

Annex E: FCoE Pre-FIP Virtual Link instantiation protocol (Informative)

David Peterson, Brocade
Erik Smith, EMC
Fred Knight, NetApp
(09-277v1)

E.1 Overview

This annex documents a pre-standard implementation of FCoE that was used in ENode to FCF direct-attach connections (i.e., an ENode is connected to an FCF via a single cable without any intermediate Ethernet bridge) before the creation of FIP as specified in clause 7.

Implementations compliant to this standard support FIP as specified in clause 7, not the methods described in this annex (i.e., this annex is not applicable for ENode and FCF implementations that support FIP as specified in clause 7). An implementation supporting this annex and not supporting FIP as specified in clause 7 is not compliant with this standard.

This annex is intended to be removed in the next version of this standard.

E.2 Protocol Summary

The pre-FIP Virtual Link instantiation protocol consists of two phases, DCBX (Data Center Ethernet Bridging Capability Exchange Protocol) and Fabric Login.

A pre-standard version of DCBX as defined in the “DCB Capability Exchange Protocol Specification” (see http://download.intel.com/technology/eedc/dcb_cep_spec.pdf) is performed first.

Once DCBX has determined that the FCoE Logical Link status is “up”, the ENode transmits a FLOGI to the FCF.

E.3 Functionality for all ENodes and FCFs

The FCoE frame format as specified in 7.7 is used.

All FCoE frames:

- a) contain an 802.1Q tag header and are priority-tagged (i.e., the VID field is set to 0);
- b) have the Priority Code Point (PCP) field in the VLAN tag set to 011b as the default value; and
- c) support only the FC-MAP value of 0EFC00h.

Per-Priority Pause (PPP) as specified in <http://www.ieee802.org/1/files/public/docs2007/new-cm-barrass-pause-proposal.pdf> is used to reduce loss in the Ethernet network due to congestion.

E.4 Functionality for ENodes

For all DCB Features TLVs, the willing bit is set to one.

The FCoE Controller in the ENode determines the FCoE and LAN Logical Link status based upon both the FCoE Logical Link Status and the LAN Logical Link Status TLVs, or the physical state of the link (i.e., up or down).

E.5 Functionality for FCFs

For all DCB Features TLVs, the willing bit is set to zero.

The FCoE Controller in the FCF indicates the FCoE and LAN Logical Link status based upon both the FCoE Logical Link Status and the LAN Logical Link Status TLVs or the physical state of the link (i.e. up or down).

E.6 Functionality for DCBX Features

At minimum, the following DCBX Feature TLVs are present in each DCBX frame:

- a) Priority Flow Control – Admin mode bit is set to one for FCoE PCP only;
- b) Application – Subtype FCoE and User Priority Map is set to 011b as the default value;
- c) Logical Link Down – Subtype FCoE Logical Link status; and
- d) Logical Link Down – Subtype LAN Logical Link status.

E.7 Ethernet destination address (DA) and source address (SA) format

When transmitting a FLOGI ELS or FDISC ELS request, the ENode sets the DA to 0EFC00FFFFFFEh and the SA to the ENode's universal MAC address (i.e., the same MAC address used for transmitting DCBX frames).

When transmitting a FLOGI LS_ACC or FDISC LS_ACC reply, the FCF sets the DA to the universal MAC address of the ENode and the SA to the FCF-MAC address.

Subsequent FCoE frames from the ENode to the FCF have:

- a) the DA set to the FCF-MAC Address;
- b) the three most significant bytes of the SA set to the FC-MAP value of 0EFC00h; and
- c) the three least significant bytes of the SA set to the value of the FC Source_ID field.

Subsequent FCoE frames from the FCF to the ENode have:

- a) the three most significant bytes of the DA set to the FC-MAP value of 0EFC00h;
- b) the three least significant bytes of the DA set to the value of the FC Destination_ID field; and
- c) the SA set to the FCF-MAC address.