Copies of this document may be purchased from:
Global Engineering, 15 Inverness Way East,
Englewood, CO 80112-5704
Phone: (800) 854-7179 or (303) 792-2181 Fax: (303) 792-2192

INCITS xxx-200x
T11/Project 1871-D/Rev 2.00

FIBRE CHANNEL

BACKBONE - 5

(FC-BB-5)

REV 2.00

INCITS working draft proposed
American National Standard
for Information Technology

June 4, 2009

Secretariat: Information Technology Industry Council

NOTE:

This is a working draft American National Standard of Accredited Standards Committee INCITS. As such this is not a completed standard. Representatives of the T11 Technical Committee may modify this document as a result of comments received anytime, or during a future public review and its eventual approval as a Standard. Use of the information contained herein is at your own risk.

POINTS OF CONTACT:

Steven Wilson (T11 Chair)
Brocade Communcations, Inc.
1745 Technology Drive
San Jose, CA 95131
Voice: 408-333-8128
swilson@brocade.com

Claudio Desanti (T11 Vice Chair)
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
Voice: 408-853-9172
cds@cisco.com

Craig W. Carlson (T11.3 Chair)
QLogic Corporation
6321 Bury Drive
Eden Prairie, MN 55346
Voice: 952-932-4064
craig.carlson@qlogic.com

Claudio DeSanti (FC-BB-5 Chair)
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
Voice: 408-853-9172
cds@cisco.com

David Peterson (FC-BB-5 Editor)
Brocade Communications, Inc.
6000 Nathan Lane North
Plymouth, MN 55442
Voice: 612-802-3299
david.peterson@brocade.com

BSR INCITS xxx-200x

American National Standard
for Information Technology

Fibre Channel —
Fibre Channel Backbone - 5 (FC-BB-5)

Secretariat

Information Technology Industry Council

Approved (not yet approved)

American National Standards Institute, Inc.

Abstract

This standard defines the functions and mappings for transporting Fibre Channel over different network technologies.

American National Standard

Approval of an American National Standard requires review by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer.

Consensus is established when, in the judgement of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that a concerted effort be made towards their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and under no circumstance gives an interpretation of any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

PATENT STATEMENT

The developers of this standard have requested that holders of patents that may be required for the implementation of the standard disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard. As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more such claims has been received. By publication of this standard, no position is taken with respect to the validity of this claim or of any rights in connection therewith. The known patent holder(s) has (have), however, filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. Details may be obtained from the publisher. No further patent search is conducted by the developer or publisher in respect to any standard it processes. No representation is made or implied that this is the only license that may be required to avoid infringement in the use of this standard.

Published by

American National Standards Institute
11 West 42nd Street, New York, NY 10036

Printed in the United States of America

Foreword (This Foreword is not part of American National Standard INCITS xxx-200x.)

This standard defines the functions and mappings for transporting Fibre Channel over different network technologies.

This standard was developed by Task Group T11.3 of Accredited Standards Organization INCITS during 2007-2009. The standards approval process started in 2008. This document includes annexes that are informative and are not considered part of the standard.

Requests for interpretation, suggestions for improvements or addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, Information Technology Industry Council, 1250 Eye Street, NW, Suite 200, Washington, DC 20005-3922.

This standard was processed and approved for submittal to ANSI by the International Committee for Information Technology Standards (INCITS). Committee approval of the standard does not necessarily imply that all committee members voted for approval.

At the time it approved this standard, INCITS had the following members:

(to be filled in by INCITS)

Technical Committee T11 on Fibre Channel Interfaces, which reviewed this standard, had the following members:

Steve Wilson, Chair
Claudio DeSanti, Vice-Chair
Bob Nixon, Secretary

Company	Name

{ To be filled in prior to submission to INCITS. }

Task Group T11.3 on Interconnection Schemes, which developed and reviewed this standard, had the following members:

Craig W. Carlson, Chair
William R. Martin, Vice-Chair
Landon Curt Noll, Secretary

Company	Name

{ To be filled in prior to submission to INCITS. }

Introduction

FC-BB-5 defines mappings for transporting Fibre Channel over different network technologies. FC-BB-5 defines four distinct Fibre Channel mappings: FC over TCP/IP, FC over GFPT, FC over MPLS, and FC over Ethernet.

The FC over ATM and FC over SONET backbone mappings are not specified in FC-BB-5. As such, FC-BB-5 is not a complete replacement of FC-BB-3 (i.e., see FC-BB-3 for the specification of the FC over ATM and FC over SONET backbone mappings).

Contents **Page**

Figure **Page**

Table **Page**

American National Standard
for Information Technology —

Fibre Channel —
Backbone - 5 (FC-BB-5)

1 Scope

This standard consists of distinct Fibre Channel mappings resulting in the following models:

- FC-BB_IP (FC over TCP/IP backbone network)
- Transparent FC-BB consisting of:
  - FC-BB_GFPT (FC over SONET/SDH/OTN/PDH backbone network using GFPT adaptation)
  - FC-BB_PW (FC over MPLS network using PW adaptation)
- FC-BB_E (FC over Ethernet)

Figure 1, figure 2, figure 3, and figure 4 illustrate the scope and the major components of the FC-BB-5 models and its relationship to the IETF, ITU-T, and IEEE standards. Table 1 shows the organization of this standard. FC-BB_IP, Transparent FC-BB, and FC-BB_E do not interoperate in any manner and are independent models.

Table 1 – FC-BB-5 organization

Model Type	Applicable Clauses and Annexes
FC-BB_IP, FC-BB_GFPT, FC-BB_PW, FC-BB_E	1, 2, 3, 4
FC-BB_IP	5
Transparent FC-BB	
FC-BB_GFPT	6, Annex A
FC-BB_PW	6
FC-BB_E	7, Annex B, Annex C, Annex D, Annex E, Annex F

The scope and components of the FC-BB_IP model is shown in figure 1.

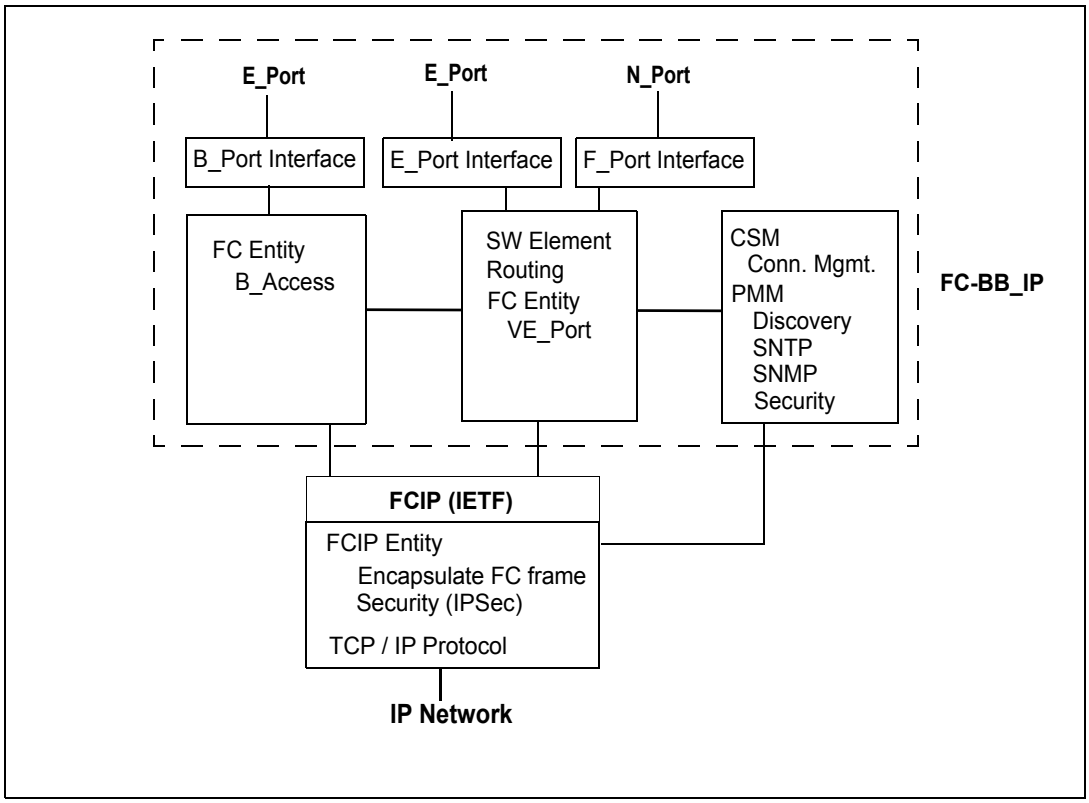

Figure 1 – Scope and components of FC-BB_IP model

The scope and components of the FC-BB_GFPT model is shown in figure 2.

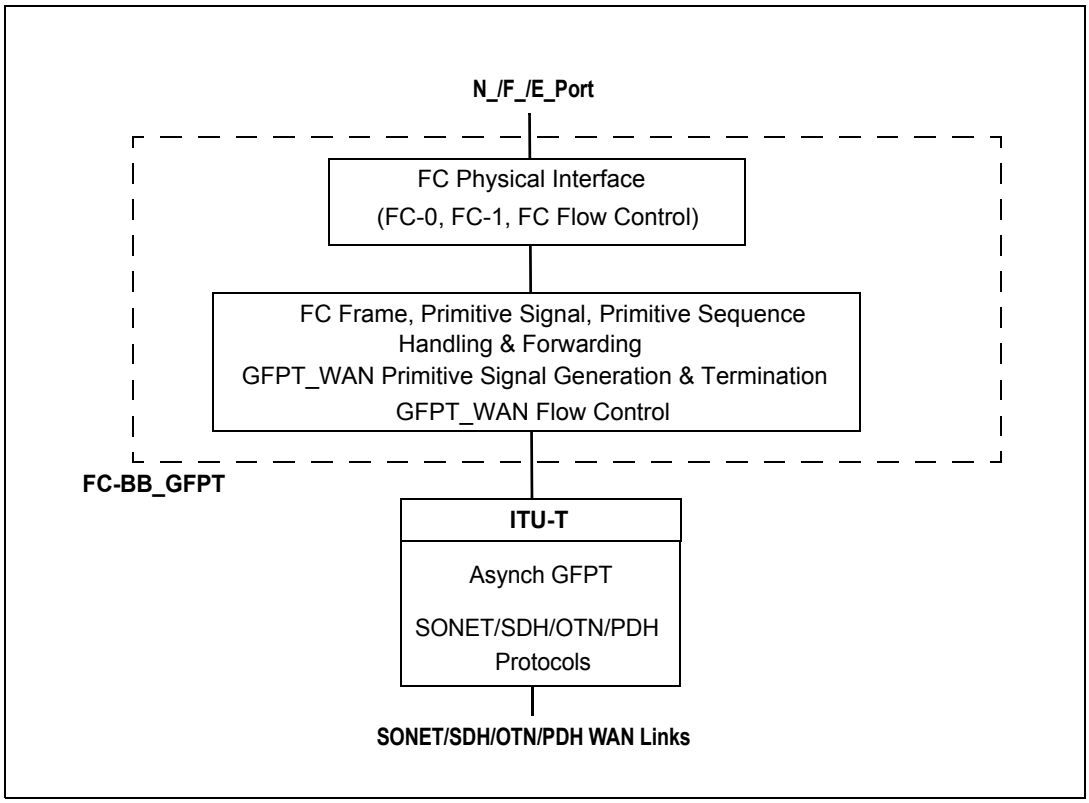

Figure 2 – Scope and components of FC-BB_GFPT model

The scope and components of the FC-BB_PW model is shown in figure 3.

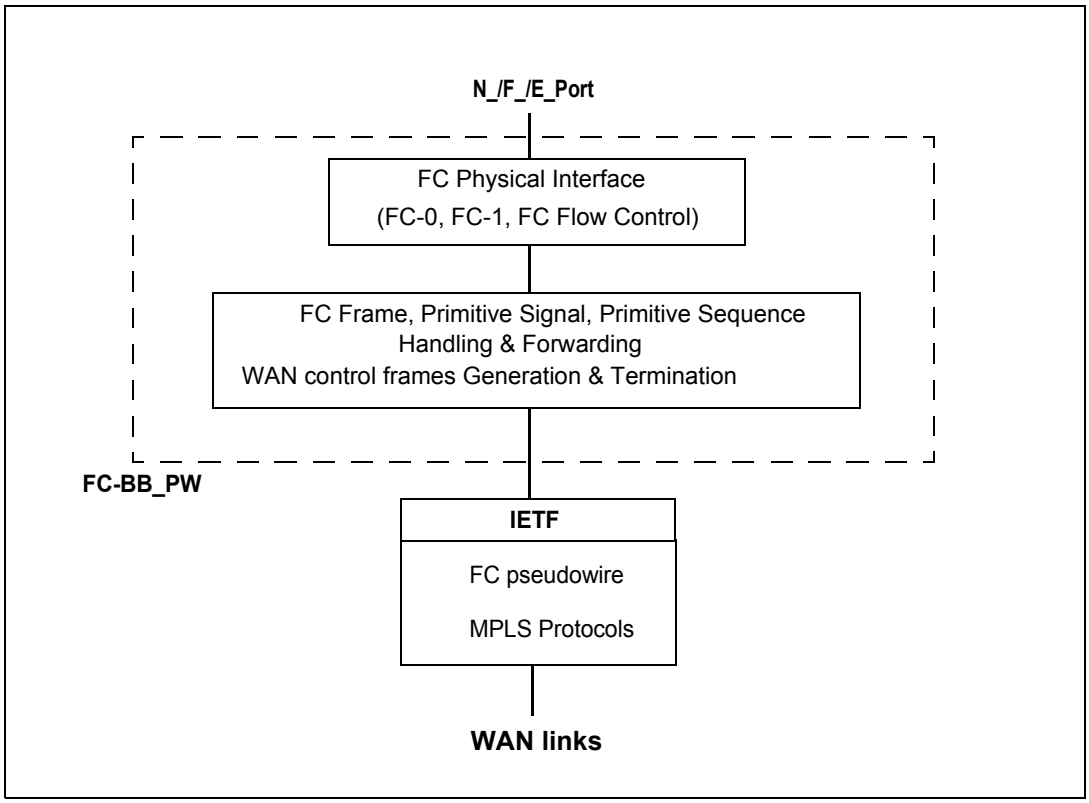

Figure 3 – Scope and components of FC-BB_PW model

The scope and components of the FC-BB_E model is shown in figure 4.

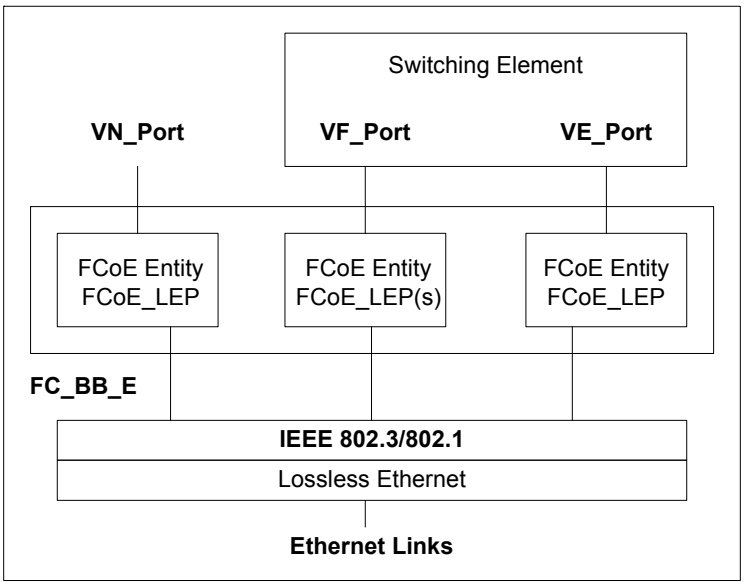

Figure 4 – Scope and components of FC-BB_E model

2 Normative References

2.1 Overview

The following standards contain provisions that, through reference in the text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards listed below.

For electronic copies of ANSI and INCITS standards, visit ANSI's Electronic Standards Store (ESS) at http://www.ansi.org. For printed versions of most standards listed here, contact Global Engineering Documents, 15 Inverness Way East, Englewood, CO; 80112-5704, (800) 854-7179.

Orders for ISO Standards and ISO publications should normally be addressed to the ISO member in your country. If that is impractical, ISO Standards and ISO publications may be ordered from ISO Central Secretariat (ISO/CS):

Phone +41 22 749 01 11
Fax +41 22 749 09 47
E-mail sales@iso.org
Post ISO, 1, rue de Varembé, CH-1211
Geneva 20, Switzerland

In order to avoid delivery errors, it is important that you accurately quote the standard's reference number given in the ISO catalogue. For standards published in several parts, you should specify the number(s) of the required part(s). If not, all parts of the standard will be provided.

Copies of the following documents may be obtained from ANSI, an ISO member organization:

a) approved ANSI standards;
b) approved and draft international and regional standards (ISO and IEC); and
c) approved foreign standards (JIS and DIN).

For further information, contact the ANSI Customer Service Department:

Phone +1 212-642-4900
Fax: +1 212-302-1286
Web: http://www.ansi.org
E-mail: ansionline@ansi.org

or the InterNational Committee for Information Technology Standards (INCITS):

Phone 202-626-5738
Web: http://www.incits.org
E-mail: incits@itic.org

Additional availability contact information is provided below as needed.

2.2 Approved references

ANSI T1.105-2001, *Synchronous Optical Network (SONET) - Basic Description Including Multiplex Structures, Rates, and Formats.*

ANSI INCITS 426-2007, *Fibre Channel - Security Protocols (FC-SP).*

ANSI INCITS 241-1994 (R1999), *Data Compression Method –Adaptive Coding with Sliding Window for Information Interchange*.

2.3 References under development

At the time of publication, the following referenced standards were still under development. For information on the current status of the documents, or regarding availability, contact the relevant standards body or other organization as indicated.

For electronic copies of references under development by INCITS T11, see www.t11.org.

T11/Project 1674-D, *Fibre Channel - Switch Fabric - 5 (FC-SW-5).*

T11/Project 2103-D, *Fibre Channel - Link Services - 2 (FC-LS-2).*

T11/Project 1861-D, *Fibre Channel - Framing and Signaling - 3 (FC-FS-3).*

For electronic copies of references under development by the Internet Engineering Task Force (IETF), see www.ietf.org.

Roth, Solomon, Tsurusawa, "Encapsulation Methods for Transport of Fibre Channel frames Over MPLS Networks", draft-ietf-pwe3-fc-encap-09.txt (*RFC reference and date to be added during standards action*).

Roth, Solomon, Tsurusawa, "Reliable Fibre Channel Transport Over MPLS Networks", draft-ietf-pwe3-fc-flow-00.txt.txt (*RFC reference and date to be added during standards action*).

For electronic copies of references under development by the Institute of Electrical and Electronics Engineers (IEEE), see www.ieee802.org.

802.1Qbb: *Virtual Bridged Local Area Networks — Amendment XX: Priority-based Flow Control*.

2.4 ITU-T references

Copies of the following approved ITU-T standards may be obtained through the ITU-T Publications department at http://www.itu.int.

ITU-T Rec. G.707/Y.1322, (2007), *Network node interface for the synchronous digital hierarchy (SDH).*

ITU-T Rec. G.7041/Y.1303, (2005), *Generic Framing Procedure (GFP).*

ITU-T Rec. G.783, (2006), *Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks.*

ITU-T Rec. G.806, (2006), *Characteristics of transport equipment - Description methodology and generic functionality*

ITU-T Rec. G.702, (1988), *Digital Hierachy Bit Rates*

2.5 IETF references

Copies of the following approved IETF standards may be obtained through the Internet Engineering Task Force (IETF) at www.ietf.org.

RFC 4330, *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI,* January 2006.

RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior),* March 2002.

RFC 3643, *Fibre Channel (FC) Frame Encapsulation,* December 2004.

RFC 3821, *Fibre Channel Over TCP/IP (FCIP)*, July 2004.

RFC 3822, *Finding Fibre Channel over TCP/IP (FCIP) Entities Using Service Location Protocol version 2 (SLPv2),* July 2004.

RFC 3031, *Multiprotocol Label Switching (MPLS) Architecture,* January 2001.

RFC 3985, *Pseudowire Emulation Edge-to-Edge (PWE3) Architecture,* March 2005.

RFC 4385, *Multiprotocol Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN,* February 2006.

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP),* April 2006.

2.6 IEEE references

Copies of the following approved IEEE standards may be obtained through the Institute of Electrical and Electronics Engineers (IEEE) at http://standards.ieee.org.

802.3-2008: *Carrier sense multiple access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.*

802.1Q-2005: *Virtual Bridged Local Area Networks*.

3 Definitions and conventions

3.1 Common definitions

3.1.1 B_Port: A Bridge Port on a device that implements FC-BB_IP and connects to an E_Port on an FC switch.

3.1.2 B_Port_Name: A Name_Identifier (see 3.1.24) that identifies a B_Port (see 3.1.1) for identification purposes. The format of the name is specified in FC-SW-5.

3.1.3 codeword: A sequence of bits of a code corresponding to a symbol.

3.1.4 E_Port: A Fabric expansion port that attaches to another E_Port to create an Inter-Switch Link (see FC-SW-5).

3.1.5 E_Port_Name: A Name_Identifier (see 3.1.24) that identifies an E_Port (see 3.1.4).

3.1.6 ELP: Exchange Link Parameters SW_ILS (see FC-SW-5).

3.1.7 F_BSY: Fabric Busy (see FC-FS-3).

3.1.8 F_Port: A port by which non-loop N_Ports are attached to a Fabric. Does not include FL_Ports (see FC-SW-5 and FC-FS-3).

3.1.9 F_Port_Name: A Name_Identifier (see 3.1.24) that identifies an F_Port (see 3.1.8)

3.1.10 Fabric Initialization: A process for configuring and building a Fabric (see FC-SW-5).

3.1.11 Fabric_Name: A Name_Identifier (see 3.1.24) associated with a Fabric (see FC-SW-5).

3.1.12 FC-BB_E: A protocol mapping defined by this standard in order to transport Fibre Channel over a Lossless Ethernet network.

3.1.13 FC-BB_GFPT: Equipment model defining gateway functionality for the interconnection of two non-Arbitrated Loop FC physical ports across a GFPT_WAN infrastructure (e.g., SONET, SDH, OTN, PDH). Supports both arbitrary-rate WAN transport and distance extension of buffer-to-buffer flow control.

3.1.14 FC-BB_IP: A model defining equipment that interfaces with a Fibre Channel switched network on one side and an IP network on the other side.

3.1.15 FC-BB_PW: Equipment model defining gateway functionality for the interconnection of two non-Arbitrated Loop FC physical ports across a PSN (see 3.4.5). Supports both arbitrary-rate WAN transport and distance extension of buffer-to-buffer flow control.

3.1.16 FC_Port: A port is capable of transmitting and receiving Fibre Channel frames (see FC-FS-3).

3.1.17 Fibre Channel Backbone link: The Transport Trail or equivalent network channel connection used for communications between two FC-BB-5 devices. This encompasses FC-BB_IP and GFPT_WAN links. A Fibre Channel Backbone link may, in some cases, be made up of more than one physical or logical connection.

3.1.18 FLOGI: Fabric Login ELS (see FC-LS-2).

3.1.19 Generic Framing Procedure (GFP): A procedure for adaptation of data (i.e., PDUs or 8B/10B encoded characters) to octet-synchronous (i.e., SONET, SDH, OTN) and bit-synchronous (i.e., PDH) Wide Area Network transport infrastructures, specified by ITU-T. See ITU-T Rec. G.7041/Y.1303.

3.1.20 ISL: Inter-Switch Link (see FC-SW-5).

3.1.21 Keep Alive Timeout Value (K_A_TOV): A timer that is used by the Link Keep Alive (LKA) ELS (see FC-LS-2) as a trigger for issuing LKA.

3.1.22 LS_ACC: Link Service Accept (see FC-LS-2).

3.1.23 LS_RJT: Link Service Reject (see FC-LS-2).

3.1.24 Name_Identifier: A 64-bit identifier, with a 60-bit value preceded with a 4-bit Network_Address_Authority Identifier, used to identify entities in Fibre Channel (e.g., N_Port, node, F_Port, or Fabric) (see FC-FS-3).

3.1.25 Node_Name: A Name_Identifier (see 3.1.24) associated with a node (see FC-FS-3).

3.1.26 N_Port: A device port that generates/terminates FC-4 channel traffic.

3.1.27 N_Port_Name: A Name_Identifier (see 3.1.24) that identifies an N_Port (see 3.1.26).

3.1.28 Ordered Set: See FC-FS-3.

3.1.29 OTN: An acronym for Optical Transport Network. OTN is a term that refers to the rates and formats specified in ITU-T G.709/Y.1331.

3.1.30 PDH: An acronym for Plesiochronous Digital Hierarchy. PDH is a term that refers to the rates and formats specified in ITU-T G.702.

3.1.31 PLOGI: N_Port Login (see FC-LS-2).

3.1.32 P_BSY: N_Port Busy (see FC-FS-3).

3.1.33 SONET: An acronym for Synchronous Optical NETwork. SONET is a term that refers to the rates and formats specified in ANSI T1.105.

3.1.34 Switch_Name: A Name_Identifier (see 3.1.24) that identifies a Switch or a Bridge device. The format of the name is specified in FC-FS-3. Each Switch and Bridge device shall provide a unique Switch_Name within the Fabric.

3.1.35 SW_ACC: Switch Fabric Internal Link Service Accept (see FC-SW-5).

3.1.36 SW_RJT: Switch Fabric Internal Link Service Reject (see FC-SW-5).

3.1.37 Transparent FC-BB: A general model consisting of the FC-BB-GFPT (see 3.1.13) and the FC-BB_PW (see 3.1.15) models.

3.1.38 WAN interface: An interface that connects to a Wide Area Network. May be physical (e.g., SONET) or logical (e.g., GFPT_WAN).

3.1.39 Synchronous Digital Hierarchy (SDH): A term that refers to the rates and formats specified in ITU-T G.707/Y.1322.

3.2 FC-BB_IP definitions

3.2.1 B_Access: A component of the FC Entity (see 3.2.8) that interfaces with the FCIP_LEP (see 3.2.20) component of the FCIP Entity (see 3.2.15) on one side and the B_Port on the other side.

3.2.2 B_Access_Name: The Name_Identifier (see 3.1.24) of the B_Access portal.

3.2.3 B_Access Virtual ISL: A Virtual ISL (see 3.2.27) that connects two B_Access portals.

3.2.4 Control and Service Module (CSM): A control component of the FC-BB_IP interface that mainly handles connection management. CSM interfaces with the PMM (see 3.2.22).

3.2.5 encapsulated FC frame: An SOF/EOF delimited FC frame prefixed with a 28-byte FC frame Encapsulation Header (see RFC 3643).

3.2.6 Encapsulated Frame Receiver Portal: The TCP access point through which an encapsulated FC frame (see 3.2.5) is received from the IP network by an FCIP_DE (see 3.2.14).

3.2.7 Encapsulated Frame Transmitter Portal: The TCP access point through which an encapsulated FC frame (see 3.2.5) is transmitted to the IP network by the FCIP_DE (see 3.2.14).

3.2.8 FC Entity: The FC Entity is the principal interface point to the FC switched network on one side and in combination with the FCIP Entity to the IP network on the other side. It is the data forwarding component of the FC-BB_IP interface consisting of VE_Port(s) (see 3.2.24) and/or B_Access (see 3.2.1) portals.

3.2.9 FC Entity Protocol Layer: The protocol layer that lies between the Fibre Channel level FC-2 and the FCIP Entity Protocol Layer (see 3.2.16). Its primary function is to support one or more Virtual E_Ports (see 3.4.24) or B_Access (see 3.2.1) portals and to communicate with the FCIP Entity (see 3.2.8).

3.2.10 FC Receiver Portal: The access point through which an FC frame and timestamp enters an FCIP_DE (see 3.2.14) from the VE_Port/B_Access (see 3.2.24/3.2.1).

3.2.11 FC Transmitter Portal: The access point through which an FC frame and timestamp leaves an FCIP_DE (see 3.2.14) to the VE_Port/B_Access (see 3.2.24/3.2.1).

3.2.12 FC-BB_IP device: A device defined by the FC-BB_IP model.

3.2.13 FC-BB_IP interface: The point that has interfaces to the FC switched network on one side and the IP network on the other side. It consists of a Switching Element, FC/FCIP Entity pair(s), the CSM, and the PMM.

3.2.14 FCIP Data Engine (FCIP_DE): The data forwarding component of the FCIP Entity's (see 3.2.15) FCIP_LEP (see 3.2.20) that handles FC frame encapsulation, de-encapsulation, and transmission of encapsulated frames through a single TCP connection.

3.2.15 FCIP Entity: The data forwarding component of the FC-BB_IP interface consisting of the FCIP_LEP (see 3.2.20) and is the principal interface point to the IP network on one side and in combination with the FC Entity (see 3.2.8) to the FC switched network on the other side. Its primary

function is formatting, encapsulating, and forwarding encapsulated FC frames (see 3.2.5) across the IP network interface.

3.2.16 FCIP Entity Protocol Layer: The protocol layer that lies between the FC Entity (see 3.2.8) layer and the TCP layer.

3.2.17 FCIP frame: The FCIP term for an encapsulated FC frame (see 3.2.5).

3.2.18 FCIP Link: A virtual link that connects an FCIP_LEP (see 3.2.20) in one FC-BB_IP device (see 3.2.12) with another. It consists of one or more TCP connections.

3.2.19 FCIP Link Originator and Acceptor: The FC-BB_IP FCIP_LEP (see 3.2.20) that originates an FCIP Link is defined as the FCIP Link Originator. The corresponding FCIP_LEP that accepts this link is defined as the FCIP Link Acceptor.

3.2.20 FCIP Link Endpoint (FCIP_LEP): The component of an FCIP Entity (see 3.2.15) that contains one or more FCIP_DEs (see 3.2.14).

3.2.21 FCIP Transit Time (FTT): The total transit time of an encapsulated Fibre Channel frame in the IP network.

3.2.22 Platform Management Module (PMM): A management component of the FC-BB_IP interface that handles time synchronization, discovery, and security. It interfaces with the CSM (see 3.2.4).

3.2.23 Request For Comment (RFC): A document at one stage of the IETF standardization process. Documents that are RFCs are the final draft of a specification intended to be approved as a standard or permanent document and are usually treated by industry as equivalent to a standard.

3.2.24 Virtual E_Port (VE_Port): The data forwarding component of the FC Entity (see 3.2.8) that emulates an E_Port (see 3.1.4). The term virtual indicates the use of a non Fibre Channel link connecting the VE_Ports. In the case of the FC-BB_IP model, a VE_Port interfaces with the FCIP_LEP component (see 3.2.20) of the FCIP Entity (see 3.2.15) on one side and a Fibre Channel Switching Element on the other side.

3.2.25 VE_Port_Name: The Name_Identifier (see 3.1.24) of the VE_Port (see 3.2.24).

3.2.26 VE_Port Virtual ISL: A Virtual ISL (see 3.2.27) that connects two VE_Ports (see 3.2.24).

3.2.27 Virtual ISL: An ISL that connects two VE_Ports (see 3.2.24) or two B_Access portals (see 3.2.1) across a non-FC link.

3.3 FC-BB_GFPT definitions

3.3.1 ASFC_PAUSE: The GFPT_WAN Primitive Signal used to pause flow on a GFPT_WAN link (see 6.4.4). ASFC_PAUSE is never transmitted to, or expected from, FC_Ports.

3.3.2 ASFC_RESUME: The GFPT_WAN Primitive Signal used to resume flow on a GFPT_WAN link (see 6.4.4). ASFC_RESUME is never transmitted to, or expected from, FC_Ports.

3.3.3 GFP Server: Generic Framing Procedure (see 3.1.19) adaptation/de-adaptation engine.

3.3.4 GFPT: (Asynchronous) Transparent Generic Framing Procedure (see 3.1.19).

3.3.5 GFPT_WAN interface: Transport network-side interface, on an FC-BB_GFPT device, corresponding to one GFPT_WAN facility (see 3.3.7), and to one Transport Trail (see 3.3.14). May or may not correspond to the full SONET/SDH/OTN/PDH access facility/bandwidth.

3.3.6 GFPT_WAN link: Transport Trail (see 3.3.14) assigned to one GFPT_WAN facility (see 3.3.7).

3.3.7 GFPT_WAN facility: Transport Trail (see 3.3.14), GFP Server (see 3.3.3), FC-BB_GFPT devices, and their respective GFPT_WAN interfaces (see 3.3.5), corresponding to one interconnected FC_Port pair.

3.3.8 inbound: Sent from the FC-BB_GFPT device to the attached FC_Port.

3.3.9 LEM: Login Exchange Monitor (see 6.4.3).

3.3.10 outbound: Sent from the attached FC_Port to the FC-BB_GFPT device.

3.3.11 PING: The GFPT_WAN Primitive Signal used to initiate latency measurement on a GFPT_WAN link. PING is never transmitted to, or expected from, FC_Ports.

3.3.12 PING_ACK: The GFPT_WAN Primitive Signal used to reply to a PING and complete round-trip latency measurement on a GFPT_WAN link. PING_ACK is never transmitted to, or expected from, FC_Ports.

3.3.13 RPSC ELS: Report Port Speed Capabilities ELS (see FC-LS-2).

3.3.14 Transport Trail: A contiguously or virtually-concatenated signal group (see T1.105-2001) made up of one or more standardized SONET/SDH/OTN/PDH synchronous transport signals.

3.3.15 WAN Primitive Signal: An ASFC_PAUSE (see 3.3.1), ASFC_RESUME (see 3.3.2), PING (see 3.3.11), or PING_ACK (see 3.3.12) Primitive Signal. These Primitive Signals are always generated and terminated by FC-BB_GFPT devices and transmitted only between FC-BB_GFPT devices. They are never transmitted to nor received from FC_Ports.

3.3.16 WAN_HOLDOFF_TOV: A time-out value, specific to FC-BB_GFPT devices, which defines the period that elapses, following detection/indication of a GFPT_WAN link failure, before a GFPT_WAN Down condition is declared for the purposes of the state machine described in 6.4.2. The criteria for such detection are WAN-specific and outside the scope of this standard.

3.4 FC-BB_PW definitions

3.4.1 Customer Edge (CE): A device where one end of a service originates and/or terminates. The CE is not aware is it using an emulated service rather than a native service (see RFC3985).

3.4.2 Multiprotocol Label Switching (MPLS): A data-carrying mechanism that belongs to the family of packet switched networks (see 3.4.5).

3.4.3 Provider Edge (PE): A device that provides PWE3 (see 3.4.4) to a CE (see 3.4.1).

3.4.4 Pseudowire Emulation Edge-to-Edge (PWE3): A mechanism that emulates the essential elements of an emulated service from one PE (see 3.4.3) to one or more PEs over a PSN (see 3.4.5).

3.4.5 Packet Switched Network (PSN): In the context of PWE3, a network using IP or MPLS as the mechanism for packet forwarding.

3.5 FC-BB_E definitions

3.5.1 ENode (FCoE Node): A Fiber Channel node (see FC-FS-3) that is able to transmit FCoE frames using one or more ENode MACs.

3.5.2 ENode MAC: A Lossless Ethernet MAC coupled with an FCoE Controller in an ENode.

3.5.3 ENode MAC address: The MAC address used by the FCoE Controller on an ENode MAC for the FCoE Initialization Protocol (FIP).

3.5.4 Fabric Provided MAC Address (FPMA): A MAC address that is assigned by an FCF to a single ENode MAC, and is not assigned to any other MAC within the same Ethernet VLAN. A Fabric Provided MAC Address is associated with a single VN_Port at that ENode MAC.

3.5.5 FC-MAP (Mapped Address Prefix): In a Fabric Provided MAC Address, the required value for the upper 24 bits of a MAC address assigned to a VN_Port.

3.5.6 FCF (FCoE Forwarder): A Fibre Channel Switching Element (see FC-SW-5) that is able to forward FCoE frames across one or more FCF-MACs, and that optionally includes one or more Lossless Ethernet bridging elements and/or a Fibre Channel Fabric interface.

3.5.7 FCF-MAC: A Lossless Ethernet MAC coupled with an FCoE Controller in an FCF.

3.5.8 FCF-MAC address: The MAC address of an FCF-MAC.

3.5.9 FCoE Controller: A functional entity, coupled with a Lossless Ethernet MAC, instantiating and de-instantiating VE_Ports, VF_Ports, VN_Ports, and/or FCoE_LEPs.

3.5.10 FCoE Entity: The interface, containing one or more FCoE_LEPs, between a VN_Port, a VF_Port, or a VE_Port, and a Lossless Ethernet MAC.

3.5.11 FCoE frame: An Ethernet frame (see IEEE 802.3-2008) that contains an FCoE PDU (see 3.5.13).

3.5.12 FCoE_LEP (FCoE Link End-Point): 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.

3.5.13 FCoE PDU: A PDU identified by the FCoE Ethernet Type that encapsulates a byte-encoded FC frame (see 7.7).

3.5.14 FIP frame: An Ethernet frame (see IEEE 802.3-2008) containing a FIP PDU.

3.5.15 FIP PDU: A PDU identified by the FIP Ethernet Type that encapsulates one or more FIP operations (see 7.8.6).

3.5.16 Lossless Ethernet bridging element: An Ethernet bridging function operating across Lossless Ethernet MACs.

3.5.17 Lossless Ethernet MAC: A full duplex Ethernet MAC implementing extensions to avoid Ethernet frame loss due to congestion (e.g., the PAUSE mechanism (see IEEE 802.3-2008) or the Priority-based Flow Control mechanism (see IEEE 802.1Qbb)).

3.5.18 Lossless Ethernet network: An Ethernet network composed only of full duplex links, Lossless Ethernet MACs, and Lossless Ethernet bridging elements (see 4.4.4).

3.5.19 Multicast MAC address: A MAC address associated with a group of logically related Ethernet stations on an Ethernet network and called a Multicast-Group Address in IEEE 802.3-2008.

3.5.20 PE_Port (Physical E_Port): The LCF within the Fabric that attaches to another PE_Port through a native FC link (see FC-SW-5).

3.5.21 PF_Port (Physical F_Port): The LCF within the Fabric that attaches to a PN_Port through a native FC link (see FC-SW-5).

3.5.22 PN_Port (Physical N_Port): An LCF that supports only VN_Ports (see FC-FS-3).

3.5.23 Server Provided MAC Address (SPMA): A MAC address that is assigned by an ENode to a single one of its ENode MACs, and is not assigned to any other MAC within the same Ethernet VLAN. A Server Provided MAC Address may be associated with more than one VN_Port at that ENode MAC.

3.5.24 Unicast MAC address: A MAC address associated with a particular Ethernet station on an Ethernet network and called an Individual Address in IEEE 802.3-2008.

3.5.25 Vendor_ID: An 8-byte ASCII string, the value of which shall be assigned by INCITS Technical Committee T10 (see http://www.t10.org/lists/2vid.htm), used to uniquely identify an organizational entity.

3.5.26 VE_Port (Virtual E_Port): An instance of the FC-2V sublevel of Fibre Channel that communicates with another VE_Port (see FC-SW-5) and that is dynamically instantiated on successful completion of a FIP ELP Exchange.

3.5.27 VE_Port/FCoE_LEP pair: A VE_Port and its associated FCoE_LEP.

3.5.28 VF_Port (Virtual F_Port): An instance of the FC-2V sublevel of Fibre Channel that communicates with one or more VN_Ports (see FC-SW-5) and that is dynamically instantiated on successful completion of a FIP FLOGI Exchange.

3.5.29 VF_Port/FCoE_LEP pair: A VF_Port and one of its associated FCoE_LEPs.

3.5.30 VF_Port/FCoE_LEP pair: A VN_Port and its associated FCoE_LEP.

3.5.31 Virtual Link: The logical link connecting two FCoE_LEPs (see 7.5).

3.5.32 VN_Port (Virtual N_Port): An instance of the FC-2V sublevel of Fibre Channel that operates as an N_Port (see FC-FS-3) and is dynamically instantiated on successful completion of a FIP FLOGI or FIP NPIV FDISC Exchange.

3.5.33 VN_Port MAC address: The MAC address used by an ENode for a particular VN_Port.

3.6 Editorial Conventions

In FC-BB-5, a number of conditions, mechanisms, sequences, parameters, events, states, or similar terms are printed with the first letter of each word in uppercase and the rest lowercase (e.g., Exchange, Sequence). Any lowercase uses of these words have the normal technical English meanings.

Lists sequenced by letters (e.g., a-red, b-blue, c-green) show no ordering relationship between the listed items. Numbered lists (e.g., 1-red, 2-blue, 3-green) show an ordering relationship between the listed items.

In case of any conflict between figure, table, and text, the text, then tables, and finally figures take precedence. Exceptions to this convention are indicated in the appropriate clauses.

In all of the figures, tables, and text of this document, the most significant bit of a binary quantity is shown on the left side. Exceptions to this convention are indicated in the appropriate clauses.

Data structures in this standard are displayed in Fibre Channel format (i.e., “big-endian”), while specifications originating in IEEE and IETF may display data structures in Ethernet format (i.e., “little-endian”).

When the value of the bit or field is not relevant, x or xx appears in place of a specific value. If a field or a control bit in a frame is specified as not meaningful, the entity that receives the frame shall not check that field or control bit.

Numbers that are not immediately followed by lower-case b or h are decimal values.

Numbers immediately followed by lower-case b (xxb) are binary values.

Numbers or upper case letters immediately followed by lower-case h (xxh) are hexadecimal values.

3.7 List of commonly used acronyms and abbreviations

Abbreviations and acronyms applicable to this standard are listed. Definitions of several of these items are included in clause 3.

3.7.1 General

BB Backbone
EBP Exchange B_Access Parameters
ELP Exchange Link Parameters
EOF End of Frame
EOFni End of Frame Invalid
ESC Exchange Switch Capabilities
F_BSY Fabric Busy
FCS Frame Check Sequence
FC-FS-3 Fibre Channel - Framing and Signaling - 3
FC-LS-2 Fibre Channel - Link Services - 2
FC-SP Fibre Channel - Security Protocol
FC-SW-5 Fibre Channel - Switched Fabric - 5
FLOGI Fabric Login
ISL Inter-switch Link
ITU-T International Telecomm. Union - Telecommunication Standardization Section
K_A_TOV Keep Alive Timeout value
LKA Link Keep Alive
LS_ACC Link Service Accept Reply Frame
LS_RJT Link Service Reject Reply Frame
LSB least significant byte
MSB most significant byte
P_BSY N_Port Busy
PDU Protocol Data Unit

PLOGI N_Port Login
SOF Start of Frame
SW_ACC Switch Fabric Internal Link Service Accept
SW_ILS Switch Fabric Internal Link Services
SW_RJT Switch Fabric Internal Link Service Reject
WAN Wide Area Network

3.7.2 FC-BB_IP

B_Access B_Access Portals
CSM Control and Service Module
FCIP FC over TCP/IP
FCIP_DE FCIP Data Engine
FCIP_LEP FCIP Link Endpoint
IETF IETF Internet Engineering Task Force (www.ietf.org)
PMM Platform Management Module
RFC Request For Comment
VE_Port Virtual E_Port

3.7.3 FC-BB_GFPT

ASFC Alternate Simple Flow Control
GFP Generic Framing Procedure
GFPT **(**Asynchronous) Transparent Generic Framing Procedure
GFPT_WAN GFPT Wide Area Network
LEM Login Exchange Monitor

3.7.4 FC-BB_PW

EF PHB Expedited Forwarding Per-Hop Behavior
PE Provider Edge
PW pseudowire
MPLS/PW Multiprotocol Label Switching pseudowire

3.7.5 FC-BB_E

ACE Access Control Entry
ACL Access Control List
D_A_TOV Discovery Advertisement Timeout Value
ENode FCoE Node
FC-MAP FCoE Mapped Address Prefix
FCF FCoE Forwarder
FCF-MAC FCoE Forwarder Media Access Control
FCoE_LEP FCoE Link Endpoint
FIP FCoE Initialization Protocol
FPMA Fabric Provided MAC Address
LAN Local Area Network
MAC Media Access Control
PHY Physical Layer
SPMA Server Provided MAC Address
TLV Type, Length, Value
VE_Port Virtual E_Port
VF_Port Virtual F_Port
VLAN Virtual Local Area Network

VN_Port Virtual N_Port

3.8 Symbols

Unless indicated otherwise, the following symbol has the listed meaning.

!= not equal

3.9 Keywords

3.9.1 ignored: A keyword used to describe an unused bit, byte, word, field or code value. The contents or value of an ignored bit, byte, word, field or code value shall not be examined by the receiving device and may be set to any value by the transmitting device.

3.9.2 invalid: A keyword used to describe an illegal or unsupported bit, byte, word, field or code value. Receipt of an invalid bit, byte, word, field or code value shall be reported as an error.

3.9.3 mandatory: A keyword indicating an item that is required to be implemented as defined in this standard.

3.9.4 may: A keyword that indicates flexibility of choice with no implied preference (equivalent to “may or may not”).

3.9.5 may not: A keyword that indicates flexibility of choice with no implied preference (equivalent to “may or may not”).

3.9.6 optional: A keyword that describes features that are not required to be implemented by this standard. However, if any optional feature defined by this standards is implemented, then it shall be implemented as defined in this standard.

3.9.7 reserved: A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients are not required to check reserved bits, bytes, words or fields for zero values. Receipt of reserved code values in defined fields shall be reported as an error.

3.9.8 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

3.9.9 should: A keyword indicating flexibility of choice with a strongly preferred alternative; equivalent to the phrase “it is strongly recommended”.

3.9.10 x or **xx:** The value of the bit or field is not relevant.

4 FC-BB-5 Structure and Concepts

4.1 FC-BB-5 backbone mappings

FC-BB-5 models (i.e., FC-BB_IP, FC-BB_GFPT, FC-BB_PW, and FC-BB_E), specified in this standard, define mappings for transporting Fibre Channel over different network technologies.

The FC-BB_IP model uses TCP connections over IP networks. The FC-BB_GFPT model makes use of the Asynchronous Transparent Generic Framing Procedure (GFPT) (see ITU-T Rec. G.7041/Y.1303). GFPT may be used for adaptation to different transport facilities including SONET, SDH, OTN and PDH. Details regarding the mapping of GFPT-adapted traffic to such transport facilities are elaborated in various ITU-T standards (see 2.4). The FC-BB_PW model uses PW connections over MPLS networks. The FC-BB_E model uses Lossless Ethernet links (i.e., full duplex Ethernet links extended to become lossless, see 4.4.4).

A second important distinction among the mappings discussed in this standard relates to supported architectures (i.e., network and/or link topologies) and the place of the defined devices within them. FC-BB_IP defines an FC Entity that consists of:

a) one or more Virtual E_Port (VE_Port) virtual ISL(s) that interconnect the E_Ports of external FC switches and N_Ports of external end devices; or
b) one or more B_Access virtual ISL(s) that interconnect the E_Ports of external FC switches.

The B_Access virtual ISL(s) are part Fabric bridge device with resident switch-facing interfaces called B_Ports. B_Ports have selected Fabric functions (see FC-SW-5). B_Ports are Fabric ports, and FC-BB_IP devices are components of an FC Fabric. FC-BB_IP also defines support for FC bridge devices with Fabric-facing B_Ports. However, FC-BB_IP also supports functional integration within an FC Switch. Thus FC-BB_IP devices may also have E_Ports and F_Ports.

FC-BB_GFPT and FC_BB_PW define a device that is not a component of a Fabric, and supports no Fabric functionality. Instead, it interconnects two Fibre Channel physical ports (i.e., attached FC_Ports), appearing architecturally as a wire to those ports.

FC-BB_E defines end devices (i.e., ENodes) and Fabric devices (i.e., FCFs). ENodes are Fibre Channel nodes (see FC-FS-3) that are able to transport Fibre Channel over Lossless Ethernet. FCFs are Fibre Channel Switching Elements (see FC-SW-5) that are able to transport Fibre Channel over Lossless Ethernet.

4.2 FC-BB-5 reference models

FC-BB-5 defines reference models corresponding to the FC-BB_IP, FC-BB_GFPT, FC-BB_PW, and FC-BB_E models. These reference models are shown in figure 5, figure 6, figure 7, and figure 8 respectively.

The FC-BB_IP model supports the attachment of FC switches (i.e., E_Ports) via one or more B_Ports or E_Ports and the attachment of N_Ports via one or more F_Ports. The FC-BB_GFPT and FC-BB_PW models support the attachment of N_Ports, F_Ports, and E_Ports, and the following Fibre Channel port interconnections:

a) N_Port to N_Port;
b) N_Port to F_Port; and
c) E_Port to E_Port.

The FC-BB_E model supports the operation of VN_Ports (see FC-FS-3) in ENodes and of VF_Ports and VE_Ports (see FC-SW-5) in FCFs.

Table 2 summarizes the resident FC_Port types for the different FC-BB-5 models.

Table 2 – Models and resident FC_Port types

	Reference Model			
	FC-BB_IP	**FC-BB_GFPT**	**FC-BB_PW**	**FC-BB_E[a]**
Resident FC_Port Type(s)	B_Port, E_Port, F_Port	None (FC Physical Interface)	None (FC Physical Interface)	N_Port, E_Port, F_Port
a) Resident FC_Ports are optional for FC-BB_E.				

In figure 5 (i.e., FC-BB_IP), frames destined for a remote FC network enter a B_Port, an E_Port, or an F_Port, and are forwarded on the IP network to their destination.

