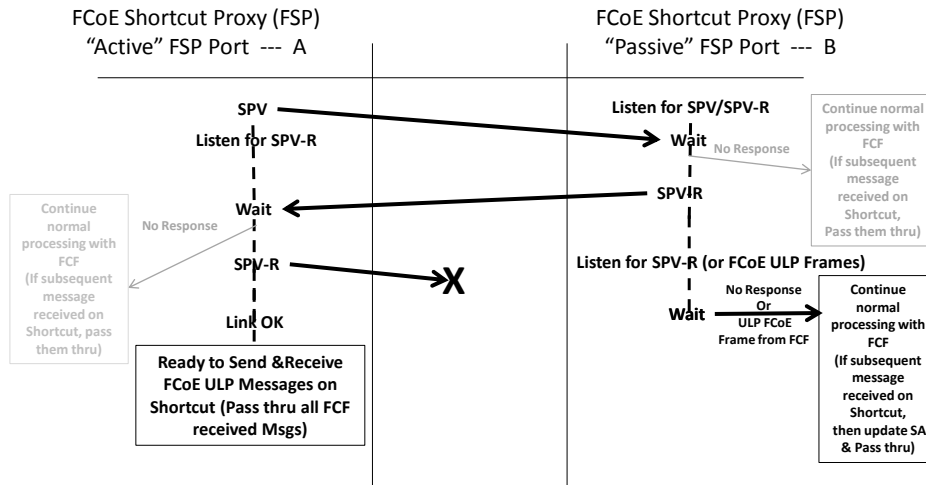
SPV=Shortcut_Path_Verification SPV-R=Shortcut_Path_Verification Response

Fig. 7 – FSP Flow Ladder for a Shortcut Path Establishment Protocol in the presences of an error Example #1

Referring to Fig. 7 – if the PASSIVE FSP port’s 1st Wait (for the “Shortcut_Path_Verification” FIP message) is longer than an implementation’s time-out value the PASSIVE FSP port should continue FCoE processing using only the path to the FCF’s corresponding VF_Port by passing the FIP and FCoE frames to (and from) the FCF’s VF_Port without any modification. (It should act like a normal Ethernet switch.) Then, if any subsequent message is received on the Shortcut path, it is probably some sort of processing error and should be handled as a normal message and passed through to the local CNA/VN_Port without any modification. (It should act like a normal Ethernet switch.) It is expected that, according to the FC-BB-5 standard, these messages that arrived via the Shortcut path will be discarded and ignored by the Local non shortcut enabled CNA/VN_Port. (That is because the arriving Frame will not have an SA of the corresponding PASSIVE side FCF VF_Port.)

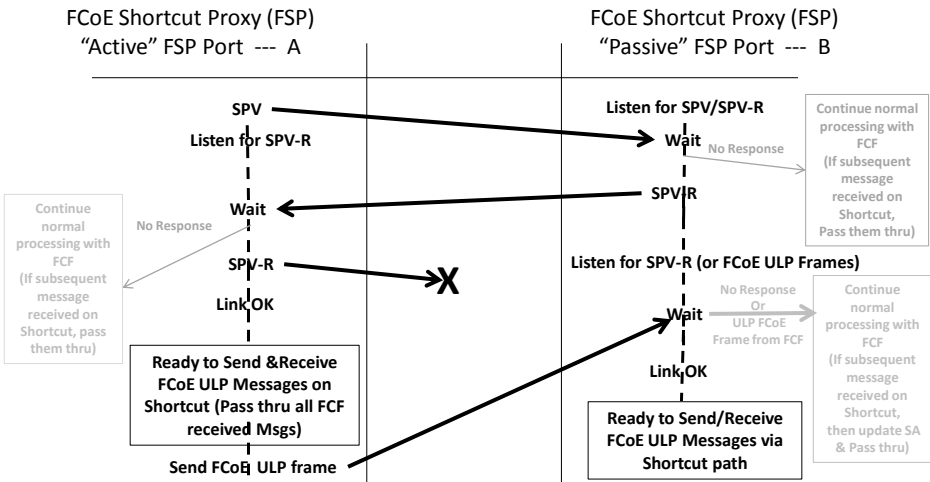
Note: referring to Fig.7, these are the same actions that are taken by an FSP Port when the other side is a non Shortcut CNA, and has no FSP.

that, according to the FC-BB-5 standard, these messages which arrived via the Shortcut path will be discarded and ignored by the Local non shortcut enabled CNA/VN_Port. (That is because the arriving Frame will not have an SA of the corresponding ACTIVE side FCF VF_Port.)



SPV=Shortcut_Path_Verification SPV-R=Shortcut_Path_Verification Response
 Note: If FCoE ULP Msg arrives on both paths (FCF & SC) pass it through w/o any SA update
Fig. 9a – FSP Flow Ladder for a Shortcut Path Establishment Protocol in the presences of an error Example #3

Fig 9a and 9b show how an FCS error on the Shortcut_Path_Verification Response message, which is sent by the ACTIVE FSP port, may be handled. All the same responses have been detailed above, but the triggered actions shown in Fig 9a and 9b may be slightly different. In Fig 9a, the error occurs when the ACTIVE FSP port sends the Shortcut_Path_Verification Response FIP message. The ACTIVE FSP port will not know the error occurred and will continue to prepare for shortcut processing. However, in Fig. 9a, the ACTIVE FSP port will not have an FCoE ULP message ready to send before the PASSIVE FSP port peer’s Wait times out. At this point the PASSIVE FSP port will follow the processes detailed in the description regarding Fig 8, above, regarding the PASSIVE FSP port’s 2nd Wait’s Time-Out. That is, the PASSIVE FSP port should continue FCoE processing using only the path to the FCF’s corresponding VF_Port and passing on all FIP and FCoE frames (outgoing or incoming) without any modification (i.e. It should act like a normal Ethernet switch). Then, if any subsequent message is received by the PASSIVE FSP port, on the Shortcut path, headed for the PASSIVE CNA/VN_Port, its SA should be set to the corresponding FCF VF_Port's MAC Address and then the Frame should be passed on to the Destination.



SPV=Shortcut_Path_Verification SPV-R=Shortcut_Path_Verification Response
 Note: If FCoE ULP Msg arrives on both paths (FCF & SC) pass it through w/o any SA update
Fig. 9b – FSP Flow Ladder for a Shortcut Path Establishment Protocol in the presences of an error Example #4

In Fig 9b, in spite of the previous error in the Shortcut_Path_Verification Response FIP Message, the ACTIVE FSP port does not know that the Shortcut_Path_Verification Response FIP Message that it sent to the PASSIVE FSP port did not arrive. Therefore, it will continue its processing and send any new (or Queued) FCoE ULP messages to the PASSIVE FSP port. In this example, the FCoE ULP message is received before the PASSIVE side's Wait times out. At that point both sides continue processing in Shortcut mode.

Anytime the FSP detects that it is receiving FCoE ULP messages on both the Shortcut and FCF paths headed for the same VN_Port peers the SA modification of FCoE ULP message arriving on the Shortcut path should cease and the messages should just be passed on like would happen in a normal Ethernet switch. It is expected that the CNA will discard those FCoE ULP frames since they will not have the MAC Address of the related FCF VF_Port in its SA field.