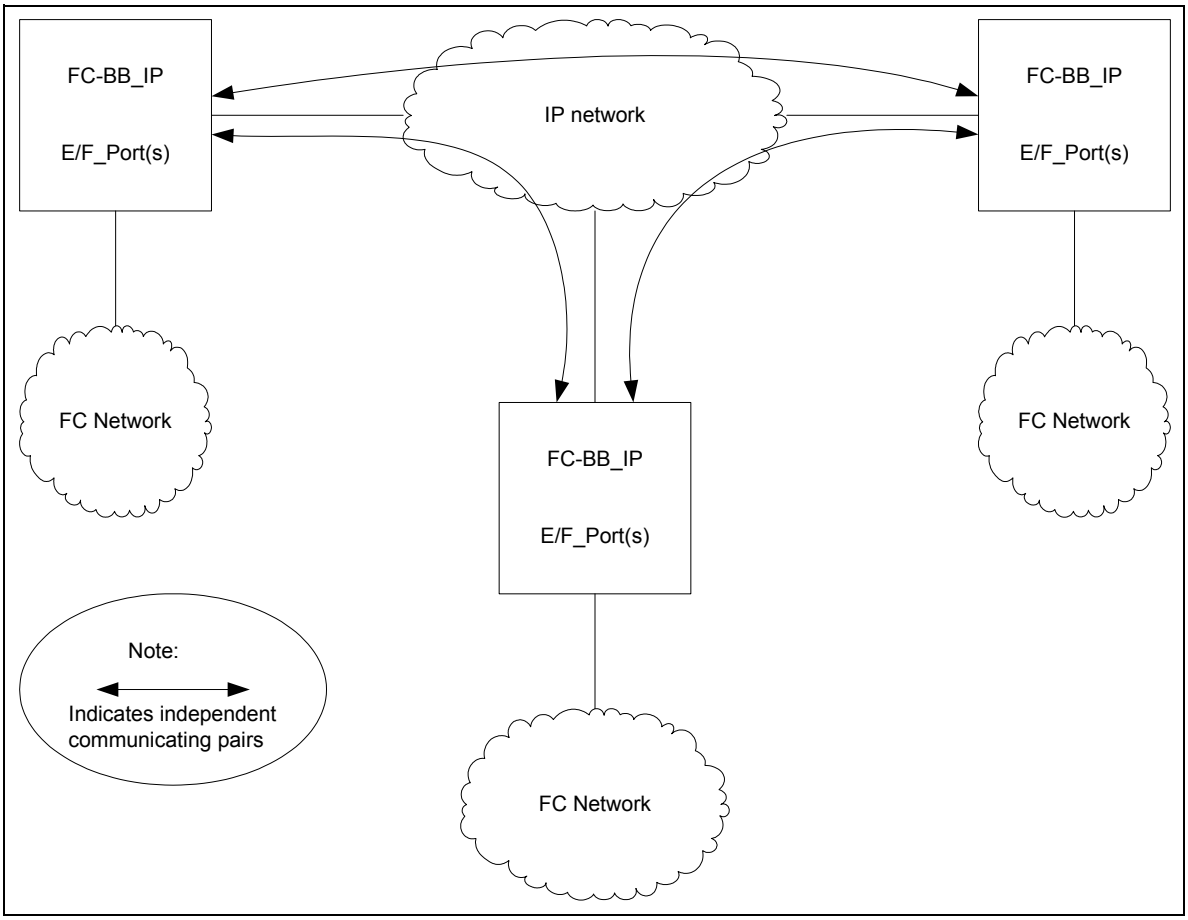

Figure 5 – FC-BB_IP reference model

In figure 6 (i.e., FC-BB_GFPT), FC physical signals (i.e., relevant 8B/10B codewords) enter an FC physical port on an FC-BB_GFPT device, and are forwarded on the SONET/SDH/OTN/PDH network to their destination.

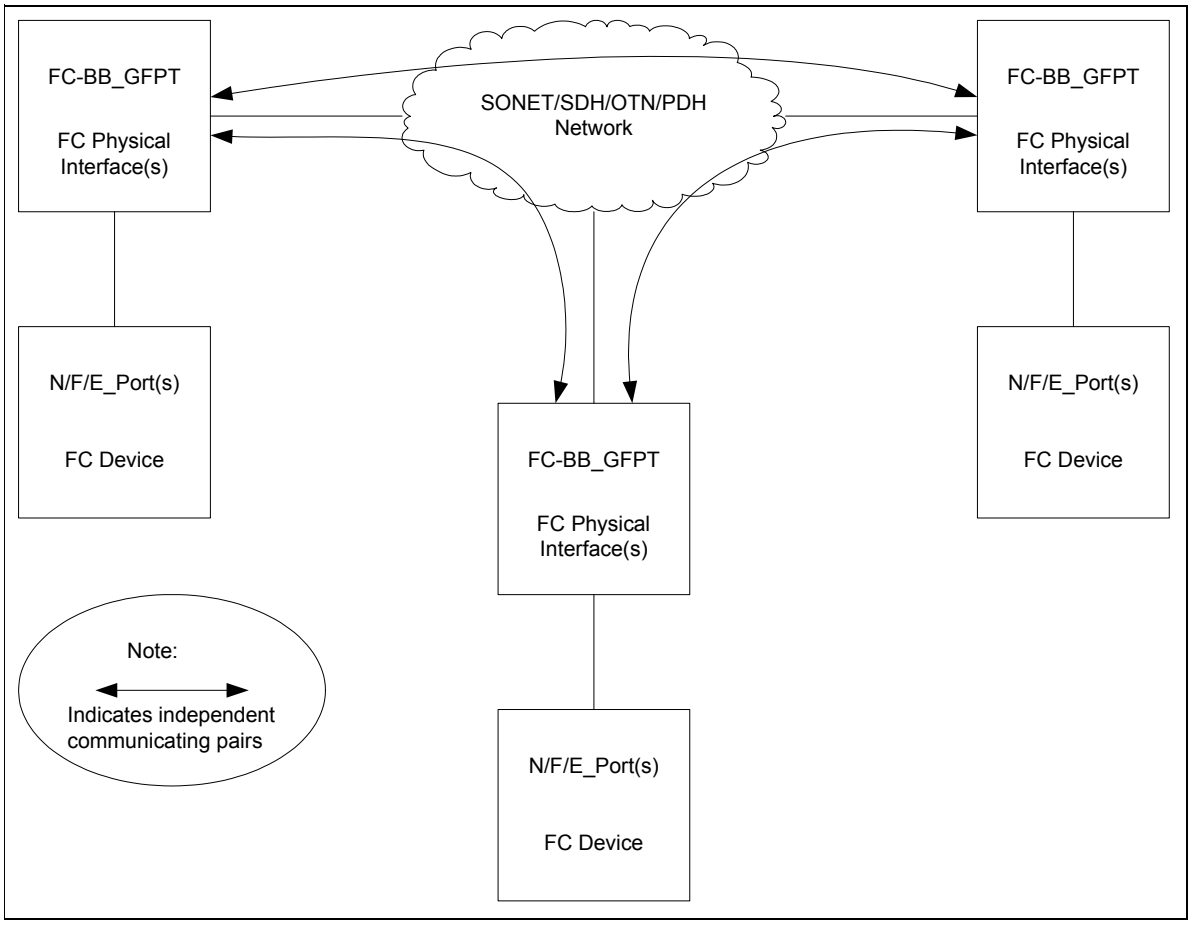

Figure 6 – FC-BB_GFPT reference model

In figure 7 (i.e., FC-BB_PW), FC physical signals enter an FC physical port on an FC-BB_PW device, and are forwarded on the PSN network to their destination.

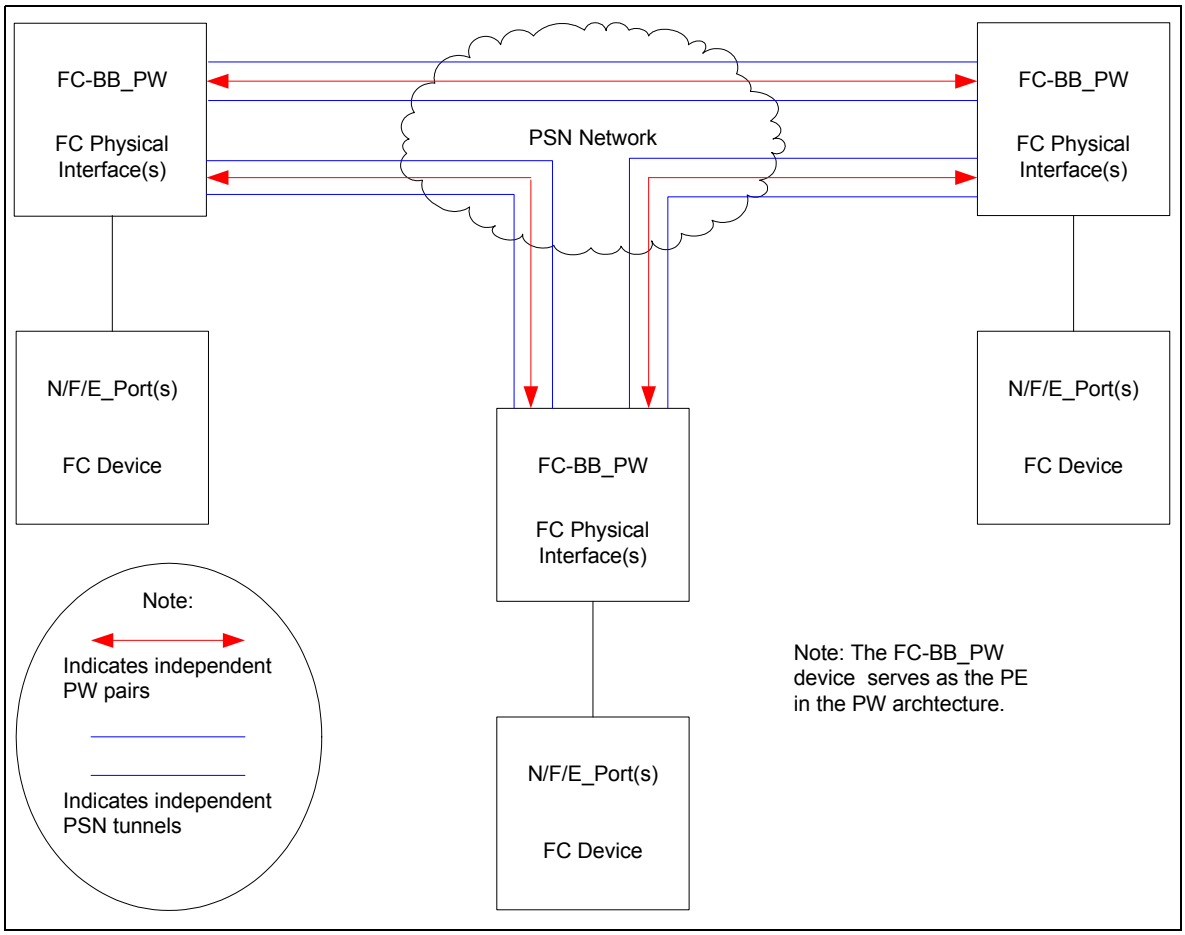

Figure 7 – FC-BB_PW reference model

Figure 8 (i.e., FC-BB_E) shows the VN_Port to VF_Port reference model and the VE_Port to VE_Port reference model. FCoE frames originated by a VN_Port are transported over the Lossless Ethernet network to the VF_Port the VN_Port is logged in with. FCoE frames originated by a VE_Port

are transported over the Lossless Ethernet network to the VE_Port that the VE_Port is logically connected to.

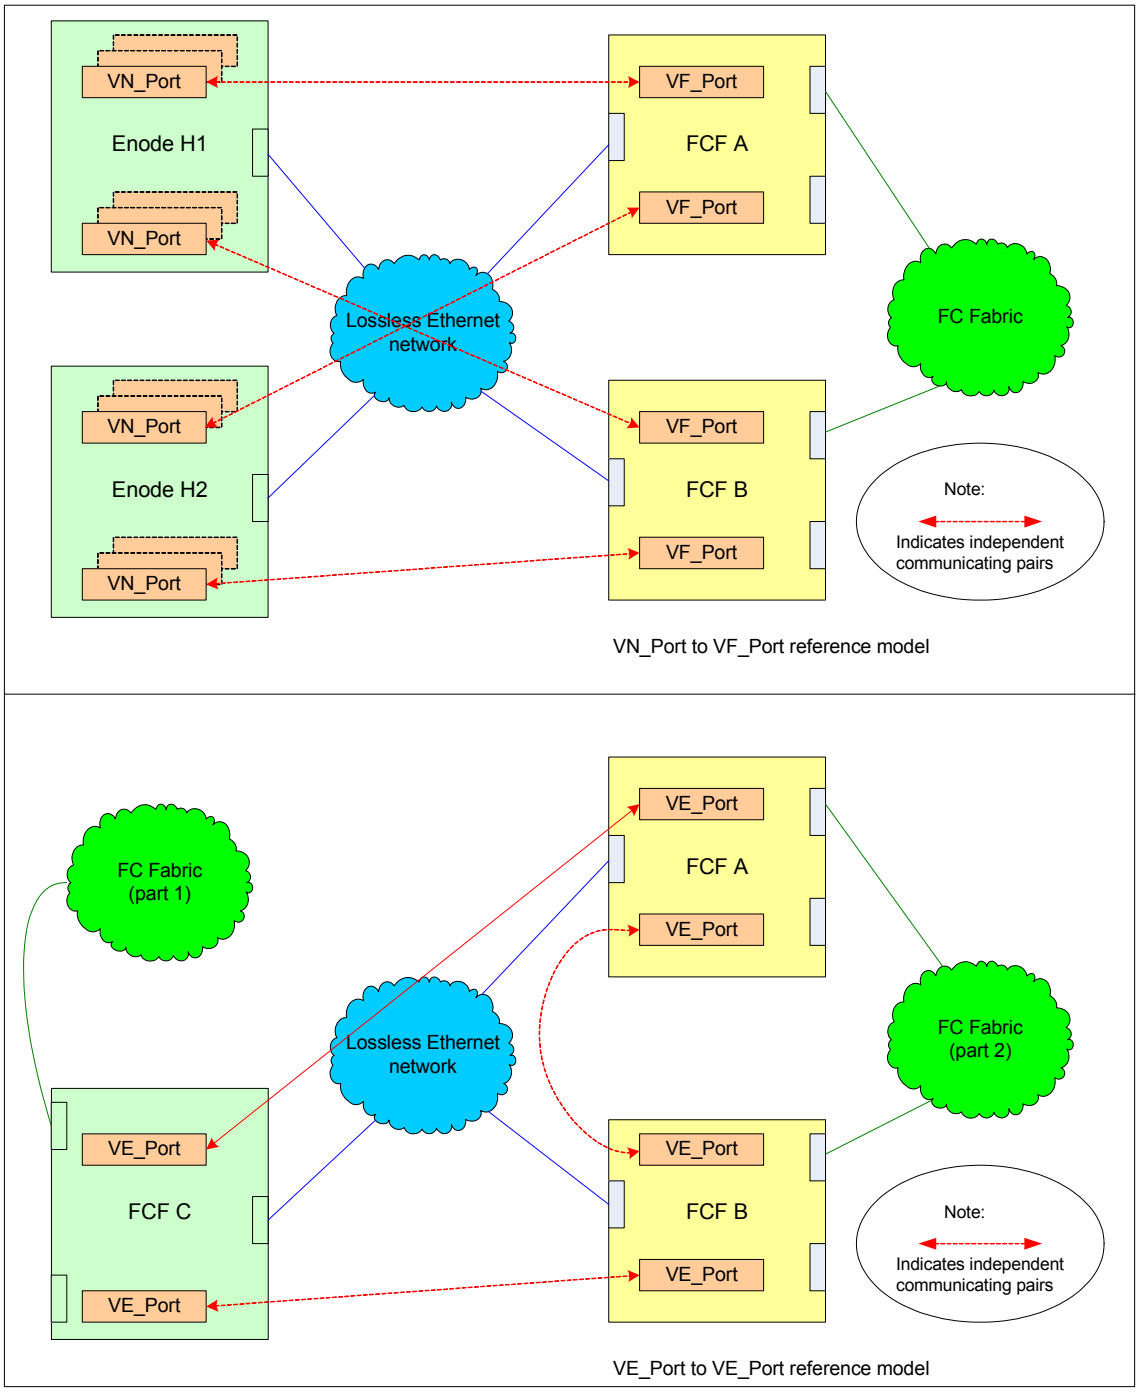

Figure 8 – FC-BB_E reference model

4.3 FC-BB-5 models overview

4.3.1 FC-BB_IP

The FC-BB_IP model defines the means by which Fibre Channel networks interface with and connect across an IP network. FC-BB_IP makes use of the FCIP standard (see RFC 3821) to define the mapping and control required by the TCP/IP protocol and the FC frame encapsulation standard (see RFC 3643) to define the encapsulation. FC-BB_IP also defines the connection management, addressing, time synchronization, discovery, security, switching, routing, and error recovery required to support Fibre Channel over TCP/IP. FC-BB_IP is independant of the underlying physical technology that exists beneath the IP layer. In this sense, the IP network could use SONET, Gigabit Ethernet, or any other link-level technology below it.

FC-BB_IP encapsulates byte-encoded Class 2, 3, or F Fibre Channel frames into a suitable format (i.e., encapsulated FC frames*)* for carriage over the IP network. Subclause 5.5 describes encapsulated FC frames in detail. The TCP/IP protocol suite provides a reliable transport of frames over the IP network. TCP provides flow control and error recovery.

The FC-BB_IP protocol provides mechanisms to create VE_Port or B_Access connectivity over the IP network (see 5).

The FC-BB_IP device interfaces to attached FC_Ports are FC physical interfaces operating at standard rates. The FC-BB_IP devices may support automatic WAN link speed negotiation with the attached IP networks.

4.3.2 FC-BB_GFPT

The FC-BB_GFPT model defines the means by which FC physical links may be extended over any WAN transport infrastructure for which GFP mapping is defined. FC-BB_GFPT supports the interconnection of arbitrary, legal, non-Arbitrated Loop FC_Port combinations, imposing no requirements, and making no suppositions regarding the topology, or even the presence of an FC Fabric. FC-BB_GFPT supports efficient transport of FC data over transport facilities of arbitrary bandwidths and potentially large distances. FC-BB_GFPT supports Class 2, 3, and F traffic.

FC-BB_GFPT devices do not generate FC frames and do not directly participate in port initialization or other Exchanges. FC-BB_GFPT devices are exempt from any requirements regarding FC_Port authentication (see FC-SP), and they do not impede or interfere with any such processes that may occur between the attached and interconnected FC_Ports. FC-BB_GFPT devices have no FC identity or visibility, and administratively they may be kept separate and distinct from FC Fabrics and ports.

FC_Ports are interconnected pair-wise over SONET/SDH/OTN/PDH networks, via FC-BB_GFPT devices, in a point-to-point fashion. Although multiple FC_Ports may interface with a single FC-BB_GFPT device, each opposing FC_Port pair is connected via a dedicated Transport Trail (e.g., a contiguously or virtually-concatenated group). Since trail and access section configurations may differ, FC-BB_GFPT devices have both physical interfaces to the transport network, and individual FC-BB_GFPT devices may have more than one such physical interface, as well as logical interfaces associated with individual circuits. Logical interfaces are referred to as GFPT_WAN interfaces. A GFPT_WAN interface corresponds to a specific Transport Trail, and always to a single attached FC_Port pair. Governance of the relationship of GFPT_WAN interfaces to physical SONET/SDH/OTN/PDH interfaces, and of any changes of such relationships (e.g., as may occur during network protection events), is specified in the appropriate ITU-T and ANSI-T1 standards (see clause 2), and is therefore outside the scope of this standard. Multiple GFPT_WAN links originating on one FC-BB_GFPT device may be terminated on different, and geographically disparate, FC-

BB_GFPT devices. The routing and provisioning of network facilities underlying GFPT_WAN links is outside the scope of this standard.

The FC-BB_GFPT device interfaces to attached FC_Ports are FC physical interfaces operating at standard rates. The FC physical interfaces on FC-BB_GFPT devices may support link speed negotiation with the attached FC_Ports.

4.3.3 FC-BB_PW

The FC-BB_PW model defines the means by which FC physical links may be extended over a wide area MPLS network. FC-BB_PW is specified in conjunction with draft-ietf-pwe3-fc-encap-09 to define the mapping and control required by the MPLS/PW protocol. According to this model the FC-BB_PW device serves as a PE network element in the PW architecture. FC-BB_PW is independant of the physical-level and link-level technologies that exist beneath the MPLS layer.

FC-BB_PW supports the interconnection of FC_Ports transporting Class 2, 3, and F traffic. FC-BB_PW encapsulates byte-encoded FC frames and a selected set of Primitive Signals and Primitive Sequences into PW PDUs for transport over the MPLS network. FC-BB_PW utilizes reliable transport of FC traffic over the MPLS network provided by the PW termination layer as specified in draft-ietf-pwe3-fc-encap-09.

The FC-BB_PW entity does not generate FC frames and does not directly participate in port initialization or other Exchanges. The FC-BB_PW entity is exempt from any requirement regarding FC_Port authentication (see FC-SP), and it does not impede or interfere with any such processes that may occur between the attached and interconnected FC_Ports. The FC-BB_PW entity has no FC identity or visibility, and administratively it may be kept separate and distinct from FC Fabrics and ports.

4.3.4 FC-BB_E

The FC-BB_E model defines the means by which Fibre Channel frames are transported over a Lossless Ethernet network (see 4.4.4). Although a generic Ethernet network may lose frames due to congestion, a proper implementation of appropriate Ethernet extensions (e.g., the PAUSE mechanism defined in IEEE 802.3-2008) allows a full duplex Ethernet link to provide a lossless behavior similar to the one provided by the buffer-to-buffer mechanism (see FC-FS-3). 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 transporting FCoE frames).

FC-BB_E encapsulates byte-encoded Class 2, 3, or F Fibre Channel frames into a suitable format (i.e., FCoE frames) for transport over the Lossless Ethernet network. See 7.7.

The FC-BB_E protocol provides mechanisms to create VN_Port to VF_Port virtual links and to create VE_Port to VE_Port virtual links.

4.4 FC-BB-5 requirements

4.4.1 Fibre Channel Class support

Class F shall be supported between FC-BB_IP, FC-BB_GFPT, and FC-BB_PW devices. Class 2 and Class 3 may be supported between FC-BB_IP, FC-BB_GFPT, and FC-BB_PW devices.

Class 3 shall be supported between FC-BB_E devices. Class 2 and Class F may be supported between FC-BB_E devices.

NOTE 1 – A Lossless Ethernet network configuration that includes Lossless Ethernet bridges may not provide sufficient functionality to support all aspects of Class 2 service. FCFs and ENodes may indicate support for Class 2 service in the presence of Lossless Ethernet bridges, but such a configuration should be managed accordingly.

4.4.2 Payload transparency

4.4.2.1 FC-BB_IP

Arriving Class 2, 3, and F Fibre Channel frames from an FC network and destined to a remote FC network shall be encapsulated using the FC-BB_IP defined mechanisms and transmitted to the appropriate FC-BB_IP device.

Arriving encapsulated frames received from remote FC-BB_IP device shall be de-encapsulated and sent to an FC network.

Primitive Signals and Primitive Sequences shall not be transported between FC-BB_IP devices.

4.4.2.2 Transparent FC-BB (FC-BB_GFPT and FC-BB_PW)

FC frames inbound from one attached FC_Port shall be delivered to the remote FC_Port in native form (i.e., without further encapsulation) across the transport network according to the adaptation processes described in 6.4.8.1 for FC-BB_GFPT devices and 6.4.8.2 for FC-BB_PW devices. Frames received from remote Transparent FC-BB devices shall be forwarded to the attached FC_Port. Selected ELP, SW_ACC, FLOGI, PLOGI, and LS_ACC frames may be subject to inspection and/or minor modifications, in transiting one or the other Transparent FC-BB device, as described in 6.4.4.

Primitive Signals transmitted by an attached FC_Port are forwarded across the transport network for delivery to the remote FC_Port according to the rules described in clause 6. Primitive Signals are forwarded by an FC-BB_GFPT device in native form (i.e., without further encapsulation) according to the adaptation processes described in 6.4.8.1. Primitive Signals are forwarded by an FC-BB_PW device encapsulated within control frames as specified in draft-ietf-pwe3-fc-encap-09.

Primitive Sequences transmitted by an attached FC_Port are forwarded across the transport network for delivery to the remote FC_Port according to the rules described in clause 6. Primitive Sequences are forwarded by an FC-BB_GFPT device in native form (i.e., without further encapsulation) according to the adaptation (including rate adaptation) processes described in 6.4.8.1. Primitive Sequences are forwarded by an FC-BB_PW device encapsulated within control frames as specified in draft-ietf-pwe3-fc-encap-09.

4.4.2.3 FC-BB_E

Class 2, 3, and F Fibre Channel frames arriving from a VN_Port, a VF_Port, or a VE_Port shall be encapsulated in FCoE frames and transmitted to the appropriate FC-BB_E device.

FCoE frames received from a remote FC-BB_E device shall be de-encapsulated and sent to the appropriate VN_Port, VF_Port, or VE_Port.

Primitive Signals and Primitive Sequences shall not be forwarded between FC-BB_E devices.

4.4.3 Latency delay and timeout value

FC-BB_IP shall ensure that the incoming encapsulated FC frames whose FCIP Transit Time (FTT) exceeds 1/2 E_D_TOV shall be discarded and not admitted into the FC network. Fibre Channel timeout values shall be administratively set to accommodate the FTT.

FC-BB_IP shall allow Class F encapsulated FC frames to be transmitted with a zero timestamp value.

Transparent FC-BB requires that the latency between two Transparent FC-BB devices be:

a) no more than one-half of the E_D_TOV value of the attached devices for frames;
b) no more than one-half of the R_T_TOV value of the attached devices for Primitive Sequences; and
c) within the R_A_TOV values of the attached Fabric(s), if any.

The Lossless Ethernet bridges used to build a Lossless Ethernet network suitable for FC-BB_E usage should enforce a maximum bridge transit delay of 500 ms (see IEEE 802.1Q-2005), according to the best practice used by Fibre Channel switches.

4.4.4 QoS and bandwidth

FC-BB_IP recommends that some form of preferential QoS be used for the FCIP traffic in the IP network to minimize latency and packet drops although no particular form of QoS is recommended. See RFC 3821.

FC-BB_GFPT has no specific transport service requirements.

FC-BB_PW recommends that Primitive Sequences are carried with low latency and no loss over the MPLS network. In addition to these properties, FC data traffic should be provided with assurance of some amount of bandwidth, however no specific recommendation is made in this standard. The Differentiated Services EF PHB (see RFC 3246) is an example of a mechanism that may be used for FC-BB_PW traffic management.

FC-BB_E is intended to operate over an Ethernet network that does not discard frames in the presence of congestion. Such an Ethernet network is called Lossless Ethernet in this standard. Lossless Ethernet may be implemented through the use of some Ethernet extensions. A possible Ethernet extension to implement Lossless Ethernet is the PAUSE mechanism defined in IEEE 802.3-2008. Another possible Ethernet extension to implement Lossless Ethernet is the Priority-based Flow Control (PFC) mechanism defined in IEEE 802.1Qbb. When PFC is used to implement Lossless Ethernet, FCoE frames shall use a lossless priority (see IEEE 802.1Qbb).

4.4.5 In-order delivery

FC-BB_IP devices shall guarantee in-order delivery of frames within the scope of any TCP connection.

FC-BB_GFPT devices shall provide in-order delivery within each provisioned Transport Trail for all transmitted data (i.e., frames, Primitive Signals and Primitive Sequences), as discussed and with the exceptions detailed in clause 6.

FC-BB_PW devices shall provide in-order delivery within each provisioned PW for all transmitted data (i.e., frames, Primitive Signals and Primitive Sequences), as discussed and with the exceptions detailed in clause 6.

FC-BB_E devices shall guarantee in-order delivery of FCoE frames within the Lossless Ethernet network.

4.4.6 Flow control

FC-BB_IP devices shall ensure that TCP flow control and error recovery acts in proper concert with the Fibre Channel BB_Credit flow control mechanism.

Flow control at E_Ports, F_Ports, VE_Ports, and B_Ports shall operate as defined in FC-SW-5.

The Alternate Simple Flow Control (ASFC) mechanism (see 6.4.4) shall be used between FC-BB_GFPT devices. Flow control on FC_Port-facing links is specified in clause 6.

Flow control at FC-BB_GFPT and FC-BB_PW physical interfaces to attached FC devices shall operate as defined in FC-SW-5 or FC-FS-3, as appropriate.

Flow control between two FC-BB_PW devices shall operate as specified in draft-ietf-pwe3-fc-encap-09.

FC-BB_E devices rely on Ethernet extensions for proper flow control of FCoE frames. Ethernet extensions achieve the same level of service provided by the Fibre Channel buffer-to-buffer flow control (see FC-FS-3). Suitable Ethernet extension for FC-BB_E usage are the PAUSE mechanism defined in IEEE 802.3-2008 and the Priority-based Flow Control mechanism defined in IEEE 802.1Qbb.

4.5 FC-BB-5 SW_ILS codes

Table 3 shows the SW_ILS codes allocated for FC-BB-5 use.

Table 3 – FC-BB-5 SW_ILS codes

Encoded Value (hex)	Description	Abbr.	Reference
28 03 00 00	Authentication Special Frame Request	ASF	5.6.2.3.2
28 01 00 00	Exchange B_Access Parameter	EBP	5.3.3.3.1

Table 4 shows the ELS codes allocated for FC-BB-5 use.

Table 4 – FC-BB-5 ELS codes

Encoded Value (hex)	Description	Abbr.	Reference
80 00 00 00	Link Keep Alive Request	LKA	FC-LS-2

5 FC-BB_IP Structure and Concepts

5.1 FC-BB_IP overview

Figure 9 shows a network configuration consisting of three FC-BB_IP devices. FC-BB_IP is a Fibre Channel backbone transport protocol that tunnels encapsulated FC frames across the IP network. An FC-BB_IP device has interfaces to both the IP and the FC network. The FC network interface supports multiple E_Ports/F_Ports (see figure 10) or multiple B_Ports (see figure 14).

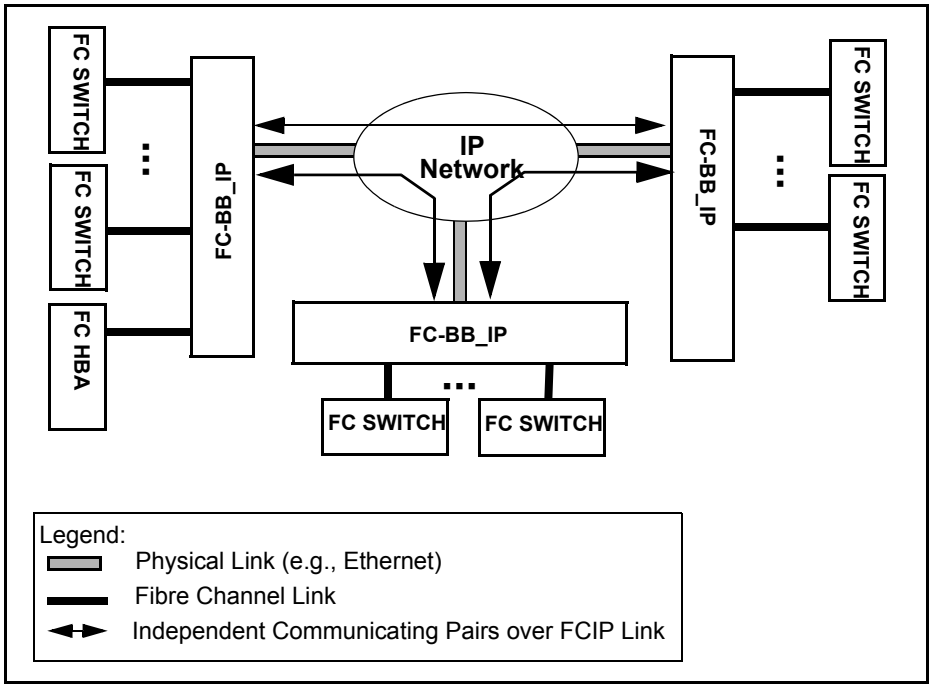

Figure 9 – FC-BB_IP network configuration

Only FC-BB_IP devices that support E_Ports or F_Ports require FC switching.

The FC-BB_IP protocol provides mechanisms to create VE_Port or B_Access connectivity over the IP network. The FC-BB_IP protocol communication occurs between pairs of FC-BB_IP devices over virtual constructs (i.e., FCIP Links) that are described in 5.2.4.2.4. Although the communication occurs between pairs of FC-BB_IP devices, a single FC-BB_IP device may communicate with more than one device at the same time (see figure 17).

NOTE 2 – Although the current scheme allows an FC-BB_IP device to independently connect to more than one FC-BB_IP device, it does not specify a point-to-multipoint connection.

The FC-BB_IP protocol uses encapsulated FC frames created by prefixing a 28-byte FC Encapsulation Header to the incoming SOF/EOF delimited FC frame (see RFC 3643). FC-BB_IP devices are not required to interpret the data content of the FC frames other than capturing and recording their SOF/EOF identities in the encapsulated FC frame. As such, FC Sequences and Exchanges are not visible to the FC-BB_IP protocol. All encapsulated FC frames are transparently transported over the IP network

FC-BB_IP devices also exchange SW_ILS control information using Class F FC frames (see figure 12 and figure 15). These FC frames are encapsulated and tunneled in the same way as the incoming FC frames.

Encapsulated FC frames join the TCP byte stream in order (see figure 18). TCP segments are created from TCP byte streams without any visibility or regard to encapsulated FC frame boundaries.

TCP flow control between two FC-BB_IP devices provides a reliable transport of encapsulated FC frames across the IP network. The only delivery order guarantee provided by TCP with respect to the FCIP protocol is the correctly ordered delivery of encapsulated FC frames within a single TCP connection. The FC Entity is expected to specify and handle all other FC frame delivery ordering requirements.

5.2 VE_Port functional model

5.2.1 FC-BB_IP interface protocol layers

Figure 10 shows the VE_Port functional model of an FC-BB_IP device that consists of the E_Port/F_Port FC interface, the FC-BB_IP interface, and the IP network interface.

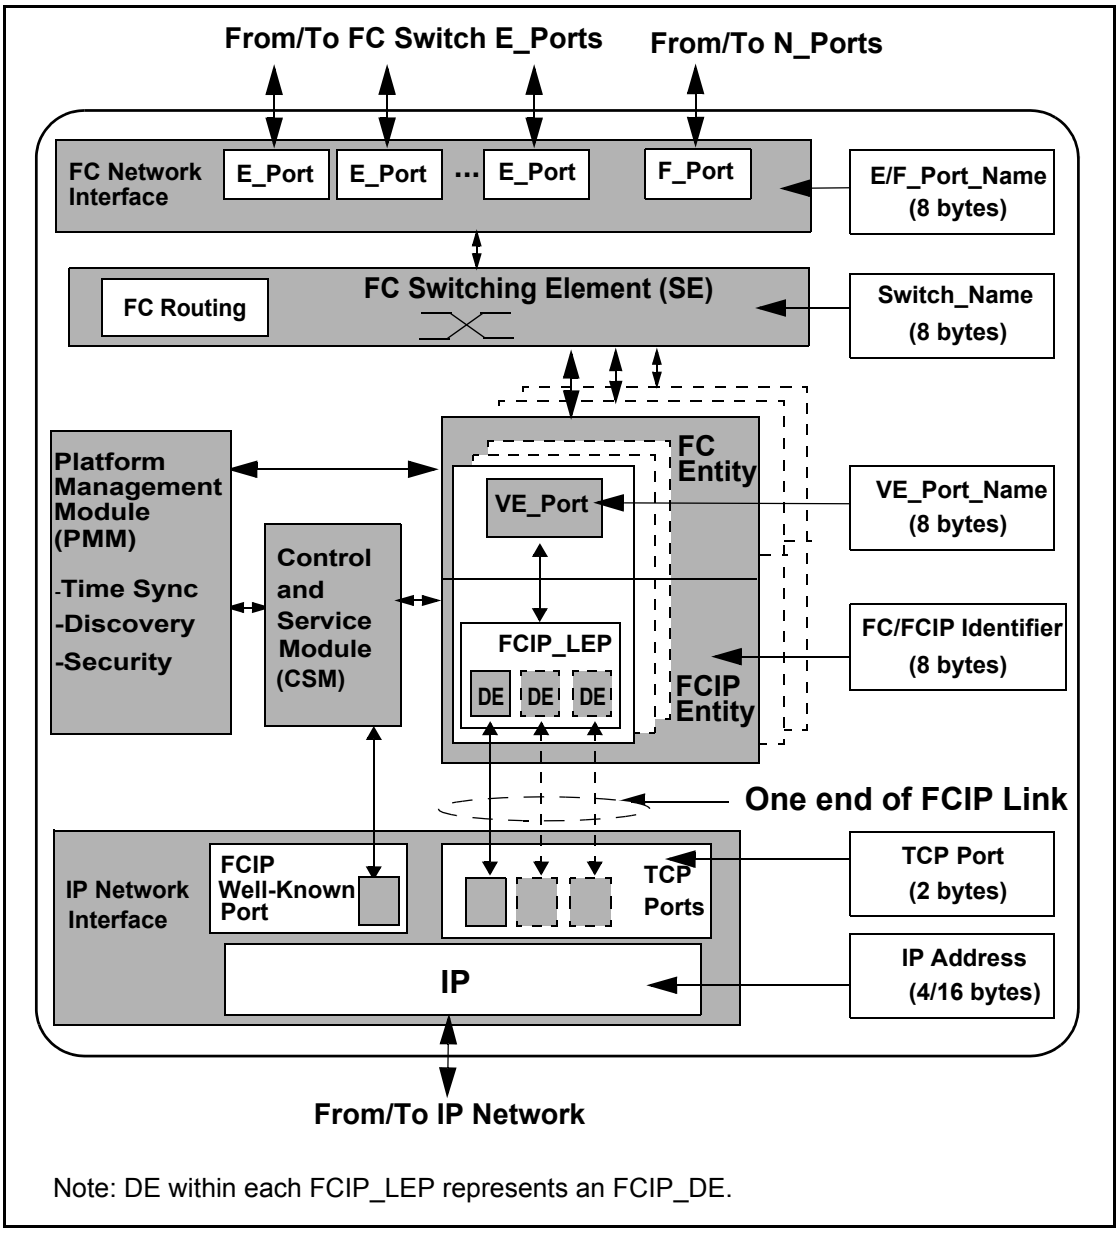

Note: DE within each FCIP_LEP represents an FCIP_DE.

Figure 10 – FC-BB_IP VE_Port functional model

The protocol layers at these interfaces are:

a) an E_Port/F_Port FC interface: FC-0, FC-1, and FC-2 levels;
b) an FC Switching Element (SE) with FC routing;
c) an FC-BB_IP protocol interface: FC Entity and FCIP Entity protocol layers; and
d) an IP network interface: TCP and IP layers.

Figure 11 illustrates the protocol layers across these interfaces.

5.2.2 E_Port/F_Port FC interface

The FC-BB_IP FC interface supports one or more E_Ports or F_Ports thus requiring the support of the FC-0, FC-1, and FC-2 levels. The E_Ports in general connect to different external FC switches, but connectivity to the same FC switch is also allowed. The data emerging from the FC-2 level is fed into an FC switching element.

The initialization of any generic E_Port or F_Port is described in FC-SW-5. An E_Port indicates its support for the ELP/ESC parameters using the ELP/ESC SW_ILS that is capable of parameter negotiation. Since FC-BB-5 does not support Class 1, the Class 1 Port Parameter VAL bit in the ELP shall be set to 0 (i.e., invalid). An ELP received at an E_Port may be rejected using SW_RJT for many reasons (see FC-SW-5).

An E_Port/F_Port is uniquely identified by an 8-byte E_Port_Name/F_Port_Name.

5.2.3 FC Switching Element (SE) with FC routing

The FC Switching Element (SE) switches and routes the incoming FC frames from the E_Port or F_Port to the proper VE_Port (see FC-SW-5). Routing is accomplished with the support of the FSPF routing protocol. Similarly, the FC SE switches and routes the data arriving from a VE_Port to the proper E_Port or F_Port.

The SE is uniquely identified by an 8-byte Switch_Name.

5.2.4 FC-BB_IP protocol interface

5.2.4.1 Major components

The FC-BB_IP protocol has interfaces to the FC network on one side and the IP network on the other. In addition to the two network interfaces, it consists of the following major components:

a) FC and FCIP Entities;
b) Control and Service Module (CSM); and
c) Platform Management Module (PMM).

FC routing occurs at a higher level than IP routing. FC/FCIP Entities themselves do not actively participate in FC frame routing. FC routing uses the FSPF protocol described in FC-SW-5. FSPF routes are mapped onto the FCIP Links interconnecting FC-BB_IP devices. An FC frame's FSPF route decides the selection of the VE_Port/FCIP_LEP pair within a selected FC/FCIP Entity pair, when multiple pairs are in use. When multiple DEs within an FCIP_LEP are in use, the selection of which FCIP_DE to use is described in 5.6.3.5.

5.2.4.2 FC and FCIP Entities

5.2.4.2.1 Function

The FC Entity is the principal interface point to the FC network on one side and, in combination with the FCIP Entity, to the IP network on the other side. The primary functions of the FC Entity are to support one or more VE_Ports and to communicate with the FCIP Entity. The FC Entity layer lies between the FC-2 FC level and the FCIP Entity layer as shown in figure 11.

The FCIP Entity is the principal interface point to the IP network on one side, and in combination with the FC Entity, to the FC network on the other. The primary function of the FCIP Entity is formatting, encapsulating, and forwarding encapsulated FC frames across the IP network interface.

The FC/FCIP Entity pair interfaces with the CSM and the PMM through a vendor-specific mechanism.

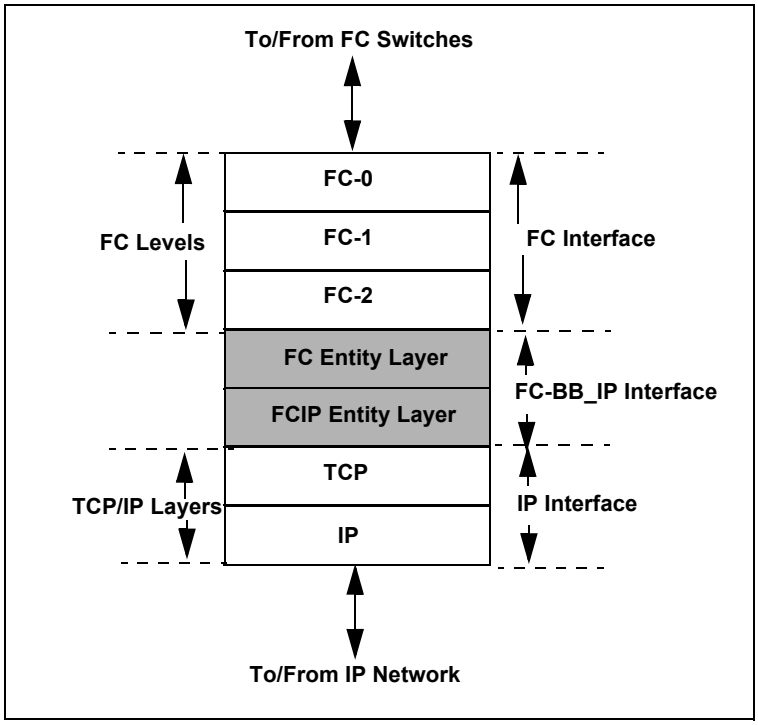

Figure 11 – FC-BB_IP Protocol Layers

5.2.4.2.2 FC Entity

The FC-BB_IP interface may support multiple instances of the FC/FCIP Entity pair. Each instance of the FC/FCIP Entity pair consists of one or more VE_Port/FCIP_LEP pairs. A VE_Port emulates an E_Port and interfaces with the FCIP_LEP component of the FCIP Entity. The term “Virtual” in VE_Port indicates the use of a non Fibre Channel link connecting the VE_Ports.

The VE_Port receives FC frames from the FC side and sends them to the FCIP_LEP for encapsulation and transmission on the IP network. The VE_Port may also exchange Class F control frames with the remote VE_Port via the LEPs. There is a one-to-one relationship between a VE_Port and an FCIP_LEP. VE_Ports communicate via VE_Port Virtual ISLs (see 5.2.4.2.4).

NOTE 3 – The term Virtual ISL when used unqualified refers to both a VE_Port Virtual ISL and a B_Access Virtual ISL.

A VE_Port is uniquely identified by an 8-byte VE_Port_Name.

Within an FC-BB_IP device, each FC/FCIP Entity pair instance is uniquely identified by an 8-byte identifier called the FC/FCIP identifier. The FC/FCIP identifier uses the Name_Identifier format.

Initialization at the FC-BB-5 protocol interface occurs between VE_Ports in a manner identical to standard E_Ports and is described in 5.2.4.3.

5.2.4.2.3 FCIP Entity

The FCIP_LEP is a component of the FCIP Entity that formats, encapsulates, and forwards encapsulated FC frames. Encapsulated FC frames are sent as TCP segments over the IP network.

The FCIP_LEP receives byte-encoded SOF/EOF delimited FC frames and a timestamp (see 5.2.4.5.2.2) from its VE_Port. The FCIP Data Engine (FCIP_DE) is the data forwarding component of the FCIP_LEP. The FCIP_DE handles all encapsulation and de-encapsulation, and transmission and reception of the encapsulated FC frames on the FCIP Link. The FCIP_LEP contains one or more FCIP_DEs, each corresponding to a TCP connection.

The FCIP_DE has four interface points (see RFC 3821):

a) **FC Receiver Portal:** The access point through which a byte-encoded SOF/EOF delimited FC frame and timestamp enters an FCIP_DE from the VE_Port;
b) **FC Transmitter Portal:** The access point through which a reconstituted byte-encoded SOF/EOF delimited FC frame and timestamp leaves an FCIP_DE to the VE_Port;
c) **Encapsulated Frame Receiver Portal:** The TCP access point through which an encapsulated FC frame is received from the IP network by the FCIP_DE; and
d) **Encapsulated Frame Transmitter Portal:** The TCP access point through which an encapsulated FC frame is transmitted to the IP network by the FCIP_DE.

5.2.4.2.4 VE_Port Virtual ISL and FCIP Link

The FC/FCIP Entity pair provides a data forwarding path between itself and a remote FC/FCIP Entity pair via virtual constructs. Two types of virtual constructs are defined:

a) a VE_Port Virtual ISL is a logical construct that is created between two FC Entity VE_Ports for the explicit purpose of sending and receiving byte-encoded SOF/EOF delimited FC frames via the FCIP Entity. Conceptually, communication between two VE_Ports is similar to communication between E_Ports; and
b) an FCIP Link is a logical construct that is created between two FCIP Entity LEPs for the explicit purpose of sending and receiving encapsulated FC frames and encapsulated FCIP control information. Conceptually, communication between two LEPs is similar to the communication between two instances of a TCP application.

There is a one-to-one mapping between a VE_Port Virtual ISL and an FCIP Link. Each FCIP Link consists of one or more TCP connections, all between the same two FC-BB_IP devices. Although more than one FCIP Link may be formed between a pair of FC-BB_IP devices, a typical configuration may only consist of a single FCIP Link. See figure 17 for some examples of allowed network topologies.

The FCIP_LEP that originates an FCIP Link is defined as the FCIP Link Originator. The corresponding FCIP_LEP that accepts this link is defined as the FCIP Link Acceptor. An FCIP Link is fully characterized by its FCIP Link Originator and FCIP Link Acceptor identities. An FCIP Link Originator or FCIP Link Acceptor is fully identified by all of the following:

a) an 8-byte Switch_Name;
b) an 8-byte VE_Port_Name; and
c) an 8-byte FC/FCIP Entity identifier.

To uniquely identify an FCIP Link, all of the following are required:

a) the 8-byte Switch_Name of the FCIP Link Originator;
b) the 8-byte VE_Port_Name of the FCIP Link Originator;
c) the 8-byte FC/FCIP Entity identifier of the FCIP Link Originator; and
d) the 8-byte Switch_Name of the FCIP Link Acceptor.

The FCIP Link Acceptor's 8-byte FC/FCIP Entity identifier and its VE_Port_Name provide additional information about an FCIP Link but are not required to uniquely identify it.

5.2.4.3 VE_Port Virtual ISL exchanges

5.2.4.3.1 SW_ILS exchanges

VE_Ports exchange SW_ILSs on the VE_Port Virtual ISL. The SW_ILSs that occur on the VE_Port Virtual ISL are the standard E_Port SW_ILSs (e.g., ELP, ESC, EFP, etc.), and in addition the LKA ELS (see FC-LS-2). Figure 12 shows the scope of the VE_Port Virtual ISLs.

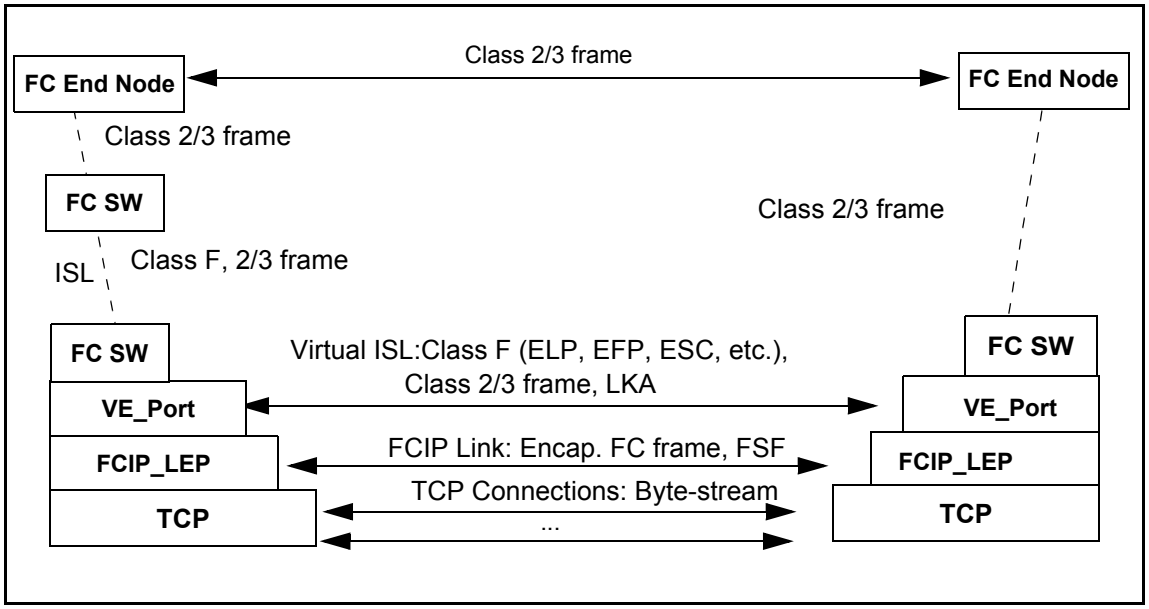

Figure 12 – Scope of VE_Port Virtual ISL

5.2.4.4 Control and Service Module (CSM)

The CSM is a control component of the FC-BB_IP interface that mainly deals with connection management. The CSM creates the FC/FCIP Entity pair during the Virtual ISL/FCIP Link setup. The CSM processes all requests for a link setup via the FCIP-registered TCP Port 3225 or optionally another TCP Port. The CSM processes requests to add additional TCP connections over the same FCIP Link. The CSM is also responsible for tearing down existing FCIP Links and TCP connections and deleting the FC/FCIP Entity pair.

NOTE 4 – Some aspects of the CSM functions are discussed only in RFC 3821.

5.2.4.5 Platform Management Module (PMM)

5.2.4.5.1 Function

The PMM is a management component of the FC-BB_IP interface that handles time synchronization, discovery, and security. The PMM is also the intended component for any miscellaneous housekeeping functions such as maintenance of event logs (see 5.7.3.5)

5.2.4.5.2 Time synchronization

5.2.4.5.2.1 FCIP Transit Time (FTT)

FCIP Transit Time (FTT) is defined as the total transit time of an encapsulated Fibre Channel frame in the IP network. The FCIP Transit Time is calculated by subtracting the timestamp value in the arriving encapsulated FC frame from the synchronized time in the FCIP Entity.

5.2.4.5.2.2 Building outgoing FC frame encapsulation headers

The FC Entity shall establish and maintain a synchronized time value in Simple Network Time Protocol (SNTP) Version 4 format (see RFC 2030) for use in computing the IP network transit times. The FC Entity shall use suitable internal clocks and one of the following mechanisms to establish and maintain the synchronized time value:

a) Fibre Channel time services; or
b) IP network SNTP server(s).

Each byte-encoded SOF/EOF delimited FC frame that the FC Entity delivers to the FCIP_DE through the FC receiver portal shall be accompanied by a timestamp value obtained from the synchronized time service. The FCIP_DE places the timestamp in the encapsulation header part of the encapsulated FC frame that carries the FC frame (see RFC 3643). If no synchronized timestamp value is available to accompany an entering Class 2 or Class 3 FC frame, the frame should not be delivered to the FCIP_DE. However, FC-BB_IP shall allow any class F encapsulated FC frames to be transmitted with a zero timestamp value.

5.2.4.5.2.3 Checking IP network transit times in incoming FC frame encapsulation headers.

Each byte-encoded SOF/EOF delimited FC frame delivered to the FC Entity through the FCIP_DE FC transmitter portal is to be accompanied by the timestamp value taken from the encapsulation header of the encapsulated FC frame. As noted in 5.2.4.5.2.2, the timestamp may be zero indicating that no valid timestamp was supplied by the sending FC Entity. Any frame other than a Class F frame whose timestamp is zero shall be discarded. A Class F frame whose timestamp is zero shall be processed as if it met all Fibre Channel timeout requirements.

When the timestamp is non-zero, the FTT of the arriving encapsulated Fibre Channel frame shall be compared to 1/2 E_D_TOV. If the FTT exceeds 1/2 E_D_TOV, then the frame shall be discarded. Otherwise the frame shall be processed normally. Fibre Channel timeout values shall be administratively set to accommodate the FTT.

5.2.4.5.3 Discovery

Discovery of FC-BB_IP devices is handled in accordance with the procedures outlined in 5.6.2.2 and in RFC 3821 and RFC 3822.

5.2.4.5.4 Security

Security in FC-BB_IP is defined at two levels, FC and FCIP. The FC level is secured through FC-SP mechanisms that are extended by FC-BB_IP. The FCIP level is secured through IPSec mechanisms (see RFC 3821). Figure 13 illustrates the scope of the two security mechanisms.

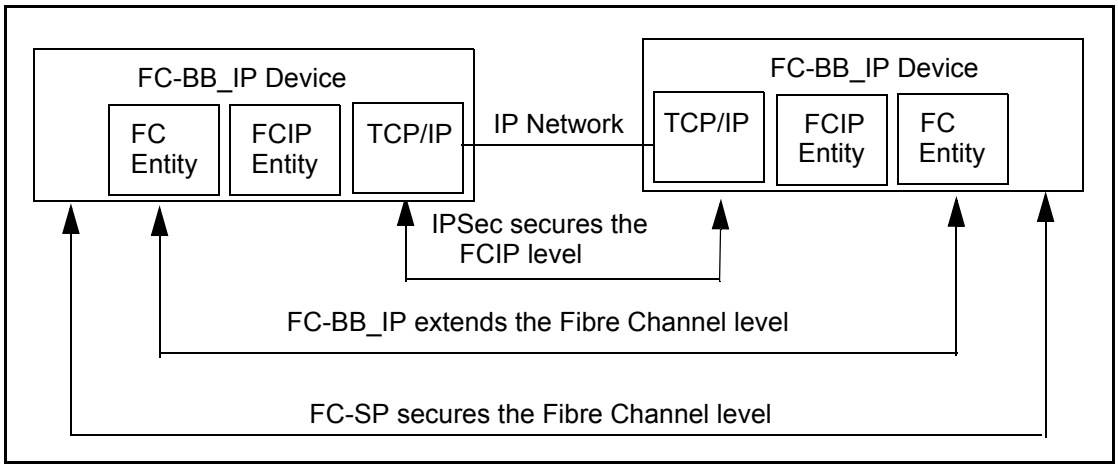

Figure 13 – Security layers

In most cases, the security requirements of an FC/FCIP Entity pair are satisfied outside the scope of this standard as follows:

a) security for the FC Fabric is provided by the FC-SP capabilities (e.g., switch-to-switch authentication, frame authentication and confidentiality); and
b) security for the TCP connections used to transit the IP network is provided by the security features described in FCIP (e.g., IPSec packet authentication and confidentiality).

Depending on the security requirements of a given configuration, any or all of the security capabilities described in other standards may be enabled or disabled. However, it is important to note that the public IP network is subject to a large variety of security attacks, meaning that serious consideration should be given to enabling the full suite of security features described in RFC 3821 whenever the public IP network is to be used to transmit FCIP frames.

An FC/FCIP entity pair has a potential security vulnerability where interactions may not be fully secured by either the FC-SP or FCIP security features. This vulnerability occurs when two or more TCP connections are aggregated in a single FCIP Link. The first TCP connection in an FCIP Link and its associated Virtual ISL may be authenticated using the FC-SP mechanisms. However, no such authentication is defined for subsequent TCP connections, since to FC they all appear to be part of an already authenticated Virtual ISL.

To prevent attacking entities in the IP network from forging additional invalid TCP connections, the FC-BB_IP mechanism described in 5.6.2.3 extends the protection of FC-SP authentication to subsequently-added TCP connections. The extension to FC-SP authentication described in 5.6.2.3.2 is based on the exchange of Class F requests and responses between FC Entities. This mechanism works in concert with the FC-SP Virtual ISL authentication mechanism, handling the Class F requests and responses over a previously authenticated TCP connection. In some configurations, this overhead may be unnecessary. However, in cases where Fabric entities are capable of being authenticated without having their behavior fully trusted, the extension to FC-SP authentication should be used in combination with other FC-SP and FCIP security mechanisms to assure trustworthy formation of FCIP Links and Virtual ISLs.

5.2.5 IP network interface

The FC-BB_IP VE_Port reference model supports one logical IP interface and allows sharing a 4-byte IPv4 or 16-byte IPv6 address in the following ways:

a) a single IP address per FC-BB_IP device (i.e., a single IP address shared by all FC/FCIP Entity pairs);
b) multiple IP addresses per FC-BB_IP device (i.e., a single IP address per FC/FCIP Entity pair);
c) multiple IP addresses per FC/FCIP Entity pair (i.e., single IP address per VE_Port/FCIP_LEP pair); and
d) multiple IP Addresses per FCIP Link (i.e., a single IP address per TCP Port).

Use of different IP address schemes at the two ends of an FCIP Link is not expected to cause interoperability problems.

As shown in figure 11, the IP network interface consists of the TCP and IP layers. The encapsulated FC frame emerging from the FCIP_DE interfaces with the TCP layer. The IP layer interfaces with the TCP layer above it and the IP network below it. The TCP layer supports multiple TCP connections, each corresponding to an FCIP_DE. Each TCP connection within an FCIP Link is assigned a unique local TCP Port Number. Either the FCIP well-known TCP Port 3225 or optionally another TCP Port is used for accepting connection requests. These ports interface with the CSM through a vendor-specific mechanism.

IP routing occurs inside the IP network. Within the IP network, the route taken by an encapsulated FC frame is determined by the normal routing procedures of the IP network.

5.3 B_Access functional model

5.3.1 FC-BB_IP interface protocol layers

Figure 14 shows the functional model of an FC-BB_IP device that consists of the B_Port FC interface, the FC-BB_IP protocol interface, and the IP network interface. Figure 11 shows the details of the protocol layers across these interfaces.

NOTE 5 – Because of the similarity between the E_Port and B_Port functional models this subclause only describes any unique definitions for the B_Access. Other definitions and descriptions from 5.2 apply equally well and remain unchanged.

5.3.2 B_Port FC interface

The FC-BB_IP FC network interface supports one or more B_Ports thus requiring the support of the FC-0, FC-1, and FC-2 levels. B_Ports in general connect to different external FC switches, but connectivity to the same FC switch is allowed.

B_Ports are uniquely identified by an 8-byte B_Port_Name.

5.3.3 FC-BB_IP protocol interface

5.3.3.1 Major components

The B_Port FC-BB_IP interface consists of all the components of the VE_Port functional model (see 5.2.4.1) except the FC Switching Element with FC routing.

5.3.3.2 FC and FCIP Entities

5.3.3.2.1 Function

The primary function of the FC Entity is to support one or more B_Access portals and to communicate with the FCIP Entity.

The function of the FCIP Entity is identical to its function in the VE_Port functional model described in 5.2.4.2.3.

The FC/FCIP Entity pair interfaces with the CSM and the PMM through a vendor-specific mechanism.

5.3.3.2.2 FC Entity

The FC-BB_IP interface may support multiple instances of the FC/FCIP Entity pairs. Each instance of the FC/FCIP Entity pair consists of one or more B_Access/FCIP_LEP pairs. A B_Access portal is a component of the FC Entity that interfaces with the FCIP_LEP component of the FCIP Entity. The B_Access portal receives FC frames from the B_Port and sends them to the FCIP_LEP for encapsulation and transmission on the IP network. The B_Access portal may also exchange Class F control frames with the remote B_Access portal via the LEPs. There is a one-to-one relationship between a B_Access portal and an FCIP_LEP. B_Access portals communicate via B_Access Virtual ISLs (see 5.3.3.2.4).

There is no switching and routing required in the case of the B_Port functional model. However, the forwarding of FC frames across the B_Access/FCIP_LEP pair is still required. When multiple DEs within an FCIP_LEP are in use, the selection of which FCIP_DE to use is described in 5.6.3.5.

Initialization at the FC-BB-5 protocol interface occurs with EBP SW_ILS exchanges between B_Access portals in a manner identical to standard E_Ports and is described in 5.3.3.2.4. The B_Access initialization state machine is described in 5.3.3.3.1.1.

5.3.3.2.3 FCIP Entity

The FCIP_LEP receives byte-encoded SOF/EOF delimited FC frames and a timestamp from its B_Access portals. All other functions are identical to the functions of the FCIP Entity in the VE_Port functional model (see 5.2.4.2.3).

5.3.3.2.4 B_Access Virtual ISL and FCIP Links

A B_Access Virtual ISL is a logical construct that is created between two FC Entity B_Access portals for the explicit purpose of sending and receiving byte-encoded SOF/EOF delimited FC frames via the FCIP Entity. Conceptually, communication between two B_Access portals is similar to communication between two VE_Ports.

There is a one-to-one mapping between a B_Access Virtual ISL and an FCIP Link.

An FCIP Link Originator or FCIP Link Acceptor is fully identified by all of the following:

a) an 8-byte Fabric_Name;
b) an 8-byte B_Access_Name; and
c) an 8-byte FC/FCIP Entity identifier.

To uniquely identify an FCIP Link, the following items are required:

a) the 8-byte Fabric_Name of the FCIP Link Originator;
b) the 8-byte B_Access_Name of the FCIP Link Originator;
c) the 8-byte FC/FCIP Entity identifier of the FCIP Link Originator; and
d) the 8-byte Fabric_Name of the FCIP Link Acceptor.

The FCIP Link Acceptor's 8-byte FC/FCIP Entity identifier and its B_Access_Name provide additional information about an FCIP Link but are not required to uniquely identify it.

The FC-BB_IP B_Access functional model is shown in figure 14.

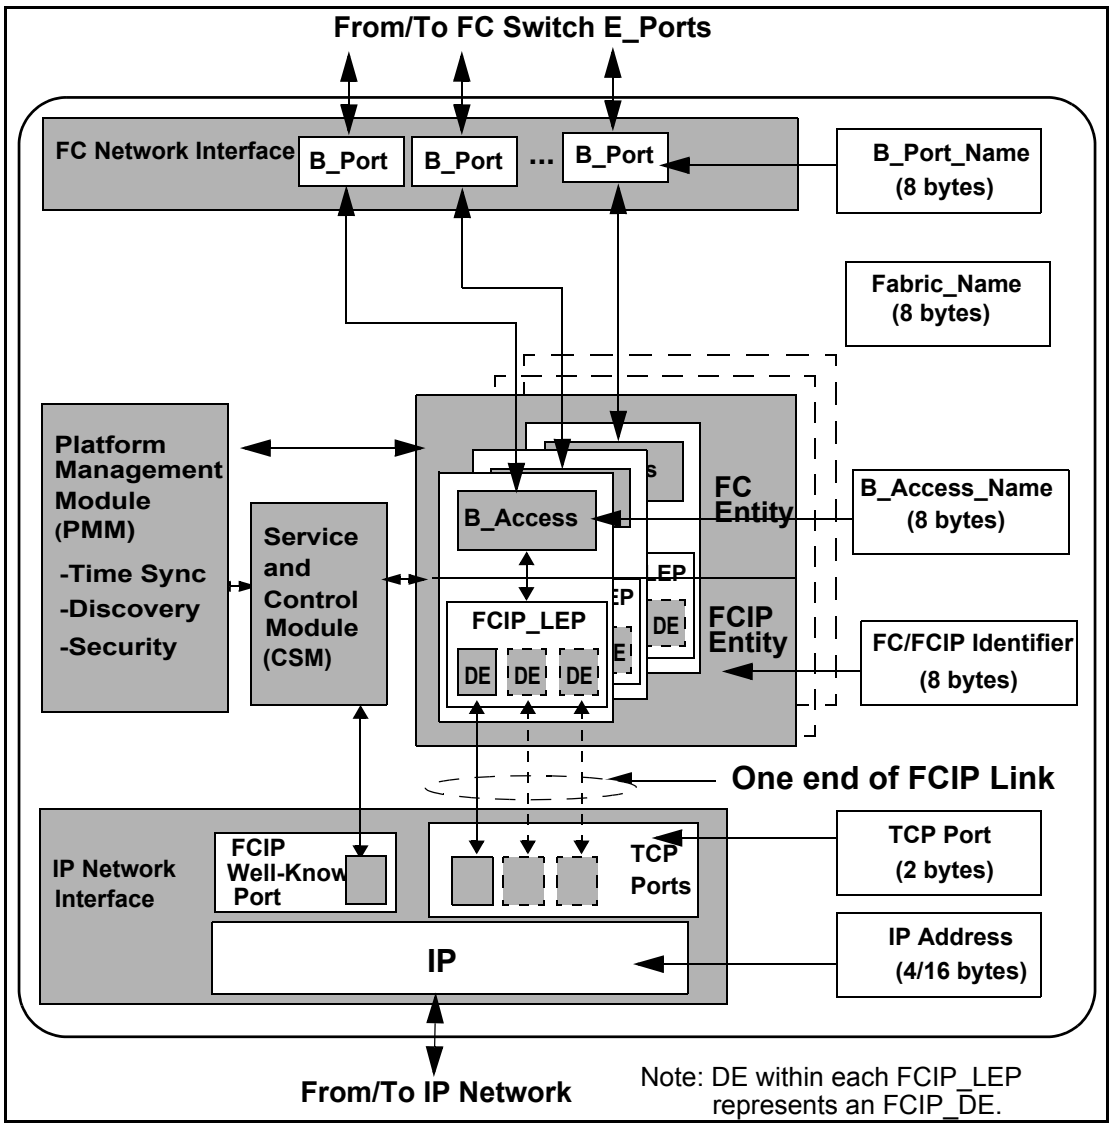

Figure 14 – FC-BB_IP B_Access functional model

5.3.3.3 B_Access Virtual ISL exchanges

5.3.3.3.1 Exchange B_Access Parameters (EBP) SW_ILS

B_Access portals exchange SW_ILSs on the B_Access Virtual ISL. The SW_ILSs that occur on the B_Access Virtual ISL are the EBP SW_ILS and the LKA ELS (see FC-LS-2). Figure 15 shows the scope of the B_Access Virtual ISLs.

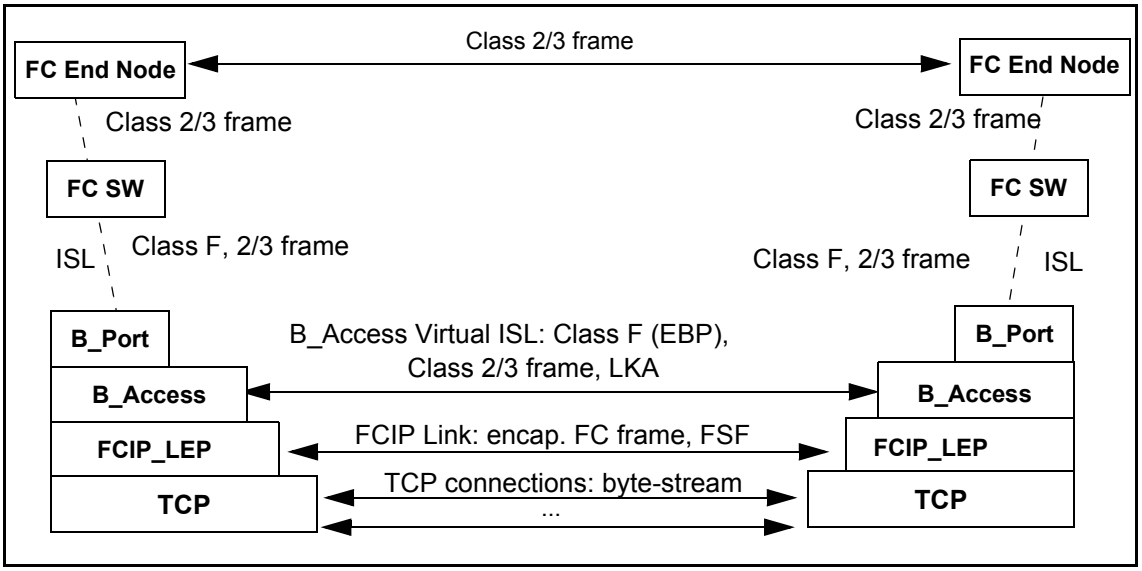

Figure 15 – Scope of B_Access Virtual ISL

The Exchange B_Access Parameters (EBP) Switch Fabric Internal Link Service (SW_ILS) is sent by a B_Access portal to a remote B_Access portal in order to establish operating link parameters and port capabilities for the B_Access Virtual ISL formed by the two B_Access portal peers. Successful acceptance of EBP SW_ILS shall be completed before the B_Ports begin switch port mode initialization.

Protocol:

a) Exchange B_Access Parameters (EBP) request Sequence; and
b) Reply Switch Fabric Internal Link Service Sequence.

Addressing: For use in switch port configuration, the S_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating B_Access portal. The D_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination B_Access portal.

Payload: The format of the EBP request payload is shown in table 5.

Table 5 – EBP request payload

Item	Size Bytes	Remarks
28 01 00 00h	4	
R_A_TOV	4	Value in milliseconds
E_D_TOV	4	Value in milliseconds
K_A_TOV	4	Value in milliseconds
Requester B_Access_Name	8	
Class F Service Parameters	16	

Requester B_Access_Name: This field shall contain the B_Access_Name of the device that originated the EBP request.

R_A_TOV: This field shall be set to the value, in milliseconds, of R_A_TOV required by the FC-BB_IP device.

E_D_TOV: This field shall be set to the value, in milliseconds, of E_D_TOV required by the FC-BB_IP device.

K_A_TOV: This field shall be set to the value, in milliseconds, of K_A_TOV required by the FC-BB_IP device.

Class F Service Parameters: This field shall contain the B_Access Class F Service Parameters and its format is identical with its use in the ELP SW_ILS (see FC-SW-5).

Reply Switch Fabric Internal Link Service Sequence:

Service Reject (SW_RJT)

Accept (SW_ACC)

Accept payload

Payload: The format of the EBP Accept payload is shown in table 6.

Table 6 – EBP accept payload

Item	Size Bytes	Remarks
02 00 00 00h	4	
R_A_TOV	4	Value in milliseconds
E_D_TOV	4	Value in milliseconds
K_A_TOV	4	Value in milliseconds
Responder B_Access_Name	8	
Class F Service Parameters	16	

The fields in table 6 are the same as defined for table 5 except for the Responder B_Access_Name field.

Responder B_Access_Name: This field shall contain the B_Access_Name of the remote device that responds to the EBP request.

The SW_RJT Reply payload format is given in FC-SW-5. The EBP reject reason code explanation is shown in table 7.

Table 7 – EBP reject reason code explanation

Encoded Value (Bits 23-16)	Description
0000 0000	No additional explanation
0000 0001	Class F Service Parameter error
0000 0010	Invalid B_Access_Name
0000 0011	K_A_TOV mismatch
0000 0100	E_D_TOV mismatch
0000 0101	R_A_TOV mismatch
others	Reserved

5.3.3.3.1.1 B_Access initialization state machine

The B_Access initialization state machine is shown in figure 16.

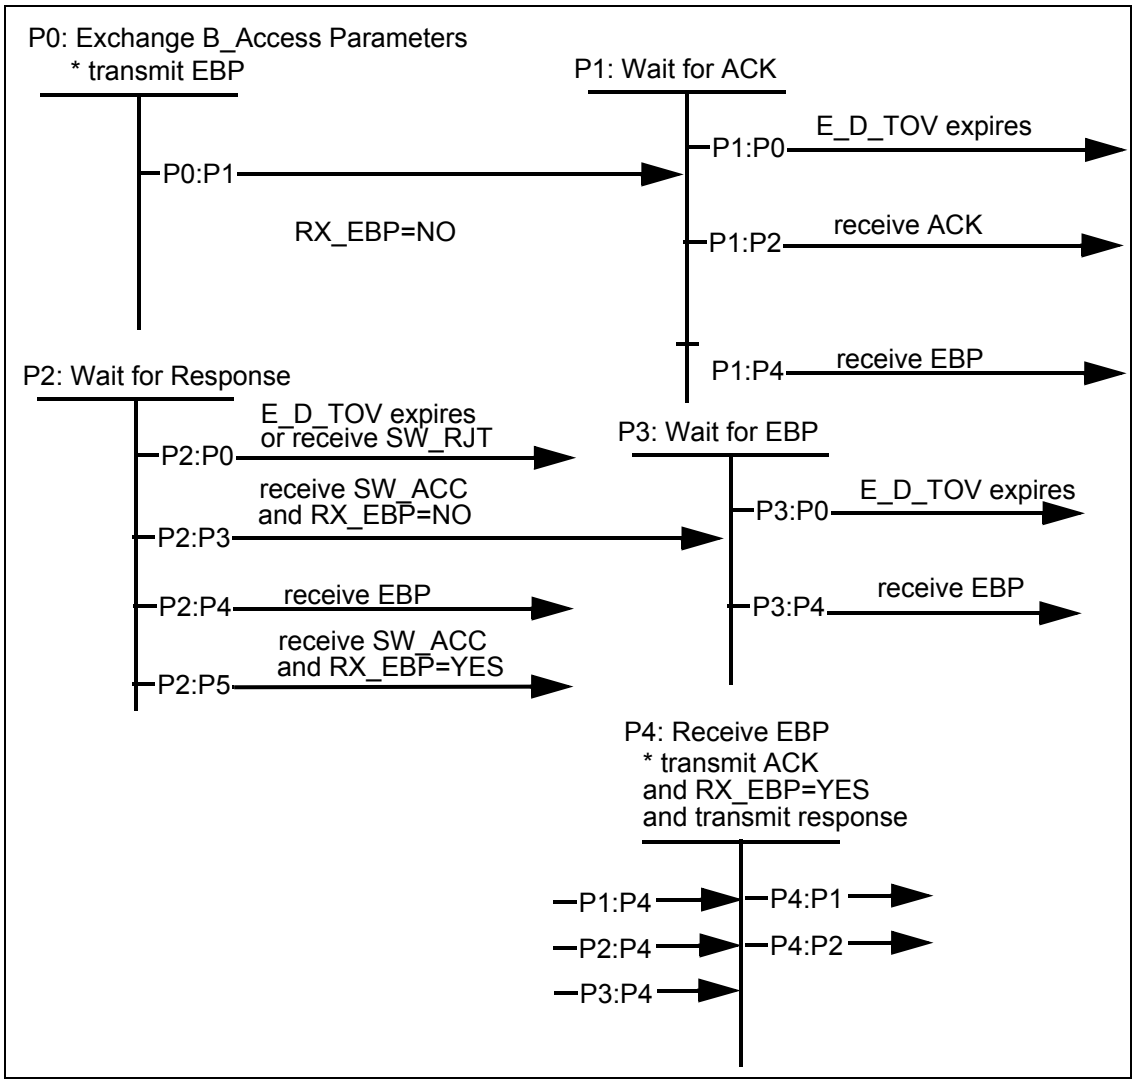

Figure 16 – B_Access initialization state machine

State P0: Exchange B_Access Parameters. This state marks the beginning of the B_Access initialization. Activity other than that described within the state machine is suspended until initialization is complete.

Transition P0:P1. The B_Access resets the RX_EBP flag.

State P1: Wait for ACK. In this state the B_Access waits until an ACK for the B_Access's transmitted EBP is received.

Transition P1:P0. This transition occurs when the B_Access has not received an ACK within E_D_TOV after the transmission of an EBP.

Transition P1:P2. This transition occurs when the B_Access receives an ACK before E_D_TOV expires.

Transition P1:P4. This transition occurs when the B_Access receives an EBP while waiting for an ACK.

State P2: Wait for Response. In this state the B_Access has received an ACK for its EBP and is waiting for a response.

Transition P2:P0. This transition occurs when the B_Access has not received a response within E_D_TOV after the transmission of an EBP or receives an SW_RJT.

Transition P2:P3. This transition occurs when the B_Access receives an SW_ACC and has not received an EBP.

Transition P2:P4. This transition occurs when the B_Access receives an EBP while waiting for a response.

Transition P2:P5. This transition occurs when the B_Access receives an SW_ACC and has received an EBP.

State P3: Wait for EBP. in this state the B_Access has received an ACK for its EBP and is waiting for an EBP.

Transition P3:P0. This transition occurs when the B_Access has not received an EBP within E_D_TOV of the transmission of an EBP.

Transition P3:P4. This transition occurs when a B_Access receives an EBP while waiting for a response.

State P4: Receive EBP. In this state the B_Access has received an EBP. The B_Access responds with an ACK and transmits an SW_ACC or SW_RJT depending upon whether or not the received configuration parameters contained within the EBP are acceptable. The B_Access sets RX_EBP to indicate an EBP has been received and is accepted.

Transition P4:P1. This transition occurs when a B_Access receives an EBP from its peer yet hasn't received an ACK for a previously transmitted EBP.

Transition P4:P2. This transition occurs when a B_Access receives an EBP from its peer yet hasn't received a response for a previously transmitted EBP.

5.3.3.4 B_Port Control and Service Module (CSM)

The B_Port CSM is identical to the E_Port CSM described in 5.2.4.4.

5.3.3.5 B_Port Platform Management Module (PMM)

The B_Port PMM is identical to the E_Port PMM described in 5.2.4.5.

5.3.4 IP Network Interface

The B_Port IP network interface is identical to the E_Port IP network interface described in 5.2.5 with a change in item (c), where a single IP address is per B_Access/FCIP_LEP pair.

5.4 FC-BB_IP Network Topologies

Figure 17 shows some example FC-BB_IP network topologies that exists between three FC-BB_IP sites:

a) FCIP Link 1 connects Sites 1 and 2 and consists of three TCP connections;
b) FCIP Link 2 connects Sites 1 and 2 and consists of two TCP connections. FCIP Link 2 however is distinct from Link 1 although it exists between the same two FC/FCIP Entity pairs (i.e., FC/FCIP_Entity_1 and FC/FCIP_Entity_2);
c) FCIP Link 3 connects Sites 1 and 3 and consists of two TCP connections. FCIP Link 3 exists between FC/FCIP_Entity_3 and FC/FCIP_Entity_5; and
d) FCIP Link 4 connects Sites 2 and 3 and consists of one TCP connection. FCIP Link 4 exists between FC/FCIP_Entity_4 and FC/FCIP_Entity_6.

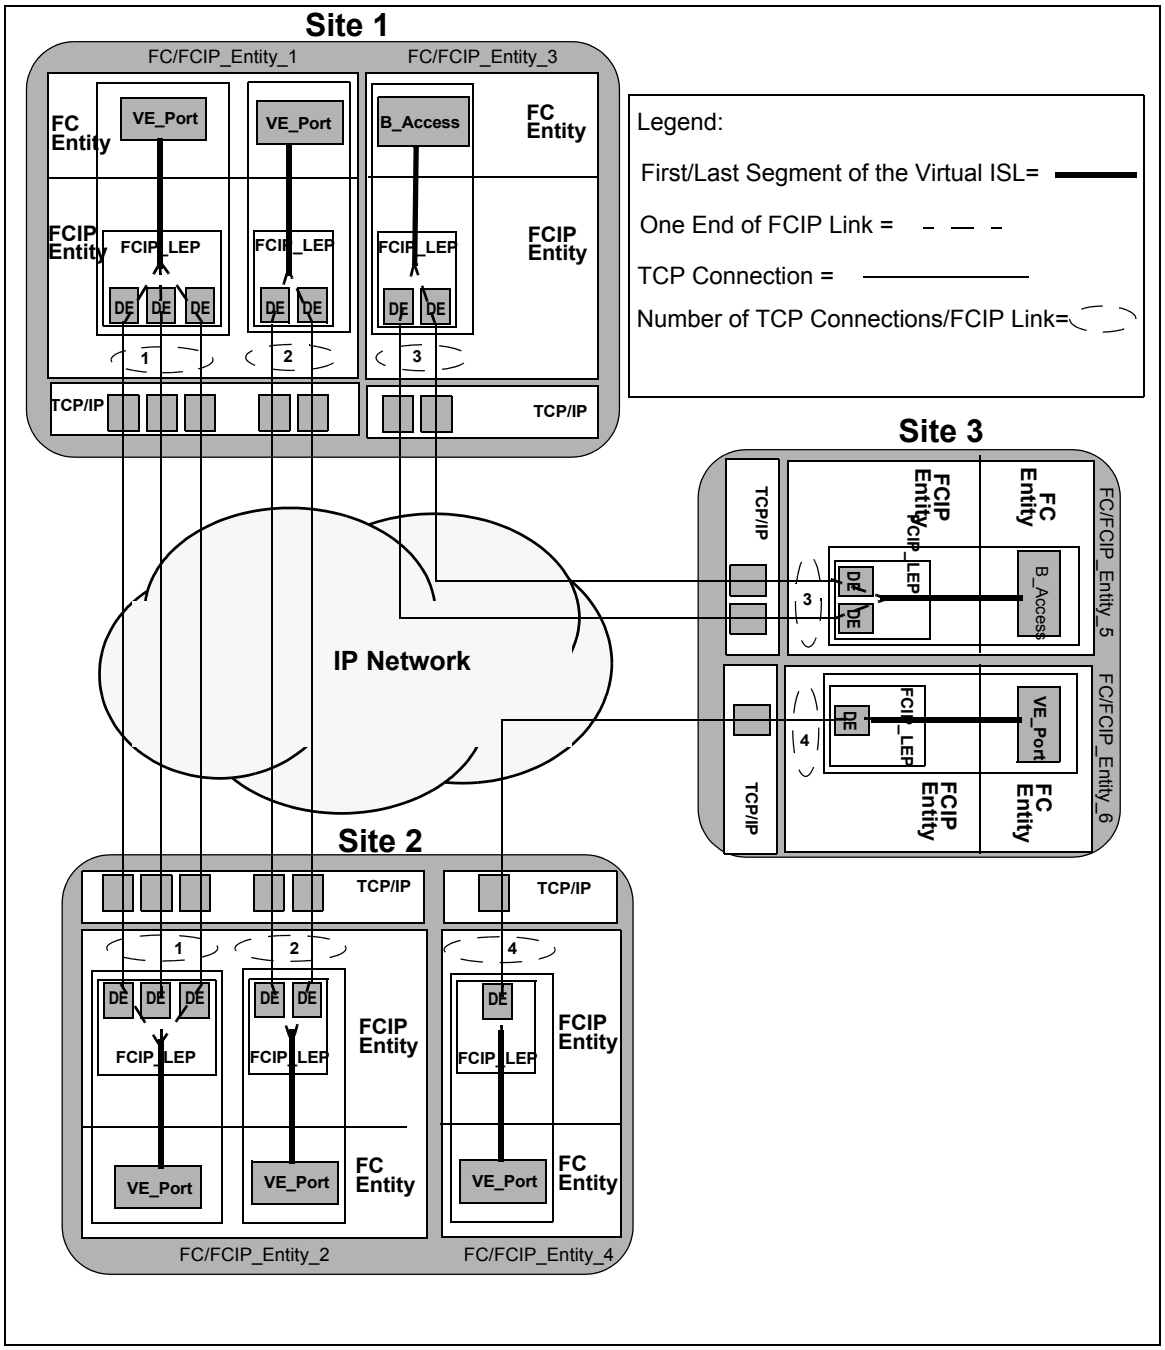

Figure 17 – FC-BB_IP network topologies

5.5 Mapping and message encapsulation using TCP/IP

5.5.1 Encapsulated frame structures

5.5.1.1 FC frame encapsulation structure

An encapsulated FC frame is carried in one or more TCP segments as shown in figure 18. Each segment's format is shown in table 8. The structure of each encapsulated FC frame is shown in table 9 and consists of a FC Encapsulation Header and a byte-encoded SOF/EOF delimited Class 2, 3, or F FC frame as described in RFC 3643.

Table 8 – TCP/IP Segment structure carrying encapsulated FC frame

Field	Sub-field	Size (Bytes)
IP Header		Min:20 Max:40
TCP Header		Min:20 Max:40
TCP payload	one or more portions of encapsulated FC frames	Min:64 Max:2176

Table 9 – Encapsulated FC frame structure

Field	Size (Bytes)
FC Encapsulation Header	28
SOF	4
FC-Header	24
FC frame payload (includes optional header)	Min: 0 Max:2112
CRC	4
EOF	4

FC frame encapsulation (see RFC 3643) describes the structures of the 4-byte SOF/EOF values fields and the FC Encapsulation Header. The FC Encapsulation Header consists of several fields: Protocol#, Version, pFlags, Flags, Frame Length, Timestamp, and CRC.

Protocol#: indicates the FCIP protocol.

Version: indicates the version number.

pFlags: defines flag bits FSF and Ch that distinguish encapsulated FC frames from FCIP originated or echoed control frames.

Flags: the CRCV bit value indicates if the contents of the CRC field are valid or invalid. For FC-BB_IP protocol the CRCV bit shall be zero (i.e., invalid).

Frame Length: contains the length of the entire FC encapsulated frame including the FC Encapsulation Header and the FC frame, including SOF and EOF words.

Timestamp: contains the time at which the FC encapsulated frame was sent as known to the sender. The format of integer and fraction timestamp word values is specified in Simple Network Time Protocol (SNTP) Version 4 (see RFC 2030). The contents of the timestamp integer and timestamp fraction words shall be set as described in 5.2.4.5.2.

CRC: for FC-BB_IP protocol the CRC field shall be set to zero.

5.5.1.2 Encapsulated FCIP Special Frame (FSF) structure

An encapsulated FCIP Special Frame (FSF) is carried as a TCP segment as shown in table 10. The structure of an encapsulated FSF is shown in table 11 and consists of an FC Encapsulation Header and an FCIP Special Frame.

Table 10 – TCP/IP Segment structure carrying encapsulated FSF

Field	Sub-field	Size (Bytes)
IP Header		Min:20 Max:40
TCP Header		Min:20 Max:40
TCP payload	encapsulated FSF	76

Table 11 – Encapsulated FSF structure

Field	Size (Bytes)
FC Encapsulation Header	28
FCIP Special Frame (FSF)	48

See 5.5.1.1 for a description of the FC Encapsulation Header structure and format.

The FSF structure is defined in FCIP (see RFC 3821) and consists of several fields: Source FC Fabric_Name, Source FC/FCIP Entity identifier, Connection Nonce, Connection Usage Flags, Connection Usage Code, Destination FC Fabric_Name, and K_A_TOV.

Source FC Fabric_Name: the identifier for the FC Fabric associated with the FC/FCIP Entity pair that generates the FCIP Special Frame. If the FC Fabric is an FC switch, then the field contains the Switch_Name.

Source FC/FCIP Entity identifier: a unique identifier for the FC/FCIP Entity pair that generates the FSF. The value is assigned by the FC Fabric whose name appears in the Source FC Fabric_Name field.

Connection Nonce: contains a 64-bit random number generated to uniquely identify a single TCP connect request in order to provide sufficient security for the nonce, the randomness recommendations described in RFC 3821 should be followed.

Connection Usage Flags: identifies the types of SOF values to be carried on the connection. All or none of the bits corresponding to Class F, 2, or 3 may be set to one. If all of the bits are zero, then the types of FC frames intended to be carried on the connection has no specific relationship to SOF code.

Connection Usage Code: contains Fibre Channel defined information regarding the intended usage of the connection. The FCIP Entity uses the contents of the Connection Usage Flags and the Connection Usage Code fields to locate appropriate QoS settings in the shared database of TCP connection information and apply those settings to a newly formed connection. All values for this field are reserved.

Destination FC Fabric_Name: may contain the Fibre Channel identifier for the FC Fabric associated with the FC/FCIP Entity pair that echoes, as opposed to generates, the FSF.

K_A_TOV: contains the FC Keep Alive Timeout value to be applied to the new TCP connection.

5.5.2 TCP/IP encapsulation

Figure 18 illustrates the TCP/IP encapsulation of an encapsulated FC frame. The TCP/IP encapsulation of an encapsulated FSF is similar.

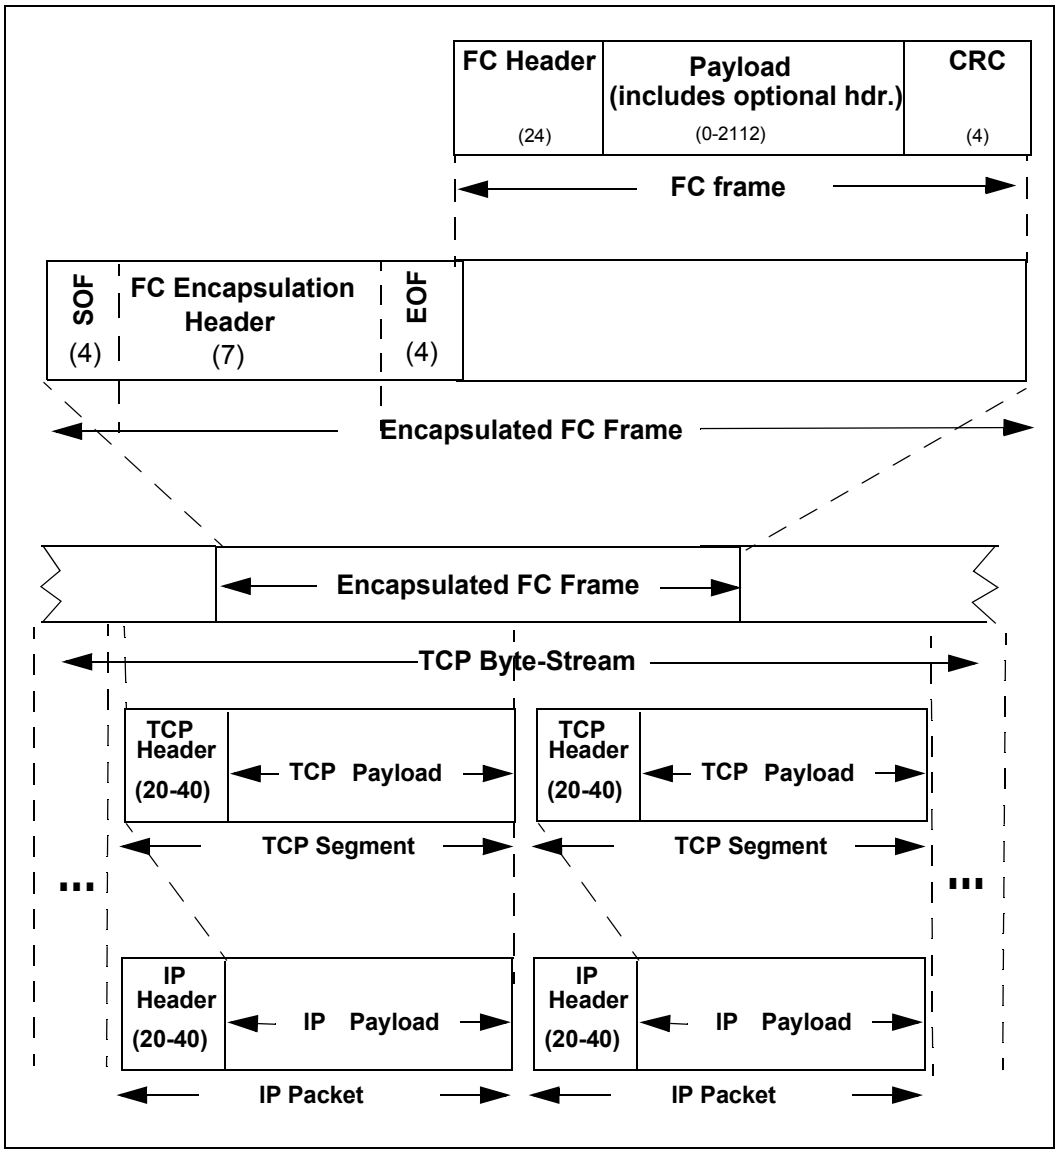

Figure 18 – TCP/IP encapsulation of an encapsulated FC frame

5.6 FC-BB_IP Protocol Procedures

5.6.1 Overview

This subclause describes the FC-BB_IP protocol procedures for platform management (see 5.6.2), connection management (see 5.6.3), and error detection and recovery (see 5.6.4). There are no specific procedures defined for housekeeping functions such as maintenance of error or event logs.

5.6.2 Procedures for platform management

5.6.2.1 Function

Platform management has three main functions, discovery, security, and time synchronization.

5.6.2.2 Procedures for discovery

Device discovery is one of the functions of the Platform Management Module (PMM). Each FC-BB_IP device is statically or dynamically configured with a list of IP addresses and other identifiers (e.g., N_Port_Names) corresponding to participating FC/FCIP Entities. If dynamic discovery of participating FC-BB_IP devices is supported, the function is performed using Service Location Protocol version 2 (see RFC 3822).

FC/FCIP Entities themselves do not actively participate in the discovery of FC source and destination identifiers. Discovery of FC addresses accessible via the FC/FCIP Entity is provided by techniques and protocols within the FC architecture as described in FC-FS-3 and FC-SW-5.

5.6.2.3 Procedures for extending FC-SP security

5.6.2.3.1 Authentication mechanisms

The Platform Management Module (PMM) is responsible for extending security at the FC level.

Entity authentication occurs at the FCIP and FC levels as illustrated in figure 13. Authentication mechanisms at the FCIP level are defined in FCIP (see RFC 3821). Authentication mechanisms at the FC level are defined in FC-SP.

During initialization of a Virtual ISL, each switch may authenticate the other switch with FC-SP authentication mechanisms. FC-BB_IP provides for extending the protection of FC-SP authentication to subsequently added TCP connections via either the ASF SW_ILS described in 5.6.2.3.2 or vendor-specific configuration information.

NOTE 6 – The unqualified use of the term Virtual ISL refers to both VE_Port Virtual ISL and B_Access Virtual ISL.

When an FCIP Entity receives a TCP connect request for an additional TCP connection to an existing FCIP Link to which FC-SP authentication has been applied, the FCIP Entity generates a request to the FC Entity to authenticate the additional TCP connection including at least the following information:

a) Connection Nonce;
b) Destination FC Fabric_Name;
c) Connection Usage Flags; and
d) Connection Usage Code.

If FC-SP authentication procedures are not being applied to the Virtual ISL, the FC Entity shall respond to the FCIP Entity indicating that the new TCP is authentic.

NOTE 7 – If the first TCP connection in a Virtual ISL is not authenticated using the applicable FC-SP procedures, no security is gained by authenticating other TCP connections.
NOTE 8 – The preferred security mechanism for the Public Internet IP network is the success or failure of an ASF SW_ILS.

5.6.2.3.2 Authenticate Special Frame (ASF)

The Authenticate Special Frame (ASF) Switch Fabric Internal Link Service (SW_ILS) is used by an FC Entity to authenticate additional TCP connections on existing FCIP Links. To authenticate a new TCP connection using the ASF SW_ILS, the FC Entity shall use the information provided by the FCIP Entity to transmit an ASF request on the Virtual ISL to which the new TCP connection is being added using a TCP connection in the Virtual ISL that has already been authenticated.

The FC Entity shall use the information from the new FSF request to populate the fields in the ASF request. The fields are the same as defined for FSF (see 5.5.1.2). The format of the ASF request payload is specified in table 12.

The FC Entity shall transmit the ASF over the previously authenticated TCP connection. This piggybacking technique authenticates additional TCP connections by riding on top of previously authenticated TCP connections.

An FC Entity that receives an ASF SW_ILS shall verify that the information in the request payload identifies a TCP connection initiated by that FC/FCIP Entity pair. If it verifies that this information is sound then the FC Entity shall respond with an SW_ACC (see table 13), otherwise it shall respond with an SW_RJT with a reason code of “Unable to perform command request” and a reason code explanation of “Class F Service Parameter error”.

Protocol:

a) Authenticate Special Frame (ASF) request Sequence; and

b) Reply Switch Fabric Internal Link Service Sequence.

Addressing: The S_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating FC Entity. The D_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the receiving FC Entity.

Payload: The format of the ASF request payload is specified in table 12.

Table 12 – ASF request payload

Item	Size Bytes
28 03 00 00h	4
Destination FC Fabric_Name	8
Connection Nonce	8
Connection Usage Flags	1
Reserved	1
Connection Usage Code	2
Reserved	4

Destination FC Fabric_Name: This field is the Fabric_Name of the destination switch and is the Source FC Fabric_Name from the FSF frame.

Connection Nonce: This field is the Connection Nonce from the FSF request.

Connection Usage Flags: This field is the Connection Usage Flags from the FSF request and signifies the acceptance of these flags.

Connection Usage Code: This field is the Connection Usage Code from the FSF request and signifies the acceptance of these codes.

Reply Switch Fabric Internal Link Service Sequence:

a) Service Reject (SW_RJT); or
b) Accept (SW_ACC).

Accept payload.

Payload: The format of the ASF accept payload is specified in table 13.

Table 13 – ASF accept response payload

Item	Size Bytes
02 00 00 00h	4

5.6.3 Procedures for connection management

5.6.3.1 Function

The primary function of the Control and Services Module (CSM) is managing connections.

5.6.3.2 Procedures for link setup

In order to realize a Virtual ISL/FCIP Link between two FC-BB_IP endpoints, an FC-BB_IP device establishes TCP connection(s) with its peer FC-BB_IP device.

NOTE 9 – A Virtual ISL exists between two VE_Ports or two B_Access portals and an FCIP Link exists between two FCIP_LEPs. Conceptually, the procedures for establishing these two are identical.

It may also be useful to assign a pool of connections for transmission of high priority and control frames (e.g., Class F) on connections so they do not encounter head-of-line blocking behind Class 2, or Class 3 traffic. The use of multiple connections and policies for distributing frames on these connections is described in 5.6.3.5.

A Virtual ISL/FCIP Link and the two FC-BB_IP device endpoints that are involved become operational only after the first TCP connection is established. The sequence of operations performed in order to establish a Virtual ISL/FCIP Link is as follows:

1) the FC-BB_IP device initializes its local resources to enable it to listen to TCP connection requests;
2) the FC-BB_IP device discovers the FC-BB_IP device endpoints to which it is able to establish a Virtual ISL/FCIP Link. The result of the discovery shall be, at the minimum, the IP address and

the TCP port of the peer endpoint. The discovery process may rely on administrative configuration or on services such as SLPv2 as described in 5.6.2.2;

3) the processes defined by FCIP are used to establish TCP connections. FC level authentication of the first TCP connection is accomplished using the mechanisms and management controls described in FC-SP. To extend FC-SP authentication to additional TCP connections the mechanisms described in 5.6.2.3 shall be followed;
4) at this point, both endpoints have their respective VE_Port/FCIP_LEP pairs or B_Access/FCIP_LEP pairs established;
5) after connection establishment, the FC-BB_IP device constructs the encapsulated FC frames according to the methods described in 5.2.4.5.2.2;
6) at this point the Virtual ISL endpoints shall exchange FC virtual port initialization frames to enable and identify port operation. The E_Port port mode initialization state machine is described in FC-SW-5 and the B_Access portal initialization state machine is described in 5.3.3.3.1.1. Switch-to-switch authentication shall use FC-SP authentication mechanisms;
7) an FC-BB_IP device operates in E_Port or B_Port mode. When operating in E_Port mode, normal FSPF messages are exchanged and the switch port becomes operational. When operating in B_Port mode, it is expected that the external E_Ports may exchange FSPF messages over the Virtual ISL which result in the link becoming operational;
8) link costs are implementation-defined;
9) in certain deployments, a single FC-BB_IP device may establish Virtual ISLs/FCIP Links with multiple FC-BB_IP device endpoints. In this situation, the FC-BB_IP device shall manage TCP operational parameters independently for each Virtual ISL or FCIP Link. Also, the FC-BB_IP device VE_Port may perform the E_Port initialization independently, for each Virtual ISL/FCIP Link. The B_Access also may perform initialization independently, for each Virtual ISL/FCIP Link; and
10) the FC Entity may participate in determining allowed TCP connections, TCP connection parameters, quality of service usage, and security usage by modifying interactions with the FCIP Entity that are modeled as a shared database. See RFC 3821.

5.6.3.3 Procedures for data transfer

The procedures for data transfer are as follows:

a) the sending FC Entity shall deliver FC frames to the correct FCIP_LEP/FCIP_DE in the correct FCIP Entity;
b) each FC frame delivered to the FCIP_DE shall be accompanied by a time value synchronized with the clock maintained by the FC Entity at the other end of the FCIP Link (see 5.2.4.5.2.2); and
c) when FC frames exit FCIP_DE(s) via the FC Transmitter Portal(s), the FC Entity should forward them to the FC Fabric. However, before forwarding the FC frame the FC Entity shall verify the end-to-end transit time as described in 5.2.4.5.2.3.

5.6.3.4 Procedures for FCIP Link disconnection

The FC Entity may require the FCIP Entity to perform TCP close requests (e.g., to perform a controlled shutdown of a link or to respond to high link error rates). If the FC Entity requests the closure of all TCP connections in an FCIP Link, the FCIP Link is disconnected.

When the FCIP Link is disconnected, notification of the disconnection shall be accomplished according to the procedures in 5.7.3.5.

5.6.3.5 Procedures for multiple connection management

A pair of FC-BB_IP device endpoints may establish a number of TCP connections between them. Since a Virtual ISL potentially maps a fairly large number of FC flows, where a flow is defined as a pair of Fibre Channel S_ID/D_ID addresses, it may not be practical to establish a separate TCP connection for each FC flow. However, once an FC flow is assigned to an FCIP_DE within the Virtual ISL, all FC frames of that flow shall be sent on that same FCIP_DE. This rule is in place to honor any in-order delivery guarantees that may have been made between the two end points of the FC flow.

When a TCP connect request is received and that request adds a new TCP connection to an existing FCIP_LEP, the procedures described in 5.6.2.3.1 shall be followed.

5.6.4 Procedures for error detection recovery

5.6.4.1 Procedures for handling invalid FC frames

Data corruption is detected at two different levels, TCP checksum and FC frame encapsulation errors. Data corruption detected at the TCP level shall be recovered via TCP data recovery mechanisms. The recovery for FC frame errors is described below. The TCP and FC frame recovery operations are performed independently.

Fibre Channel frame errors and the expected resolution of those errors are described in RFC 3821 and summarized below:

NOTE 10 – The behavior given below is that of the FCIP Entity.

a) all incoming frames on the FC receiver port are verified for correct header, proper format, valid length and valid CRC. A frame having an incorrect header or CRC shall be discarded or processed in accordance with the rules for the particular type of FC_Port;
b) all frames transmitted by the encapsulated frame tansmitter are valid FC encapsulations of valid FC frames with correct TCP check sums on the correct TCP/IP connection;
c) the FC frames contained in incoming encapsulated frames on the encapsulated frame receiver port are verified for a valid header, proper content, proper SOF and EOF values, and valid length. FC frames that are not valid according to those checks are managed according to the following rules:
   A) the frame may be discarded; or
   B) the frame may be transmitted in whole or in part by the FC transmitter port and ended with an EOF indicating that the content of the frame is invalid; and
d) if there is any discrepancy between statements in this subclause and RFC 3821, then RFC 3821 shall prevail.

5.6.4.2 Procedures for error recovery

The FC Entity shall recover from events that the FCIP Entity is unable to handle, such as:

a) loss of synchronization with FCIP frame headers from the encapsulated frame receiver portal requiring resetting the TCP connection; and
b) recovering from FCIP frames that are discarded as a result of synchronization problems.

The FC Entity may recover from connection failures.

Since FC Primitive Signals and Primitive Sequences are not exchanged between FCIP devices, there may be times when an FC frame is lost within the IP network. When this event occurs it is the responsibility of the communicating FC devices to detect and correct the errors based on the features defined in FC-FS-3.

In order to facilitate faster detection of loss of link connectivity, FC Entities shall make use of the Link Keep Alive (LKA) ELS (see FC-LS-2). The LKA ELS is exchanged across the Virtual ISL as shown in figure 12 (i.e., E_Port implementation) or figure 15 (i.e., B_Port implementation). The exact number of lost LKA heartbeats that forces the FC Entity to mark the link down is a configurable parameter with a default value of 2. Once the link has been marked down, the FC Entity shall attempt to re-establish the link via the FCIP Entity.

5.6.5 FC-BB_IP system parameters

5.6.5.1 FC timers

FC has two important timeouts, E_D_TOV and R_A_TOV.

E_D_TOV determines the life of an individual FC frame in any particular Fabric element. The effects of E_D_TOV on the Fabric as a whole are typically cumulative since each Fabric element contains its own E_D_TOV timers for any frame received.

R_A_TOV determines the life of an individual FC frame in the Fabric as a whole. For a Fabric, R_A_TOV implies that no particular frame shall remain in, and thus be emitted from, the Fabric after the timer expires.

K_A_TOV is a timer defined in this standard that is used by the Link Keep Alive (LKA) ELS (see FC-LS-2) as a trigger for issuing LKA. The LKA should be sent at least every K_A_TOV if no traffic has been sent and/or received on the connection. The default value for K_A_TOV is 1/2 E_D_TOV.

5.6.5.2 TCP timers

Given the multitude of current and probable TCP implementations, IETF Requests For Comments related to TCP, applications network requirements, etc., it is impossible to provide even rudimentary guidance in suggesting values for the tunable parameters associated with TCP.

5.6.5.3 Maximum number of attempts to complete an encapsulated FC frame transmission

This is an unspecified parameter and is implementation specific.

5.6.5.4 Maximum number of outstanding encapsulated FC frames

This is an unspecified parameter and is implementation specific.

5.7 FC-BB_IP service considerations

5.7.1 Latency delay

The time required for a frame to pass from one FC-BB_IP device to another across the IP network is variable and beyond the direct control of the FC/FCIP Entity pair. However, the IP network transit time affects the FC Entity's ability to meet FC timeout requirements (e.g., the R_A_TOV requirements of the Fabric). Therefore, the FC Entity is required to use facilities provided by the FCIP Entity to compute the IP network transit time for frames. See 5.2.4.5.2.

Class F frames may be excepted from IP network transit time checking, however, all other classes of frames shall have their IP network transit time computed and checked. If a frame is found to have an IP network transit time that causes the frame's lifetime in the Fabric to exceed FC requirements, the FC Entity shall discard the frame.

5.7.2 Throughput

5.7.2.1 How timeouts affect throughput

Both FC and TCP timeouts affect throughput as follows:

a) small R_A_TOV values may cause encapsulated FC frames to be discarded frequently in the FCIP_DE necessitating FC end-node retransmissions;
b) large TCP timeouts may result in encapsulated FC frames becoming stale in the IP network, leading the FCIP_DE to discard them again necessitating FC end-node retransmissions; and
c) discarding encapsulated FC frames due to improper settings of timeout values and errors in the IP network lowers the effective throughput.

The FC/FCIP Entities have little or no control over TCP timeouts. The FC/FCIP Entities never initiate retransmissions, that is done either by TCP or by the FC end nodes.

5.7.2.2 How loss affects throughput

TCP retransmissions occur due to loss or corruption of TCP segments. If TCP retransmissions cause the allowed transit time to exceed a threshold, then encapsulated FC frames shall be discarded. Either case is likely to cause the effective throughput to be reduced.

5.7.2.3 Other factors that affect throughput

Throughput may be affected by a mismatch in the effective rates of data transfer across the FC and the IP network interfaces. This mismatch may occur due to differences in the physical line speeds at the FC network and the IP network interfaces or due to the fundamental difference in the two flow control mechanisms.

FC uses BB_Credit flow control and TCP uses a sliding window based flow control. FC-BB_IP does not specify the mechanism that aligns the two flow control schemes, although it is thought that performance may be affected if this aspect is not considered. The FC-BB_IP device needs to ensure that the TCP connections are able to handle the frame arrival rate from the FC Fabric. The FC Entity shall work cooperatively with the FCIP Entity to manage flow control problems in either the IP network or FC Fabric.

In order to achieve better TCP aggregate throughput properties in the face of packet losses, a pair of peer FC-BB_IP devices may use multiple DEs between them, and use appropriate policies for mapping FC frames to these connections.

5.7.3 Reliability

5.7.3.1 Loss of connectivity

The FC-BB_IP device has the capability of detecting loss of connectivity with its remote peer (see 5.6.4.2). Upon detecting a loss of connectivity, an FC-BB_IP device establishes a new connection, or uses an existing TCP connection to the same FC-BB_IP device endpoint. An FC-BB_IP device shall not retransmit an encapsulated FC frame on the new connection. This is to ensure exactly-once delivery semantics to the FC endpoint.

The FC Entity may test for failed TCP connections. Should such a test detect a failed TCP connection, the FC Entity shall disconnect that connection following the procedures in 5.6.3.4.

5.7.3.2 Loss of synchronization

The FC Entity shall recover from events that the FCIP Entity is unable to handle, such as:

a) loss of synchronization with FC-BB_IP encapsulated FC frame headers from the Encapsulated Frame Receiver Portal requiring resetting the TCP connection; and
b) recovering from FC-BB_IP encapsulated FC frames that are discarded as a result of synchronization problems (see RFC 3821).

5.7.3.3 Loss or corruption of TCP segments

TCP flow control and error control has mechanisms to detect lost or corrupted TCP segments. TCP retransmits the TCP segments that were lost or corrupted.

TCP flow control provides the ability to regulate the flow of data on the IP network interface based on the perceived IP network congestion conditions, potentially avoiding large losses of data.

5.7.3.4 Loss or corruption of FC frames

The FC interface of the FC-BB_IP device has no mechanisms to detect lost data but only to detect corrupted frames. Corrupted frames detected prior to transmission into the IP network, are discarded and not sent over the IP network.

FC BB_Credit flow control provides the ability to regulate the flow of data on the FC network interface with no loss.

5.7.3.5 FCIP error reporting

The FC Entity receives notifications from the FCIP Entity due to a number of errors detected by the FCIP Entity. As a result, the E_Port implementation of the FC Entity shall report those errors to the local FC switch element via the local VE_Port (see figure 10). Similarly the B_Port implementation shall report the error to the local B_Access (see figure 14). In addition, the FC Entity may pass these error reports to the local PMM for inclusion in a local event log.

The FC Entity shall convert the error message received from the FCIP Entity into a Registered Link Incident Report (RLIR) (see FC-LS-2). It is the RLIR that is forwarded from the FC Entity to either the VE_Port (see figure 10) or B_Access (see figure 14).

On receipt of the message from the FC Entity, VE_Port or B_Access shall immediately forward the RLIR to the Domain Controller of the Switch.

As a minimum the FC Entity shall accept the following information from the FCIP Entity:

a) loss of FC frame synchronization (see RFC 3821);
b) failure to setup TCP connection (see RFC 3821);
c) duplicate connect request (see RFC 3821) ;
d) TCP connect request timeout (see RFC 3821) ;
e) successful completion of FC Entity request to close TCP connection (see RFC 3821);
f) loss of TCP connectivity (see RFC 3821);
g) excessive number of dropped datagrams (see RFC 3821);
h) any confidentiality violations (see RFC 3821);
i) SA parameter mis-match (see RFC 3821); and
j) LKA timeout notification (see FC-LS-2).

The FC Entity shall generate and forward an RLIR to the management server for the following:

a) loss of FC frame synchronization (see RFC 3821);
b) failure to setup additional TCP connection (see RFC 3821); and
c) additional duplicate TCP connect request (see RFC 3821).

5.7.4 Quality of Service (QoS)

The FC-BB_IP protocol may use TCP/IP QoS features to support FC capabilities.

5.7.5 Delivery order

Each VE_Port/FCIP_LEP pair defines a separate FCIP Link. FCIP_DEs within an FCIP_LEP share the FCIP Link. Multiple FCIP_DEs between FCIP_LEPs introduce multiple traffic paths (e.g., Class F, Class 2/3). The order in which the FCIP_DEs are serviced on the FCIP Link is not specified. One possibility is providing different priority levels to each traffic path changing the overall delivery order.

The only delivery order guarantee provided by TCP is correctly ordered delivery of FC-BB_IP encapsulated FC frames between a pair of FCIP_DEs. The FC Entity is expected to specify and handle all other FC frame delivery ordering requirements.

NOTE 11 – The order of the FC frames sent by the encapsulated frame transmitter may not be the same as the order sent by the source FC end node. This is due to the fact that some types of FC login allow FC frames to be re-ordered in the FC Fabric before reaching the FC receiver port.

5.7.6 IP multicast and broadcast

An FC-BB_IP device shall not make use of IP multicast and broadcast.

5.7.7 Security and authentication

The IETF security standards referenced by RFC 3821 provide numerous mechanisms for securing TCP connections between FC/FCIP Entity pairs (e.g., IPsec packet authentication and confidentiality). It is important to note that the public Internet IP network is subject to a large variety of security attacks, meaning that serious consideration should be given to enabling the full suite of security features described in FCIP whenever the public Internet IP network is to be used to transmit FCIP frames.

The TCP connection authentication mechanism described in 5.6.2.3.1 provides FC-BB-5 specific authentication for the second, third, etc., TCP connections in an FCIP Link and its associated Virtual ISL as long as the first TCP connection is authenticated using the mechanisms described in 5.6.2.3 and FC-SP.

6 Transparent FC-BB (FC-BB_GFPT and FC-BB_PW) Structure and Concepts

6.1 Applicability

Clause 4 discussed the FC-BB_GFPT and FC-BB_PW reference models. This clause discusses the FC-BB_GFPT and FC-BB_PW functional models.

6.2 FC-BB_GFPT overview

This clause discusses further aspects of FC-BB_GFPT operation, including initialization, flow control, and procedures for adaptation of FC information for transport using the Asynchronous Transparent Generic Framing Procedure (GFPT). Mapping FC-BB_GFPT into Asynchronous GFPT allows for applications where the WAN transport rate is less that the attached FC_Port data rate.

Figure 19 illustrates the protocol levels and layers involved in FC-BB_GFPT processes and devices.

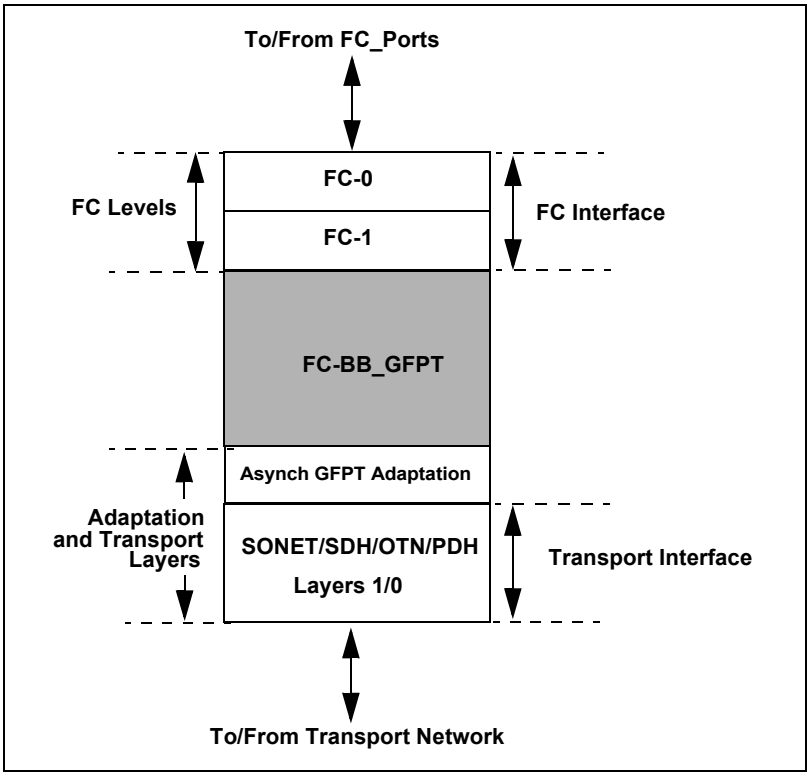

Figure 19 – FC-BB_GFPT protocol levels and layers

The FC-BB_GFPT level, the functions of which are defined by this standard, produces a stream of 8B/10B codewords in either direction. In the FC interface direction, it produces a full FC rate, synchronous codeword stream, constructed from the FC Ordered Sets (i.e., frames, Primitive Signals and Primitive Sequences) delivered by the Asynchronous GFPT adaptation level, and emulates a standard FC-2 level interfacing to a standard FC-1 level. In the WAN network direction, the FC-BB_GFPT level produces a filtered stream of codewords that includes all FC frames, selected Primitive Signals and selected Ordered Sets of Primitive Sequences that are forwarded by the FC-1 level, as well as GFPT_WAN Primitive Signals that are used for WAN flow control and management purposes. The Asynchronous GFPT adaptation level and the transport (i.e., SONET/SDH/OTN/PDH) layers are described by ITU-T standards (see 2.4).

If the FC-BB_GFPT device supports frame compression (see ANSI INCITS 241-1994 (R1999)), the FC Frames in the WAN network direction may be encoded and compressed to allow greater throughput for cases where the provisioned WAN bandwidth is less than the attached FC_Port data rate.

Buffering of selected codewords and Ordered Sets, in both directions of propagation (i.e., from the FC-1 level, and from the GFPT adaptation level), is one function of the FC-BB_GFPT level components.

6.3 FC-BB_PW overview

This subclause discusses further aspects of FC-BB_PW operation, including initialization and procedures for adaptation of FC information for transport over MPLS networks (see RFC 3031).

The FC-BB_PW functionality is separated into two layers as shown in figure 20.

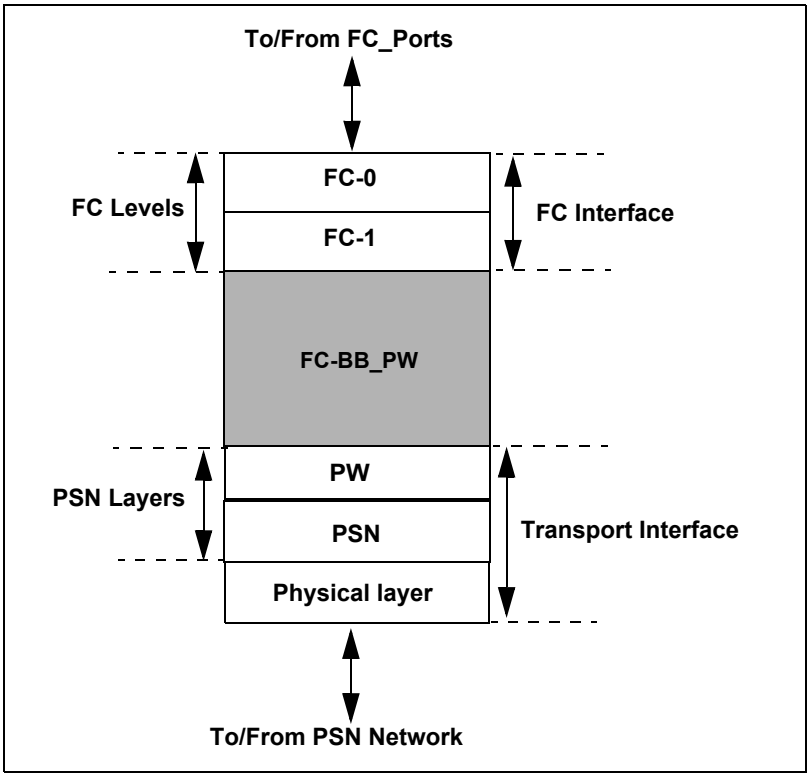

Figure 20 – FC-BB_PW protocol levels and layers

The FC-BB_PW level, specified in this standard, produces a full FC rate, synchronous stream of 8B/10B codewords in the FC interface direction. The stream is constructed from FC frames, Primitive Signals, and Primitive Sequences delivered by the PW Termination layer. This layer emulates a standard FC-2 layer to the underlying standard FC-1 layer. In the WAN direction, the FC-BB_PW layer produces a filtered stream of bytes that include all FC frames, selected Primitive Signals, and selected Primitive Sequences that are delivered by the FC-1 layer. The FC-BB_PW Termination layer is described in draft-ietf-pwe3-fc-encap-06, and the MPLS layer is described in RFC 3031.

6.4 Transparent FC-BB functional model

6.4.1 Transparent FC-BB initialization

Transparent FC-BB devices do not directly participate in FC link initialization or FC_Port initialization. FC link initialization and FC_Port initialization occurs between the attached FC_Port and the remote FC_Port. This is a key distinction between the Transparent FC-BB device model and other FC-BB-5 models.

Transparent FC-BB devices monitor and transport the FC_Port initialization Exchanges (i.e., ELP, FLOGI, and PLOGI) that take place between the attached FC_Port and the remote FC_Port, but only to capture and potentially modify the parameters that are relevant to link-level flow control.

If an FC-BB_GFPT device hosts more than one GFPT_WAN facility, then a separate state machine operates on each GFPT_WAN facility. The state machines in each FC-BB_GFPT device on the same GFPT_WAN facility operate independently of one another.

If an FC-BB_PW device hosts more than one PW facility, then a seperate state machine operates on each PW facility. The state machines in each FC-BB_PW device on the same PW facility operate independently of one another.

6.4.2 Transparent FC-BB initialization state machine

6.4.2.1 Initialization state machine keywords

The keywords used in the Transparent FC-BB initialization state machine diagram (see figure 21) are specified in table 14.

Table 14 – Transparent FC-BB initialization state machine keywords

Keyword	Description
WAN	Received from remote Transparent FC-BB device.
FC	Received from attached FC_Port.
WAN-Error	For FC-BB_GFPT, WAN-Error indicates the reception of 10B_ERR (i.e., a GFPT-defined character used to represent an FC illegal codeword or running disparity error at ingress, or an irresolvable character error produced during GFPT_WAN transmission. Interpretation of these characters is described in 6.4.8.1).) For FC-BB_PW, WAN-Error indicates the reception of an errored FC frame as specified in FC-FS-3 (e.g., CRC error, invalid EOF), or the reception of an error indication control frame.
FC-Error	8B/10B character error or running disparity error.
PSig	Primitive Signal
PSeq	Primitive Sequence
Arb-Loop	Arbitrated Loop

6.4.2.2 Initialization state machine

The Transparent FC-BB initialization state machine is specified in figure 21.

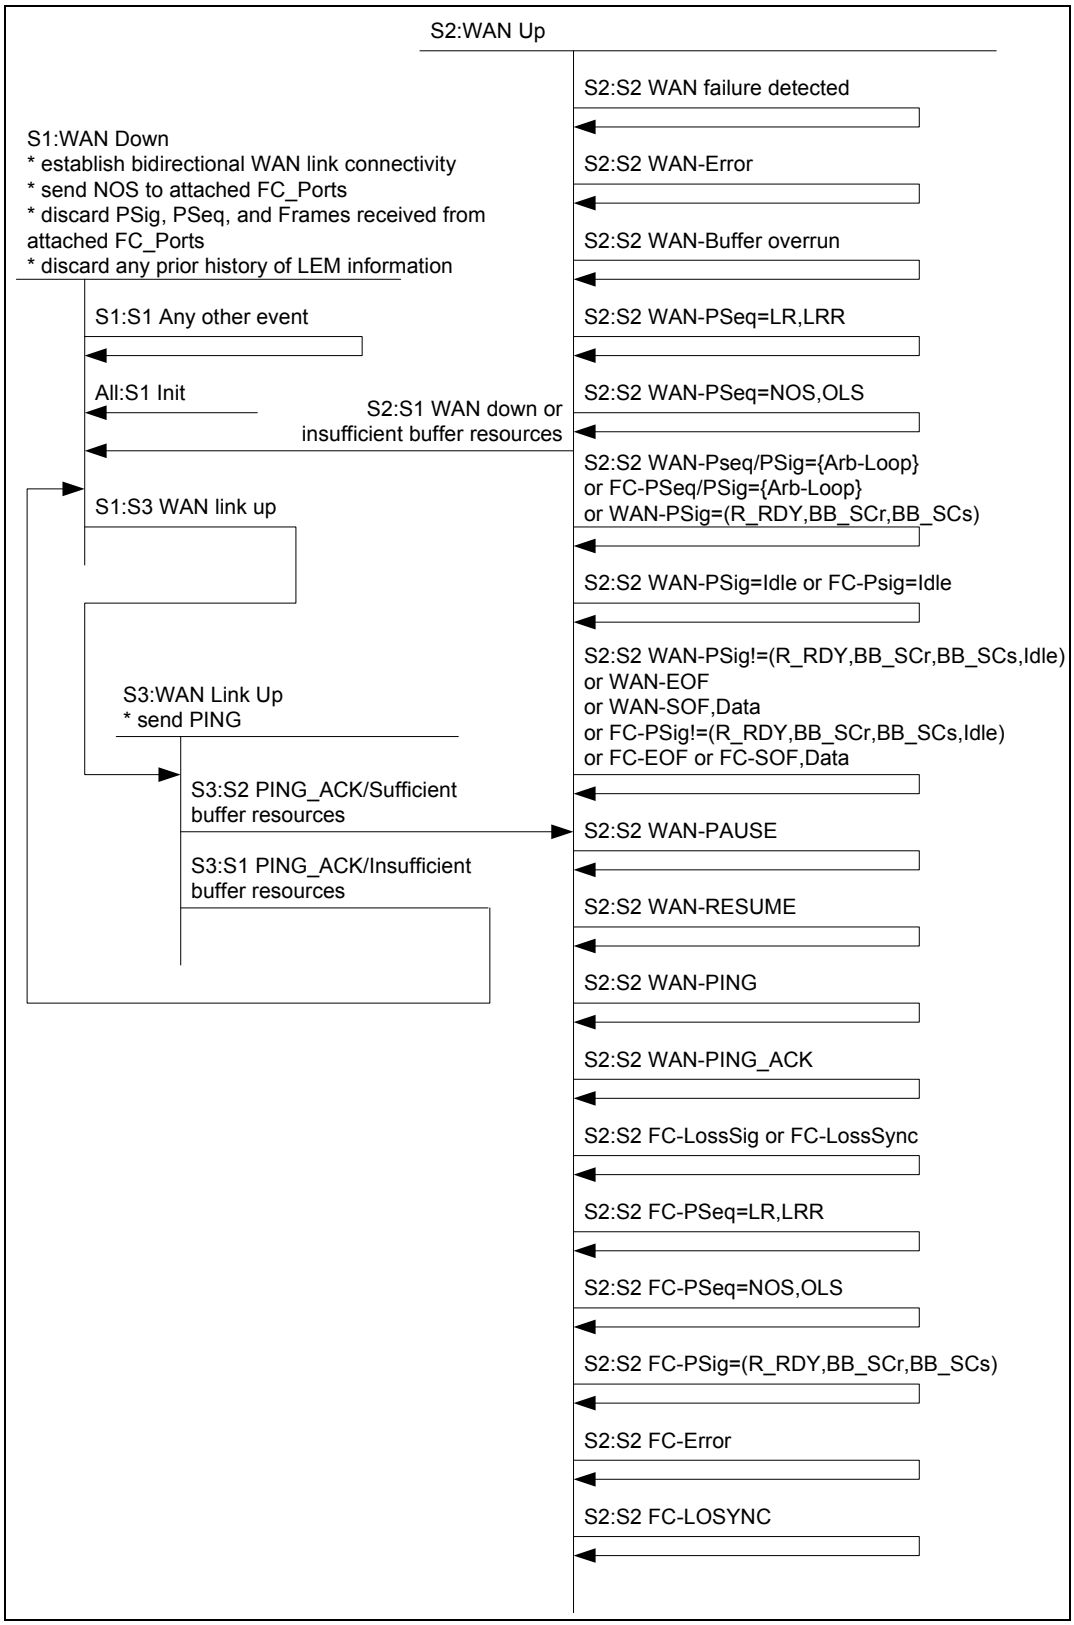

Figure 21 – Transparent FC-BB initialization state machine

Transition All:S1 Init. This transition occurs when an initialization event occurs in a state where it is not already handled. An initialization event may be:

a) a power-on reset condition;
b) outside intervention requesting an initialization event; or
c) a vendor-specific initialization event.

State S1: WAN Down. In this state, Transparent FC-BB devices shall:

a) initiate establishment of WAN links, either Transport Trail with a fully operational GFP server in both directions (see 6.4.8.1) or PW;
b) send the Not Operational (NOS) Primitive Sequence to attached FC_Ports, if any;
c) discard any Primitive Signals, Primitive Sequences, or frames received from attached FC_Ports; and
d) discard any prior history of Login Exchange Monitor (LEM) information (see 6.4.3).

Transition S1:S1 Any other event. This transition occurs when an event other than what is specified in state S1 occurs.

Transition S1:S3 WAN link up. This transition occurs when a WAN link up event is detected.

State S3: WAN Link Up. In this state, the Transparent FC-BB device shall transmit one or more PING signals to the remote Transparent FC-BB device.

Transition S3:S2 PING_ACK/Sufficient buffer resources. This transition occurs upon reception of at least one PING_ACK signal from the remote Transparent FC-BB device and sufficient buffer resources are available to prevent data loss (see 6.4.7). This transition is not coordinated between the Transparent FC-BB devices at each end of the WAN link.

Transition S3:S1 PING_ACK/Insufficient buffer resources. This transition occurs upon reception of at least one PING_ACK signal from the remote Transparent FC-BB device and WAN latency measurements indicate insufficient buffer resources are available within the device to prevent data loss (see 6.4.7).

State S2: WAN Up. In this state, the WAN link is normally (i.e., excluding WAN protection events) operational and the attached FC_Ports at either end of the link are permitted to communicate with each other.

Validated (see 6.4.8.1) Primitive Sequences received from the attached FC_Ports shall be transmitted between attached FC_Ports by the Transparent FC-BB devices according to the rate adaptation and other rules described in 6.4.8.1.

If a Primitive Sequence is received from an attached FC_Port while a Transparent FC-BB device has ceased transmitting information across the WAN link because it has received one or more PAUSEs, the Transparent FC-BB device shall:

a) flush its WAN-facing buffer;
b) clear the WAN pause condition; and
c) transmit the Primitive Sequence to the remote Transparent FC-BB device.

If a Primitive Sequence is received from a remote Transparent FC-BB device while a Transparent FC-BB device is unable to transmit frames to an attached FC_Port because it has no available BB_Credit, the Transparent FC-BB device shall:

a) flush its attached FC_Port-facing buffer; and
b) transmit the Primitive Sequence to the attached FC_Port.

When FC-BB_GFPT devices have no data to transmit to the WAN, the FC-BB_GFPT device shall follow the rules specified in ITU-T Rec. G.7041/Y.1303.

Transition S2:S1 WAN down or insufficient buffer resources. This transition occurs when:

a) a WAN down event is detected. A WAN down event occurs when GFPT_WAN recovery does not occur within WAN_HOLDOFF_TOV (see 6.4.9) of the time the GFPT_WAN failure is detected; or
b) WAN latency measurements (i.e., using PING and PING_ACK signals) indicate insufficient buffer resources are available within the device to prevent data loss (see 6.4.7).

This transition is not coordinated between the Transparent FC-BB devices at each end of the WAN link.

Transition S2:S2 WAN failure detected. This transition occurs when WAN failure is detected. WAN failure is detected according to criteria that are WAN specific and outside the scope of this standard (e.g., LOS, LOP, AIS, in the case of SONET/SDH, see ITU-T Rec. G.806 and ITU-T Rec. G.783, or RFI, LDP signaling failure in the case of PW, see RFC 4447) (see 6.4.8.1).

The Transparent FC-BB device shall start the WAN_HOLDOFF_TOV timer (see 6.4.9) upon detection of a WAN failure.

WAN link interruptions or protection events may be concealed from the attached FC equipment in a manner beyond the scope of this standard. During such interruptions, the Transparent FC-BB device may operate as though the WAN link were operational and transmit to attached FC_Ports:

a) the Primitive Sequence that they were transmitting at the time the WAN failure was detected;
b) Idles if they were not transmitting a Primitive Sequence at the time the WAN failure was detected; or
c) for an FC-BB_GFPT device, 8B/10B-encoded interpretations of 10B_ERR characters (see ITU-T Rec. G.7041/Y.1303).

Transition S2:S2 WAN-Error. For an FC-BB_GFPT device, this transition occurs when a 10B_ERR character is received from the WAN (see 6.4.8.1). For an FC-BB_PW device, this transition occurs when an errored frame (see FC-FS-3) or error indication control frame is received from the WAN (see 6.4.8.2).

Transition S2:S2 WAN-Buffer overrun. This transition occurs when a buffer overrun is detected. This transition is not expected to occur, but if it does, the number of dropped FC frames shall be counted.

Transition S2:S2 WAN-PSeq=LR,LRR. This transition occurs when a LR or LRR Primitive Sequence is received from the WAN. BB_Credit_CNT shall be set to zero, and BB_Credit shall be set to the value that was established by the last LEM (see 6.4.3).

Transition S2:S2 WAN-PSeq=NOS,OLS. This transition occurs when a NOS or OLS Primitive Sequence is received from the WAN. Any existing history of LEM information shall be discarded, BB_Credit_CNT shall be set to zero, and BB_Credit shall be set to one.

Transition S2:S2 WAN-PSeq/PSig={Arb-Loop} or FC-PSeq/PSig={Arb-Loop} or WAN-PSig=(R_RDY,BB_SCr,BB_SCs). This transition occurs when an Arbitrated Loop Primitive

Sequence or Primitive Signal is received from the WAN or the attached FC_Port, or when an R_RDY, BB_SCr, or BB_SCs Primitive Signal is received from the WAN. These Primitive Sequences and Primitive Signals are not expected and shall be discarded.

Transition S2:S2 WAN-PSig=Idle or FC-PSig=Idle. This transition occurs when an Idle Primitive Signal is received from the WAN or from the attached FC_Port.

Idle Primitive Signals shall be handled in the same manner as Primitive Sequences when they are received by a Transparent FC-BB device in a direction of propagation that has most recently seen reception of a Primitive Sequence, as opposed to a frame or FC Primitive Signal. At all other times (i.e., when they appear in a direction of propagation that has most recently seen reception of a frame or FC Primitive Signal, as opposed to a Primitive Sequence), Idle Primitive Signals shall be

a) discarded by FC-BB_PW devices; or
b) discarded or transported end-to-end by FC-BB_GFPT devices without impeding the forwarding of FC frames and other Primitive Signals, which have higher priority.

In all circumstances, Idle Primitive Signals shall be generated as required for forwarding to attached FC_Ports, consistent with the rules defined in FC-FS-3.

Transition S2:S2 WAN-PSig!=(R_RDY,BB_SCr,BB_SCs,Idle) or WAN-EOF or WAN-SOF,Data or FC-PSig!=(R_RDY,BB_SCr,BB_SCs,Idle) or FC-EOF or FC-SOF,Data. This transition occurs when a Primitive Signal other than R_RDY, BB_SCr, BB_SCs, or Idle, an EOF, or SOF with subsequent data is received from the WAN or a Primitive Signal other than R_RDY, BB_SCr, BB_SCs, or Idle, an EOF, or SOF with subsequent data is received from the attached FC_Port.

These Primitive Signals, Ordered Sets, and data shall be transported end-to-end without modification in the order in which they were received, except when flow control conditions require the flushing of buffers. The Ordered Sets shall be transported using the rate adaptation rules described in 6.4.8.1.

If a valid ELP, FLOGI, or PLOGI request frame (see FC-SW-5 and FC-LS-2) is received from the WAN or the attached FC_Port, the Transparent FC-BB device shall follow the Login Exchange Monitor (LEM) rules specified in 6.4.3.

Valid ELP, FLOGI, PLOGI, SW_ACC, and LS_ACC frames, may be modified by Transparent FC-BB devices according to the rules specified in 6.4.4, but only when such frames open or close a LEM as specified in 6.4.3.

Transition S2:S2 WAN-PAUSE. This transition occurs when a PAUSE signal is received from the WAN. Transmission of frames and Primitive Signals over the WAN shall be paused. See 6.4.7.

Transition S2:S2 WAN-RESUME. This transition occurs when a RESUME signal is received from the WAN. Transmission of frames and Primitive Signals over the WAN shall resume. See 6.4.7.

Transition S2:S2 WAN-PING. This transition occurs when a PING signal is received from the WAN. The Transparent FC-BB device shall transmit a PING_ACK to the remote Transparent FC-BB device. See 6.4.7.

Transition S2:S2 WAN-PING_ACK. This transition occurs when a PING_ACK signal is received from the WAN. See 6.4.7.

Transition S2:S2 FC-LossSig or FC-LossSync. This transition occurs when Loss of Signal or Loss of Synchronization (see FC-FS-3) is detected on the link to the attached FC_Port for a period of time greater than R_T_TOV. The Transparent FC-BB device shall transmit the NOS Primitive Sequence

for the duration of the Loss of Signal or Loss of Synchronization condition to the remote Transparent FC-BB device to initiate the Link Failure protocol (see FC-FS-3).

Transition S2:S2 FC-PSeq=LR,LRR. This transition occurs when a LR or LRR Primitive Sequence is received from the attached FC_Port. Any existing WAN Pause shall be cleared, BB_Credit_CNT shall be set to zero, and BB_Credit shall be set to the value that was established by the last LEM (see 6.4.3).

Transition S2:S2 FC-PSeq=NOS,OLS. This transition occurs when a NOS or OLS Primitive Sequence is received from the attached FC_Port. Any existing WAN Pause shall be cleared, any existing history of LEM information shall be discarded, BB_Credit_CNT shall be set to zero, and BB_Credit shall be set to one.

Transition S2:S2 FC-PSig=(R_RDY,BB_SCr,BB_SCs). This transition occurs when an R_RDY, BB_SCr, or BB_SCs Primitive Signal is received from the attached FC_Port. These Primitive Signals shall be terminated and shall not be forwarded to the remote Transparent FC-BB device. The R_RDY Primitive Signal is used for BB_Credit management with the attached FC_Port. The BB_SCr and BB_SCs Primitive Signals are used for BB_Credit Recovery with the attached FC_Port. See 6.4.5.

Transition S2:S2 FC-Error. This transition occurs when an 8B/10B character error or running disparity error is detected on the link to the attached FC_Port. See 6.4.8.1 for FC-BB_GFPT and 6.4.8.2 for FC-BB_PW.

Transition S2:S2 FC-LOSYNC. This transition occurs when a Loss of Synchronization is detected on the link to the attached FC_Port for a period of time less than R_T_TOV.

An FC-BB_GFPT device shall transmit the 10B_ERR character to the remote FC-BB_GFPT device for the duration of the Loss of Synchronization condition.

An FC-BB_PW shall indicate Loss of Synchronization to the remote FC-BB_PW device by transmitting to the WAN:

a) an FC frame with EOFni if it was in the middle of transmitting a frame at the time the FC-LOSYNC was detected; or
b) an error indication control frame (see 6.4.8.2.2) (i.e., if it was transmitting either Primitive Sequences or Idles at the time the FC-LOSYNC was detected). In this case the FC-BB_PW device shall continue transmitting to the WAN:
   A) the Primitive Sequence that it was transmitting at the time the FC-LOSYNC was detected; or
   B) Idles if it was not transmitting a Primitive Sequence at the time the FC-LOSYNC was detected.

6.4.3 Login Exchange Monitors

Login Exchange Monitors (LEMs) identify frames belonging to the ELP, FLOGI, and PLOGI Exchanges that take place between the attached FC_Ports that are directly interconnected by two Transparent FC-BB devices across a WAN link. The function of a LEM is to determine, and possibly modify, the BB_Credit and BB_SC_N values that are negotiated between the attached FC_Ports at each end of the link. LEMs operate according to a state machine that is described in table 15.

Six different LEMs are defined and designated using the notation LEM(rank, direction), where rank is either ELP, FLOGI, or PLOGI and direction is either outbound (see 3.3.10) or inbound (see 3.3.8).

The outbound direction is used when an ELP, FLOGI, or PLOGI request frame has been received from the attached FC_Port. The inbound direction is used when an ELP, FLOGI, or PLOGI request frame has been received from the remote Transparent FC-BB device.

Candidate LEM frames (see 6.4.8.1) shall be valid frames (see FC-FS-3).

The three LEM ranks are hierarchical as follows:

1) ELP (i.e., highest rank);
2) FLOGI; and
3) PLOGI (i.e., lowest rank).

All six LEMs have two states, open, and closed. LEMs may transition between opened and closed during WAN initialization in state S2 (see 6.4.2). LEMs have no significance in state S1 (see 6.4.2).

If a valid ELP, FLOGI, or PLOGI request frame (see FC-SW-5 and FC-LS-2) is observed in either direction, then a corresponding, direction-specific LEM is opened:

a) if no history has been saved for a successfully closed LEM in the specified direction; or
b) if the rank of the most recent successfully closed LEM in the specified direction is less than the rank of the valid ELP, FLOG, or PLOGI request frame being processed.

LEMs are opened separately and independently according to both rank and direction of propagation.

One or more LEMs per direction of propagation (i.e., outbound or inbound) may be open at one time. If a LEM is already in the open state, the observation of an ELP, FLOGI, or PLOGI request frame in the same direction (i.e., that may open a LEM corresponding to the same direction) either:

a) forces closure of the open LEM, and opening of the LEM corresponding to the observed ELP, FLOGI, or PLOGI request frame, if the new LEM has the same or higher rank than the already open LEM; or
b) if the new LEM has a lower rank than the already open LEM, the ELP, FLOGI, or PLOGI request frame and subsequent Exchange are saved (i.e., multiple LEM ranks may be open).

An open LEM is successfully closed by:

a) observation of a valid ELP, FLOGI, or PLOGI request frame in the direction associated with the LEM and of the same or higher rank (see 6.4.3). In this case, a new LEM is also opened; or
b) observation of a corresponding valid SW_ACC or LS_ACC frame in the direction opposite to the one associated with the LEM with the N_Port/F_Port bit set appropriately (i.e., zero for LS_ACC to PLOGI and one for LS_ACC to FLOGI).

Transparent FC-BB devices shall store the rank (i.e., ELP, FLOGI, or PLOGI) of the successful LEM closure. The successful closure of a LEM, in either direction, forces unsuccessful closure of all other LEMs in both directions.

When a LEM closes successfully, BB_Credit is initialized as per the newly negotiated parameters (see 6.4.5). LEM closure and BB_Credit initialization are not directly coordinated between the Transparent FC-BB devices at each end of the WAN link.

An open LEM is not successfully closed by:

a) observation of a corresponding valid SW_RJT, LS_RJT, F_RJT, P_RJT, F_BSY, P_BSY, or ABTS frame in the direction opposite to the one associated with the LEM; or

b) observation of a corresponding valid SW_ACC or LS_ACC frame in the direction opposite to the one associated with the LEM with the N_Port/F_Port bit set inappropriately (i.e., one for LS_ACC to PLOGI and zero for LS_ACC to FLOGI).

Table 15 provides a current state/next state description of the LEM state machine. For convenience of presentation, the six LEMs are collapsed into a single generic one. Every observation of candidate LEM frames shall be evaluated with respect to all six combinations of rank and direction. In most cases, the applicable row of table 15 is different among the six LEMs.

When successful LEM closure occurs in response to input event (5), input event (6) shall be applied to the other five LEMs.

Table 15 – Login Exchange Monitor (LEM) state machine

	Current State			
	LEM(Rank=A, Dir=X) Closed		LEM(Rank=A, Dir=X) Open	
Input Event:	**Next State(s)**	**Comments**	**Next State(s)**	**Comments**
1) Observation of valid ELP/FLOGI/PLOGI request frame, where rank of frame > A, in direction X	LEM(A,X) Closed	No action	LEM(A,X) Closed	Unsuccessful LEM closure
2) Observation of valid ELP/FLOGI/PLOGI request frame, where rank of frame = A, in direction X	LEM(A,X) Open	New LEM opening	LEM(A,X) Open	LEM "re-opening" - i.e., new information
3a) Observation of a valid FLOGI/PLOGI request frame, where rank of LEM = ELP, in direction X	LEM(A,X) Closed	No action	LEM(A,X) Open	No action
3b) Observation of a valid PLOGI request frame, where rank of LEM = FLOGI, in direction X	LEM(A,X) Closed	No action	LEM(A,X) Open	Open PLOGI LEM
4) Observation of valid ELP/FLOGI/PLOGI request frame, of any rank, in direction opposite to X	LEM(A,X) Closed	No action	LEM(A,X) Open	No action
5) Observation of valid SW_ACC/LS_ACC frame, corresponding to LEM(A,X)	LEM(A,X) Closed	Not expected - disregard	LEM(A,X) Closed	Successful LEM closure: rank = A; apply input event (6) against other LEMs
6) Successful closure of a LEM other than LEM(A,X)	LEM(A,X) Closed	No action	LEM(A,X) Closed	Unsuccessful LEM closure
7) Observation of valid SW_RJT/LS_RJT/F_BSY/P_BSY/F_RJT/P_RJT frame corresponding to LEM(A,X) or ABTS frame in the direction opposite to the one associated with LEM(A,X)	LEM(A,X) Closed	Not expected - disregard	LEM(A,X) Closed	Unsuccessful LEM closure
8) Primitive Sequence (LR,LRR) received in either direction	LEM(A,X) Closed	No action	LEM(A,X) Closed	Unsuccessful LEM closure
9) Primitive Sequence (OLS,NOS) received in either direction	LEM(PLOGI(credit=1), both directions) closed	No action	LEM(PLOGI(credit=1),both directions) closed	Unsuccessful LEM closure

Actions associated with LEM openings and closings are of two types, those associated with BB_Credit initialization, and those associated with the capture or modification of parameters that are

relevant to link-level flow control in frames processed by LEMs. These actions are discussed in detail, in 6.4.2 and 6.4.4 respectively, and are not considered in table 15.

6.4.4 Port initialization parameter observation and modification

Transparent FC-BB devices shall note the ISL Flow Control Mode value on inbound (see 3.3.8) and outbound (see 3.3.10) ELP request/reply SW_ACC frames of open LEMs. When the LEM is successfully closed and flow control other than either R_RDY flow control or VC_RDY flow control has been established between the interconnected FC_Ports, the Transparent FC-BB devices behavior is outside the scope of this standard.

If the use of R_RDY or VC_RDY flow control has been negotiated between the interconnected FC_Ports, the FC-BB_GFPT devices shall operate using unchannelized ASFC (see 6.4.7) and FC-BB_PW devices shall operate using Selective Retransmission mode as specified in draft-ietf-pwe3-fc-encap-06. Transparent FC-BB devices shall note the BB_Credit value on outbound ELP/FLOGI/PLOGI request/reply SW_ACC/LS_ACC frames of open LEMs, for subsequent use by BB_Credit flow control management. Similarly, Transparent FC-BB devices shall note the BB_Credit value on inbound ELP/FLOGI/PLOGI request/reply SW_ACC/LS_ACC frames of valid or open LEMs. If the BB_Credit value exceeds that corresponding to the size of the Transparent FC-BB device's outbound buffer for the attached FC_Port, then the BB_Credit value shall be overwritten with the value that corresponds to the outbound buffer size. Transparent FC-BB devices are responsible for recalculating the FC frame CRC in such cases. The number of BB_Credits to be used by flow control management at any time, and the governance of (re-)initialization of BB_Credit values, are specified in 6.4.2 and 6.4.5.

The use of channelized ASFC (see 6.4.7) when VC_RDY flow control has been negotiated between the interconnected FC_Ports is not defined in this standard.

Transparent FC-BB devices shall note the BB_Credit Management bit value on outbound FLOGI/PLOGI request/reply LS_ACC frames of open LEMs, for appropriate use by BB_Credit flow control management. When the use of alternate BB_Credit management has been indicated by an attached FC_Port and the Transparent FC-BB device does not support such management, the Transparent FC-BB device behavior is outside the scope of this standard.

Transparent FC-BB devices shall support BB_Credit Recovery (see 6.4.5). Transparent FC-BB devices shall note the BB_SC_N value in both inbound and outbound ELP/FLOGI/PLOGI reply SW_ACC/LS_ACC frames of open LEMs. The larger such value shall be used for BB_Credit Recovery purposes, unless either such value is zero, in which case BB_Credit Recovery shall not operate (see FC-FS-3).

The use of the RPSC (Report Port Speed Capabilities) ELS, by either attached FC_Port on a WAN facility, may in principle result in a mismatch of negotiated and reported port speed capabilities. The interception and local handling of RPSC ELSs by Transparent FC-BB devices, to prevent such mismatch, is not possible in general, since RPSC ELSs may be secured, depending on the security relationships that may be established between the communicating FC_Ports. However, any such reported mismatch should be strictly a second-order management issue and should not in any way impact the functioning of the FC_Ports and devices, the Transparent FC-BB devices, or the end-to-end data link.

6.4.5 Handling of BB_SCs, BB_SCr, and R_RDY Primitive Signals and BB_Credit initialization

Reception of BB_SCs and BB_SCr Primitive Signals from attached FC_Ports shall be noted for use in BB_Credit Recovery at the FC physical interface on the Transparent FC-BB device. Transparent

FC-BB devices shall generate BB_SCs and BB_SCr Primitive Signals, for transmission to attached FC_Ports, as required to properly perform BB_Credit Recovery with those ports. Adjustment of BB_Credit as specified in FC-FS-3 may be required on reception of a BB_SCs or BB_SCr Primitive Signal.

Reception of R_RDY Primitive Signals from attached FC_Ports shall be noted for credit/buffer management purposes. R_RDY Primitive Signals shall be generated by Transparent FC-BB devices, for transmission to attached FC_Ports, under the direction of credit/buffer managers.

The BB_Credit values assumed by credit/buffer managers, significant only in state S2, are generally those established by the most recently successfully closed LEM, and credit values/counts are initialized to these values immediately following LEM closure or upon transit of LR or LRR Primitive Sequences in either direction. In state S2, credit/buffer managers shall use a BB_Credit value of 1 in the absence of a history of successful LEM closure, and they shall also initialize credit values/counts to 1 upon transit in either direction of a Primitive Sequence other than LR or LRR.

The reception of R_RDY Primitive Signals, by Transparent FC-BB devices, from attached FC_Ports shall increment the credit available for transmission of frames to those ports, subject to the limitation that such credit counts/values shall not be incremented beyond the values of BB_Credit established by the most recently successfully closed LEM. The transmission of frames to attached FC_Ports shall decrement the credit available for transmission of frames to those ports, subject to the limitation that such credit counts/values shall not be reduced to a value less than zero. Frames shall be transmitted to attached FC_Ports only when current credit values are positive.

The generation of R_RDY Primitive Signals, by Transparent FC-BB devices, shall be managed such that buffer overflow and frame loss within the Transparent FC-BB devices, is precluded.

An example of a “start-up” sequence, showing the normal progression of state machine transitions, as well as R_RDY and BB_Credit handling, is shown in figure 22.

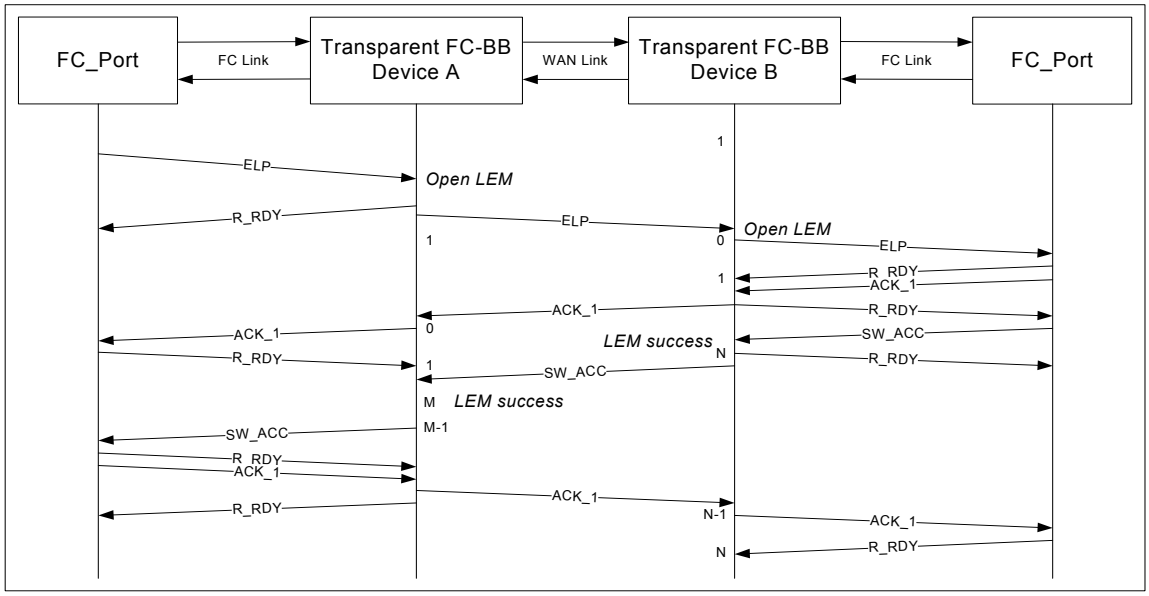

Figure 22 – Example port initialization process

Figure 22 illustrates an example port initialization process for ISL extension using Transparent FC-BB devices (i.e., attached port initialization is ELP). Simultaneous initiation of ELP by both ports represents a straightforward extension of the illustrated process. Only one of the ELP Exchange leads to successful LEM closure in both Transparent FC-BB devices. Both Transparent FC-BB devices are shown as beginning in state S2 and link initialization is assumed to have been completed between the attached E_Ports. The LEM opening and closing (i.e., success) events are shown for both devices. The evolution of credit counts used by Transparent FC-BB devices, to pace the delivery of frames to attached devices are also shown. LEM closure initializes the active values of those credit counts to M and N for Transparent FC-BB Device A and Transparent FC-BB Device B, respectively. M and N correspond to the BB_Credit values reported by the E_Ports in the ELP and SW_ACC frames, respectively. Corruption/loss of frames and/or R_RDY Primitive Signals, at any point in the end-to-end system, produces situations not qualitatively different, or different in terms of method of recovery, from corresponding events on standard ISLs.

If BB_Credit_CNT remains at zero for an interval of E_D_TOV or greater while the LEM is successfully closed, the Transparent FC-BB device shall transmit NOS to the local attached FC_Port and the remote attached FC_Port for 100ms +/- 25ms.

6.4.6 Transparent FC-BB Primitive Signals

Transparent FC-BB WAN signals are generated and terminated by Transparent FC-BB devices. Transparent FC-BB WAN signals are not sent to FC_Ports, nor do Transparent FC-BB devices expect to receive them from FC_Ports.

FC-BB_GFPT WAN signals are Ordered Sets defined as GFPT_WAN Primitive Signals. The most significant byte of GFPT_WAN Primitive Signals is K28.1. This character is chosen to prevent any possibility of conflict with Primitive Signals or other Ordered Sets that may be defined in any FC standard, for generation by any FC_Port.

FC-BB_PW WAN signals are control frames transmitted between FC-BB_PW devices.

6.4.7 Transparent FC-BB flow control

6.4.7.1 Overview

Alternate Simple Flow Control (ASFC) (see 6.4.7.2) shall be supported between FC-BB_GFPT devices.

FC-BB_PW devices shall support Selective Retransmission as specified in draft-ietf-pwe3-fc-encap-06.

6.4.7.2 FC-BB_GFPT Alternate Simple Flow Control (ASFC)

ASFC defines two Primitive Signals, ASFC_PAUSE and ASFC_RESUME (see table 16), that may be sent by one FC-BB_GFPT device to a remote FC-BB_GFPT device.

Table 16 specifies the Ordered Set values for the ASFC_PAUSE and ASFC_RESUME FC-BB_GFPT WAN Primitive Signals.

Table 16 – Values of FC-BB_GFPT ASFC_PAUSE and ASFC_RESUME Primitive Signals

FC-BB_GFPT Primitive Signal	Ordered Set (Value)
ASFC_PAUSE	K28.1 D00.0 D00.0 D00.0 - unchannelized
ASFC_RESUME	K28.1 D04.0 D00.0 D00.0 - unchannelized

An FC-BB_GFPT device sends ASFC_PAUSE to a remote FC-BB_GFPT device to direct that device to cease forwarding FC frames and non-GFPT_WAN Primitive Signals. Reception of a single ASFC_PAUSE Primitive Signal suffices to trigger such cessation, which shall take effect:

a) within 25 microseconds of reception of the ASFC_PAUSE Primitive Signal; or
b) immediately upon completion of the forwarding of a frame or Primitive Signal if forwarding of a frame or non-GFPT_WAN Primitive Signal is in progress when the ASFC_PAUSE Primitive Signal is received and the time required for such completion is greater than 25 microseconds.

When FC frame and non-GFPT_WAN Primitive Signal forwarding has been suspended due to reception of an ASFC_PAUSE Primitive Signal, it shall remain suspended until at least one ASFC_RESUME Primitive Signal has been received. FC-BB_GFPT devices shall:

a) not insert ASFC_PAUSE or ASFC_RESUME Primitive Signals within any FC Order Set;
b) insert ASFC_PAUSE or ASFC_RESUME Primitive Signals on transmission word boundries; and
c) not insert more than two ASFC_PAUSE and ASFC_RESUME Primitive Signals within the boundaries of any frame.

Otherwise, ASFC_PAUSE and ASFC_RESUME Primitive Signals are sent at times, in numbers, and at frequencies at the discretion of the sending FC-BB_GFPT device.

6.4.7.3 PING and PING_ACK signals

6.4.7.3.1 Overview

To assist in buffer management, as well as for other purposes, it may be useful to measure the latency on the WAN link. To facilitate this, two additional Transparent FC-BB signals are defined, PING and PING_ACK. For FC-BB_GFPT devices these signals are defined as Primitive Signals (see 6.4.7.3.2). For FC-BB_PW devices these signals are defined as control frames (see 6.4.7.3.3). As well as facilitating latency measurements, PING and PING_ACK signals are also used to communicate selected operational information between Transparent FC-BB devices.

The PING signal may be sent at any time by either Transparent FC-BB device, but not more frequently than once every 100 milliseconds. A single PING_ACK signal shall be transmitted by an Transparent FC-BB device in response to each received PING signal, within 25 microseconds of reception of the latter. FC-BB_GFPT devices shall respect FC Ordered Set delineation when inserting PING and PING_ACK Primitive Signals.

Transparent FC-BB devices shall begin transmission of PING signals as soon as possible upon establishment of WAN connectivity in state S1, and shall not exit state S1 until at least one PING_ACK signal has been received from remote Transparent FC-BB devices. This ensures that communication of support for channelized ASFC capabilities takes place between FC-BB_GFPT

devices before they exit state S1. Transmission of PING signals is otherwise undertaken at the discretion of individual Transparent FC-BB devices, within the frequency limits specified above.

6.4.7.3.2 FC-BB_GFPT PING and PING_ACK

Table 17 specifies the Ordered Set values for the PING and PING_ACK FC-BB_GFPT WAN Primitive Signals.

Table 17 – FC-BB_GFPT PING and PING_ACK Primitive Signal values

Primitive Signal	Ordered Set (Value)
PING	K28.1 D16.0 D(XNCCCRRRb) D(YZZZZZZZb)
PING_ACK	K28.1 D00.4 D(XNCCCRRRb) D(YZZZZZZZb)

The nomenclature D(MMMMMMMMb), MSB left to LSB right, is interpreted as the 8B/10B D-character equivalent of the binary value MMMMMMMMb.

The X bit indicates support (X=1b) or lack of support (X=0b) of channelized ASFC capability in the FC-BB_GFPT device transmitting the Primitive Signal. Since channelized ASFC is not defined by this standard, the value of the X bit bit shall be ignored by FC-BB_GFPT devices receiving PING and PING_ACK Primitive Signals.

The N bit may be set to 0b or 1b by devices transmitting PING, and the PING_ACK generated in response shall use the same value of N received in the associated PING.

Table 18 specifies the CCC bit field values.

Table 18 – FC-BB_GFPT PING and PING_ACK CCC bit field values

Value	Description	Reference
000b	No compression	
001b	LZS compression	6.4.10.3
010b-111b	Reserved	

The values of the R bits are reserved and shall be set to 0b.

The Y bit specifies vendor-specific use (i.e., Y=1b), or standard use (i.e., Y=0b), of the remaining bits (i.e., bits denoted Z) of the Primitive Signal. If Y=0b, the values of the bits denoted Z are reserved and shall be set to 0b. If Y=1b, the values assigned of the bits denoted Z are vendor specific.

6.4.7.3.3 FC-BB_PW PING and PING_ACK

The FC-BB_PW PING and PING_ACK control frame format is specified in figure 23.

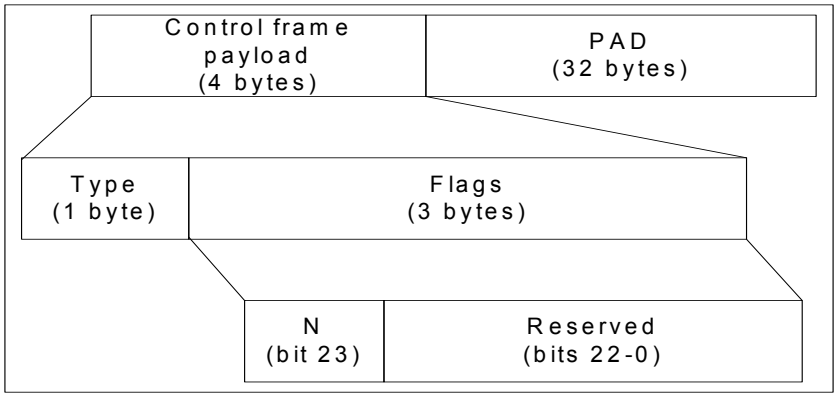

Figure 23 – FC-BB_PW PING and PING_ACK control frame format

The FC-BB_PW PING and PING_ACK control frame payload shall be set as specified in table 19.

Table 19 – FC-BB_PW PING and PING_ACK control frame payload values

Field	Value	Description
Type	0x81 0x82	PING PING_ACK
N	0b or 1b	The N bit is set to alternating 0b and 1b. For PING_ACK the N bit is set to the same value as the associated PING.
PAD	Zeroes	Set to all zeroes.

This standard specifies the use of Type field value 0x11 for transmission of error indications (see 6.4.8.2.2).

The FC-BB_PW PING and PING_ACK control frames are encapsulated and formatted for transmission as specified in draft-ietf-pwe3-fc-encap-06.

6.4.8 Adaptation of FC information for Transparent FC-BB

6.4.8.1 Adaptation of FC information for GFPT transport in FC-BB_GFPT

Adaptation and de-adaptation of attached FC_Port data in FC-BB_GFPT devices shall use the Asynchronous Transparent Generic Framing Procedure (GFPT) methodology (see ITU-T Rec. G.7041/Y.1303). Arbitrary concatenations of 8B/10B characters may be transported using this methodology, so that frame or message-oriented encapsulation of attached FC_Port data is not required. As such:

a) Ordered Sets (i.e., selected Primitive Signals and sub-rated Primitive Sequences) are directly adapted for transport using Asynchronous GFPT;
b) with no compression (see table 18), FC frames are directly adapted for transport using Asynchronous GFPT;

c) with compression (see table 18), at the ingress FC-BB_GFPT device, FC frames are encoded into GFPT_WAN compressed FC frames before entering the GFPT layer and at the egress FC-BB_GFPT device GFPT_WAN compressed frames are decoded back into FC frames; and
d) a GFPT_WAN compressed FC frame is composed of a special K28.7 SOF delimiter, compressed frame payload, and an EOF delimiter (6.4.10.1).

Asynchronous GFPT allows for adaptation to transport facilities of arbitrary bandwidths. Characters selected for transport are delivered in order from the de-adaptation process, irrespective of the GFPT_WAN link bandwidth. In this sense, the adaptation of FC frames and Primitive Signals is transparent.

Primitive Sequence adaptation is not purely transparent since the GFPT_WAN link may be sub-rate. Valid Primitive Sequences received from an attached FC_Port shall be adapted for transport by presenting at least three instances of the corresponding FC Ordered Set to the Asynchronous GFPT adaptation engine. The maximum rate of presentation of the corresponding FC Ordered Set to the Asynchronous GFPT adaptation engine shall not exceed the rate at which the transport facility is able to transport them. Primitive Sequences received from the GFPT_WAN link shall be recognized as valid when three Ordered Sets, corresponding to the Primitive Sequence, have been received in succession, excluding both 65B_Idle characters (see ITU-T Rec. G.7041/Y.1303) and GFPT_WAN Primitive Signals. If a valid Primitive Sequence is received from the GFPT_WAN link, the corresponding full-rate Primitive Sequence shall be constructed for forwarding to the FC_Port. Such forwarding shall continue until superceded by reception, from the remote FC-BB_GFPT device, of a different, valid Primitive Sequence, or a non-GFPT_WAN Primitive Signal or FC frame.

The handling of Idle Primitive Signals is described in 6.4.2.

In general, FC frames shall be forwarded without checking or modification of CRC values, and FC-BB_GFPT devices shall not discard or terminate any FC frames, except as dictated by 6.4.2 (i.e., when in state S1). Buffering that is otherwise necessary for CRC checking and potential frame discard based on the results of such checking, is therefore not required and should be avoided. FC frames should be transferred directly through the FC-BB_GFPT device without intermediate buffering to minimize overall latency, except as required for compression (see table 18) which may require temporary intermediate storage for the frame being compressed or decompressed. FC frames with invalid CRC values shall not be candidates for LEMs (see 6.4.3). The transit of candidate LEM frames through an FC-BB_GFPT device may be delayed by a maximum of 100 milliseconds while CRC checking, and operations specified in 6.4.2, 6.4.3, and 6.4.4 are performed. Such delays shall not affect the ordering of FC frames, Primitive Signals and Primitive Sequences that transit the FC-BB_GFPT device. All applicable rules (see FC-FS-3) regarding running disparity shall be followed by FC-BB_GFPT devices, with respect to the forwarding of FC frames, Primitive Signals, and Primitive Sequences received from a remote FC-BB_GFPT device to attached FC_Ports.

A Transport Trail that is in use as a GFPT_WAN link between FC-BB_GFPT devices is identified by the presence of the appropriately set User Payload Identifier (UPI) value (i.e., 0Ch) in the Payload Header of the GFP frames borne on the facility (see ITU-T Rec. G.7041/Y.1303).

Ingress FC character errors, ingress FC running disparity errors, or character errors resulting from bit errors occurring during transport on the GFPT_WAN link facility, are represented by GFPT (de-)adaptation engines using the 10B_ERR character (see ITU-T Rec. G.7041/Y.1303). FC-BB_GFPT devices shall interpret the 10B_ERR character, when GFPT de-adaptation engines present it, as (i.e., they shall generate) an unrecognized 8B/10B neutral disparity codeword, depending on beginning running disparity (RD+ or RD-), as follows:

a) either 001111 0001 (RD-) or 110000 1110 (RD+); or

b) some other invalid transmission characters, provided the invalid transmission characters also meet all 8B/10B coding rules, are of neutral disparity, and contain a minimum of one transition within both the first four bits and the last four bits of the codeword.

Synchronous and Asynchronous GFPT has significant, inherent transport error detection and correction capabilities. All detected but uncorrected codeword errors lead to the generation of 10B_ERR characters (i.e., to the generation of 8B/10B-encoded characters at egress). This functionality is sufficient to prevent the possibility of ambiguity in Primitive Signal or Ordered Set interpretation due to bit errors produced on the transport facility. But it may also result in the effective reception of incomplete Ordered Sets (i.e., Ordered Sets consisting partially of 10B_ERR characters) by an FC-BB_GFPT device from the GFPT_WAN. Corrupted FC frames, possibly including corrupted SOF and/or EOF Ordered Sets, thus may be produced. Such cases are indistinguishable, at the remote FC-BB_GFPT device, from FC frames that were received corrupted from FC_Ports and propagated to the remote FC-BB_GFPT device. In either case, no frame delimiter clean-up actions are mandated by this standard, nor are any vendor-specific implementations of such actions precluded. Otherwise, corrupted Ordered Sets received from remote FC-BB_GFPT devices shall be propagated through to attached FC_Ports, including during transmission of Primitive Sequences. Corrupted Ordered Sets shall not be interpreted as valid GFPT_WAN Primitive Signals.

GFPT error correction, if enabled, provides for minimization of such Ordered Set corruption during GFPT_WAN transport, and for improved channel performance in the face of an imperfect GFPT_WAN facility.

Communication of Loss of Client Signal and Loss of Client Synchronization using GFP Client Management Frames is not required by FC-BB_GFPT. If generated, these Client Management Frames shall not be used for FC-BB_GFPT state machine management or other purposes as specified in 6.4.2. FC-BB_GFPT devices are not required to generate or interpret these Client Management Frames, which may be used for administration and management purposes (e.g., to distinguish between Loss of Client Signal and Loss of Client Synchronization events).

6.4.8.2 Adaptation of FC information for PW transport in FC-BB_PW

6.4.8.2.1 Adaptation of FC information for PW transport in FC-BB_PW overview

Adaptation of FC data in FC-BB_PW equipment shall be performed by transparently encapsulating FC frames within PW packets. FC Ordered Sets (i.e., selected Primitive Signals and sub-rated Primitive Sequences) are encapsulated within FC-BB_PW control frames. The encapsulation method is outside the scope of this standard and is specified in draft-ietf-pwe3-fc-encap-06.

FC Frames and Ordered Sets selected for transport are delivered in order, irrespective of the PW bandwidth. In this sense, the adaptation of FC frames and Primitive Signals is transparent.

Primitive Sequence adaptation is not purely transparent since the PW may be sub-rate. Valid Primtive Sequences received from an attached FC_Port shall be adapted for transport by generating a low rate stream of control frames corresponding to the FC Ordered Set at the PW termination engine. If a control frame corresponding to a FC Primitive Sequence is received from the PW, the corresponding full-rate Primitive Sequence shall be constructed for forwarding to the FC_Port. Such forwarding shall continue until superceded by reception, from the remote FC-BB_PW device, of a different control frame, an FC Primitive Signal or an FC frame.

The handling of Idle Primitive Signals is described in 6.4.2.

In general, FC frames shall be forwarded without checking or modification of CRC values, and FC-BB_PW devices shall not discard any FC frames, except as dictated by 6.4.2 (i.e., when in state S1).

Buffering that is otherwise necessary for CRC checking and potential frame discard based on the results of such checking, is therefore not required and should be avoided. FC frames should be transferred directly through the FC-BB_PW device without intermediate buffering to minimize overall latency. FC frames with invalid CRC values shall not be candidates for LEMs (see 6.4.3). The transit of candidate LEM frames through an FC-BB_PW device may be delayed by a maximum of 100 milliseconds while CRC checking, and operations specified in 6.4.2, 6.4.3, and 6.4.4 are performed. Such delays shall not affect the ordering of FC frames, Primitive Signals and Primitive Sequences that transit the FC-BB_PW device. All applicable rules (see FC-FS-3) regarding running disparity shall be followed by FC-BB_PW devices, with respect to the forwarding of FC frames, Primitive Signals, and Primitive Sequences received from a remote FC-BB_PW device to attached FC_Ports.

When the FC-BB_PW device detects ingress FC character errors, ingress FC running disparity errors, or character errors resulting from bit errors it shall transmit to the WAN:

a) an FC frame with EOFni if it was in the middle of transmitting a frame at the time the error was detected (i.e., an error has occured in at least one of the characters of the FC frame); or
b) an error indication control frame (see 6.4.8.2.2) (i.e., if it was transmitting either Primitive Sequences or Idles at the time the character or running disparity error was detected). In this case the FC-BB_PW shall continue transmitting to the WAN:
   A) the Primitive Sequence that was transmitting at the time the error was detected; or
   B) Idles if it was not transmitting a Primitive Sequence at the time the error was detected.

If an FC-BB_PW device recieves an error indication control frame from the WAN it shall generate a single unrecognized 8B/10B neutral disparity codeword, depending on beginning running disparity (RD+ or RD-), as follows:

a) either 001111 0001 (RD-) or 110000 1110 (RD+); or
b) some other invalid transmission characters, provided the invalid transmission characters also meet all 8B/10B coding rules, are of neutral disparity, and contain a minimum of one transition within both the first four bits and the last four bits of the codeword.

After transmitting this codeword the FC-BB_PW device shall resume normal operation (i.e., transmitting FC frames, Primitive Sequences or Idel Primitive Signals as necessary).

Layer 2 technologies used to transport PW/MPLS frames have significant, inherent transport error detection capabilities. If a codeword error is detected in an FC frame, then the associated EOF Ordered Set shall be replaced with an EOFni Ordered Set. If a codeword error is detected in an FC-BB_PW control frame, then the control frame shall be discarded. This functionality is sufficient to prevent the possibility of errors in Primitive Signal or Ordered Set generation at egress due to bit errors produced on the transport facility. Corrupted FC frames, including an EOFni, may be transmitted. Such cases are indistinguishable, at the remote FC-BB_PW device, from FC frames that were received corrupted from FC_Ports and propagated to the remote FC-BB_PW device with an EOFni. In either case, no frame delimiter error correction actions are mandated by this standard, nor are any vendor-specific implementations of such actions precluded.

6.4.8.2.2 FC-BB_PW error indication control frame

The FC-BB_PW error indication control frame format is specified in figure 24.

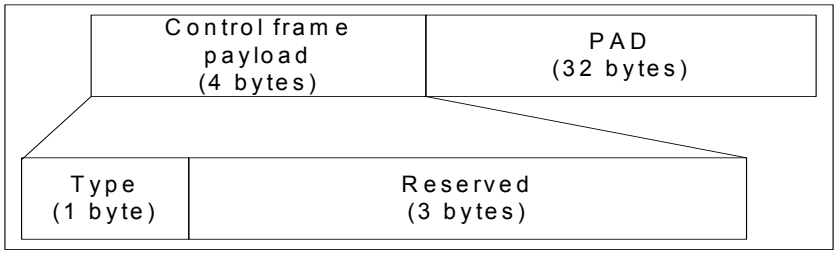

Figure 24 – FC-BB_PW error indication control frame format

The FC-BB_PW error indication control frame payload shall be set as specified in table 20.

Table 20 – FC-BB_PW error indication control frame payload values

Field	Value	Description
Type	0x11	FC-ERROR - indicates a received character error or Loss of Synchronization.
PAD	Zeroes	Set to all zeroes.

This standard specifies the use of Type field values 0x81 and 0x82 for transmission of FC-BB_PW PING and PING_ACK (see 6.4.7.3.3).

The FC-BB_PW error indication control frame is encapsulated and formatted for transmission as specified in draft-ietf-pwe3-fc-encap-06.

6.4.9 WAN Holdoff Timeout Value (WAN_HOLDOFF_TOV)

The WAN_HOLDOFF_TOV is the maximum time between detection of a GFPT_WAN link failure and state machine transition from S2 to S1 (see 6.4.2). The default value of WAN_HOLDOFF_TOV is 100 ms (i.e., 2 times the normal SONET/SDH Trail restoration time limit).

6.4.10 Transparent FC-BB frame compression encoding

6.4.10.1 FC-BB_GFPT FC frame compression

If an ingress FC-BB_GFPT device is operating with compression enabled (see 6.4.7.3.2), the ingress FC-BB_GFPT device:

a) shall send errored FC frames as FC frames; and
b) for error-free FC frames, optionally encode them into GFPT_WAN compressed FC frames as follows:
   1) replace the K28.5 character of the SOF with a K28.7 character and leave the remaining three characters of the SOF unchanged;
   2) send all the data words in the FC frame between the SOF and the EOF through the compression algorithm;
   3) pad the end of the encoded data stream to a transmission word boundary prior to the EOF; and

4) append the original EOF.

NOTE 12 – A compression algorithm may not compress some frames due to throughput performance or compression performance reasons.

After encoding, the stream is converted using Asynchronous GFPT.

If an egress FC-BB_GFPT device is operating with compression allowed (see 6.4.7.3.2), the egress FC-BB_GFPT device shall process a frame as follows:

a) Convert the entire stream from Asynchronous GFPT encoding into FC encoding;
b) if the frame begins with a K28.7 SOF:
   1) translate the FC frame K28.7 SOF back to the original K28.5 SOF;
   2) send all the data words in the FC frame between the SOF and the EOF through the decompression algorithm; and
   3) append the original EOF; and
c) forward the data stream to the FC layer for further processing (e.g., LEM).

6.4.10.2 FC-BB_PW FC frame compression

An FC-BB_PW device does not use frame compression.

6.4.10.3 LZS compression algorithm

The LZS algorithm (see ANSI INCITS 241-1994 (R1999)) is a general purpose lossless compression algorithm for use with a wide variety of data types.

The sender shall reset the LZS decompression history prior to processing each compressed FC frame. This ensures that each FC frame may be decompressed independently of any other and prevents errors in one frame from causing decompression errors in any subsequent frames that the receiver decompresses.

The decompression process of each compressed frame shall be is performed as specified in ANSI INCITS 241-1994 (R1999).

7 FC-BB_E Structure and Concepts

7.1 Applicability

Clause 4 discussed the FC-BB_E reference model. This clause discusses the FC-BB_E functional models.

7.2 FC-BB_E overview

This clause discusses aspects of the FC-BB_E mapping, including initialization and procedures for the mapping of Fibre Channel frames over Ethernet.

Figure 25 shows how FC-BB_E maps the Fibre Channel levels and sublevels over IEEE 802.3 layers.

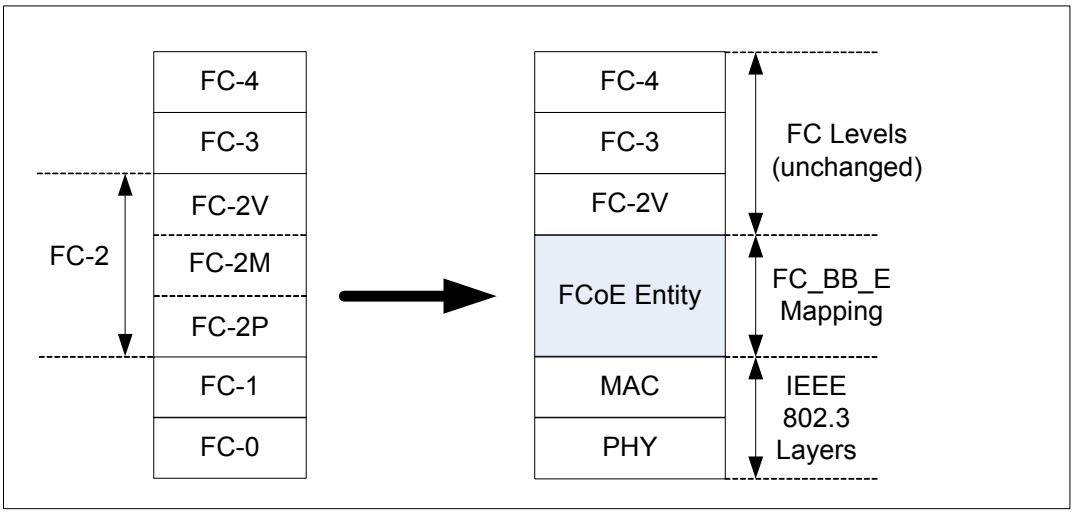

Figure 25 – FC-BB_E mapping

Figure 26 shows how the FC-BB_E mapping applies to FCoE Forwarders (FCF) and FCoE Nodes (ENodes).

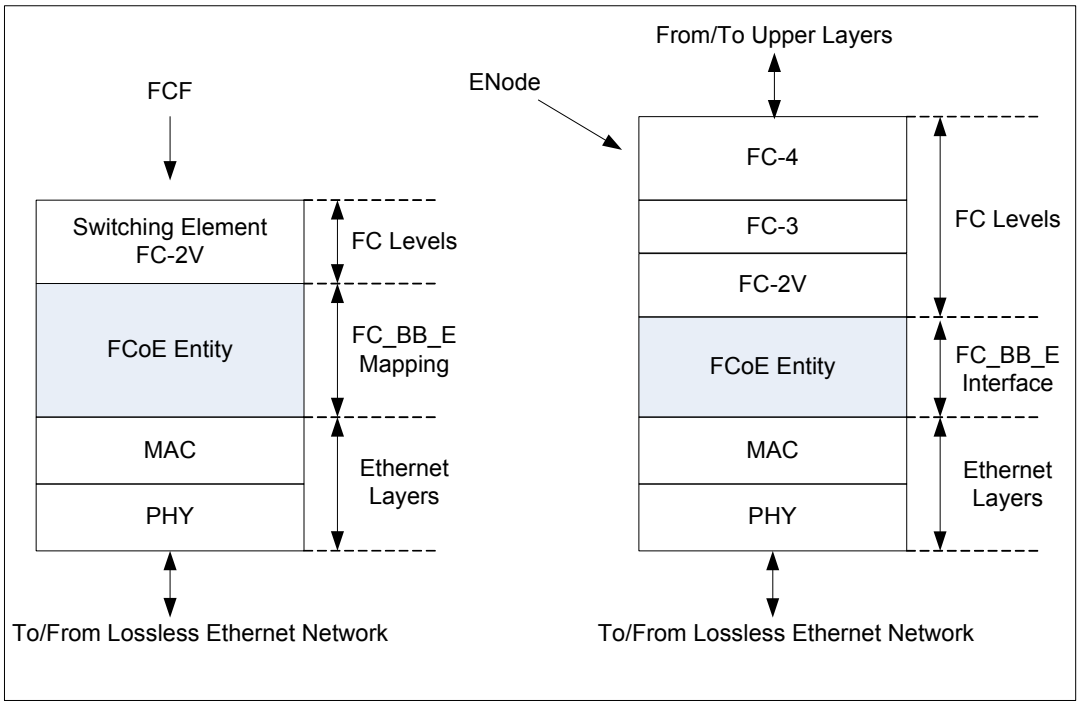

Figure 26 – FC-BB_E protocol levels and layers

FC-BB_E defines a direct mapping of Fibre Channel over Ethernet (FCoE). Although a generic Ethernet network may lose frames due to congestion, a proper implementation of appropriate Ethernet extensions (e.g., the PAUSE mechanism defined in IEEE 802.3-2008) allows a full duplex Ethernet link to provide a lossless behavior equivalent to the one provided by the buffer-to-buffer credit mechanism (see FC-FS-3). The protocol mapping defined by FC-BB_E is referred to as Fibre Channel over Ethernet (FCoE) and shall use an underlying Ethernet layer (i.e., composed only of full duplex links and providing a lossless behavior when carrying FCoE frames (see 4.4.4)). The Lossless Ethernet layer provides sequential delivery of FCoE frames.

In native Fibre Channel, Fibre Channel Nodes (see FC-FS-3) and Switches (see FC-SW-5) communicate through FC_Ports. Fibre Channel links connect PN_Ports to PF_Ports and PE_Ports to PE_Ports.

In Fibre Channel over Ethernet, FCoE Nodes (ENodes) and FCoE Forwarders (FCFs) communicate through Ethernet ports over a Lossless Ethernet network. FCoE Virtual Links replace the physical Fibre Channel links by encapsulating FC frames in Ethernet frames. FCoE supports VE_Port to VE_Port Virtual Links and VN_Port to VF_Port Virtual Links. A VE_Port to VE_Port Virtual Link is identified by the pair of MAC addresses of the two link end-points. A VN_Port to VF_Port Virtual Link is identified in general by the pair of MAC addresses of the two link end-points and by the N_Port_ID assigned to the VN_Port. When VN_Port MAC addresses are unique per VN_Port (e.g., when FPMAs are used (see 7.6)), the pair of MAC addresses of the two link end-points is enough to identify a VN_Port to VF_Port Virtual Link.

Figure 27 shows an example FCoE VN_Port to VF_Port network configuration.

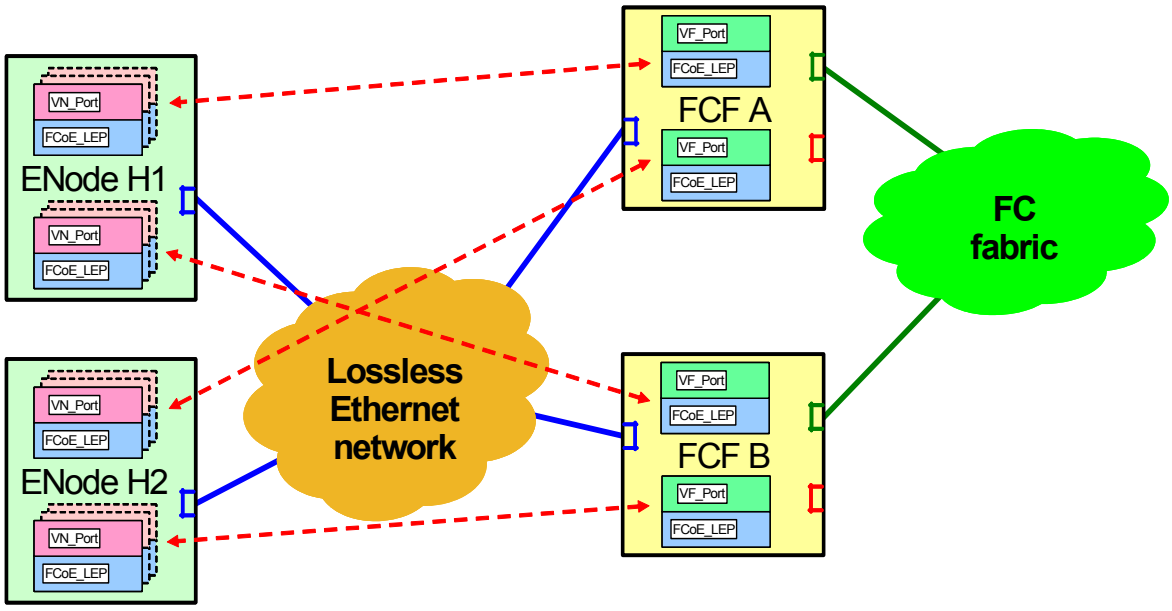

Figure 27 – FCoE VN_Port to VF_Port network configuration example

Each of the two ENodes H1 and H2 depicted in figure 27 has a single physical Ethernet connection to the Lossless Ethernet network. Each of the two FCFs, FCF A and B, has a single physical Ethernet connection to the Lossless Ethernet network. Each ENode may instantiate multiple VN_Ports, connected to VF_Ports instantiated by the FCFs through FCoE Virtual Links. The dotted lines in figure 27 depict possible VN_Port to VF_Port Virtual Links. In this case, a Lossless Ethernet network is reduced by FCoE to a set of point-to-point VN_Port to VF_Port Virtual Links where the VN_Port to VF_Port Fibre Channel protocols are able to operate.

Figure 28 shows an FCoE VE_Port to VE_Port network configuration.

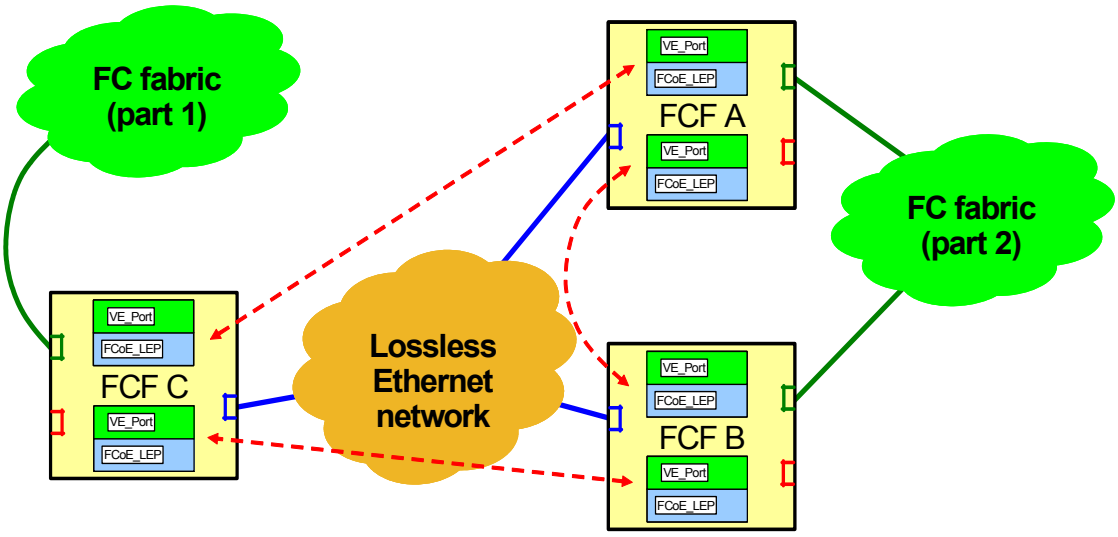

Figure 28 – FCoE VE_Port to VE_Port network configuration example

Each of the three FCFs A, B, and C depicted in figure 28 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 28 depict possible VE_Port to VE_Port Virtual Links. In this case, a Lossless Ethernet network is reduced by FCoE to a set of point-to-point VE_Port to VE_Port Virtual Links where the VE_Port to VE_Port Fibre Channel protocols are able to operate.

As shown in figure 27 and figure 28, Fibre Channel over Ethernet enables some additional features in respect to native Fibre Channel:

a) an ENode may establish VN_Port to VF_Port Virtual Links (i.e., perform Fabric Login) with multiple FCFs through a single Lossless Ethernet MAC;
b) an FCF may establish VN_Port to VF_Port Virtual Links (i.e., accept Fabric Login) with multiple ENodes through a single Lossless Ethernet MAC; and
c) an FCF may establish VE_Port to VE_Port Virtual Links with multiple other FCFs through a single Lossless Ethernet MAC.

7.3 ENode functional model

Figure 29 shows the functional model of an ENode, where the bracketed functional components are optional. An ENode is functionally composed of at least one Lossless Ethernet MAC (i.e., the ENode MAC), and an FCoE Controller function for each ENode MAC.

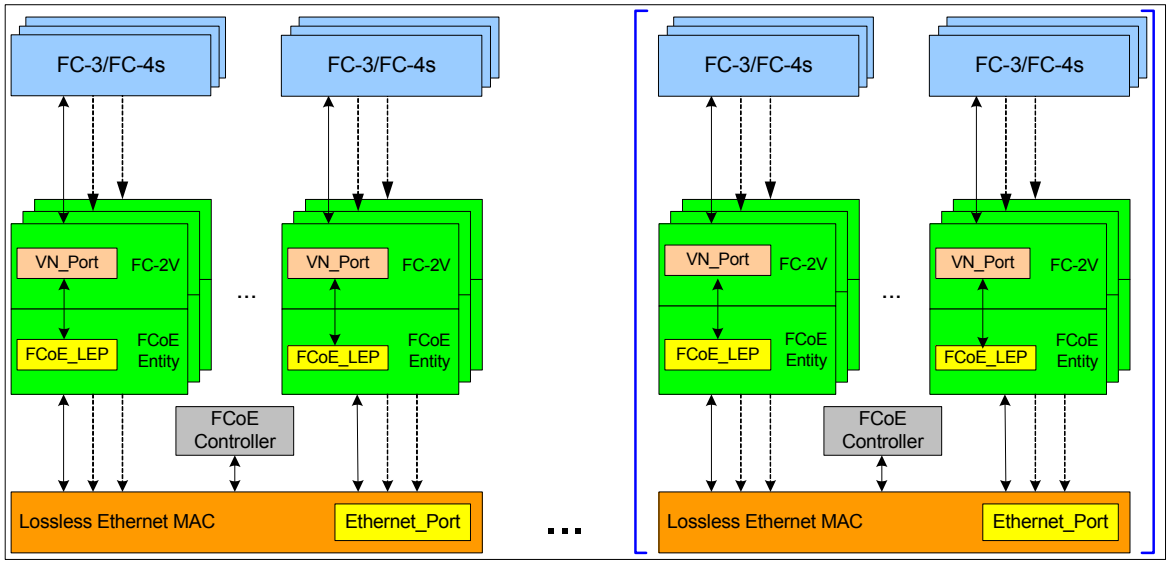

Figure 29 – ENode functional model

The FCoE Controller associated with an ENode MAC shall support the instantiation of VN_Port/FCoE_LEP pairs.

The FCoE Controller is the functional entity that performs the FCoE Initialization Protocol (FIP) and instantiates or de-instantiates VN_Port/FCoE_LEP pairs as needed.

For an ENode's MAC, the FCoE Controller:

a) optionally initiates the FIP VLAN discovery protocol to discover FCoE VLANs;

b) should initiate the FIP discovery protocol in order to discover VF_Port capable FCF-MACs connected to the same Lossless Ethernet network;
c) initiates FIP FLOGI Exchanges and instantiates a VN_Port/FCoE_LEP pair on successful completion of each FIP FLOGI Exchange with a VF_Port capable FCF-MAC;
d) optionally initiates FIP NPIV FDISC Exchanges and instantiates a VN_Port/FCoE_LEP pair on successful completion of each FIP NPIV FDISC Exchange with a VF_Port capable FCF-MAC;
e) de-instantiates a VN_Port/FCoE_LEP pair when that VN_Port is logged out from the Fabric;
f) initiates FIP LOGO Exchanges when explicit Fabric logout is needed;
g) de-instantiates the indicated VN_Port/FCoE_LEP pairs on receiving FIP Clear Virtual Link requests;
h) transmits periodic FIP Keep Alive frames on behalf of the ENode MAC every FKA_ADV_PERIOD, unless the D bit is set to one in received Discovery Advertisements (see 7.8.5.2);
i) monitors the status of instantiated VN_Port/FCoE_LEP pairs and transmits periodic FIP Keep Alive frames on their behalf every FKA_VN_PERIOD, unless the D bit is set to one in received Discovery Advertisements (see 7.8.5.2); and
j) monitors the status of VF_Ports to which the instantiated VN_Port/FCoE_LEP pairs are logged in by maintaining timers and verifying periodic FIP Discovery Advertisements are received within every FKA_ADV_PERIOD, unless the D bit is set to one in received Discovery Advertisements (see 7.8.5.2).

The FCoE Controller of an ENode MAC may perform FIP FLOGIs and FIP NPIV FDISCs with multiple VF_Port capable FCF-MACs. VN_Ports instantiated by the FCoE Controller of an ENode MAC on successful completion of FIP NPIV FDISC Exchanges with a VF_Port capable FCF-MAC are all associated with the same VF_Port, instantiated by the FCoE Controller of that VF_Port capable FCF-MAC on successful completion of a FIP FLOGI Exchange. In figure 29, each stack of VN_Port/FCoE_LEP pairs represents an association with a different VF_Port capable FCF-MAC.

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 an FCoE_LEP of an ENode 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 its MAC address.

A VN_Port is an instance of the FC-2V sublevel of Fibre Channel that operates as an N_Port (see FC-FS-3) and is dynamically instantiated together with its FCoE_LEP 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 an N_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-2 and FC-FS-3, with the exception that 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. When receiving FC frames from its FCoE_LEP, a VN_Port shall verify that the D_ID of the received FC frame is equal to its address identifier. If the check fails the FC frame shall be discarded.

NOTE 13 – The receive checks performed by the VN_Port/FCoE_LEP pair verify that the correct MAC destination address, MAC source address, and D_ID are present in a received FCoE frame (see D.5).

7.4 FCF functional model

Figure 30 shows the functional model of an FCF, where the bracketed functional components are optional. An FCF is functionally composed of a Fibre Channel Switching Element (see FC-SW-5) 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 Lossless 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.

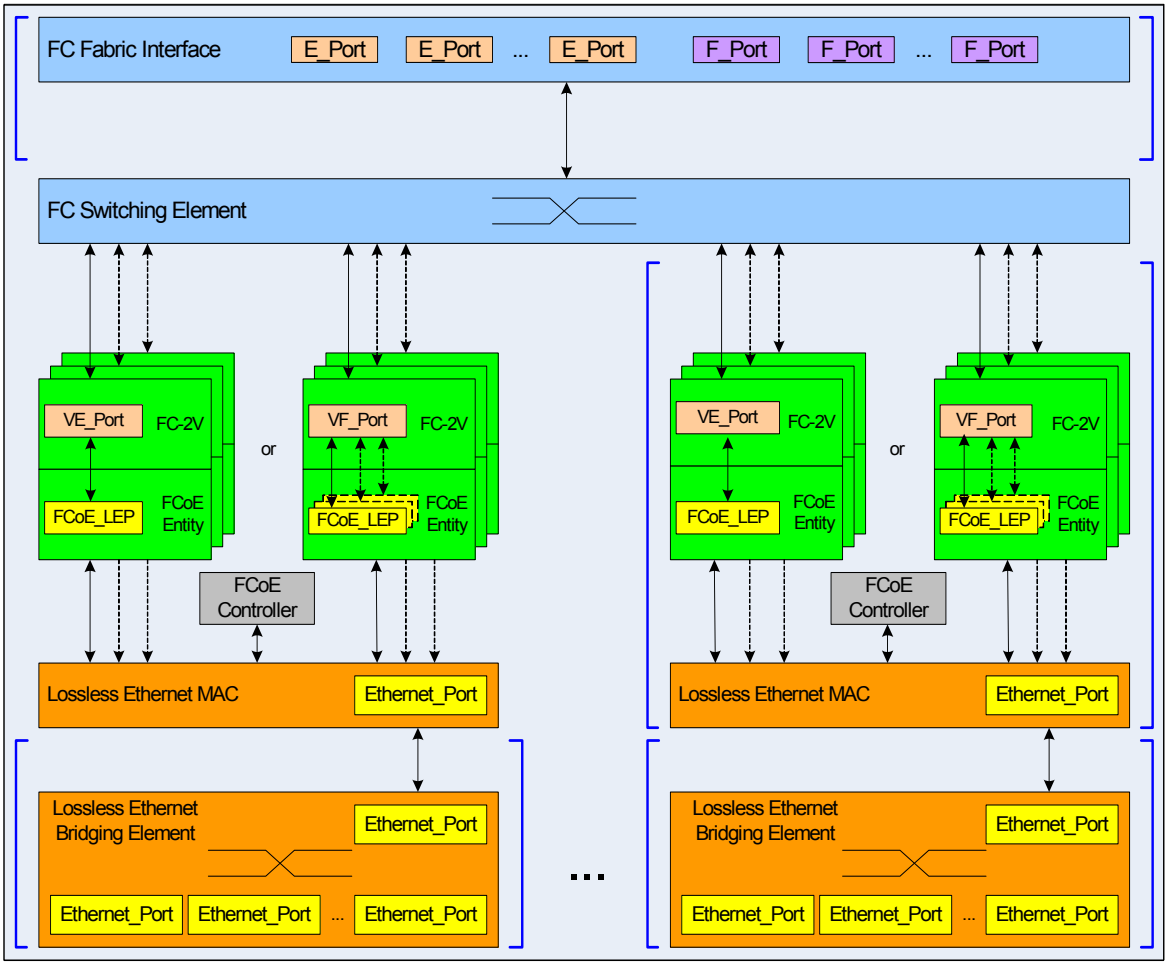

Figure 30 – FCF functional model

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

When an FCF includes Lossless Ethernet bridging elements, an FCF-MAC address may be used by multiple Ethernet ports of the FCF.

The FCoE Controller associated with an FCF-MAC shall support the instantiation of VE_Port/FCoE_LEP pairs or VF_Port/FCoE_LEP pairs. An FCF-MAC supporting the instantiation of VE_Port/FCoE_LEP pairs is referred to as a VE_Port capable FCF-MAC. An FCF-MAC supporting the instantiation of VF_Port/FCoE_LEP pairs is referred to as a VF_Port capable FCF-MAC. Support for both VE_Port/FCoE_LEP pairs and VF_Port/FCoE_LEP pairs on the same FCF-MAC is prohibited.

MAC addresses used by FCFs for FCF-MACs shall be different than MAC addresses used by ENodes for ENode MACs.

The FCoE Controller is the functional entity that performs the FCoE Initialization Protocol (FIP) and instantiates or de-instantiates VE_Port/FCoE_LEP pairs or VF_Port/FCoE_LEP pairs as needed.

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

a) optionally performs the FIP VLAN discovery protocol to discover FCoE VLANs;
b) discovers other VE_Port capable FCF-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 FCF-MAC;
d) de-instantiates a VE_Port/FCoE_LEP pair on receiving a FIP Clear Virtual Link request;
e) monitors the status of the instantiated VE_Port/FCoE_LEP pairs;
f) initiates FIP Clear Virtual Link requests as needed to terminate Virtual Links to other VE_Ports;
g) transmits periodic FIP Discovery Advertisements to the All-FCF-MACs address every FKA_ADV_PERIOD; and
h) monitors the status of remote VE_Ports by maintaining timers and verifying that periodic FIP Discovery Advertisements are received within every FKA_ADV_PERIOD.

For a VF_Port capable FCF-MAC, the FCoE Controller:

a) participates in the FIP VLAN discovery protocol initiated by an ENode MAC;
b) participates in the FIP discovery protocol initiated by an ENode MAC;
c) instantiates a VF_Port/FCoE_LEP pair on successful completion of each FIP FLOGI Exchange initiated by an ENode MAC;
d) instantiates an additional FCoE_LEP on successful completion of each FIP NPIV FDISC Exchange initiated by an already logged in ENode MAC. The resulting VF_Port/FCoE_LEP pair shares the VF_Port with the VF_Port/FCoE_LEP pair instantiated on successful completion of the FIP FLOGI Exchange;
e) when a VN_Port is logged out, de-instantiates the FCoE_LEP associated to that VN_Port and de-instantiates the corresponding VF_Port if that FCoE_LEP was the only one associated with that VF_Port;
f) initiates FIP Clear Virtual Link requests as needed to terminate Virtual Links to VN_Ports;
g) monitors the status of the instantiated VF_Port/FCoE_LEP pairs;
h) transmits periodic FIP Discovery Advertisements to the All-ENode-MACs address every FKA_ADV_PERIOD;
i) monitors the status of the logged in ENode MACs by verifying that periodic FIP Keep Alive frames are received within FKA_ADV_PERIOD, unless the D bit is set to one in received Discovery Advertisements (see 7.8.5.2); and
j) monitors the status of the logged in VN_Ports by maintaining timers and verifying that periodic FIP Keep Alive frames are received within FKA_VN_PERIOD, unless the D bit is set to one in received Discovery Advertisements (see 7.8.5.2).

The FCoE Controller of an ENode MAC may perform FIP FLOGIs and FIP NPIV FDISCs with multiple VF_Port capable FCF-MACs. VN_Ports instantiated by the FCoE Controller of an ENode MAC on successful completion of FIP NPIV FDISC Exchanges with a VF_Port capable FCF-MAC are all associated with the same VF_Port, instantiated by the FCoE Controller of that 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 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 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.

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 its MAC address.

A VE_Port is an instance of the FC-2V sublevel of Fibre Channel that operates as an E_Port in accordance with FC-SW-5 and is dynamically instantiated together with its FCoE_LEP 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 an E_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-5, with the exception that 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 an instance of the FC-2V sublevel of Fibre Channel that operates as an F_Port in accordance with FC-SW-5 and is dynamically instantiated together with its FCoE_LEP 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 Fibre Channel Switching element. A VF_Port is uniquely identified by an F_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-2 and FC-FS-3, with the exception that 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. When receiving FC frames from one of its FCoE_LEPs, a VF_Port shall verify that the S_ID of the received FC frame is equal to the address identifier of the VN_Port associated to that FCoE_LEP. If the check fails the FC frame shall be discarded.

NOTE 15 – The receive checks performed by the VF_Port/FCoE_LEP pair verify that the correct MAC destination address, MAC source address, and D_ID are present in a received FCoE frame (see D.5).

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-5.

7.5 FCoE Virtual Links

Figure 31 shows how the model defined in 7.4 models VE_Port to VE_Port Virtual Links.

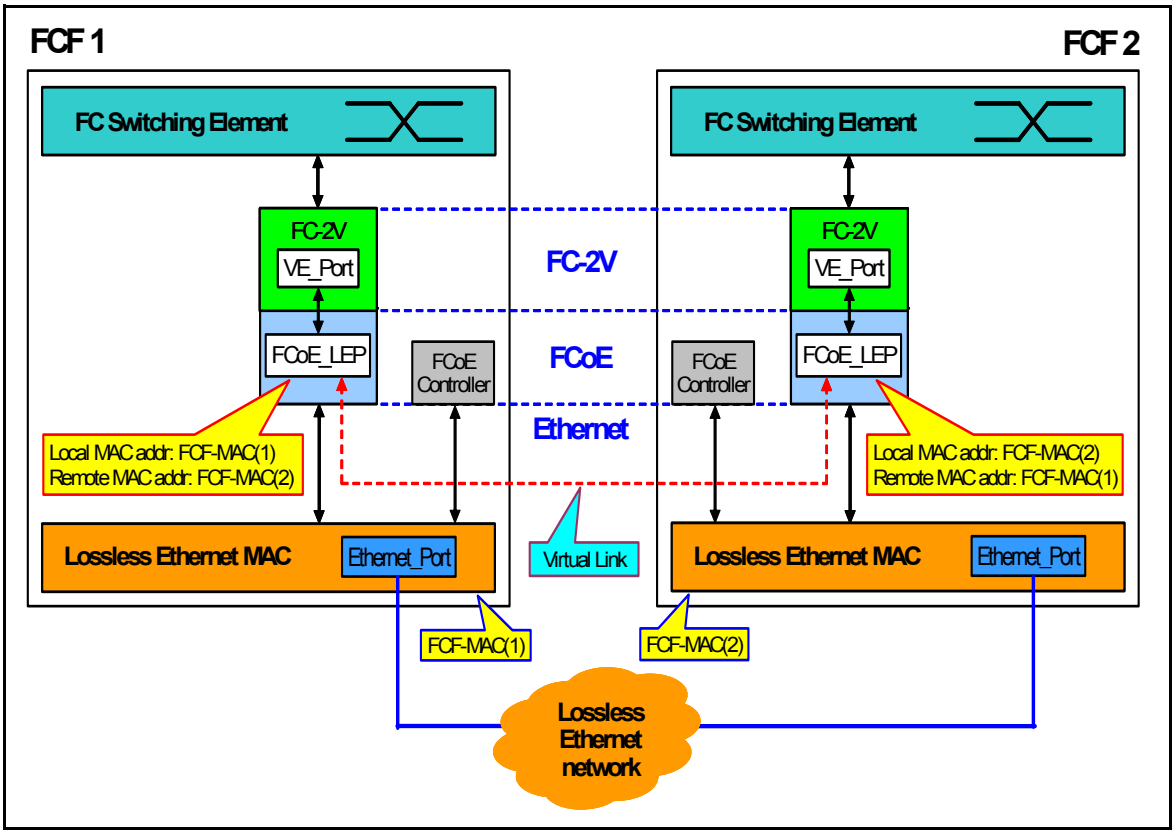

Figure 31 – VE_Port to VE_Port Virtual Links example

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 31 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)).

Figure 32 shows how the models defined in 7.3 and 7.4 model VN_Port to VF_Port Virtual Links.

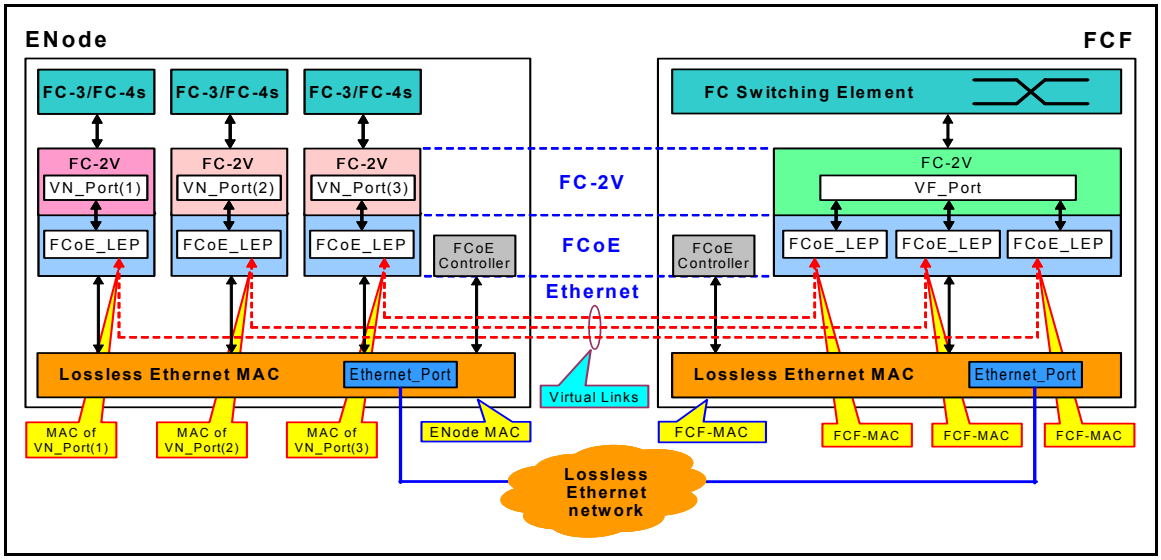

Figure 32 – VN_Port to VF_Port Virtual Links example

On successful completion of a FIP FLOGI Exchange, the FCoE Controller for an ENode MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(1) in figure 32) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates a VF_Port/FCoE_LEP pair.

On successful completion of a FIP NPIV FDISC Exchange, the FCoE Controller for an ENode MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(2) in figure 32) 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 NPIV FDISC Exchange, the FCoE Controller of an ENode MAC instantiates a VN_Port/FCoE_LEP pair (VN_Port(3) in figure 32) and the FCoE Controller of a VF_Port capable FCF-MAC instantiates an additional FCoE_LEP to the instantiated VF_Port.

Figure 32 shows the Virtual Links end-points, that are the MAC addresses used by the VN_Ports (i.e., MAC of VN_Port(1), MAC of VN_Port(2), and MAC of VN_Port(3)), and the FCF-MAC address.

When SPMAs are used, multiple VN_Ports associated with an ENode MAC may use the same local MAC address to establish Virtual Links to the same VF_Port capable FCF-MAC. In this case, the Fibre Channel addressing information is needed to identify a specific VN_Port to VF_Port Virtual Link in addition to the pair of MAC addresses of the two link end-points. This case is modeled on an ENode MAC with multiple VN_Port/FCoE_LEP pairs, in which the FCoE_LEPs operate using the same local MAC address/remote MAC address pair. Incoming FCoE frames are delivered to the proper VN_Port/FCoE_LEP pair on the basis of the D_ID field. This case is modeled on an FCF-MAC with multiple FCoE_LEPs associated with a VF_Port, in which the FCoE_LEPs operate using the same local MAC address/remote MAC address pair.

7.6 VN_Port MAC addresses

ENodes may use Fabric Provided MAC Addresses (FPMAs) and/or Server Provided MAC Addresses (SPMAs) as VN_Port MAC addresses. The FIP protocol is used to negotiate between ENodes and FCFs which type of VN_Port MAC addresses are used (see 7.8.4.1).

FPMAs are assigned by FCFs while assigning an N_Port_ID to a VN_Port (see 7.8.4.1). A properly formed FPMA is one in which the 24 most significant bits equal the Fabric's FC-MAP value and the least significant 24 bits equal the N_Port_ID assigned to the VN_Port by the FCF. This guarantees that FPMAs are unique within the Fabric. The FC-MAP value is checked by the FIP discovery protocol (see 7.8.3) to ensure it is consistent across the Fabric. FPMAs should not be used for other protocols.

If the FC-MAP value is not administratively configured, then the FC-MAP value shall be set to DEFAULT_FC-MAP (see table 47). If the FC-MAP value is administratively configured, then the FC-MAP value should be in the range 0EFC00h to 0EFCFFh.

SPMAs are assigned by ENodes and validated by FCFs. SPMAs should be globally assigned, not locally generated (i.e., they should have the U/L bit set to zero, see IEEE 802.3-2008). SPMAs used for FCoE and FIP traffic should not be used for other protocols.

ENodes and FCFs shall implement support for FPMAs as VN_Port MAC addresses. ENodes and FCFs may implement support for SPMAs as VN_Port MAC addresses.

7.7 FCoE frame format

An FCoE frame is an Ethernet frame (see IEEE 802.3-2008) containing an FCoE PDU. FCoE frames shall be formatted in accordance with IEEE 802.3-2008 and the MAC Client Data field within the Ethernet frame shall contain an FCoE PDU (see table 21). The use of an IEEE 802.1Q tag header is optional and additional IEEE 802.1 defined tags may be present in an FCoE frame. See Annex B for examples of FCoE frames.

The format of an FCoE PDU is specified in table 21.

Table 21 – FCoE PDU format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = FCoE_TYPE																Version				Reserved											
1	Reserved																															
2	Reserved																															
3	Reserved																								SOF							
4	(MSB)																															
n+3	Encapsulated FC Frame (n words)																														(LSB)	
n+4	EOF								Reserved																							

The Type field in the Ethernet header shall be set to FCoE_TYPE (see table 47). The Type field in the Ethernet header is not part of the FCoE PDU.

The Version field shall be set to FCoE_FRAME_VER (see table 47).

The SOF field specifies the SOF delimiter for the encapsulated FC frame. The value of the SOF field shall be as specified in table 22.

Table 22 – FCoE SOF field

Value	Description	Reference
28h	SOFf	FC-FS-3
2Dh	SOFi2	FC-FS-3
35h	SOFn2	FC-FS-3
2Eh	SOFi3	FC-FS-3
36h	SOFn3	FC-FS-3

The Encapsulated FC Frame field shall contain:

a) FC Extended_Header(s) (see FC-FS-3), if any;
b) the FC Frame_Header (see FC-FS-3);
c) the FC Data_Field (see FC-FS-3); and
d) the FC CRC (see FC-FS-3).

The EOF field specifies the EOF delimiter for the encapsulated FC frame. The value of the EOF field shall be set as specified in table 23.

Table 23 – FCoE EOF field

Value	Description	Reference
41h	EOFn	FC-FS-3
42h	EOFt	FC-FS-3
49h	EOFni	FC-FS-3
50h	EOFa	FC-FS-3

7.8 FC-BB_E device initialization

7.8.1 FCoE Initialization Protocol (FIP) overview

The FCoE Initialization Protocol (FIP) is used to perform the functions of FC-BB_E device discovery, initialization, and maintenance. To perform these functions, encapsulated FIP operations (see 7.8.6.2) are specified.

The FIP Ethernet Type (see 7.8.6.1) is different than the FCoE Ethernet Type (see 7.7) to enable the distinction of discovery, initialization, and maintenance traffic from other FCoE traffic.

FIP frames are used to perform the following protocols:

a) FIP VLAN discovery (see 7.8.2);
b) FIP discovery (see 7.8.3);
c) FCoE Virtual Link instantiation (see 7.8.4); and
d) FCoE Virtual Link maintenance (see 7.8.5).

All FIP protocols are performed on a per-VLAN basis. It is recommended to use the FIP VLAN Discovery protocol on the default VLAN (see IEEE 802.1Q-2005). All other FIP protocols shall be performed in each VLAN that provides FC-BB_E services.

In order to provide FC-BB_E services on a VLAN, FCoE and FIP protocols, other than FIP VLAN discovery, shall both be performed on that VLAN. Support for multiple Fabrics per VLAN is outside the scope of this standard.

NOTE 16 – The security provisions of this standard are not sufficient to ensure that FCoE frames remain associated with the correct Fabric if multiple Fabrics are used on the same VLAN.

On ENodes, the ENode MAC address shall be used for all FIP frames, except the VN_Port FIP Keep Alive frame (see 7.8.7.5). On FCFs, the FCF-MAC address shall be used for all FIP frames.

ENode MACs shall listen to the All-ENode-MACs group address, FCF-MACs shall listen to the All-FCF-MACs group address, and both ENode MACs and FCF-MACs shall listen to the All-FCoE-MACs group address.

An ENode MAC shall discard a FIP message destined to an address other than its ENode MAC address or the All-ENode-MACs address.

7.8.2 FIP VLAN discovery protocol

When becoming operational, an ENode MAC or an FCF-MAC may invoke the FIP VLAN discovery protocol to discover the VLANs in the Lossless Ethernet network that provide FC-BB_E services. The FIP VLAN discovery protocol is not needed if these VLANs are already known or if VLANs are not used.

An ENode MAC may send a FIP VLAN Request frame to the All-FCF-MACs MAC address over an available VLAN (e.g., the port VLAN). VF_Port capable FCF-MACs that receive a FIP VLAN Request frame shall respond with a unicast FIP VLAN Notification frame over the same VLAN. The FIP VLAN Notification frame should provide the list of VLAN IDs over which the originating FCF offers FC-BB_E services. The ENode MAC that received a FIP VLAN Notification frame may enable one or more of these VLANs for subsequent operations. VF_Port capable FCF-MACs may limit the number of VLAN IDs listed in a FIP VLAN Notification frame on a per-requester basis.

A VF_Port capable FCF-MAC shall discard a multicast FIP VLAN Request frame that has a source address equal to its FCF-MAC address, and the FIP VLAN Request frame should be reported in a vendor specific way as an indication of a MAC address duplication.

If the configuration of VLANs on which a VF_Port capable FCF-MAC supports FC-BB_E services changes, that FCF-MAC should send a unicast FIP VLAN Notification frame to each ENode MAC address with which that FCF-MAC has established VN_Port to VF_Port Virtual Links. The unicast FIP VLAN Notification frame shall carry the revised list of VLAN IDs over which the originating VF_Port capable FCF-MAC offers FC-BB_E services.

A VE_Port capable FCF-MAC may send a FIP VLAN Request frame to the MAC address All- FCF-MACs over an available VLAN (e.g., the default VLAN). VE_Port capable FCF-MACs that receive a FIP VLAN Request frame shall respond with a unicast FIP VLAN Notification frame over the same VLAN. The FIP VLAN Notification frame carries the list of VLAN IDs over which the originating FCF offers FC-BB_E services. The VE_Port capable FCF-MAC that received a FIP VLAN Notification frame may enable one or more of these VLANs for subsequent operations.

A VE_Port capable FCF-MAC shall discard a multicast VLAN Request frame that has a source address equal to its FCF-MAC address. Such a VLAN Request frame should be reported in a vendor specific way as an indication of a MAC address duplication.

FCF-MACs shall listen to the All-FCF-MACs group address in the default VLAN and in other VLANs that ENodes or FCFs may use to invoke this protocol.

If the configuration of VLANs on which a VE_Port capable FCF-MAC supports FC-BB_E services changes, that FCF-MAC should send a unicast FIP VLAN Notification frame to each FCF-MAC address with which that FCF-MAC has established VE_Port to VE_Port Virtual Links. The unicast FIP VLAN Notification frame shall specify the revised list of VLAN IDs over which the originating VE_Port capable FCF-MAC offers FC-BB_E services.

7.8.3 FIP discovery protocol

7.8.3.1 Overview

On a network deploying multiple VLANs, the FIP discovery protocol is performed in the VLANs where FC-BB_E services are offered when these VLANs are known (e.g., upon performing the FIP VLAN discovery protocol (see 7.8.2)).

7.8.3.2 ENode/FCF discovery

The FCoE Controller of a VF_Port capable FCF-MAC shall periodically transmit multicast Discovery Advertisements (see 7.8.7.3) to the All-ENode-MACs group address every FKA_ADV_PERIOD. The FKA_ADV_PERIOD period shall be randomized by adding a random delay uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network. The FCoE Controller of a VF_Port capable FCF-MAC should begin transmitting unsolicited multicast Discovery Advertisements on completion of Fabric configuration (see FC-SW-5).

On receiving Discovery Advertisements, the FCoE Controller of an ENode MAC shall verify the VN_Port addressing capabilities of the advertising FCF-MAC (i.e., the values of the FP and SP flags, see table 27) against its VN_Port addressing capabilities. The FCoE Controller of an ENode MAC shall discard incompatible Discovery Advertisements and shall create an entry for each compatible FCF-MAC in an internal FCF list.

NOTE 17 – The internal data structures used to describe this protocol are a model to express the behavior, not an implementation requirement.

Each entry in the FCF list has the following flags:

a) ‘Max FCoE Size Verified’ - set to zero for entries created from unsolicited multicast Discovery Advertisements, set to one when a solicited unicast Discovery Advertisement is received; and
b) ‘Available for Login’ - reflects the value of the A bit provided by the most recently received Discovery Advertisement from that VF_Port capable FCF-MAC.

The FCoE Controller of an ENode MAC selects for login a subset of the FCF-MACs in the FCF list having the ‘Available for Login’ flag set to one (i.e., the FCF Login Set) on the basis of a local policy that should default to selecting the one(s) with higher priority (i.e., lower priority value) in the absence of explicit configuration of other selection criteria. A FIP FLOGI may be performed with an FCF-MAC in the FCF Login Set only if its ‘Max FCoE Size Verified’ flag is set to one. In order to perform a FIP FLOGI with an FCF-MAC in the FCF Login Set with the ‘Max FCoE Size Verified’ flag set to zero, the FCoE Controller of an ENode MAC shall transmit a unicast Discovery Solicitation (see 7.8.7.2) to that FCF-MAC address and receive a solicited unicast Discovery Advertisement in response.

The periodic reception of unsolicited multicast Discovery Advertisements allows the FCoE Controller of ENode MACs to continuously verify FCF-MAC connectivity. The Available for Login (A) bit in received Discovery Advertisements provides the information that the transmitting FCF-MAC is available for FIP FLOGI/FDISC, and this information is updated in the FCF list and FCF Login Set on reception of Discovery Advertisements. The A bit is informational and shall have no effect on existing logins.

When the FCoE Controller of an ENode MAC becomes operational it should discover VF_Port capable FCF-MACs with which it may perform FIP FLOGI by transmitting a multicast Discovery Solicitation to the All-FCF-MACs group address. In response to a Discovery Solicitation from an ENode MAC, a VF_Port capable FCF-MAC shall transmit a solicited unicast Discovery Advertisement to the soliciting ENode MAC if its VN_Port addressing modes are compatible with the modes of the ENode MAC (see table 27) and if it is configured to allow a FIP FLOGI from that ENode. The solicited unicast Discovery Advertisement shall be transmitted to the MAC address specified in the MAC address descriptor in the received Discovery Solicitation. The solicited unicast Discovery Advertisement shall be transmitted within ADV_TOV (see table 47) upon reception of the Discovery Solicitation. Discovery Advertisements transmitted in response to a multicast Discovery Solicitation should be delayed by a random time uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network. This delay should not be applied to solicited unicast Discovery Advertisements sent in response to unicast Discovery Solicitations. Solicited unicast Discovery Advertisements should not be transmitted until Fabric configuration (see FC-SW-5) is completed.

NOTE 18 – An ENode MAC may also wait to receive unsolicited multicast Discovery Advertisements and then send unicast Discovery Solicitations to the FCF-MACs selected for login from the FCF Login set.

A Discovery Solicitation shall carry in the Max FCoE Size descriptor the maximum FCoE PDU size the ENode MAC intends to use for FCoE traffic (see 7.8.6.3.7). The FIP PDU (see table 24) in a solicited unicast Discovery Advertisement shall be extended to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.8.7.3).

NOTE 19 – If an ENode transmits an FCoE frame with an FCoE PDU size that is greater than its maximum FCoE PDU size, the network may not deliver the FCoE frame.

An ENode MAC may generate multiple Discovery Solicitations.

NOTE 20 – As an example, an ENode MAC that does not receive a solicited unicast Discovery Advertisement in response to a Discovery Solicitation may transmit additional Discovery Solicitations, unicast or multicast.

Reception of a solicited unicast Discovery Advertisement from an FCF-MAC shall set the ‘Max FCoE Size Verified’ flag to one in the entry for that FCF-MAC in the FCF Login Set of an ENode MAC.

It is possible for an FCF to receive a multicast Discovery Solicitation from the same ENode MAC on multiple FCF-MACs. In this case, a separate solicited unicast Discovery Advertisement shall be transmitted by each of the FCF-MACs that received the Discovery Solicitation. The ENode MAC that transmitted the multicast Discovery Solicitation is able to determine that it received multiple solicited unicast Discovery Advertisements from the same FCF since the value of the Name_Identifer field in the Name_Identifier descriptor is the same in each of the solicited unicast Discovery Advertisements (see 7.8.7.3). In this case the ENode MAC should select the FCF-MAC for Fabric login with that FCF based on the value of the Priority descriptor in the Discovery Advertisements.

It is possible for an ENode MAC to receive multiple unsolicited multicast Discovery Advertisements from multiple FCF-MACs of the same FCF. The ENode MAC is able to determine that those

unsolicited multicast Discovery Advertisements are from the same FCF since the value of the Name_Identifer field in the Name_Identifier descriptor is the same in each of the unsolicited multicast Discovery Advertisements (see 7.8.7.3). In this case the ENode MAC should select the FCF-MAC for Fabric login with that FCF based on the value of the Priority descriptor in the Discovery Advertisements.

An ENode MAC shall discard any received Discovery Solicitation. A VF_Port capable FCF-MAC shall discard any Discovery Solicitation originated by a VE_Port capable FCF-MAC (i.e., having the F bit set to one (see 7.8.6.2)).

An ENode MAC shall discard an unsolicited multicast Discovery Advertisement that has a source address equal to its ENode MAC address. Such a Discovery Advertisement should be reported in a vendor specific way as an indication of a MAC address duplication.

A VF_Port capable FCF-MAC shall discard a multicast Discovery Solicitation that has a source address equal to its FCF-MAC address. Such a Discovery Solicitation should be reported in a vendor specific way as an indication of a MAC address duplication.

Reception of Discovery Advertisements for more than one Fabric on the same VLAN should be reported by an ENode MAC in a vendor specific manner and no subsequent VN_Port to VF_Port Virtual Links should be instantiated.

7.8.3.3 FCF/FCF discovery

The FCoE Controller of a VE_Port capable FCF-MAC shall periodically transmit multicast Discovery Advertisements (see 7.8.7.3) to the All-FCF-MACs group address every FKA_ADV_PERIOD. The FKA_ADV_PERIOD period shall be randomized by adding a random delay uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network.

On receiving Discovery Advertisements, the FCoE Controller of a VE_Port capable FCF-MAC shall create an entry per FCF-MAC in an internal FCF list.

NOTE 21 – The internal data structures used to describe this protocol are a model to express the behavior, not an implementation requirement.

Each entry in the FCF list has the following flags:

a) ‘Max FCoE Size Verified’ - set to zero for entries created from unsolicited multicast Discovery Advertisements, set to one when a solicited unicast Discovery Advertisement is received; and
b) ‘Available for ELP’ - reflects the value of the A bit provided by the most recently received Discovery Advertisement from that VE_Port capable FCF-MAC.

A FIP ELP may be performed with an FCF-MAC in the FCF list only if its ‘Max FCoE Size Verified’ flag is set to one. In order to perform a FIP ELP with an FCF-MAC in the FCF list with the ‘Max FCoE Size Verified’ flag set to zero, the FCoE Controller of a VE_Port capable FCF-MAC shall transmit a unicast Discovery Solicitation (see 7.8.7.2) to that FCF-MAC address and receive a solicited unicast Discovery Advertisement in response.

The periodic reception of unsolicited multicast Discovery Advertisements allow the FCoE Controller of VE_Port capable FCF-MACs to continuously verify the FCF-MACs connectivity. The ‘Available for Login’ (A) bit in received Discovery Advertisements provides the information that the transmitting FCF-MAC is available for FIP ELP, and this information is updated in the FCF list on reception of Advertisements. The A bit is informational and shall have no effect on existing VE_Port to VE_Port Virtual Links.

When the FCoE Controller for a VE_Port capable FCF-MAC becomes operational it should discover other VE_Port capable FCF-MACs by transmitting a multicast Discovery Solicitation to the All-FCF-MACs group address. In response to a Discovery Solicitation from an FCF-MAC, a VE_Port capable FCF-MAC shall transmit a solicited unicast Discovery Advertisement to the soliciting FCF-MAC if the FC-MAP value in the Discovery Solicitation is compatible with the FC-MAP configured on the FCF and if it is configured to allow a Virtual Link with that FCF.

The solicited unicast Discovery Advertisement shall be transmitted to the MAC address specified in the MAC address descriptor in the received Discovery Solicitation. The solicited unicast Discovery Advertisement shall be transmitted within ADV_TOV (see table 47) upon reception of the Discovery Solicitation. Discovery Advertisements transmitted in response to a multicast Discovery Solicitation should be delayed by a random time uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network. This delay should not be applied to solicited unicast Discovery Advertisements sent in response to unicast Discovery Solicitations.

NOTE 22 – A VE_Port capable FCF-MAC may also wait to receive unsolicited multicast Discovery Advertisements and then send unicast Discovery Solicitations to the FCF-MACs in the FCF list.

A Discovery Solicitation shall specify in the Max FCoE Size descriptor (see 7.8.6.3.7) the maximum FCoE PDU size the VE_Port capable FCF-MAC intends to use for FCoE traffic. The FIP PDU (see table 24) in a solicited unicast Discovery Advertisement shall be extended to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.8.7.3).

NOTE 23 – If a VE_Port capable FCF-MAC transmits an FCoE frame with an FCoE PDU size that is greater than its maximum FcoE PDU size, the network may not deliver the FCoE frame.

A VE_Port capable FCF-MAC may generate multiple Discovery Solicitations.

NOTE 24 – As an example, a VE_Port capable FCF-MAC that does not receive a solicited unicast Discovery Advertisement in response to a Discovery Solicitation may transmit additional Discovery Solicitations.

Reception of a solicited unicast Discovery Advertisement from an FCF-MAC shall set the ‘Max FCoE Size Verified’ flag to one in the entry for that FCF-MAC in the FCF list of the receiving VE_Port capable FCF-MAC.

It is possible for an FCF to receive multicast Discovery Solicitations from the same VE_Port capable FCF-MAC on multiple FCF-MACs. In this case, a separate solicited unicast Discovery Advertisement shall be transmitted by each of the FCF-MACs that received the Discovery Solicitation. The VE_Port capable FCF-MAC that transmitted the multicast Discovery Solicitation is able to determine that it received multiple solicited unicast Discovery Advertisements from the same FCF since the value of the Name_Identifer field in the Name_Identifier descriptor is the same in each of the solicited unicast Discovery Advertisements (see 7.8.7.3).

It is possible for a VE_Port capable FCF-MAC to receive multiple unsolicited multicast Discovery Advertisements from multiple FCF-MACs of the same FCF. The VE_Port capable FCF-MAC is able to determine that those unsolicited multicast Discovery Advertisements are from the same FCF since the value of the Name_Identifer field in the Name_Identifier descriptor is the same in each of the unsolicited multicast Discovery Advertisements (see 7.8.7.3).

After receiving a Discovery Solicitation originated by an FCF (i.e., the F bit is set to one), an FCF-MAC shall perform the following verification checks:

a) the Name_Identifier field value in the Discovery Solicitation is different than the Switch_Name of the recipient FCF (see 7.8.7.2.2); and
b) either:
   A) the FP bit is set to one (see 7.8.6.2) and the FC-MAP value in the FC-MAP descriptor in the Discovery Solicitation is the same as the FC-MAP value of the recipient FCF; or
   B) the FP bit is set to zero, the SP bit is set to one (see 7.8.6.2), and the FC-MAP value in the FC-MAP descriptor in the Discovery Solicitation is zero.

If any verification check is false, then the Discovery Solicitation shall be discarded.

After receiving a Discovery Advertisement, an FCF-MAC shall perform the following verification checks:

a) the Name_Identifier field value in the Discovery Advertisement is different than the Switch_Name of the recipient FCF (see 7.8.7.3); and
b) either:
   A) the FP bit is set to one (see 7.8.6.2) and the FC-MAP value in the Fabric descriptor in the Discovery Advertisement is the same as the FC-MAP value of the recipient FCF; or
   B) the FP bit is set to zero, the SP bit is set to one (see 7.8.6.2), and the FC-MAP value in the Fabric descriptor in the Discovery Advertisement is zero.

If any verification check is false, then the Discovery Advertisement shall be discarded.

NOTE 25 – It is possible for an FCF to receive a multicast Discovery Solicitation or a multicast Discovery Advertisement that it originated because FIP frames sent to the All-FCF-MACs group address may be forwarded to other ports on the same FCF by intermediate Ethernet bridges.

A VE_Port capable FCF-MAC shall discard any Discovery Solicitation originated by an ENode (i.e., having the F bit set to zero (see 7.8.6.2)).

A VE_Port capable FCF-MAC shall discard a multicast Discovery Solicitation that has a source address equal to its FCF-MAC address. Such a Discovery Solicitation should be reported in a vendor specific way as an indication of a MAC address duplication.

Reception of Discovery Advertisements for more than one Fabric on the same VLAN should be reported by VE_Port capable FCF-MAC in a vendor specific manner and no subsequent VE_Port to VE_Port Virtual Links should be instantiated.

7.8.4 FCoE Virtual Link instantiation protocol

7.8.4.1 VN_Port to VF_Port Virtual Links

The FCoE Controller of an ENode MAC instantiates VN_Port to VF_Port Virtual Links on successful completion of a FIP Fabric login request. Fabric login (i.e., FLOGI, NPIV FDISC) shall be performed using FIP frames (see table 24) and the associated FIP descriptor type (see table 29). Fabric login (i.e., FLOGI, NPIV FDISC) shall not be performed using FCoE frames .

In addition to providing Fabric login, the FIP Fabric login provides a method to assign a MAC address for the VN_Port (see 7.8.7.4.2).

When the FCoE Controller of an ENode MAC transmits a FIP FLOGI Request or FIP NPIV FDISC Request it shall indicate the addressing mode it intends to use (i.e., FPMA, SPMA, or both (see table 27)). The MAC address returned by the FCF in a FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC shall be used as the VN_Port MAC address (see 7.6).

If the SP bit is set to one in a FIP FLOGI Request or FIP NPIV FDISC Request (see table 27) and the FCF selects to use SPMA, the FCF shall return the MAC address specified in the FIP FLOGI Request or FIP NPIV FDISC Request in the FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC, if that MAC address is valid (see 7.8.6.3.3).

If the FP bit is set to one in a FIP FLOGI Request or FIP NPIV FDISC Request (see table 27) and the FCF selects to use FPMA, the FCF shall return a properly formed FPMA MAC address in the FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC (see 7.6).

If both the FCF and ENode support both SPMA and FPMA, the FCF shall select a type and shall return a MAC address for the selected type.

Explicit VN_Port to VF_Port Virtual Link de-instantation is performed by an ENode MAC by performing Fabric logout. Fabric logout (i.e., Fabric LOGO) shall be performed by an ENode using FIP frames (see table 24) and the associated FIP descriptor type (see table 29). Fabric logout shall not be performed using FCoE frames.

In addition to providing Fabric logout, the FIP Fabric logout provides a method to de-assign a MAC address for the VN_Port (see 7.8.7.4.3).

7.8.4.2 VE_Port to VE_Port Virtual Links

The FCoE Controller of a VE_Port capable FCF-MAC instantiates VE_Port to VE_Port Virtual Links on successful completion of a FIP ELP request. ELP shall be performed using FIP frames (see table 24) and the associated FIP descriptor type (see table 29). ELP shall not be performed using FCoE frames.

In addition to providing ELP, the FIP ELP provides a method to communicate the MAC address for the VE_Port (see 7.8.7.4.4).

7.8.5 FCoE Virtual Link maintenance protocol

7.8.5.1 Virtual Link maintenance protocol overview

VN_Port to VF_Port Virtual Links (see figure 27) and VE_Port to VE_Port Virtual Links (see figure 28) overlay over a Lossless Ethernet network. The Virtual Link maintenance protocol specifies how to deal with faults that may occur in a Lossless Ethernet network.

7.8.5.2 VN_Port to VF_Port Virtual Link maintenance protocol

To deal with local physical layer faults, an ENode MAC shall de-instantiate all its VN_Ports to VF_Port Virtual Links upon detecting that its physical layer is not operational. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports. A VF_Port capable FCF-MAC shall de-instantiate all its VF_Ports upon detecting that its physical layer is not operational.

To deal with non-local faults, the FCoE Controllers of an ENode MAC and of a VF_Port capable FCF-MAC shall continuously verify the state of the VN_Port to VF_Port Virtual Link by transmitting appropriate FIP frames and by verifying received FIP frames. This behavior may be disabled by VF_Port capable FCF-MACs under administrative control by setting to one the D bit in the FKA_ADV_Period descriptor in Discovery Advertisements (see 7.8.6.3.13). The D bit in the FKA_ADV_Period descriptor may be set to one only in a direct-attach topology (i.e., when an ENode is directly connected to an FCF without any intermediate Ethernet bridges).

The FCoE Controller of an ENode MAC shall transmit a unicast FIP Keep Alive frame on behalf of the ENode MAC (i.e., with the ENode MAC address as source MAC address and without a Vx_Port Identification descriptor in the FIP Descriptor list (see table 45)) to each VF_Port capable FCF-MAC with which it has VN_Ports logged in. This ENode FIP Keep Alive frame shall be transmitted every FKA_ADV_PERIOD. The FKA_ADV_PERIOD is obtained from the Discovery Advertisements received from the VF_Port capable FCF-MACs with which the ENode MAC has VN_Ports logged in.

In addition, the FCoE Controller of an ENode MAC shall transmit a unicast FIP Keep Alive frame on behalf of each VN_Port (i.e., with the VN_Port MAC address as source MAC address and containing a Vx_Port Identification descriptor for that VN_Port in the FIP Descriptor list, see table 45) to the VF_Port capable FCF-MAC with which that VN_Port is logged in. This VN_Port FIP Keep Alive frame shall be transmitted every FKA_VN_PERIOD.

The FCoE Controller of an ENode MAC shall monitor the status of a VF_Port with which it has VN_Ports logged in by verifying reception of unsolicited multicast Discovery Advertisements from that VF_Port capable FCF-MAC. Unsolicited multicast Discovery Advertisements are expected to be received every FKA_ADV_PERIOD. If unsolicited multicast Discovery Advertisements are not received within 2.5 * FKA_ADV_PERIOD, all the VN_Port to VF_Port Virtual Links with that VF_Port shall be implicitly de-instantiated. This condition should be counted as a Virtual Link failure and shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports. That FCF-MAC shall be removed from the FCF Login Set (see 7.8.3.2). A subsequent FIP Fabric Login may be performed with an FCF-MAC in the current FCF Login Set as specified in see 7.8.3.2.

The FCoE Controller of a VF_Port capable FCF-MAC shall transmit an unsolicited multicast Discovery Advertisement to the All-ENode-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller of a VF_Port capable FCF-MAC shall monitor the status of an ENode MAC with which it has active VN_Port to VF_Port Virtual Links by verifying the reception of FIP Keep Alive frames from that ENode MAC and its VN_Ports. VN_Port FIP Keep Alive frames (i.e., those containing a Vx_Port Identification descriptor) are expected to be received every FKA_VN_PERIOD and ENode FIP Keep Alive frames (i.e., those not containing a Vx_Port Identification descriptor) are expected to be received every FKA_ADV_PERIOD.

If VN_Port FIP Keep Alive frames are not received within 2.5 * FKA_VN_PERIOD, the associated VN_Port to VF_Port Virtual Link shall be explicitly de-instantiated (i.e., a FIP Clear Virtual Links frame listing the unreachable VN_Port shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Port. If ENode FIP Keep Alive frames are not received within 2.5 * FKA_ADV_PERIOD, all associated VN_Port to VF_Port Virtual Links shall be explicitly de-instantiated (i.e., a FIP Clear Virtual Links frame listing all the unreachable VN_Ports shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

NOTE 26 – The use of a faster FKA_ADV_PERIOD rate for ENode FIP Keep Alive frames allows fast response times against loss of connectivity in the Ethernet realm with limited overhead. The use of VN_Port FIP Keep Alive frames transmitted at a slower FKA_VN_PERIOD rate allows the clearing of state associated with each individual VN_Port when that VN_Port becomes not operational.

Explicit VN_Port to VF_Port Virtual Link de-instantation is invoked by a VF_Port capable FCF-MAC by transmitting a FIP Clear Virtual Links frame (see 7.8.7.6). A FIP Clear Virtual Links frame transmitted to an ENode MAC with logged in VN_Ports provides the list of VN_Ports to be removed. An ENode MAC shall de-instantiate the VN_Ports listed in a FIP Clear Virtual Link frame upon reception of the FIP frame. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

The size of a FIP Clear Virtual Links frame shall not exceed the standard Ethernet MAC Client Data size (i.e., 1 500 bytes for basic frames and 1 504 bytes for Q-tagged frames, see IEEE 802.3-2008). If the list of VN_Ports to be removed does not fit in a single FIP frame, multiple FIP frames should be transmitted to convey the entire list.

On receiving a VN_Port FIP Keep Alive frame coming from a VN_Port that is not logged, the FCoE Controller of a VF_Port capable FCF-MAC shall transmit a FIP Clear Virtual Links frame listing that VN_Port.

On receiving an ENode FIP Keep Alive frame coming from an ENode MAC that is not logged in, the FCoE Controller of a VF_Port capable FCF-MAC shall transmit a FIP Clear Virtual Links frame listing no VN_Ports. A FIP Clear Virtual Links frame listing no VN_Ports shall be handled by an ENode MAC by de-instantiating all VN_Port to VF_Port Virtual Links with that VF_Port capable FCF-MAC. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

FIP Clear Virtual Links frames may be generated by FCFs whenever appropriate to speed-up fault recovery.

NOTE 27 – As an example, in certain topologies an FCF may generate a FIP Clear Virtual Links frame to de-instantiate the VN_Port to VF_Port Virtual Links affected by a local physical layer fault on other ports upon detection of that fault.

The FKA_ADV_PERIOD value (see table 47) may be changed on a FCF under administrative control. When this happens, each VF_Port capable FCF-MAC of the FCF shall advertise the updated FKA_ADV_PERIOD in subsequent unsolicited multicast Discovery Advertisements.

When the FKA_ADV_PERIOD value is decreased, a VF_Port capable FCF-MAC shall transmit unsolicited multicast Discovery Advertisements at the interval specified by the updated value, but shall not use the updated value for detection of missing ENode FIP Keep Alives until five times the old value has elapsed since the transmission of the first updated unsolicited multicast Discovery Advertisement.

When the FKA_ADV_PERIOD value is increased, a VF_Port capable FCF-MAC shall transmit unsolicited multicast Discovery Advertisements at the interval specified by the old value until five times the updated value has elapsed since the transmission of the first updated unsolicited multicast Discovery Advertisement, but shall use the updated value for detection of missing ENode FIP Keep Alives.

On detecting the updated value, an ENode having VN_Port to VF_Port Virtual Links instantiated with that FCF shall transmit ENode FIP Keep Alive frames at the interval specified by the updated FKA_ADV_PERIOD value and shall use the updated value for detection of missing unsolicited multicast Discovery Advertisements.

7.8.5.3 VE_Port to VE_Port Virtual Link maintenance protocol

To deal with local physical layer faults, a VE_Port capable FCF-MAC shall de-instantiate all its VE_Port to VE_Port Virtual Links upon detecting that its physical layer is not operational.

To deal with non-local faults, the FCoE Controllers for VE_Port capable FCF-MACs shall continuously verify the state of a VE_Port to VE_Port Virtual Link by transmitting unsolicited multicast Discovery Advertisements and by verifying received unsolicited multicast Discovery Advertisements.

The FCoE Controller for a VE_Port capable FCF-MAC shall transmit a Discovery Advertisement to the All-FCF-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller for a VE_Port capable FCF-MAC shall monitor the status of a VE_Port to VE_Port Virtual Link by verifying the reception of unsolicited multicast Discovery Advertisements. Unsolicited multicast Discovery Advertisements are expected to be received every FKA_ADV_PERIOD. If unsolicited multicast Discovery Advertisements are not received within 2.5 * FKA_ADV_PERIOD, the VE_Port to VE_Port Virtual Link associated with that FCF-MAC shall be explicitly de-instantiated along with the associated VE_Port. That FCF-MAC shall be removed from the FCF List (see 7.8.3.3).

Explicit VE_Port to VE_Port Virtual Link de-instantation is performed by transmitting a FIP Clear Virtual Links frame from a VE_Port capable FCF-MAC. A FIP Clear Virtual Links frame transmitted to a VE_Port capable FCF-MAC provides the destination FCF-MAC address in the Vx_Port Identification descriptor. A VE_Port capable FCF-MAC shall de-instantiate the VE_Port listed in a FIP Clear Virtual Link frame upon reception of the FIP frame.

The FKA_ADV_PERIOD value (see table 47) may be changed on a FCF under administrative control. When this happens, each VE_Port capable FCF-MAC of the FCF shall advertise the updated FKA_ADV_PERIOD in subsequent unsolicited multicast Discovery Advertisements and shall transmit them at the interval specified by the updated value.

On detecting the updated value in received unsolicited multicast Discovery Advertisements, a VE_Port capable FCF-MAC shall use the updated value for detection of missing unsolicited multicast Discovery Advertisements.

7.8.6 FIP frames

7.8.6.1 FIP frame format

A FIP frame is an Ethernet frame (see IEEE 802.3-2008) containing a FIP PDU. FIP frames shall be formatted in accordance with IEEE 802.3-2008 and the MAC Client Data field within the Ethernet frame shall contain a FIP PDU (see table 24). The use of an 802.1Q tag header is optional and additional IEEE 802.1 defined tags may be present in a FIP frame. See Annex B for examples of FIP frames.

The format of a FIP PDU is specified in table 24.

Table 24 – FIP PDU format

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = FIP_TYPE																Version				Reserved											
1	(MSB)																															
	Encapsulated FIP operation (n words)																															
n																																(LSB)

The Type field in the Ethernet header shall be set to FIP_TYPE (see table 47). The Type field is not part of the FIP PDU.

The Version field shall be set to FIP_FRAME_VER (see table 47).

The Encapsulated FIP operation field shall be as specified in 7.8.6.2.

A received FIP frame containing no encapsulated FIP operation shall be discarded.

7.8.6.2 Encapsulated FIP operation

The format of the encapsulated FIP operation is specified in table 25.

Table 25 – Encapsulated FIP operation format

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	FIP Protocol Code																Reserved								FIP Subcode							
1	FIP Descriptor List Length																FP	SP	Reserved												A	S	F
2	FIP Descriptor List (m words)																															
n-1	FIP_Pad (n-m-2 words)																															

The FIP Protocol Code and FIP Subcode field values and operations are specified in table 26.

Table 26 – FIP Protocol Code and FIP Subcode field values

FIP Protocol Code	FIP Subcode	FIP Operation	Reference
0001h	01h	Discovery Solicitation	7.8.7.2
	02h	Discovery Advertisement	7.8.7.3
0002h	01h	Virtual Link Instantiation Request	7.8.7.4
	02h	Virtual Link Instantiation Reply	7.8.7.4
0003h	01h	FIP Keep Alive	7.8.7.5
	02h	FIP Clear Virtual Links	7.8.7.6
0004h	01h	FIP VLAN Request	7.8.7.7
	02h	FIP VLAN Notification	7.8.7.8
FFF8h to FFFEh	00h to FFh	Vendor Specific	7.8.7.9
All others	All others	Reserved	

The FIP Descriptor List Length field shall be set to the length in words of all FIP descriptor(s) that follow. A receiver shall ignore any data beyond the length specified in this field.

The Fabric Provided (FP) bit and Server Provided (SP) bit settings are dependent on the FIP operation and shall be set as specified in table 27

Table 27 – FP bit and SP bit setting

Bit	FIP Protocol	FIP Operation (see table 45)	Setting
FP	Discovery	Discovery Solicitation[a]	Set to one if the originating device supports FPMA. Set to zero if the originating device does not support FPMA.
		Discovery Advertisement[a]	Set to one if the originating device supports FPMA. Set to zero if the originating device does not support FPMA.
	Virtual Link Instantiation	FIP FLOGI Request[b]	Set to one if FPMA is requested. Set to zero if FPMA is not requested.
		FIP NPIV FDISC Request[b]	Set to one if FPMA is requested. Set to zero if FPMA is not requested.
		FIP FLOGI LS_ACC[c]	Set to one if FPMA is granted. Set to zero if FPMA is not granted.
		FIP NPIV FDISC LS_ACC[c]	Set to one if FPMA is granted. Set to zero if FPMA is not granted.
	All others	All others	Reserved
SP	Discovery	Discovery Solicitation[a]	Set to one if the originating device supports SPMA. Set to zero if the originating device does not support SPMA.
		Discovery Advertisement[a]	Set to one if the originating device supports SPMA. Set to zero if the originating device does not support SPMA.
	Virtual Link Instantiation	FIP FLOGI Request[b]	Set to one if SPMA is requested. Set to zero if SPMA is not requested.
		FIP NPIV FDISC Request[b]	Set to one if SPMA is requested. Set to zero if SPMA is not requested.
		FIP FLOGI LS_ACC[c]	Set to one if SPMA is granted. Set to zero if SPMA is not granted.
		FIP NPIV FDISC LS_ACC[c]	Set to one if SPMA is granted. Set to zero if SPMA is not granted.
	All others	All others	Reserved

a Solicitation or Advertisement frames with the FP and SP bits both set to zero should not be transmitted and such frames shall be ignored on reception.
b Both the FP bit and SP bit may be set to one in a FIP FLOGI Request or FIP NPIV FDISC Request, but at least one of the bits shall be set to one.
c Only one of the FP bit and SP bit shall be set to one in a FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC. They shall not have the same value.

The Solicited (S) bit shall be set to one in solicited unicast Discovery Advertisements (i.e., Discovery Advertisements transmitted in response to a Discovery Solicitation). The S bit shall be set to zero in unsolicited multicast Discovery Advertisements (i.e., Discovery Advertisements not transmitted in response to a Discovery Solicitation). The S bit is reserved for all other FIP operations.

The FCF (F) bit shall be set to one in a Discovery Solicitation, Discovery Advertisement, or FIP VLAN Request if the originating device is an FCF. The F bit shall be set to zero in a Discovery Solicitation, Discovery Advertisement, or in a FIP VLAN Request if the originating device is not an FCF. The F bit is reserved for all other FIP operations.

The Available for Login (A) bit shall be set to one in a Discovery Advertisement if the originating FCF is available to process FIP FLOGI, FIP NPIV FDISC, or FIP ELP Requests (see 7.8.3). The A bit shall be set to zero in a Discovery Advertisement if the originating FCF is not available to process FIP FLOGI, FIP NPIV FDISC, or FIP ELP Requests. The A bit is reserved for all other FIP operations.

The FIP Descriptor List field shall contain one or more FIP descriptors (see 7.8.6.3).

The FIP_Pad field shall be used in solicited unicast Discovery Advertisements to extend the FIP PDU (see table 24) to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.8.7.3). The FIP_Pad field shall be of zero length (i.e., not present) for all other FIP operations.

Received FIP frames shall be checked for correct formatting before any FIP descriptor processing occurs. A malformed FIP frame shall be discarded and should be reported in a vendor specific way. The checks for correct formatting include:

a) the FIP Descriptor List Length value matches the sum of the descriptors' lengths in the FIP Descriptor List;
b) the FIP Protocol Code field and FIP Subcode field are valid (see table 26);
c) the critical descriptors (see 7.8.6.3.1) required by the FIP Protocol Code and FIP Subcode are present;
d) no critical descriptors other than the ones required by the FIP Protocol Code and FIP Subcode are present; and
e) descriptors use valid values for MAC addresses (see 7.8.6.3.3), FKA_ADV_PERIOD (see table 47) and VLAN IDs (see 7.8.6.3.15).

7.8.6.3 FIP descriptors

7.8.6.3.1 FIP descriptor overview

FIP descriptors are specified using a TLV format (i.e., Type, Length, Value). The length field value shall be specified as the number of words in the FIP descriptor including the TLV format. FIP descriptor type values are split into two ranges, critical and non-critical, as specified in table 28.

Table 28 – FIP descriptor type value ranges

Range	Value	Description
Critical	0 to 127	An FCoE Controller that receives a FIP frame with an unknown critical descriptor shall discard the FIP frame.
Non-critical	128 to 255	An FCoE Controller that receives a FIP frame with one or more unknown non-critical descriptors shall ignore the unknown descriptors and continue to process the FIP frame.

A descriptor with an invalid length value shall be considered invalid. If that descriptor is critical, the entire FIP operation shall be discarded. If that descriptor is not critical, the descriptor shall be ignored.

The FIP descriptor types are specified in table 29. See table 45 for how FIP descriptors are used.

Table 29 – FIP descriptor types

Range	Type	FIP Descriptor	Reference
Critical	0	Reserved	
	1	Priority	7.8.6.3.2
	2	MAC address	7.8.6.3.3
	3	FC-MAP	7.8.6.3.4
	4	Name_Identifier	7.8.6.3.5
	5	Fabric	7.8.6.3.6
	6	Max FCoE Size	7.8.6.3.7
	7	FLOGI[a]	7.8.6.3.8
	8	NPIV FDISC[a]	7.8.6.3.9
	9	LOGO[a]	7.8.6.3.10
	10	ELP[a]	7.8.6.3.11
	11	Vx_Port Identification	7.8.6.3.12
	12	FKA_ADV_Period	7.8.6.3.13
	13	Vendor_ID	7.8.6.3.14
	14	VLAN	7.8.6.3.15
	15 to 127	Reserved	
Non-critical	128 to 240	Reserved	
	241 to 254	Vendor Specific	7.8.6.3.16
	255	Reserved	
a The FC CRC, SOF, and EOF shall not be included in the FIP descriptor.			

7.8.6.3.2 FIP Priority descriptor

The FIP Priority descriptor is used in FIP operations as specified in table 45. An ENode may use the value provided in the Priority descriptor of received Discovery Advertisements to select the FCF-MAC to which to perform FIP FLOGI. The default value for the Priority field is DEFAULT_FIP_PRIORITY (see table 47). The highest priority value is 0 and the lowest priority value is 255 (i.e., lower numerical values indicate higher priorities).

The FIP Priority descriptor format shall be as specified in table 30.

Table 30 – FIP Priority descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 01h								Length = 01h								Reserved								Priority							

Priority: the Priority value associated with the sending FCF-MAC.

7.8.6.3.3 FIP MAC address descriptor

The FIP MAC address descriptor is used in FIP operations as specified in table 45.

The FIP MAC address descriptor format shall be as specified in table 31.

Table 31 – FIP MAC address descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 02h								Length = 02h								(MSB)															
1	MAC address																														(LSB)	

MAC address: the MAC address contained in the descriptor. Valid MAC addresses are unicast addresses.

7.8.6.3.4 FIP FC-MAP descriptor

The FIP FC-MAP descriptor is used in FIP operations as specified in table 45.

The FIP FC-MAP descriptor format shall be as specified in table 32.

Table 32 – FIP FC-MAP descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 03h								Length = 02h								Reserved															
1	Reserved								(MSB)								FC-MAP														(LSB)	

FC-MAP: the value to be used as the most significant 24 bits in FPMAs (see 7.6).

7.8.6.3.5 FIP Name_Identifier descriptor

The FIP Name_Identifier descriptor is used in FIP operations as specified in table 45.

The FIP Name_Identifier descriptor format shall be as specified in table 33.

Table 33 – FIP Name_Identifier descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 04h								Length = 03h								Reserved															
1	(MSB)								Name_Identifier																							
2																																(LSB)

Name_Identifier: the Name_Identifier (see FC-FS-3) contained in the descriptor.

7.8.6.3.6 FIP Fabric descriptor

The FIP Fabric descriptor is used in FIP operations as specified in table 45.

The FIP Fabric descriptor format shall be as specified in table 34.

Table 34 – FIP Fabric descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 05h								Length = 04h								VF_ID															
1	Reserved								(MSB)								FC-MAP															(LSB)
2	(MSB)								Fabric_Name																							
3																																(LSB)

VF_ID: the VF_ID (see see FC-FS-3) associated with the Fabric, if any.

FC-MAP: the value to be used as the most significant 24 bits in FPMAs (see 7.6).

Fabric_Name: the Fabric_Name (see FC-FS-3) identifying the Fabric.

7.8.6.3.7 FIP Max FCoE Size descriptor

The FIP Max FCoE Size descriptor is used in FIP operations as specified in table 45.

The FIP Max FCoE Size descriptor format shall be as specified in table 35.

Table 35 – FIP Max FCoE Size descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 06h								Length = 01h								Max_FCoE_Size															

Max_FCoE_Size: the size in bytes that the FIP PDU (see table 24) in a solicited unicast Discovery Advertisement is requested to be extended to.

7.8.6.3.8 FIP FLOGI descriptor

The FIP FLOGI descriptor is used in FIP operations as specified in table 45.

The FIP FLOGI descriptor format shall be as specified in table 36

Table 36 – FIP FLOGI descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 07h								Length								Reserved															
1	(MSB)								FLOGI Request or FLOGI LS_ACC/LS_RJT																							
n																															(LSB)	

Length: shall be set to 36 for a FLOGI Request and FLOGI LS_ACC, or to 9 for a FLOGI LS_RJT.

FLOGI Request or FLOGI LS_ACC/LS_RJT: an encapsulated FLOGI Request, FLOGI LS_ACC, or FLOGI LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field. In an FLOGI Request or FLOGI LS_ACC, the Payload bit shall be set to zero (see FC-LS-2).

7.8.6.3.9 FIP NPIV FDISC descriptor

The FIP NPIV FDISC descriptor is used in FIP operations as specified in table 45.

The FIP NPIV FDISC descriptor format shall be as specified in table 37

Table 37 – FIP NPIV FDISC descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 08h								Length								Reserved															
1	(MSB)								NPIV FDISC Request or NPIV FDISC LS_ACC/LS_RJT																							
n																															(LSB)	

Length: shall be set to 36 for an FDISC Request and FDISC LS_ACC, or to 9 for an FDISC LS_RJT.

NPIV FDISC Request or NPIV FDISC LS_ACC/LS_RJT: an encapsulated FDISC Request, FDISC LS_ACC, or FLOGI LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field. In an FLOGI Request or FLOGI LS_ACC, the Payload bit shall be set to zero (see FC-LS-2).

7.8.6.3.10 FIP LOGO descriptor

The FIP LOGO descriptor is used in FIP operations as specified in table 45.

The FIP LOGO descriptor format shall be as specified in table 38

Table 38 – FIP LOGO descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 09h								Length								Reserved															
1	(MSB)								LOGO Request or LOGO LS_ACC/LS_RJT																							
n																															(LSB)	

Length: shall be set to 11 for a LOGO Request, 8 for a LOGO LS_ACC, or to 9 for a LOGO LS_RJT.

LOGO Request or LOGO LS_ACC/LS_RJT: an encapsulated LOGO Request, LOGO LS_ACC, or LOGO LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field (see FC-LS-2).

7.8.6.3.11 FIP ELP descriptor

The FIP ELP descriptor is used in FIP operations as specified in table 45.

The FIP ELP descriptor format shall be as specified in table 39.

Table 39 – FIP ELP descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 0Ah								Length								Reserved															
1	(MSB)								ELP Request or ELP SW_ACC/SW_RJT																							
n																																(LSB)

Length: shall be set to 33 for an ELP Request and ELP SW_ACC, or to 9 for an ELP SW_RJT.

ELP Request or ELP SW_ACC/SW_RJT: an encapsulated ELP Request, ELP SW_ACC, or ELP SW_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an SW_ILS payload but without the CRC field (see FC-SW-5).

7.8.6.3.12 FIP Vx_Port Identification descriptor

The FIP Vx_Port Identification descriptor is used in FIP operations as specified in table 45.

The FIP Vx_Port Identification descriptor format shall be as specified in table 40.

Table 40 – FIP Vx_Port Identification descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 0Bh								Length = 05h								(MSB)															
1	MAC address																															(LSB)
2	Reserved								(MSB)								Address Identifier															(LSB)
3	(MSB)								Port_Name																							
4																																(LSB)

MAC address: the MAC address associated with the referred VN_Port or VE_Port. Valid MAC addresses are unicast addresses.

Address Identifier: the address identifier associated with the referred VN_Port or the value FFFFFDh for a VE_Port.

Port_Name: the N_Port_Name of the referred VN_Port or the E_Port_Name of the referred VE_Port.

7.8.6.3.13 FIP FKA_ADV_Period descriptor

The FIP FKA_ADV_Period descriptor is used in FIP operations as specified in table 45.

The FIP FKA_ADV_Period descriptor format shall be as specified in table 41.

Table 41 – FIP FKA_ADV_Period descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 0Ch								Length = 02h								Reserved															D
1	FKA_ADV_PERIOD																															

D: the value of the D bit shall be administratively configurable on FCFs. The D bit shall be set to zero unless administratively set to one. VE_Port capable FCF-MACs shall ignore the value of the D bit provided in received Discovery Advertisements. ENode MACs shall process the value of the D bit provided in received Discovery Advertisements.

When the D bit is set to zero, the receiving ENode MAC shall verify periodic reception of Discovery Advertisements and shall transmit periodic ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.8.5.2. The VF_Port capable FCF-MAC shall verify periodic reception of ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.8.5.2.

When the D bit is set to one, the receiving ENode MAC may verify periodic reception of Discovery Advertisements and should not transmit periodic ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.8.5.2. The VF_Port capable FCF-MAC shall not verify periodic reception of ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.8.5.2 and should discard possible received ENode FIP Keep Alive and VN_Port FIP Keep Alive frames.

FKA_ADV_PERIOD: the value of the advertised FKA_ADV_PERIOD (see table 47). See table 47 for the range of valid values for FKA_ADV_PERIOD.

7.8.6.3.14 FIP Vendor_ID descriptor

The FIP Vendor_ID descriptor is used in FIP operations as specified in table 45.

The FIP Vendor_ID descriptor format shall be as specified in table 42.

Table 42 – FIP Vendor_ID descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 0Dh								Length = 03h								Reserved															
1	(MSB)								Vendor_ID																							
2																															(LSB)	

Vendor_ID: the vendor's Vendor_ID value.

7.8.6.3.15 FIP VLAN descriptor

The FIP VLAN descriptor is used in in FIP operations as specified in table 45.

The FIP VLAN descriptor format shall be as specified in table 43

Table 43 – FIP VLAN descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type = 0Eh								Length = 01h								Reserved				FCoE VID											

FCoE VID: the VLAN ID of a VLAN where FCoE services may be available. The range of valid values for VLAN IDs is 001h to FFEh.

7.8.6.3.16 FIP Vendor Specific descriptors

FIP Vendor Specific descriptors are non-critical and may be used in any FIP frame. An FC-BB_E device shall not require the use of any FIP Vendor Specific descriptor in order to operate in accordance with this standard.

The FIP Vendor Specific descriptor format is specified in table 44.

Table 44 – FIP Vendor Specific descriptor format

Bit / Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Type								Length								Reserved															
1	(MSB)								Vendor_ID																							
2																															(LSB)	
3	(MSB)								Vendor Specific information																							
n																															(LSB)	

Type: the FIP Vendor Specific descriptors are identified by a type value in the range 241 to 254, inclusive.

Length: shall be set to the length in words of the descriptor.

Vendor_ID: the vendor's Vendor_ID value.

Vendor Specific Information: defined by the vendor.

7.8.7 FIP operations

7.8.7.1 FIP operations overview

Table 45 specifies the FIP descriptors required in each FIP operation and the recommended order in which they should be encapsulated by a transmitting FCoE Controller. In certain cases, as indicated, strict ordering is required. A receiving FCoE Controller shall process unknown descriptors according to the criticality of the FIP descriptor (see 7.8.6.3.1). Unless otherwise specified (e.g., for a FIP FLOGI Request), a receiving FCoE Controller shall be able to process the FIP descriptors in any order.

NOTE 28 – The ability to process FIP descriptors in any order is to provide flexibility for future protocol extensions,

A FIP operation shall contain the expected critical descriptors (see table 45) and may contain additional non-critical descriptors. If some critical descriptors are missing or unexpected, the FIP operation shall be discarded and it should be reported in a vendor specific manner.

Table 45 – FIP operation descriptors and order

FIP Operation	FIP Protocol Code/Subcode[a]	Originator	Expected Descriptors and Order
Discovery Solicitation (see 7.8.7.2.1)	0001h/01h	ENode	1) MAC address 2) Name_Identifier 3) Max FCoE Size
Discovery Solicitation (see 7.8.7.2.2)	0001h/01h	FCF	1) MAC address 2) FC-MAP 3) Name_Identifier 4) Max FCoE Size
Discovery Advertisement (see 7.8.7.3)	0001h/02h	FCF	1) Priority 2) MAC address 3) Name_Identifier 4) Fabric 5) FKA_ADV_Period
FIP FLOGI Request[b] (see 7.8.7.4.2)	0002h/01h	ENode	1) FLOGI 2) MAC address
FIP FLOGI LS_ACC[b] (see 7.8.7.4.2)	0002h/02h	FCF	1) FLOGI 2) MAC address
FIP FLOGI LS_RJT[b] (see 7.8.7.4.2)	0002h/02h	FCF	1) FLOGI
FIP NPIV FDISC Request[b] (see 7.8.7.4.2)	0002h/01h	ENode	1) NPIV FDISC 2) MAC address
FIP NPIV FDISC LS_ACC[b] (see 7.8.7.4.2)	0002h/02h	FCF	1) NPIV FDISC 2) MAC address
FIP NPIV FDISC LS_RJT[b] (see 7.8.7.4.2)	0002h/02h	FCF	1) NPIV FDISC

a) Virtual Link instantiation requests are differentiated by the contained descriptors.
b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here.
c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Link frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP).
d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames.
e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor.
f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present).

Table 45 – FIP operation descriptors and order (Continued)

FIP Operation	FIP Protocol Code/Subcode[a]	Originator	Expected Descriptors and Order
FIP Fabric LOGO[b] (see 7.8.7.4.3)	0002h/01h	ENode[c]	1) LOGO 2) MAC address
FIP Fabric LOGO LS_ACC[b] (see 7.8.7.4.3)	0002h/02h	FCF[c]	1) LOGO 2) MAC address
FIP Fabric LOGO LS_RJT[b] (see 7.8.7.4.3)	0002h/02h	FCF[c]	1) LOGO
FIP ELP Request[b] (see 7.8.7.4.4)	0002h/01h	FCF	1) ELP 2) MAC address
FIP ELP SW_ACC[b] (see 7.8.7.4.4)	0002h/02h	FCF	1) ELP 2) MAC address
FIP ELP SW_RJT[b] (see 7.8.7.4.4)	0002h/02h	FCF	1) ELP
FIP Keep Alive (see 7.8.7.5)	0003h/01h	ENode	1) MAC address 2) Vx_Port Identification[d]
FIP Clear Virtual Links (see 7.8.7.6	0003h/02h	FCF	1) MAC address 2) Name_Identifier 3) Vx_Port Identification(s)[e]
FIP VLAN Request (see 7.8.7.7)	0004h/01h	ENode or FCF	1) MAC address 2) Name_Identifier[f]
FIP VLAN Notification (see 7.8.7.8)	0004h/02h	FCF	1) MAC address 2) VLAN(s)
FIP Vendor Specific (see 7.8.7.9)	FFF8h to FFFEh / 00h to FFh	ENode or FCF	1) Vendor_ID 2) Descriptor(s)

a) Virtual Link instantiation requests are differentiated by the contained descriptors.
b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here.
c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Link frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP).
d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames.
e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor.
f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present).

7.8.7.2 FIP Discovery Solicitation

7.8.7.2.1 ENode FIP Discovery Solicitation

As specified in table 45, a Discovery Solicitation operation originated by the FCoE Controller of an ENode MAC contains a MAC address descriptor (see 7.8.6.3.3), a Name_Identifier descriptor (see 7.8.6.3.5), and a Max FCoE Size descriptor (see 7.8.6.3.7).

A Discovery Solicitation frame may be unicast (i.e., addressed to a specific FCF-MAC) or multicast (i.e., addressed to the All-FCF-MACs group address).

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

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

NOTE 29 – The Name_Identifer field may be set to zero if the Node_Name is ambiguous or not yet available when the Discovery Solicitation is transmitted.

The Max_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE PDU size the ENode MAC intends to use for FCoE traffic. The Max_FCoE_Size value shall be specified as the number of octets starting with and including Version field, up to and including the Reserved field following the EOF field (see table 21).

7.8.7.2.2 FCF FIP Discovery Solicitation

As specified in table 45, a Discovery Solicitation operation originated by the FCoE Controller of a VE_Port capable FCF-MAC contains a MAC address descriptor (see 7.8.6.3.3), an FC-MAP descriptor (see 7.8.6.3.4), a Name_Identifier descriptor (see 7.8.6.3.5), and a Max FCoE Size descriptor (see 7.8.6.3.7).

A Discovery Solicitation frame may be unicast (i.e., addressed to a specific FCF-MAC) or multicast (i.e., addressed to the All-FCF-MACs group address).

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

For FCF-MACs that support FPMA, the FC-MAP field in the FC-MAP descriptor shall be set to the FC-MAP value the FCF-MAC is using. If the FC-MAP value is not administratively configured, then the FC-MAP value shall be set to DEFAULT_FC-MAP (see table 47).

For FCF-MACs that only support SPMA, the FC-MAP field in the FC-MAP descriptor shall be set to zero.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF.

The Max_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE frame size the VE_Port capable FCF-MAC intends to use for FCoE traffic. The Max_FCoE_Size value shall be specified as the number of octets starting with and including the Version field, up to and including the Reserved field following the EOF field (see table 21).

7.8.7.3 FIP Discovery Advertisements

As specified in table 45, a Discovery Advertisement operation contains a Priority descriptor (see 7.8.6.3.2), a MAC address descriptor (see 7.8.6.3.3), a Name_Identifier descriptor (see 7.8.6.3.5), a Fabric descriptor (see 7.8.6.3.6), and a FKA_ADV_Period descriptor (see 7.8.6.3.13).

When a Discovery Advertisement frame is solicited, it shall be unicast (i.e., addressed to a specific ENode MAC or FCF-MAC address). When a Discovery Advertisement frame is unsolicited, it shall be multicast (i.e., addressed to the All-ENode-MACs or to the All-FCF-MACs group addresses).

The Priority field in the Priority descriptor shall be set to the value the originating FCF-MAC is using. If the priority value is not administratively configured, then the priority value shall be set to DEFAULT_FIP_PRIORITY (see table 47).

The MAC address field in the MAC address descriptor shall be set to the originating FCF-MAC address.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the originating FCF.

Discovery Advertisements shall only contain a single Fabric descriptor. All Discovery Advertisements from an FCF in a VLAN shall contain the same single Fabric descriptor. The VF_ID field in the Fabric descriptor shall be set to the VF_ID identifying the advertised FC Fabric. If a VF_ID is not defined for the advertised FC Fabric, the VF_ID field shall be set to zero. For FCFs that support FPMA, the FC-MAP field in the Fabric descriptor shall be set to the FC-MAP value the FCF is using. If the FC-MAP value is not administratively configured, then the FC-MAP value shall be set to DEFAULT_FC-MAP (see table 47). For FCFs that only support SPMA, the FC-MAP field in the Fabric descriptor shall be set to zero. The Fabric_Name field in the Fabric descriptor shall be set to the Fabric_Name for the originating FCF.

The FKA_ADV_PERIOD field in the FKA_ADV_Period descriptor shall be set to the FKA_ADV_PERIOD value the FCF is advertising (see table 47).

The FIP_Pad field shall be used to extend the FIP PDU (see table 24) to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation to which the Discovery Advertisement is responding. The FIP_Pad field value shall be set to zero and not checked in reception. For an unsolicited Discovery Advertisement, the FIP_Pad field shall be of zero length (i.e., not present).

7.8.7.4 FIP Virtual Link Instantiation Requests and Replies

7.8.7.4.1 FIP Virtual Link Instantiation Requests and Replies overview

FIP Virtual Link Instantiation Requests and Replies encapsulates an ELS or an SW_ILS. The encapsulated ELS or SW_ILS shall be a single-frame Sequence. FIP Virtual Link Instantiation Requests and Replies are used to perform:

a) Fabric login between ENode MACs and VF_Port capable FCF-MACs (see 7.8.7.4.2);
b) Fabric logout between ENode MACs and VF_Port capable FCF-MACs (see 7.8.7.4.3); and
c) Exchange Link Parameters between VE_Port capable FCF-MACs (see 7.8.7.4.4).

7.8.7.4.2 Fabric login

As specified in table 45:

a) a FIP FLOGI Request operation contains a FLOGI descriptor (see 7.8.6.3.8) and a MAC address descriptor (see 7.8.6.3.3);
b) a FIP FLOGI LS_ACC operation contains a FLOGI descriptor (see 7.8.6.3.8) and a MAC address descriptor (see 7.8.6.3.3);
c) a FIP FLOGI LS_RJT operation contains a FLOGI descriptor (see 7.8.6.3.8).
d) a FIP NPIV FDISC Request operation contains a NPIV FDISC descriptor (see 7.8.6.3.9) and a MAC address descriptor (see 7.8.6.3.3);
e) a FIP NPIV FDISC LS_ACC operation contains a NPIV FDISC descriptor (see 7.8.6.3.9) and a MAC address descriptor (see 7.8.6.3.3); and
f) a FIP NPIV FDISC LS_RJT operation contains a NPIV FDISC descriptor (see 7.8.6.3.9).

The FLOGI or NPIV FDISC descriptor shall be the first descriptor in the operation.

NOTE 30 – In this way the encapsulated Fibre Channel ELS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in a FCoE PDU.

The FLOGI or NPIV FDISC descriptor shall contain:

a) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and an FLOGI Request payload or an NPIV FDISC Request payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP FLOGI Request or FIP NPIV FDISC Request operations. The Payload bit in the FLOGI Request payload or NPIV FDISC Request payload shall be set to zero (see FC-LS-2);
b) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and an FLOGI LS_ACC payload or an NPIV FDISC LS_ACC payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC operations. The Payload bit in the FLOGI LS_ACC payload shall be set to zero (see FC-LS-2); or
c) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and an FLOGI LS_RJT payload or an NPIV FDISC LS_RJT payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP FLOGI LS_RJT or FIP NPIV FDISC LS_RJT operations.

The MAC address field in the MAC address descriptor of a FIP FLOGI Request operation or a FIP NPIV FDISC Request operation shall contain:

a) the proposed MAC address to use as VN_Port MAC address if the ENode is requesting to use SPMA (see table 27);
b) all zeroes to indicate no MAC address is proposed if the ENode is requesting to use FPMA (see table 27); or
c) the proposed MAC address to use as VN_Port MAC address if the ENode supports both SPMA and FPMA and leaves the decision of which addressing scheme to use to the FCF (i.e., if both the FP and SP bits are set to one, see table 27).

The MAC address field in the MAC address descriptor of a FIP FLOGI LS_ACC operation or a FIP NPIV FDISC LS_ACC operation shall contain the MAC address that the FCF granted for use as VN_Port MAC address. The FP and SP bits shall be set as shown in table 28. If the FCF granted an SPMA, the granted MAC address shall be the same as the one carried in the MAC address descriptor of the corresponding FIP FLOGI Request operation or FIP NPIV FDISC Request operation. If the FCF granted an FPMA, the granted MAC address shall be a properly formed FPMA (see 7.6). An ENode shall verify that a granted FPMA address is properly formed.

A FIP FLOGI or FIP NPIV FDISC operation shall be processed respectively as the FLOGI or FDISC ELS with S_ID = 000000h specified in FC-LS-2, ignoring the buffer-to-buffer flow control parameters. If no response to a FIP FLOGI Request or FIP NPIV FDISC Request is received within the ELS

timeout value specified in FC-LS-2, the Exchange resources may be reused and the ABTS protocol (see FC-FS-3) shall not be used.

A successful FIP FLOGI operation instantiates a VF_Port, a VN_Port, and a Virtual Link between them. Subsequent FIP NPIV FDISC operations from the same ENode MAC Address as the FIP FLOGI operation associate additional VN_Ports to the same VF_Port that was instantiated by the FIP FLOGI operation. Subsequent FIP NPIV FDISC operations from the same ENode MAC address as the FIP FLOGI operation should request the same MAC address type (i.e., FPMA or SPMA) as the FIP FLOGI operation. A subsequent FIP FLOGI operation from the same ENode MAC Address is equivalent to an implicit logout of all the VN_Ports followed by an FLOGI, as specified in FC-LS-2.

NOTE 31 – A duplicate MAC address may be a cause for an implicit logout.

FCFs shall reject FIP FLOGI Requests and FIP NPIV FDISC Requests for a single addressing mode (i.e., SPMA or FPMA) that is not supported by the FCFs. FCFs supporting FPMAs shall reject FIP FLOGI and FIP NPIV FDISC Requests with the SP bit set to one when the MAC address descriptor contains a proposed MAC address in which the 24 most significant bits match the FC-MAP in use by the FCFs. In addition, FCFs shall reject FIP FLOGI Requests and FIP NPIV FDISC Requests having:

a) both the SP bit and the FP bit set to zero;
b) the FP bit set to one, the SP bit set to zero, and the MAC address descriptor not set to zero; and
c) the SP bit set to one and the MAC address descriptor containing a MAC address that is not a unicast address.

Rejections of FIP FLOGI Requests and FIP NPIV FDISC Requests are performed with the LS_RJT Reason Codes and Reason Code Explanations specified in table 46.

Table 46 – FIP Fabric login rejections

Error Condition	Reason Code (see FC-LS-2)	Reason Code Explanation (see FC-LS-2)
The MAC addressing mode in the FIP FLOGI/FDISC Request is not supported.	FIP Error (i.e., 20h)	MAC addressing mode not supported (i.e., 60h).
The MAC address proposed in the MAC address descriptor of a FIP FLOGI/FDISC Request is incorrect for the requested addressing mode.	FIP Error (i.e., 20h)	Proposed MAC address incorrectly formed (i.e., 61h).

7.8.7.4.3 Fabric logout

As specified in table 45:

a) a FIP Fabric LOGO Request operation contains a LOGO descriptor (see 7.8.6.3.10) and a MAC address descriptor (see 7.8.6.3.3);
b) a FIP Fabric LOGO LS_ACC operation contains a LOGO descriptor (see 7.8.6.3.10) and a MAC address descriptor (see 7.8.6.3.3); and
c) a FIP Fabric LOGO LS_RJT operation contains a LOGO descriptor (see 7.8.6.3.10).

The LOGO descriptor shall be the first descriptor in the operation.

NOTE 32 – In this way the encapsulated Fibre Channel ELS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in a FCoE PDU.

The LOGO descriptor shall contain:

a) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a Fabric LOGO Request payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP Fabric LOGO Request operations;
b) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a Fabric LOGO LS_ACC payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP Fabric LOGO LS_ACC operations; or
c) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a Fabric LOGO LS_RJT payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP Fabric LOGO LS_RJT operations.

The MAC address field in the MAC address descriptor of a FIP Fabric LOGO Request and a FIP Fabric LOGO LS_ACC operation shall be set to the MAC address assigned to the VN_Port that is being logged out.

A FIP Fabric LOGO operation shall be processed as the Fabric LOGO ELS specified in FC-LS-2. If no response to a FIP Fabric LOGO Request is received within the ELS timeout value specified in FC-LS-2, the Exchange resources may be reused and the ABTS protocol (see FC-FS-3) shall not be used.

7.8.7.4.4 Exchange Link Parameters

As specified in table 45:

a) a FIP ELP Request operation contains a ELP descriptor (see 7.8.6.3.11) and a MAC address descriptor (see 7.8.6.3.3);
b) a FIP ELP SW_ACC operation contains a ELP descriptor (see 7.8.6.3.11) and a MAC address descriptor (see 7.8.6.3.3); and
c) a FIP ELP SW_RJT operation contains a ELP descriptor (see 7.8.6.3.11).

The ELP descriptor shall be the first descriptor in the operation.

NOTE 33 – In this way the encapsulated Fibre Channel SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in a FCoE PDU.

The ELP descriptor shall contain:

a) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a ELP Request payload specifying R_RDY Flow Control (see FC-SW-5) but without the CRC field (see FC-FS-3) for FIP ELP Request operations;
b) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a ELP SW_ACC payload (see FC-SW-5) but without the CRC field (see FC-FS-3) for FIP ELP SW_ACC operations; or
c) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and a ELP SW_RJT payload (see FC-SW-5) but without the CRC field (see FC-FS-3) for FIP ELP SW_RJT operations.

The MAC address field in the MAC address descriptor of a FIP ELP Request and a FIP ELP SW_ACC operation shall be set to the MAC address of the destination FCF-MAC.

A FIP ELP operation shall be processed as the ELP SW_ILS specified in FC-SW-5, ignoring the buffer-to-buffer flow control parameters.

7.8.7.5 FIP Keep Alive

As shown in table 45, a FIP Keep Alive operation contains a MAC address descriptor (see 7.8.6.3.3) and zero or one Vx_Port Identification descriptor (see 7.8.6.3.12).

ENode FIP Keep Alive operations (see 7.8.5.1) contains only a MAC address descriptor. VN_Port FIP Keep Alive operations (see 7.8.5.1) contain a MAC address descriptor and a Vx_Port Identification descriptor.

ENode FIP Keep Alive frames shall have the originating ENode MAC address as the source address. The MAC address field in the MAC address descriptor shall be set to the originating ENode MAC address.

VN_Port FIP Keep Alive frames shall have the VN_Port MAC address as the source address. The MAC address field in the MAC address descriptor shall be set to the originating ENode MAC address. In the Vx_Port Identification descriptor, the MAC address field shall be set to the VN_Port MAC address, the Address Identifier field shall be set to the VN_Port N_Port_ID, and the Port_Name field shall be set to the VN_Port N_Port_Name.

7.8.7.6 FIP Clear Virtual Links

7.8.7.6.1 FIP Clear Virtual Links to an ENode

The FCoE Controller of a VF_Port capable FCF-MAC may de-instantiate one or more VN_Port to VF_Port Virtual Links by transmitting a FIP Clear Virtual Links frame to an ENode MAC. As specified in table 45, this FIP Clear Virtual Links frame shall contain one MAC address descriptor (see 7.8.6.3.3), one Name_Identifier descriptor (see 7.8.6.3.5), and a list of Vx_Port Identification descriptors (see 7.8.6.3.12), one for each VN_Port the Virtual Link with it is requested to be de-instantiated.

The MAC address field in the MAC address descriptor shall be set to the FCF-MAC address of the originating FCF-MAC. The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the originating FCF. For each Vx_Port Identification descriptor, the MAC address field shall be set to the VN_Port MAC address, the Address Identifier field shall be set to the VN_Port N_Port_ID, and the Port_Name field shall be set to the VN_Port N_Port_Name. The FCoE Controller of a receiving ENode MAC shall ignore a Vx_Port Identification descriptor that does not match any of its instantiated VN_Ports.

7.8.7.6.2 FIP Clear Virtual Links to an FCF

The FCoE Controller for a VE_Port capable FCF-MAC may de-instantiate a VE_Port to VE_Port Virtual Link by transmitting a FIP Clear Virtual Links frame to a VE_Port capable FCF-MAC. As specified in table 45, this FIP Clear Virtual Links frame shall contain one MAC address descriptor (see 7.8.6.3.3), one Name_Identifier descriptor (see 7.8.6.3.5), and one Vx_Port Identification descriptor (see 7.8.6.3.12).

The MAC address field in the MAC address descriptor shall be set to the FCF-MAC address of the originating FCF-MAC. The Name_Identifier field in the Name_Identifier descriptor shall be set to the

Switch_Name of the originating FCF. In the Vx_Port Identification descriptor, the MAC address field shall be set to the remote FCF-MAC address, the Address Identifier field shall be set to FFFFFDh, and the Port_Name field shall be set to the remote E_Port_Name.

7.8.7.7 FIP VLAN Request

As specified in table 45, a FIP VLAN Request operation contains a MAC address descriptor and optionally a Name_Identifier descriptor (see 7.8.6.3.3). A FIP VLAN Request frame may be generated by an ENode MAC or by an FCF-MAC.

When generated by an ENode MAC, the FIP VLAN Request frame shall have the F flag set to zero and the MAC address field in the MAC address descriptor shall be set to the originating ENode MAC address. If a Name_Identifier descriptor is present, the Name_Identifier field should be set to the Node_Name of the originating ENode.

When generated by an FCF-MAC, the FIP VLAN Request frame shall have the F flag set to one and the MAC address field in the MAC address descriptor shall be set to the originating FCF-MAC address. If a Name_Identifier descriptor is present, the Name_Identifier field should be set to the Switch_Name of the originating FCF.

7.8.7.8 FIP VLAN Notification

As specified in table 45, a FIP VLAN Notification operation contains a MAC address descriptor (see 7.8.6.3.3) and one or more VLAN descriptors (see 7.8.6.3.15). A FIP VLAN Notification frame is generated by an FCF-MAC.

The MAC address field in the MAC address descriptor shall be set to the originating FCF-MAC address. The FCoE VID field of each of the FIP VLAN descriptors shall be set to a VID over which the FCF-MAC is offering FC-BB_E services.

7.8.7.9 FIP Vendor Specific frames

FIP Vendor Specific frames may be transmitted by both ENodes and FCFs. As specified in table 45, a FIP Vendor Specific frame shall include a Vendor_ID descriptor (see 7.8.6.3.14) as the first descriptor, followed by one or more additional descriptors. An unknown received FIP Vendor Specific frame shall be discarded. An FC-BB_E device shall not require the use of any FIP Vendor Specific frame in order to operate in accordance with this standard.

7.9 Timers and constants

FC-BB_E timers and constants are specified in table 47.

Table 47 – FC-BB_E timers and constants

Timer/Constant	Value	Description	Reference
FIP_TYPE	8914h	The value specified in the Ethernet Type field for a FIP PDU.	7.8.6.1
FCoE_TYPE	8906h	The value specified in the Ethernet Type field for an FCoE PDU.	7.7
FIP_FRAME_VER	0001b	The value specified in the Version field for a FIP PDU.	7.8.6.1
FCoE_FRAME_VER	0000b	The value specified in the Version field for an FCoE PDU.	7.7
All-FCoE-MACs	01-10-18-01-00-00	The group address for all FCoE devices.	7.8.1
All-ENode-MACs	01-10-18-01-00-01	The group address for all ENodes.	7.8.1
All-FCF-MACs	01-10-18-01-00-02	The group address for all FCFs.	7.8.1
DEFAULT_FIP_PRIORITY	128	The default value specified in the FIP Priority descriptor.	7.8.6.3.2
DEFAULT_FC-MAP	0EFC00h	The default value for the FC-MAP field in a FIP FC-MAP descriptor.	7.8.6.3.4
ADV_TOV	2	The interval in seconds within which solicited Discovery Advertisements are transmitted, if the FCF chooses so, upon reception of a Discovery Solicitation.	7.8.3
FKA_ADV_PERIOD	-	The interval in milliseconds between periodic Discovery Advertisements and ENodes FIP Keep Alive frames. The default value is 8 000. Valid values are in the range 250 to 90 000.	7.8.3, 7.8.5
FKA_VN_PERIOD	90[a]	The interval in seconds between periodic VN_Port FIP Keep Alive frames.	7.8.5

a) This value has been chosen as appropriate to keep the forwarding tables of intermediate Lossless Ethernet bridges updated.

7.10 FC-BB_E Link Error Status Block (LESB) definition

FC-FS-3 specifies a Link Error Status Block (LESB) to monitor link error statistics that are useful for problem determination in Fibre Channel topologies. The LESB may be obtained from an FC-BB_E VN_Port using the Read Link Error Status Block (RLS) ELS request (see FC-LS-2).

An FC-BB_E VN_Port or VF_Port that supports the LESB with the FC-BB_E format (see table 21) shall provide this LESB format in response to an RLS ELS request. If an FC-BB_E VN_Port or VF_Port that is the designated FC_Port of an RLS request sequence does not support any of the fields specified in table 48, the recipient of the RLS request sequence shall reply to the RLS request with an LS_RJT specifying a reason code of "Unable to perform command request" (i.e., 09h) and should respond with a reason code explanation of "Request not supported" (i.e., 2Ch) (see FC-LS-2). A subset of the counters in the LESB may be supported and any unsupported counters shall be set to zero.

Table 48 specifies the format for the FC-BB_E LESB that shall be used in response to an RLS request by an FC-BB_E VN_Port or VF_Port.

Table 48 – FC-BB_E Link Error Status Block format

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Link Failure Count
1	Virtual Link Failure Count
2	Missing FIP Keep Alive or Discovery Advertisement Count
3	Symbol Error During Carrier Count
4	Errored Block Count
5	Frame Check Sequence Error Count

The Link Failure Count field indicates the number of link failures detected through detection of physical link transitions (i.e., the number of times that the aMediaAvailable attribute (see IEEE 802.3-2008 subclause 30.5.1.1.4) changes from the enumeration “available” to any other enumeration.

The Virtual Link Failure Count field indicates the number of virtual link failures detected by the Virtual Link maintenance protocol.

The Missing FIP Keep Alive or Discovery Advertisement field indicates the number of missing Virtual Link maintenance protocol frames. A missing Virtual Link maintenance protocol frame is detected after 1.5 times FKA_ADV_PERIOD since the reception of the last Virtual Link maintenance protocol frame. For an ENode, the Missing FIP Keep Alive or Discovery Advertisement field indicates the number of missing Discovery Advertisements. For a VF_Port, the Missing FIP Keep Alive or Discovery Advertisement field indicates the number of missing FIP Keep Alive frames from an ENode.

The Symbol Error During Carrier Count field indicates the number of reception errors at the PHY layer that occur during frame reception. The detection procedure is dependant on media and link speed (see IEEE 802.3-2008 subclause 30.3.2.1.5).

The Errored Block Count field indicates the cumulative count of the events counted by the eight-bit errored blocks counter (see IEEE 802.3-2008 subclause 45.2.3.12.4).

The Frame Check Sequence Error Count field indicates the number of Ethernet frames received that are an integral number of octets in length and do not pass the FCS check (see IEEE 802.3-2008 subclause 30.4.3.1.6).

7.11 Link Incidents definition

The Link Incident reporting procedure (see FC-LS-2) defines link incidents and corresponding incident codes that are based on the Fibre Channel physical layer. A definition of FC-BB_E Link Incidents and their respective Incident Code values (see table 49) are based on the IEEE 802.3 physical layer and the FC-BB_E Virtual Link maintenance protocol.

Bit-error rate thresholding (see FC-FS-3) defines error intervals as a time period during which Fibre Channel invalid Tranmission Words are recognized. For FC-BB_E devices an error interval is a time period during which one or more error blocks (see IEEE 802.3-2008 subclause 45.2.3.12.4) are recognized.

Table 49 specifies the Link Incidents for FC-BB_E that shall be used in a RLIR ELS request by an FC-BB_E VN_Port, VF_Port, or VE_Port.

Table 49 – FC-BB_E Link Incidents

Value	Meaning
00h	Reserved
01h	Implicit incident: A condition, caused by an event known to have occurred within the incident port, has been recognized by the incident port. The condition affects the attached link in such a way that it may cause a link incident to be recognized by the connected port.
02h	Bit-error-rate threshold exceeded: The incident port has detected that the Error Interval Count equals the Error Threshold (see FC-FS-3) where the Error Interval Count is based on errored blocks (see IEEE 802.3-2008 subclause 45.2.3.12.4).
03h	Link Failure - Loss-of-Signal: The aLoseMediaCounter has been incremented for entering an aMediaAvailable state indicating anything other than a remote fault (see IEEE 802.3-2008 subclause 30.5.1.1.4 and IEEE 802.3-2008 subclause 30.5.1.1.5).
04h	Link Failure - Remote fault: The aLoseMediaCounter has been incremented for an aMediaAvailable state being entered indicating a remote fault (see IEEE 802.3-2008 subclause 30.5.1.1.4 and IEEE 802.3-2008 subclause 30.5.1.1.5).
05h	Link Failure - Virtual Link failure: The incident port has detected a Virtual Link failure using the Virtual Link maintenance protocol.
06h to FFh	Reserved

Annex A: FC-BB_GFPT Interoperability Guidelines (Informative)

A.1 GFPT-specific interoperability guidelines

To ensure interoperability, it is recommended that GFPT-specific implementations for FC-BB_GFPT devices should include the option of:

a) using no optional Payload FCSs (see ITU-T Rec. G.7041/Y.1303);
b) using a null Extension Header (see ITU-T Rec. G.7041/Y.1303);
c) not propagating Client Signal Fail via Client Management Frames (see ITU-T Rec. G.7041/Y.1303); and
d) having the number of superblocks per GFPT frame not exceed twice the minimum number of superblocks per GFPT frame as recommended in ITU-T Rec. G.7041/Y.1303.

Annex B: FCoE and FIP Frame Examples (Informative)

B.1 Overview

An example format of an FCoE frame is specified in table B.1.

Table B.1 – FCoE frame format example

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	(MSB) Destination Address																															
1	(LSB)																(MSB)															
2	Source Address																														(LSB)	
3	IEEE 802.1Q Tag																															
4	Type=FCoE_TYPE																Version				Reserved											
5	Reserved																															
6	Reserved																															
7	Reserved																														SOF	
8	(MSB)																															
n+7	Encapsulated FC Frame (n words)																														(LSB)	
n+8	EOF																Reserved															
n+9	FCS																															

NOTE 34 – Additional IEEE 802.1 defined tags may be present in an FCoE frame.

An example format of a FIP frame is specified in table B.2.

Table B.2 – FIP frame format example

Word \ Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	(MSB) Destination Address																															
1	(LSB)																(MSB)															
2	Source Address																														(LSB)	
3	IEEE 802.1Q Tag																															
4	Type=FCoE_TYPE																Version				Reserved											
5	(MSB)																															
n+4	Encapsulated FIP operation (n words)																														(LSB)	
n+5	FCS																															

NOTE 35 – Additional IEEE 802.1 defined tags may be present in a FIP frame.

Annex C: Increasing FC-BB_E Robustness Using Access Control Lists (Informative)

C.1 Overview

In Fibre Channel Fabrics, Fibre Channel switches are generally considered trusted devices. Fibre Channel end devices log into the switch to which they are attached before they may communicate with other end devices that are attached to the Fabric. Given that Fibre Channel links are point-to-point, the Fibre Channel switch has complete control over the traffic an end device injects into the Fabric or is received from the Fabric. As a result, the switch may enforce zoning configurations, ensure end devices are using their assigned addresses, and prevent various types of anomalous behaviors, both erroneous and malicious.

FCoE provides increased flexibility, but with this flexibility new challenges arise in assuring highly robust Fabrics. Specifically, if Ethernet bridges exist between an ENode and an FCF, the point-to-point assurance between the ENode and FCF is lost. Thus the FCF does not have the complete control that a Fibre Channel switch has.

Equivalent robustness between FCoE and Fibre Channel may be achieved by ensuring that all FCoE traffic to and from an ENode passes through an FCF and that if multiple ENodes access an FCF through a single physical FCF port, those ENodes use their assigned MAC addresses. Doing so, in effect, creates the equivalent of a point-to-point link between an ENode and FCF.

Note that the above are necessary, but not sufficient, conditions to ensure equivalent robustness. See Annex D for a complete discussion on achieving equivalent robustness.

A possible method of achieving this robustness is to ensure every ENode is physically connected to an FCF with no intervening Ethernet bridges, but in many deployments this is not practical.

Ethernet bridges commonly provide a feature called Access Control Lists (ACLs). Properly configured ACLs may emulate a point-to-point link by providing the traffic enforcement previously discussed. Furthermore, the FIP protocol has been designed to enable Ethernet bridges to efficiently monitor FIP frames passing through them. This data facilitates the automatic configuration of these ACLs. In addition, the automatic configuration is possible independent of any other ACLs that may be in use in the network for other applications.

This annex discusses the ACLs, the required Access Control Entries (ACEs) within the ACL to provide FCoE with equivalent robustness as Fibre Channel, and the process of generating these ACEs automatically by examining FIP frames.

The particular set of ACEs to be used depend on the location of the port within the network and the traffic that is administratively configured to pass through it. Figure C.1 illustrates a network along with the potential different ACEs that are applicable.

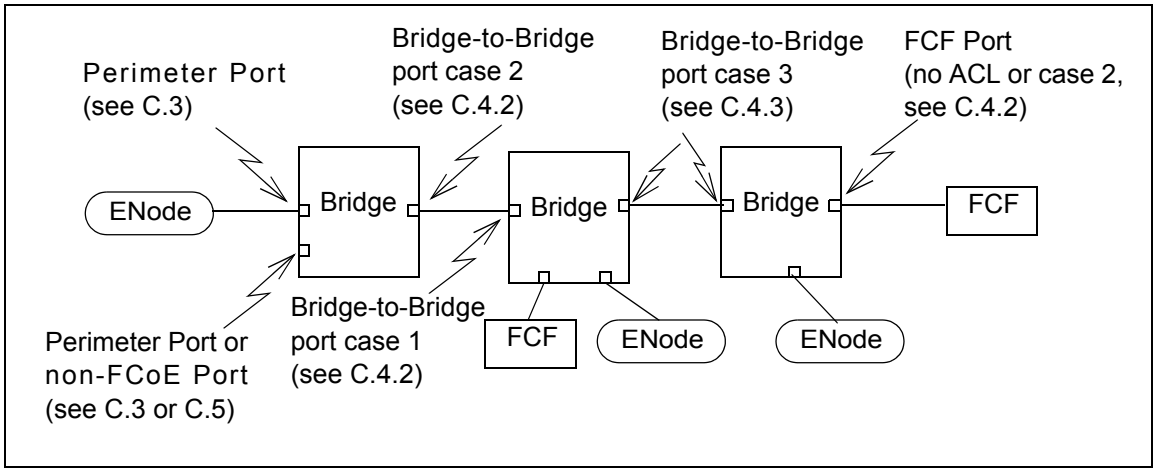

Figure C.1 – Bridge port to ACE cross reference

C.2 Access Control Lists

C.2.1 ACL overview

The implementation of ACLs is not standardized and specifics vary among Ethernet bridges. However, certain features are available from a wide variety of Ethernet bridges.

In general, an Access Control List consists of an ordered list of rules that determine whether a given frame should be forwarded (i.e.,“permit”) or discarded (i.e.,“deny”). Each rule is specified by matching bits within the received frame to a specified pattern. The pattern may require that the bits are set to one, to zero, or to don’t care. If a frame matches multiple patterns within the ACL, the first matching ACE determines whether the frame is permitted or denied. A default permit or deny may be specified in the last entry to cover the case in which no patterns match.

Most ACL implementations allow specification of ACLs per bridge port and operate on frames as they enter the bridge using the ACL specified for the ingress port (i.e., an ingress ACL). Some implementations may also apply ACLs to frames as they exit from the bridge (i.e., an egress ACL). For the purpose of this annex, all ACLs are assumed to be ingress ACLs.

It is recommended that ACL protection be applied at the edge of the network (i.e., at the ports that connect directly to the ENodes). The ACLs provided in this Annex are intended to be applied to these ports. It is also possible to construct ACLs that provide a lower level of protection on bridge-to-bridge ports. This annex also provides suggested ACLs for this purpose.

In addition, ACL implementations vary in how deeply into a frame the patterns may be applied. For the purpose of this annex, it is only necessary to examine the Source and Destination MAC address fields, the VLAN tag, and the Ethernet Type fields. Most ACL implementations are capable of this.

Implementations vary as to whether bridge learning is subject to the ingress ACL. This annex assumes that bridge learning is subject to the ACL (i.e., if a frame is denied by the ACL, its source address is not learned). For implementations that learn source addresses of denied frames, a simple extension using Static Forwarding Entries (see IEEE 802.1D-2004) is discussed to provide equivalent functionality (see C.7).

C.2.2 ACL nomenclature

The exact method of specifying ACLs is implementation specific. A generalized nomenclature is used in this annex. ACLs consist of an ordered list of access control entries. In general, an access control entry (ACE) has the form of:

[*field* = *value*],[*field* = *value*],...,**permit || deny;**

The last ACE may contain only the keyword “permit” or “deny” to cover the case that no ACEs match.

The fields used in this annex are:

a) DA: destination MAC Address (48 bits);
b) SA: source MAC Address (48 bits);
c) SApre: 24 most significant bits of the Source MAC Address (24 bits);
d) VLAN: value of the VLAN field within the VLAN tag (12 bits); and
e) Type: value of the Ethernet Type field (16 bits).

The following constants, defined in table 47, are used:

a) FIP_TYPE; and
b) FCoE_TYPE.

FC-MAP applies only if FPMAs are in use and is the 24-bit FPMA address prefix being used on the network (see 7.6).

“{FCFs}” represents the set of FCF-MAC addresses that a given ENode is allowed to connect. For simplicity, a single ACE is illustrated using this set. In general, multiple ACEs may be required to represent all the members of the set.

C.3 Perimeter ACL construction

C.3.1 Perimeter ACL construction overview

The ACL described in this subclause should be used at the perimeter of the network. This includes bridge ports connected to ENodes and unconnected bridge ports. It may also be used on bridge to bridge ports to provide additional security in depth; however, doing so in certain deployments may exceed the ACEs capacity of the bridge. In addition, deploying this ACL on bridge-to-bridge ports may limit the network ability to autonomously respond to link failures. See C.4 for additional options to address these issues on bridge-to-bridge links.

The following are the requirements of the ACL and the subclauses that describe how these requirements are met using ACEs:

a) enable transmission of FIP frames from ENodes to FCFs (see C.3.2);
b) ensure that FIP frames from ENodes may only be addressed to FCFs (see C.3.2);
c) ensure no end device uses an FCF-MAC address as its source (see C.3.3);
d) prevent transmission of all FCoE frames from an ENode prior to its successful completion of FIP FLOGI (see C.3.4);
e) after successful completion of FIP FLOGI, ensure that only the FCoE source addresses used by an ENode are the ones assigned by the FCF to that ENode (see C.3.5);
f) after successful completion of FIP FLOGI, ensure that the assigned FCoE source address is only used for FCoE traffic (see C.3.5) and FIP traffic (i.e., VN_Port FIP Keep Alive frames); and

g) after successful completion of each FIP FLOGI or FIP NPIV FDISC, ensure that FCoE frames may only be addressed to the accepting FCFs (see C.3.5).

These ACEs are constructed such that if they are inserted prior to any other non-FCoE and non-FIP related ACEs that may be in use, they do not conflict with those ACEs. In addition, these ACEs are constructed such that they do not inhibit non-FCoE and non-FIP traffic (i.e., traffic that does not contain the FCoE or FIP Ethernet Type value and does not utilize an FCoE source MAC address)

C.3.2 FIP frame transmission

An ENode is allowed to send FIP frames to FCFs, and only to FCFs. These frames may be addressed to a specific FCF, or to the All-FCF-MACS group address. ACEs that accomplish this are:

DA = All-FCF-MACs, Type = FIP_TYPE, permit;
DA = {FCFs}, Type = FIP_TYPE, permit; (see note 36)
Type = FIP_TYPE, deny;

NOTE 36 – This ACE also allows VN_Port FIP Keep Alive frames.

C.3.3 Prevention of the transmission of frames using an FCF-MAC address for the source

An ENode is not allowed to transmit frames using an FCF source address. This is necessary to prevent address learning and FCF impersonation attacks. The ACE that prevents this is:

SA = {FCFs}, deny;

C.3.4 Prevention of frames using FCoE Type or FCoE source addresses prior to successful completion of FIP FLOGI

ENodes are not permitted to send any FCoE frames prior to the successful completion of FIP FLOGI. FCoE frames are identified by the Type field being equal to FCoE_TYPE. The ACE to accomplish this is:

Type = FCoE_TYPE, deny;

C.3.5 Enabling traffic after successful completion of FIP FLOGI (or FIP NPIV FDISC)

After successful completion of FIP FLOGI, FCoE traffic between the ENode and the FCF that accepted the FLOGI using the assigned VN_Port MAC address is enabled. The following ACE accomplishes this:

SA = assigned VN_Port MAC address, DA = FCF-MAC address, Type = FCoE, permit;

For proper operation these ACEs are inserted anywhere prior to those in C.3.4 and it may be convenient to simply insert these at the top of the ACL.

C.3.6 Prevention of duplicate VN_Port MAC addresses

Duplicate VN_Port MAC addresses within a FC-BB_E network may lead to various catastrophic failures, including undetected corruption of data, denial of service, and undetected interception of data. Duplicate VN_Port MAC addresses may occur due to network configuration issues and malicious entities on the network. Duplicate VN_Port MAC addresses may be prevented with the use of FPMAs and the appropriate use of ACEs. In general, it is not practical to construct an ACE to prevent address duplication with SPMAs since doing so requires a priori knowledge of all MAC Addresses that

are being used as VN_Port MAC addresses everywhere in the network. As a result, duplicate VN_Port address prevention is beyond the scope of this annex for SPMAs.

With FPMAs, it is possible to identify all VN_Port MAC addresses. The following ACE causes a bridge port to discard any frame with a source address equal to a VN_Port MAC address:

SAPre = FC-MAP, deny;

This entry should be in place in the ACL after the entries specified in C.3.5.

C.3.7 ACL summary

Prior to receipt of any Discovery Advertisements, the initial ACL is:

DA = All-FCF-MACs, Type = FIP_TYPE, permit;
Type = FIP_TYPE, deny;
Type = FCoE_TYPE, deny;
SApre = FC-MAP, deny; --Note: applies to FPMA only
Any non-FCoE related ACEs.

After receipt of Discovery Advertisements or as the result of administrative configuration, the ACL is expanded to:

SA = {FCFs}, deny;
DA = All-FCF-MACs, Type = FIP_TYPE, permit;
DA = {FCFs}, Type = FIP_TYPE, permit;
Type = FIP_TYPE, deny;
Type = FCoE_TYPE, deny;
SApre = FC-MAP, deny; --Note: applies to FPMA only
Any non-FCoE related ACEs

For each successful FIP FLOGI (or FIP NPIV FDISC), an ACE is added prior to ‘Type=FCoE_TYPE, deny;’ of the form:

SA = assigned VN_Port MAC address, DA = FCF-MAC address, Type = FCoE, permit;

C.4 Security in depth

C.4.1 Overview

The ACL described in C.3, if properly deployed at the perimeter of the network (i.e., all bridge ports connected to all ENodes and unconnected ports), provides a high degree of network security. However, if a device is connected within this perimeter defense (e.g., a mis-configuration or an omitted port ACL), the level of security provided is diminished. Deployment of ACLs on bridge-to-bridge links provide additional defense in these situations. The ACL described in C.3 may be used for this purpose. However, as previously discussed, use of this ACL on bridge-to-bridge ports has undesirable scalability and network resiliency characteristics.

This subclause provides alternative ACLs that may be used on bridge-to-bridge links that have better scalability and network resilience characteristics. Three sets of ACLs are provided, each designed for a specific use within the network. See figure C.1 for an illustration of the following cases:

a) a bridge port, connected to a bridge-to-bridge link, receiving frames from ENode(s) destined for FCF(s) (but not the other direction);

b) a bridge port, connected to a bridge-to-bridge link or to an FCF, receiving frames from FCF(s) destined for ENode(s) (but not the other direction); and
c) a bridge port, connected to a bridge-to-bridge link, receiving frames from both FCF(s) and ENode(s).

NOTE 37 – Changes in routing (e.g., automatic spanning tree recalulation) may cause the assumptions on which the recommended ACLs are hosted are based to become invalid. This may have the result of weaker protection or Virtual Link timeouts that require relogins.

C.4.2 Bridge-to-bridge link receiving ENode frames destined to FCF(s)

In the case of a bridge-to-bridge link receiving frames from ENode(s), that by definition are destined to FCF(s), the ingress bridge port should check for basic validity of the received frames. This includes:

a) verification that FIP frames are addressed to FCFs;
b) verification that FCoE frames are addressed to FCFs;
c) verification that FCoE frames are sourced only by ENodes (see note 38); and
d) prevention of FCoE frames from one FCF destined to another FCF (see note 39).

NOTE 38 – This protection is practical with FPMA only.

NOTE 39 – Normally, FCFs are permitted to transfer frames between one another. However, this is a special case of a bridge-to-bridge link on which it is administratively known that all the traffic being received is supposed to be from ENode(s), therefore, it is proper to disallow FCF to FCF traffic. See C.4.4 to address the case in which traffic may be flowing from FCF to FCF.

A set of ACEs that accomplish this is:

DA=All-FCF-MACS, Type=FIP_TYPE, permit;
DA={FCFs}, Type=FIP_TYPE, permit;
Type=FIP_TYPE, deny;
DA={FCFs}, SApre=FC-MAP, Type=FCoE_TYPE, permit; -- FPMA only
DA={FCFs}, SA={FCFs}, Type=FCoE_TYPE, deny; -- SPMA only, but no harm for FPMA
DA={FCFs}, Type=FCoE_TYPE, permit; -- SPMA only
Type=FCoE_TYPE, deny;

C.4.3 Bridge-to-bridge link receiving FCF frames destined to ENode(s)

In the case of a bridge-to-bridge link receiving frames from FCF(s) with administrative knowledge that these frames are destined only to ENodes (i.e., there are no FCFs downstream) the ingress bridge port should check for basic validity of the received frames. This includes:

a) verification that all FIP frames are sourced from an FCFs and are not destined to other FCFs;
b) verification that all FCoE frames are sourced from an FCF;
c) verification that all FCoE frames are destined to ENodes (see note 40); and
d) prevention of FCoE frames destined for an FCF (see note 41).

NOTE 40 – This protection is practical with FPMA only.

NOTE 41 – Normally, FCFs are permitted to transfer frames between one another. However, this is a special case of a bridge-to-bridge link on which it is administratively known that all the traffic being received is supposed to be destined ENode(s), therefore, it is proper to disallow FCF to FCF traffic. See C.4.4 to address the case in which traffic may be flowing from FCF to FCF.

A set of ACEs that accomplish this is:

DA={FCFs}, Type=FIP_TYPE, deny;
DA=All-ENode-MACS, Type=FIP_TYPE, permit -- see note 42
SA={FCFs}, Type=FIP_TYPE, permit;
Type=FIP_TYPE, deny;
SA={FCFs}, DApre=FC-MAP, Type=FCoE_TYPE, permit; -- FPMA only
DA={FCFs}, SA={FCFs}, Type=FCoE_TYPE, deny; -- SPMA only, see note 43
SA={FCFs}, Type=FCoE_TYPE, permit; -- SPMA only
Type=FCoE_TYPE, deny;

NOTE 42 – This ACE should only be included if it is administratively known and trusted that only FCFs are able to inject frames destined to All-ENode-MACs onto the network. Including this ACE enables automatic population of the {FCFs} set. If this cannot be trusted, this ACE should not be included, This has the side effect of disabling automatic population of the {FCFs} set, thus requiring that set to be populated administratively.

NOTE 43 – This ACE may result in scalability issues in some deployments. Given n FCFs in a network, this expands to n squared entries, which may exceed the ACL capability of a bridge. In this case it is impractical to ensure that FCFs are not sourcing these frames when SPMAs are in use.

C.4.4 Bridge-to-bridge link receiving both FCF and ENode frames

In the case of a bridge-to-bridge link receiving frames from both FCF(s) and ENodes, including frames that are exchanged between FCFs, the ingress bridge port should check for basic validity of the received frames. This includes:

a) verify that all FIP frames are either sourced by or destined to an FCF;
b) verify that all FCoE frames are sourced by an FCF and destined to either an ENode or FCF, or sourced by an ENode and destined to an FCF (see note 44); and
c) verify that all FCoE frames or either sourced by are destined to an FCF.

NOTE 44 – This protection is practical with FPMA only.

A set of ACEs that accomplish this is:

DA=All-FCF-MACS, Type=FIP_TYPE, permit;
DA={FCFs}, Type=FIP_TYPE, permit;
SA={FCFs}, Type=FIP_TYPE, permit;
DA=All-ENode-MACs, Type=FIP_TYPE, permit -- see note 45;
Type=FIP_TYPE, deny;
DA={FCFs}, SApre=FC-MAP, Type=FCoE_TYPE, permit; -- FPMA only
SA={FCFs}, DApre=FC-MAP, Type=FCoE_TYPE, permit; -- FPMA only
DA={FCFs}, Type=FCoE_TYPE, permit; -- SPMA only
SA={FCFs}, Type=FCoE_TYPE, permit; -- SPMA only
Type=FCoE_TYPE, deny;

NOTE 45 – This ACE should only be included if it is administratively known and trusted that only FCFs are able to inject frames destined to All-ENode-MACs onto the network. Including this ACE enables automatic population of the {FCFs} set. If this cannot be trusted, this ACE should not be included. This has the side effect of disabling automatic population of the {FCFs} set, thus requiring that set to be populated administratively.

C.4.5 Additional FCF protection

The preceeding ACL recommendations have been based on the set of FCF-MAC addresses, {FCFs}, that contain all FCF-MAC addresses including MAC addresses used for VE_Ports, VF_Ports, and the FCoE Controllers of the FCF-MACs. This MAC address set may be subdivided into subsets for individual uses to create ACEs that provide greater protection against unintended FCF to FCF communication. For example, an ACE that prevents VE_Ports from communicating with VF_Ports may be constructed. While the details of achieving this are beyond the scope of this annex, this possibility may be considered for individual implementations and deployments.

C.5 Prevention of FCoE related traffic

It may be desirable to prevent the reception of all FCoE related traffic on a given bridge port (e.g., on bridge ports connected to links that are not known to be Lossless). To do this, all frames with an Ethernet Type of FCoE_TYPE or FIP_TYPE have to be denied. Additionally, it is desirable to deny all frames using any VN_Port MAC address as the source address (e.g., to prevent an attack from a rogue host or to prevent undetected data corruption due to an erroneous configuration). Denying such frames is generally practical only with FPMA. The ACEs to accomplish this are:

Type = FIP_TYPE, deny;
Type = FCoE_TYPE, deny;
SApre = FC-MAP, deny; -- Note: applies to FPMA only

C.6 Automatic configuration of ACLs

An Ethernet bridge may choose to examine FIP frames from which all, or most in the case of SPMAs, the information needed to automatically configure ACLs may be determined.

The set of FCF-MAC addresses {FCFs} may be determined by creating a list from all received FIP Discovery Advertisements. In addition, the FC-MAP may be determined from these advertisements.

It is important to note that the ACLs recommended in this annex do not prevent a rogue host from advertising itself as an FCF using the All-FCF-MACs group address. Preventing this is beyond the scope of this annex since this vulnerability exists in Fibre Channel and is not unique to FCoE. Within Fibre Channel and FCoE, this vulnerability may be addressed using FC-SP to prevent VE_Ports from being formed with such hosts.

However, bridges that examine FIP traffic to determine a list of FCF-MAC addresses include such rogue hosts in their list of valid FCFs if they consider such frames. To avoid this vulnerability, bridges should examine only the FIP advertisements addressed to the All-ENode-MACs group address. Alternatively, the list of FCF-MAC addresses may be configured administratively.

The post FLOGI/FDISC ACEs may be constructed by examining the FIP FLOGI LS_ACC/FIP NPIV FDISC LS_ACC frames transmitted by the FCFs. These frames contain the assigned VN_Port MAC address and the FCF-MAC address.

The FLOGI/FDISC ACEs may be deleted from one port and re-created on another by examining the FIP FLOGI LS_ACC/FIP NPIV FDISC LS_ACC frames transmitted by the FCFs, if it determined that an ENode has moved from one port to another.

Finally, the FIP Clear Virtual Links frame may be examined to determine that one of these ACEs may be removed from a port.

To ensure that only valid FIP frames are examined, all ports except those known to be connected to an FCF (e.g., via administrative configuration) should contain an ACE to filter FIP frames (the ACLs described in this annex contain such an entry).

The list of FCF-MAC addresses may be configured administratively. Doing so with the ACEs provided in this annex prevents a rogue host from impersonating an FCF.

The ACEs provided in this annex have soft state (i.e., a bridge should remove them if FIP Keep Alive frames and Discovery Advertisements are not seen on a periodic basis).

C.7 Ethernet bridge learning considerations

Some implementations of ACLs allow a bridge to learn a source address even if the frame is denied by the ACL. This may leave a network vulnerable to certain Ethernet learning attacks. In such implementations, Static Forwarding Entries (see IEEE 802.1Q-2005) may be used to supplement the ACL.

To accomplish this, when a post FLOGI/FDISC ACE is created, a Static Forwarding Entry for the assigned MAC address is also created. This entry should specify “forward” for the port on which the FIP FLOGI/FIP NPIV FDISC is received, and “filter” for all other ports.

C.8 VLAN considerations

It is possible for separate FCoE Fabrics to exist on separate VLANs. A common FC-MAP may be used for the entire physical infrastructure. The ACEs then need to be qualified with the appropriate VLAN ID(s).

The use of multiple FC-MAPs in a given physical infrastructure as well as the use of multiple Virtual Fabrics on a single VLAN is beyond the scope of this standard.

Annex D: FCoE Security Recommendations (Informative)

D.1 Overview

During the development of this standard, a detailed threat analysis was performed to ensure that FCoE networks may be practically deployed while maintaining security characteristics comparable to native Fibre Channel. As a result of this analysis, the deployment recommendations presented in this annex were developed. Many of the recommendations are included in other normative clauses of this standard as they are required for correct operation of FCoE in addition to providing a level of security comparable to Fibre Channel. The implementation of the remaining recommendations is not technically required for correct operation. However, to ensure the security characteristics of an FCoE network are comparable to the ones of a native Fibre Channel Fabric, these recommendations should be considered requirements.

D.2 Considerations

FCFs are assumed to be trusted devices. Therefore, a bridge known to be directly connected to an FCF is not required to perform any verification of frames received from the FCF.

For the purposes of this annex, “known” refers to knowledge gained through administrative action or from a trusted management application.

Furthermore, bridges that fully implement these recommendations provide a defensive perimeter. Therefore bridges known to be directly connected to bridges forming a defensive perimeter are not required to perform verification of frames received from the bridges forming the defensive perimeter. However, performing such verification enhances the security characteristics of the network. Doing so may come at the cost of limiting the scalability of the network and its ability to autonomously respond to faults, as well as in increased administrative complexity. These trade-offs should be considered during Fabric design and deployment.

Finally, FCoE networks may be subjected to various forms of catastrophic failures if duplication of VN_Port MAC addresses occur. These failures may include failure of the network to provide service, undetected data interception, and undetected data corruption. Addressing this issue differs depending on the addressing mode (i.e., FPMA or SPMA) being used. This is discussed in more detail in D.6.

D.3 General deployment recommendations

1) No VLAN should carry more than one Fibre Channel Virtual Fabric (applies to the LAN if VLANs are not in use).

NOTE 46 – The possibility of carrying more than one Virtual Fabric on a given VLAN was beyond the scope of the threat analysis performed during the development of this standard. Therefore, if more than one Virtual Fabric is deployed on a VLAN, this standard provides no assurance with respect to the security characteristics of the Fabric. Furthermore, it was observed that such a deployment greatly increases the possibility of a MAC address being assigned to multiple VN_Ports which may result in catastrophic Fabric failures and undetected data corruption. The probability of duplicate VN_Port MAC addresses is greatly exacerbated by the use of FPMA, but the concern applies to both FPMA and SPMA. Therefore, the deployment of multiple Virtual Fabrics on a single VLAN is strongly discouraged. This concern does not apply to the deployment of multiple Virtual Fabrics with each being deployed on an independent VLAN.

D.4 Bridge recommendations

1) All Bridge ports, except those known to be connected to FCFs, or those that are to be explicitly prohibited from carrying FCoE/FIP traffic, should implement ingress filtering that:

   I) discards all frames with a source MAC address matching that of any FCF;

   II) discards all frames with Ethernet Type = FIP_TYPE except those:

      i) addressed to the All-FCF-MACs group address; or

      ii) addressed to any FCF-MAC address

   III) discards all frames with Ethernet Type = FCoE_TYPE except those:

      i) containing a source MAC address currently assigned by an FCF to a VN_Port;

NOTE 47 – A currently assigned MAC address refers to the MAC address that has been assigned by the FCF via a FIP FLOGI/FIP NPIV FDISC LS_ACC and has not been unassigned by virtue of a corresponding FIP LOGO, FIP Clear Virtual Links, or the expiration of associated FIP Keep Alive timers.

      ii) containing a Destination MAC address equal to the FCF-MAC address of the FCF that assigned the source MAC address; and

      iii) received on the bridge port through which the FIP FLOGI/FIP NPIV FDISC LS_ACC was forwarded to the VN_Port that has been assigned the VN_Port MAC address.

Annex C provides a description of a method (i.e., the use of Access Control Lists) to satisfy this recommendation.

To ensure security characteristics comparable to native Fibre Channel, this recommendation should be applied to all bridge ports directly connected to ENodes. If this is not possible (e.g., the bridges connected to ENodes do not provide this filtering capability), enhanced security may still be obtained by applying this filtering to upstream bridge ports, but the security characteristics in this case are significantly weaker.

In many deployments, implementation of ingress filtering according to this recommendation on bridge ports connected to other bridge ports may have undesirable consequences with respect to scalability, ability of the network to autonomously respond to faults, and administrative complexity. If care is taken to ensure that the far-end bridge implements the recommendations in this annex, then this ingress filtering is not strictly required on the near-end bridge port. However, implementation of this recommendation provides additional protection (e.g., against configuration errors). Careful consideration should be given to the advantages and disadvantages of implementing this filtering on bridge ports connected to other bridges.

2) As an alternative to recommendation 1, all bridge ports known to be connected to other bridge ports should implement filtering that:

   I) if it is known that the port is to receive frames from ENodes but not FCFs:

      i) discards FIP frames that are not addressed to FCFs; and

      ii) discards FCoE frames that are not addressed to FCFs; and

iii) discards FCoE frames that are transmitted from one FCF to another FCF.

II) if it is known that the port is to receive frames from FCFs addressed to ENodes but not other FCFs:

i) discards all FIP frames that are not transmitted from an FCF;

ii) discards all FIP frames that are addressed to another FCF;

iii) discards all FCoE frames that are not transmitted from an FCF; and

iv) discards all FCoE frames that are addressed to another FCF.

III) if it is known that the port may receive frames from ENodes and FCFs that may be addressed to both FCFs and ENodes:

i) discards all FIP frames that are not either sourced by or addressed to an FCF; and

i) discards all FCoE frames that are not either sourced by or addressed to an FCF.

This recommendation provides a more scalable and resilient alternative to recommendation 1 for bridge-to-bridge ports. However, the security characteristics of this recommendation are weaker than that of recommendation 1. Consideration should be given to which recommendation best meets the overall needs of a given deployment.

Annex C provides a description of a method (i.e., the use of Access Control Lists) to satisfy this recommendation.

3) Bridges should not perform any learning function based on the source address of a frame that was discarded by recommendation 1 or recommendation 2.

4) All Ethernet bridges transmitting FCoE and FIP traffic should ensure that any VLAN carrying FCoE or FIP traffic for a given Virtual Fabric is in an independent VLAN learning set relative to all other VLANs.

5) Bridge ports intended to specifically exclude ingress of FIP and FCoE traffic should implement ingress filtering that discards all frames with an Ethernet Type equal to FIP_TYPE or FCoE_TYPE.

6) Bridges intended to transport FIP and FCoE traffic should not discard FIP and FCoE frames due to congestion.

Ethernet bridges, unlike Fibre Channel switches, do not by default provide a flow control mechanism. Therefore, Ethernet bridges discard frames when congested. Such discarding of frames may result in significant performance degradation of Fibre Channel traffic. This may be avoided by deploying mechanisms to prevent this type of packet discard. Two possible mechanisms of accomplishing this are the use of the PAUSE mechanism (see IEEE 802.3-2008) or of the Priority-based Flow Control mechanism (see IEEE 802.1Qbb) within the Ethernet bridges.

D.5 ENode and FCF recommendations

All of the recommendations in this subclause, except for recommendation 12, appear as normative requirements in this standard, and in some cases, appropriate Ethernet standards. They are repeated here to highlight their applicability to Fabric security considerations.

1) ENodes discard all received frames with an Ethernet Type equal to FIP_TYPE except:

   I) those that contain a Destination MAC address equal to All-ENode-MACs; and

   II) those that contain a Destination MAC address that equals a source MAC address used in a FIP Discovery Solicitation from the ENode.

2) ENodes discard all received frames with an Ethernet Type equal to FCoE_TYPE that:

   I) contain a destination MAC address/destination N_Port_ID pair that was not assigned by an FCF to one of the VN_Ports on the ENode; or

   II) contain a source MAC address that does not match the MAC address of the FCF that assigned the corresponding VN_Port MAC address.

In the case of SPMAs, the MAC address assigned by an FCF refers to the MAC address approved by the FCF during the FIP FLOGI/FIP NPIV FDISC process.

Using SPMAs, it is possible that multiple VN_Ports are assigned the same MAC address by one or more FCFs. This recommendation ensures that the frame is addressed from a VF_Port to its corresponding VN_Port. FPMAs uniquely address all VN_Ports.

3) FCFs discard all frames received with an Ethernet Type equal to FCoE_TYPE that:

   I) contain a destination MAC address that does not match the MAC address of one of the FCF's VE_Ports or VF_Ports;

   II) contain the source MAC address that does not match the MAC addresses that the FCF has assigned to the corresponding VN_Port or was established for the corresponding VE_Port; or

   III) in the case of a VN_Port, contains a Fibre Channel source address that does not match the one assigned to the VN_Port by the FCF.

4) On transmission, VN_Ports construct all frames with:

   I) the destination MAC address set to the MAC address of the FCF that it successfully performed a FIP FLOGI or FIP NPIV FDISC with; and

   II) the source MAC address set to the MAC address assigned to the VN_Port by the FCF as a result of the FIP FLOGI or FIP NPIV FDISC.

5) On transmission, VF_Ports construct all frames with:

   I) the destination MAC address set to the MAC address of the VN_Port as assigned by the transmitting FCF during FIP FLOGI/FIP NPIV FDISC; and

   II) the source MAC address set to the MAC address of the VF_Port (i.e., that of the FCF).

6) On transmission, VE_Ports construct all frames with:

   I) the source MAC address of the transmitting VE_Port; and

   II) the destination MAC address of the remote VE_Port.

7) The MAC Client within a FCF does not deliver:

   I) to a VE_Port or VF_Port, any frame whose Ethernet Type is not equal to FCoE_TYPE; and

   II) to the FCoE controller, any frame whose Ethernet Type is not equal to FIP_TYPE; or

   III) alternatively, VE_Ports, VF_Ports, and FCoE Controllers discard all frames that do not contain an Ethernet Type of FCoE_TYPE, FCoE_TYPE, and FIP_TYPE, respectively.

8) FCF ports that implement multiple port types (i.e., VF_Port and VE_Port) do not use the same MAC address for different port types.

9) ENodes may choose to transmit a FIP FLOGI/FIP NPIV FDISC to any FCF(s).

10) While processing a FIP FLOGI or FIP NPIV FDISC, an FCF either rejects the request or ensures that the MAC address assigned to the requesting ENode:

   I) complies with local administrative policy; and

   II) in the case of FPMA, the 24 most significant bits contain the Fabric's FC-MAP and the 24 least significant bits equal that of the assigned Fibre Channel address identifier.

For FPMAs, the fact that the assigned MAC address contains a Fabric wide unique Fibre Channel address identifier provides assurance that the MAC address itself is unique Fabric wide.

11) FCFs may chose to create or not create VE_Ports with other FCFs based on local policy information (e.g., the MAC address of other FCFs).

12) All source MAC addresses used in FIP should be globally assigned (see IEEE 802-2001 for a description of globally assigned MAC addresses).

D.6 Additional threat isolation using FPMAs

There is a class of threats related to the misuse of MAC addresses assigned to VN_Ports. If multiple VN_Ports utilize the same MAC address (e.g., through mis-configuration or other network issues), catastrophic network failures may occur including undetected corruption of data. In addition, the use of MAC addresses assigned to VN_Ports by malicious stations provides a number of attack possibilities that include denial of service attacks and undetected data interception. In addition, threats exist related to using an Fibre Channel address identifier that is not associated with the MAC address assigned to a given VN_Port.

Fabric measures to prevent such attacks using SPMAs are generally not practical. Doing so requires that all bridges with edge ports have knowledge of all MAC addresses being used for VN_Ports. The standards do not provide a mechanism where this knowledge may be reasonably obtained. Furthermore, even if these MAC addresses were known, implementations would face challenging scalability issues. Consequently, protection against these failures and attacks are accomplished by other means (e.g., careful network configuration, enforcement of strict physical security measures). Specific recommendations for protection against these failures and attacks when using SPMAs are beyond the scope of this annex.

With FPMAs it is possible to test all MAC addresses assigned to VN_Ports, both now and in the future. This may be accomplished by specifically allowing the valid VN_Port MAC addresses on a given port, which is learned by examining FIP FLOGI/NPIV FDISC LS_ACC frames egressing the bridge port, and discarding all other frames in which the 24 most significant bits of the source MAC address

match the Fabric's FC-MAP. Furthermore, it is trivial to ensure the proper association between a VN_Port MAC address and its associated Fibre Channel address identifier (i.e., the 24 least significant bits of the VN_Port MAC address match the VN_Port Fibre Channel address identifier).

To obtain the additional security capability provided by FPMA, the following recommendations are provided in addition to those previously discussed:

1) Bridge ports, except those known to be connected to FCFs, those known to be connected to other bridge ports, and those that are to be explicitly prohibited from carrying FCoE/FIP traffic, should implement ingress filtering that discards all frames containing a source MAC address in which the 24 most significant bits match the FCoE Fabric's FC-MAP and in which the source MAC address does not match a valid VN_Port MAC address assigned by an FCF to a device connected to that port. This requirement applies regardless of the Ethernet Type.

The MAC addresses currently assigned to VN_Ports that are reached through the bridge port may be obtained by examining the FIP FLOGI/FIP NPIV FDISC LS_ACC, FIP LOGO, FIP Clear Virtual Links, and FIP Keep Alive frames.

2) Bridge ports known to be connected to other bridge ports should:

   I) if it is known that the bridge port is to receive frames from ENodes but not FCFs:

      i) discard FCoE frames where the 24 most significant bits of the source MAC address do not match the Fabric's FC-MAP.

   II) if it is known that the bridge port is to receive frames from FCFs addressed to ENodes but not other FCFs:

      i) discard all FCoE frames where the 24 most significant bits of the destination MAC address do not match the Fabric's FC-MAP.

   III) if it is known that the bridge port may receive frames from ENodes and FCFs that may be adressed to both FCFs and ENodes:

      i) discard all FCoE frames that do not contain either a source MAC address matching that of an FCF or a source MAC address where the 24 most significant bits match the Fabric's FC-MAP; and

      ii) discard all FCoE frames that do not contain either a destination MAC address matching that of an FCF or a destination MAC address where the 24 most significant bits match the Fabric's FC-MAP.

3) Bridges should not perform any address learning function based on the source MAC address of a frame that was discarded by recommendation 1 or 2.

4) Bridge ports intended to specifically exclude ingress FIP and FCoE frames should implement ingress filtering that discards all frames with a source MAC address where the 24 most significant bits match the FCoE Fabric's FC-MAP.

This recommendation prevents a malicious host from injecting a packet utilizing a victim's source MAC address in an attempt to alter the bridge learning tables such that it may intercept the data destined to the victim.

5) On reception, VN_Ports verify that the destination Fibre Channel address identifier matches the 24 least significant bits of the destination MAC address.

6) On reception, VF_Ports verify that the source Fibre Channel address identifier matches the 24 least significant bits of the source MAC address.

7) On transmission, VN_Ports construct all frames such that the source Fibre Channel address identifier matches the 24 least significant bits of the source MAC address.

8) On transmission, VF_Ports construct all frames such that the destination Fibre Channel address identifier matches the 24 least significant bits of the destination MAC address.

Annex E: FCoE MIB Definition (Normative)

E.1 FCoE MIB definition

```
T11-FCoE-MIB DEFINITIONS ::= BEGIN

--
--  The FCoE MIB Module
--


IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE,
    Unsigned32, enterprises        FROM SNMPv2-SMI         -- [RFC2578]
    MODULE-COMPLIANCE,
    OBJECT-GROUP                   FROM SNMPv2-CONF        -- [RFC2580]
    RowStatus, TimeStamp,
    TruthValue, MacAddress,
    TEXTUAL-CONVENTION             FROM SNMPv2-TC          -- [RFC2578]
    SnmpAdminString                FROM SNMP-FRAMEWORK-MIB -- [RFC3411]
    InterfaceIndex,
    InterfaceIndexOrZero           FROM IF-MIB             -- [RFC2863]
    VlanIndex                      FROM Q-BRIDGE-MIB       -- [RFC4363]
    T11FabricIndex                 FROM T11-TC-MIB         -- [RFC4439]
    fcmInstanceIndex,
    fcmSwitchIndex                 FROM FC-MGMT-MIB;       -- [RFC4044]

t11 OBJECT IDENTIFIER ::= { enterprises 33317 }
t11MIBs OBJECT IDENTIFIER ::= { t11 1 }

t11FCoEMIB MODULE-IDENTITY
    LAST-UPDATED    "200905210000Z"
    ORGANIZATION    "INCITS Technical Committee T11
                     http://www.t11.org/"
    CONTACT-INFO
            "Claudio DeSanti cds@cisco.com
            Joe Pelissier jopeliss@cisco.com
            David Peterson dpeterso@brocade.com"
    DESCRIPTION
        "This MIB module is for configuring and monitoring Fibre Channel
         over Ethernet (FCoE) related entities.  This MIB defines a
         Virtual FC (VFC) Interface as an object that represents either
         a VF_Port or a VE_Port on a FCoE Forwarder (FCF).  VFC
         interfaces can be created either statically (by management
         request) or dynamically (at the time of FIP based FLOGI or ELP
         request).

         Other terminologies used in this MIB are defined by the
         T11 FCoE standard, as defined in the FC-BB-5 specification.
         See www.t11.org for more information.

         This MIB also supports configuration of the following objects:
         - Mapping of FCoE VLAN-ID used to carry traffic for a Fabric
         - FC-MAP value used by the FCF operating in FPMA mode
```

```
          - FIP snooping related objects"

    REVISION        "200905210000Z"
    DESCRIPTION
        "Initial version of this MIB module."
    ::= { t11MIBs 1 }


t11FCoEMIBObjects  OBJECT IDENTIFIER
    ::= { t11FCoEMIB 1 }

t11FCoEMIBConformance  OBJECT IDENTIFIER
    ::= { t11FCoEMIB 2 }

t11FCoEConfig  OBJECT IDENTIFIER
    ::= { t11FCoEMIBObjects 1 }

t11FCoEFIPSnoopingObjects OBJECT IDENTIFIER
    ::= { t11FCoEMIBObjects 2 }

-- Textual Conventions

T11FCoEVfcBindType ::= TEXTUAL-CONVENTION
    STATUS          current
    DESCRIPTION
        "Defines the different methods to identify (or bind to)
            - the ENode associated with a particular VF_Port VFC
            - the remote-FCF associated with a particular VE_Port VFC

            interfaceIndex(1)  - This type is used only when an ENode
                or remote-FCF is directly connected to the local FCF
                via one of the local Ethernet interfaces, in which
                case the value contains the ifIndex of said Ethernet
                interface.

            macAddress(2)  - This type is used when an ENode or
                remote-FCF is reachable from the local FCF over a
                (Layer-2) Ethernet network.  A FIP frame from an
                ENode or remote-FCF is associated to a VFC only if
                the frame's source MAC address is the same as the
                MAC Address bound on that VFC."
    SYNTAX  INTEGER {
                interfaceIndex(1),
                macAddress(2)
            }

--
-- FCoE Globals per Switch in a Fibre Channel Management Instance
-- Fibre Channel Management instance is defined in [RFC4044] as a
-- separable managed instance of Fibre Channel functionality.
-- Fibre Channel functionality may be grouped into Fibre
-- Channel management instances in whatever way is most
-- convenient for the implementation(s).
--
```

```
-- RFC4044 also defines the fcmSwitchTable as a table of
-- information about Fibre Channel switches which are managed
-- by Fibre Channel management instances.  Each Fibre Channel
-- management instance can manage one or more Fibre Channel
-- switches.  The Switch Index, fcmSwitchIndex, is
-- IMPORTed from the FC-MGMT-MIB as the index value
-- to uniquely identify a Fibre Channel switch amongst
-- those (one or more) managed by the same Fibre
-- Channel management instance.
-- In this MIB, the same fcmSwitchIndex is used to
-- identify each FCF and to distinguish it from other
-- FCFs and from other Fibre Channel switches.
--
--
t11FCoECfgTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF T11fcoeCfgEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This table facilitates configuration of FCoE parameters
         on a per Fibre Channel management instance."
    ::= { t11FCoEConfig 1 }

t11FCoECfgEntry OBJECT-TYPE
    SYNTAX          T11fcoeCfgEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "There is one entry in this table for each
         Fibre Channel management instance."
    INDEX           {
                        fcmInstanceIndex,
                        fcmSwitchIndex
                    }
    ::= { t11FCoECfgTable 1 }

T11fcoeCfgEntry ::= SEQUENCE {
        t11FCoECfgFcmap                  OCTET STRING,
        t11FCoECfgDynamicVfcCreation     TruthValue,
        t11FCoECfgDefaultFCFPriority     Unsigned32,
        t11FCoECfgDATov                  Unsigned32,
        t11FCoECfgAddressingMode         INTEGER
}

t11FCoECfgFcmap OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (3))
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "This object configures the FC-MAP value used by the FCF
         when operating in FPMA mode. The default value is 0EFC00h,
         as written in the standard."
    REFERENCE
        "Fibre Channel - Backbone - 5 (FC-BB-5),
```

```
        section 7.6 and table 47"
    DEFVAL { '0EFC00'h }
    ::= { t11FCoECfgEntry 1 }

t11FCoECfgDynamicVfcCreation OBJECT-TYPE
    SYNTAX          TruthValue
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "This object is set to 'true' to enable, or 'false' to
         disable, the dynamic creation of VFC interfaces on this FCF.
         When set to 'true', VFC interfaces are dynamically created
         as and when a FIP-based FLOGI or ELP request is received."
    DEFVAL { false }
    ::= { t11FCoECfgEntry 2 }

t11FCoECfgDefaultFCFPriority OBJECT-TYPE
    SYNTAX          Unsigned32 (0..255)
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "The FIP priority value advertised by the FCF to ENodes by
         default.  t11FCoEStaticVfcFCFPriority configured for a VFC
         interface overrides this setting for the ENode associated
         with the VFC."
    DEFVAL { 128 }
    ::= { t11FCoECfgEntry 3 }

t11FCoECfgDATov OBJECT-TYPE
    SYNTAX          Unsigned32 (1..60)
    UNITS           "seconds"
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "The Discovery_Advertisement_Timeout value configured for
         the FCF. This is used as the timeout value in seconds by
         the FCF to send periodic Discovery Advertisements."
    DEFVAL { 5 }
    ::= { t11FCoECfgEntry 4 }

t11FCoECfgAddressingMode OBJECT-TYPE
    SYNTAX          INTEGER {
                        fpma(1),
                        spma(2),
                        fpmaAndSpma(3)
                    }
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "Addressing mode(s) supported by the FCF. Implementations
         should fail SetRequests for unsupported modes."
    ::= { t11FCoECfgEntry 5 }
```

```
-- FCoE per VLAN configuration

t11FCoEVLANTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF T11fcoeVLANEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "In fabrics in which VLANs are deployed, this table facilitates
         configuration of VLAN and Virtual Fabric associations in an
         FCoE network. In such fabrics, FCoE forwarding for a fabric
         is over a VLAN in a (Layer-2) Ethernet network.  That is,
         reachability between the ENode/remote-FCF and an FCF for a
         given fabric is determined by the reachability provided by the
         Ethernet network on the corresponding VLAN.

         An active entry in this table indicates which VLAN is used
         to transport FCoE traffic for a particular Virtual Fabric.
         If VLANs are not deployed or not enabled, entries in this table
         are ignored by the bridge.

         Some implmentations may allow traffic from only one Virtual
         Fabric to be transported over a given VLAN. Such
         implementations should prevent multiple entries with the same
         VLAN-ID from being created in this table.

         Modifying existing VLAN-Virtual Fabric associations is not
         possible.  The specific row must first be deleted and then
         a new one created."
    ::= { t11FCoEConfig 2 }

t11FCoEVLANEntry OBJECT-TYPE
    SYNTAX          T11fcoeVLANEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "There is one entry in this table for each VLAN that is
         designated to transport FCoE traffic for a given Virtual
         Fabric."
    INDEX           {
                        fcmInstanceIndex,
                        fcmSwitchIndex,
                        t11FCoEVLANIndex,
                        t11FCoEFabricIndex
                    }
    ::= { t11FCoEVLANTable 1 }

T11fcoeVLANEntry ::= SEQUENCE {
        t11FCoEVLANIndex              VlanIndex,
        t11FCoEFabricIndex            T11FabricIndex,
        t11FCoEVLANOperState          INTEGER,
        t11FCoEVLANRowStatus          RowStatus
}

t11FCoEVLANIndex OBJECT-TYPE
```

```
    SYNTAX          VlanIndex
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This object identifies the VLAN-ID that the FCoE FCF function
         is being enabled for."
    ::= { t11FCoEVLANEntry 1 }

t11FCoEFabricIndex OBJECT-TYPE
    SYNTAX          T11FabricIndex
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This object identifies the Fabric Index of the Virtual Fabric
         traffic which is to be transported over the VLAN identified
         by t11FCoEVLANIndex."
    ::= { t11FCoEVLANEntry 2 }

t11FCoEVLANOperState OBJECT-TYPE
    SYNTAX          INTEGER {
                        up(1),
                        down(2)
                    }
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "Operational state of this VLAN-Virtual Fabric association
         entry. The 'up' state is achieved when both the Virtual
         Fabric and VLAN are valid."
    ::= { t11FCoEVLANEntry 3 }

t11FCoEVLANRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row. The RowStatus becomes active
        on successful creation of an entry."
    ::= { t11FCoEVLANEntry 4 }


-- Static Virtual FC interface Table

t11FCoEStaticVfcTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF T11fcoeStaticVfcEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This table facilitates the creation and deletion of static
        VFC interfaces. While VFCs can be dynamically created based on
        FIP FLOGI/ELP requests, operators may want to associate certain
        pre-configured policy for a particular ENode or a remote-FCF.
        In such cases static VFC creation becomes necessary. In addition
        to being creating, a static VFC also needs to be associated to
```

```
        an ENode or remote-FCF. The VFC binding provides such an
        associaton.  The binding does not need to be specified when
        the row for a VFC is created, but may be specified later."
    ::= { t11FCoEConfig 3 }

t11FCoEStaticVfcEntry OBJECT-TYPE
    SYNTAX          T11fcoeStaticVfcEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "There is one entry in this table for each statically created
         VFC Interface."
    INDEX           {
                        fcmInstanceIndex,
                        fcmSwitchIndex,
                        t11FCoEStaticVfcIndex
                    }
    ::= { t11FCoEStaticVfcTable 1 }

T11fcoeStaticVfcEntry ::= SEQUENCE {
        t11FCoEStaticVfcIndex             Unsigned32,
        t11FCoEStaticVfcFCFPriority       Unsigned32,
        t11FCoEStaticVfcBindType          T11FCoEVfcBindType,
        t11FCoEStaticVfcBindPortWWN       OCTET STRING,
        t11FCoEStaticVfcBindIfIndex       InterfaceIndexOrZero,
        t11FCoEStaticVfcBindMACAddress    MacAddress,
        t11FCoEStaticVfcBindVLANIndex     VlanIndex,
        t11FCoEStaticVfcIfIndex           InterfaceIndex,
        t11FCoEStaticVfcCreationTime      TimeStamp,
        t11FCoEStaticVfcFailureCause      SnmpAdminString,
        t11FCoEStaticVfcRowStatus         RowStatus
}

t11FCoEStaticVfcIndex OBJECT-TYPE
    SYNTAX          Unsigned32 (1..65535)
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This index uniquely identifies a static VFC entry in this
         table."
    ::= { t11FCoEStaticVfcEntry 1 }

t11FCoEStaticVfcFCFPriority OBJECT-TYPE
    SYNTAX          Unsigned32 (0..255)
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "If this VFC is for a VF_Port this object is used to configure
         FCF priority to be advertised to the ENode associated with the
         VFC."
   DEFVAL { 128 }
    ::= { t11FCoEStaticVfcEntry 2 }

t11FCoEStaticVfcBindType OBJECT-TYPE
```

```
    SYNTAX          T11FCoEVfcBindType
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The mechanism to identify the ENode associated with this VFC
         if it is of type VF_Port or to identify the remote-FCF
         associated with this VFC if it is of type VE_Port."
    ::= { t11FCoEStaticVfcEntry 3 }

t11FCoEStaticVfcBindPortWWN OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (16))
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "This object contains the WWN of the remote port."
    ::= { t11FCoEStaticVfcEntry 4 }

t11FCoEStaticVfcBindIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndexOrZero
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "This object is applicable only when the local FCF is
         directly connected to an ENode or remote-FCF over a
         specific Ethernet interface, in which case this object
         contains the ifIndex of said Ethernet interface.
         If the ENode or remote-FCF is not directly connected
         to the FCF, this value of this object is zero."
    ::= { t11FCoEStaticVfcEntry 5 }

t11FCoEStaticVfcBindMACAddress OBJECT-TYPE
    SYNTAX          MacAddress
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "This object is applicable when the ENode or remote-FCF to
         which the local FCF is connected is identified by a MAC
         address.  A FIP frame from a ENode or remote-FCF is
         associated with this VFC only if the source MAC address
         in the frame is the same as the value of this object."
    ::= { t11FCoEStaticVfcEntry 6 }

t11FCoEStaticVfcBindVLANIndex OBJECT-TYPE
    SYNTAX          VlanIndex
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "This object is applicable when the ENode or remote-FCF to
         which the local FCF is connected is identified by a MAC
         address.  A FIP frame from a ENode or remote-FCF is
         associated with this VFC only if the source MAC address
         in the frame is the same as the value of this object."
    ::= { t11FCoEStaticVfcEntry 7 }
```

```
t11FCoEStaticVfcIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndex
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The ifIndex of this Virtual FC interface."
    ::= { t11FCoEStaticVfcEntry 8 }

t11FCoEStaticVfcCreationTime OBJECT-TYPE
    SYNTAX          TimeStamp
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The timestamp of this entry's creation time."
    ::= { t11FCoEStaticVfcEntry 9 }

t11FCoEStaticVfcFailureCause OBJECT-TYPE
    SYNTAX          SnmpAdminString
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The cause of failure for the last bind operation. This
         object will be zero length if and only if the bind is
         successful."
    ::= { t11FCoEStaticVfcEntry 10 }

t11FCoEStaticVfcRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row.  The RowStatus becomes
         active on successful creation of a VFC.  The VFC does not
         need to be bound for the row to be active, but the VFC must
         be bound before becoming operational."
    ::= { t11FCoEStaticVfcEntry 11 }

--------------------

-- Dynamic Virtual FC interface Table

t11FCoEDynamicVfcTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF T11fcoeDynamicVfcEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This table facilitates obtaining the status of dynamic
        VFC interfaces. Dynamic VFCs are created based on
        FIP FLOGI/ELP requests. In such cases dynamic VFC creation
        becomes necessary. In addition to being created, a dynamic
        VFC also needs to be associated to an ENode or remote-FCF.
        The VFC binding provides such an associaton."
    ::= { t11FCoEConfig 4 }
```

```
t11FCoEDynamicVfcEntry OBJECT-TYPE
    SYNTAX          T11fcoeDynamicVfcEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "There is one entry in this table for each dynamically created
         VFC Interface."
    INDEX           {
                        fcmInstanceIndex,
                        fcmSwitchIndex,
                        t11FCoEDynamicVfcIndex
                    }
    ::= { t11FCoEDynamicVfcTable 1 }

T11fcoeDynamicVfcEntry ::= SEQUENCE {
        t11FCoEDynamicVfcIndex             Unsigned32,
        t11FCoEDynamicVfcBindPortWWN       OCTET STRING,
        t11FCoEDynamicVfcBindIfIndex       InterfaceIndexOrZero,
        t11FCoEDynamicVfcBindMACAddress    MacAddress,
        t11FCoEDynamicVfcBindVLANIndex     VlanIndex,
        t11FCoEDynamicVfcIfIndex           InterfaceIndex,
        t11FCoEDynamicVfcCreationTime      TimeStamp,
        t11FCoEDynamicVfcRowStatus         RowStatus
}

t11FCoEDynamicVfcIndex OBJECT-TYPE
    SYNTAX          Unsigned32 (1..65535)
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This index uniquely identifies a dynamic VFC entry in this
         table."
    ::= { t11FCoEDynamicVfcEntry 1 }

t11FCoEDynamicVfcBindPortWWN OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (16))
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "This object contains the WWN of the remote port."
    ::= { t11FCoEDynamicVfcEntry 2 }

t11FCoEDynamicVfcBindIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndexOrZero
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "This object is applicable only when the local FCF is
         directly connected to an ENode or remote-FCF over a
         specific Ethernet interface, in which case this object
         contains the ifIndex of said Ethernet interface.
         If the ENode or remote-FCF is not directly connected
         to the FCF, this value of this object is zero."
    ::= { t11FCoEDynamicVfcEntry 3 }
```

```
t11FCoEDynamicVfcBindMACAddress OBJECT-TYPE
    SYNTAX          MacAddress
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "This object is applicable when the ENode or remote-FCF to
         which the local FCF is connected is identified by a MAC
         address.  A FIP frame from a ENode or remote-FCF is
         associated with this VFC only if the source MAC address
         in the frame is the same as the value of this object."
    ::= { t11FCoEDynamicVfcEntry 4 }

t11FCoEDynamicVfcBindVLANIndex OBJECT-TYPE
    SYNTAX          VlanIndex
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "A FIP frame from a ENode or remote-FCF is
         associated with this VFC only if it arrives on the
         VLAN identified by this object."
    ::= { t11FCoEDynamicVfcEntry 5 }

t11FCoEDynamicVfcIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndex
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The ifIndex of this Virtual FC interface."
    ::= { t11FCoEDynamicVfcEntry 6 }

t11FCoEDynamicVfcCreationTime OBJECT-TYPE
    SYNTAX          TimeStamp
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The timestamp of this entry's creation time."
    ::= { t11FCoEDynamicVfcEntry 7 }

t11FCoEDynamicVfcRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row.  The RowStatus becomes
         active on successful creation of a VFC.  The VFC does not
         need to be bound for the row to be active, but the VFC must
         be bound before becoming operational."
    ::= { t11FCoEDynamicVfcEntry 8 }

---
---
--- VNPort Interface Table
---
```

```
t11FCoEVNPortTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF T11fcoeVNPortEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This table facilitates obtaining the status of VN_Port
        interfaces."
    ::= { t11FCoEConfig 5 }

t11FCoEVNPortEntry OBJECT-TYPE
    SYNTAX          T11fcoeVNPortEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "There is one entry in this table for each
         VN_Port Interface."
    INDEX           {
                        t11FCoEVNPortIndex
                    }
    ::= { t11FCoEVNPortTable 1 }

T11fcoeVNPortEntry ::= SEQUENCE {
        t11FCoEVNPortIndex                Unsigned32,
        t11FCoEVNPortAddressingMode       INTEGER,
        t11FCoEVNPortFcmap                OCTET STRING,
        t11FCoEVNPortLocalMACAddress      MacAddress,
        t11FCoEVNPortLocalPortWWN         OCTET STRING,
        t11FCoEVNPortBindMACAddress       MacAddress,
        t11FCoEVNPortBindPortWWN          OCTET STRING,
        t11FCoEVNPortBindVLANIndex        VlanIndex,
        t11FCoEVNPortIfIndex              InterfaceIndex,
        t11FCoEVNPortDATov                Unsigned32,
        t11FCoEVNPortCreationTime         TimeStamp,
        t11FCoEVNPortRowStatus            RowStatus
}

t11FCoEVNPortIndex OBJECT-TYPE
    SYNTAX          Unsigned32 (1..65535)
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This index uniquely identifies a VN_Port entry in this
         table."
    ::= { t11FCoEVNPortEntry 1 }

t11FCoEVNPortAddressingMode OBJECT-TYPE
    SYNTAX          INTEGER {
                        fpma(1),
                        spma(2)
                    }
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
```

```
        "Addressing mode in use by the VN_Port."
    ::= { t11FCoEVNPortEntry 2 }

t11FCoEVNPortFcmap OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (3))
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "FC-MAP value in use by the VN_Port.  Applies only if
        FPMA is in use, i.e. t11FCoEVNPortAddressingMode==fpma."
    REFERENCE
        "Fibre Channel - Backbone - 5 (FC-BB-5),
         section 7.6 and table 47"
    ::= { t11FCoEVNPortEntry 3 }

t11FCoEVNPortLocalMACAddress OBJECT-TYPE
    SYNTAX          MacAddress
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The MAC address being used by this VN_Port."
    ::= { t11FCoEVNPortEntry 4 }

t11FCoEVNPortLocalPortWWN OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (16))
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "This object contains the WWN of the local port."
    ::= { t11FCoEVNPortEntry 5 }

t11FCoEVNPortBindMACAddress OBJECT-TYPE
    SYNTAX          MacAddress
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The MAC address of the remote VF_Port."
    ::= { t11FCoEVNPortEntry 6 }

t11FCoEVNPortBindPortWWN OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (16))
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "This object contains the WWN of the remote VF_Port port."
    ::= { t11FCoEVNPortEntry 7 }

t11FCoEVNPortBindVLANIndex OBJECT-TYPE
    SYNTAX          VlanIndex
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The VLAN in use by this VN_Port / VF_Port port association."
    ::= { t11FCoEVNPortEntry 8 }
```

```
t11FCoEVNPortIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndex
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The ifIndex of this VN_Port interface."
    ::= { t11FCoEVNPortEntry 9 }

t11FCoEVNPortDATov OBJECT-TYPE
    SYNTAX          Unsigned32 (1..60)
    UNITS           "seconds"
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The Discovery_Advertisement_Timeout value in use
         by this VN_Port."
    ::= { t11FCoEVNPortEntry 10 }

t11FCoEVNPortCreationTime OBJECT-TYPE
    SYNTAX          TimeStamp
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The timestamp of this entry's creation time."
    ::= { t11FCoEVNPortEntry 11 }

t11FCoEVNPortRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row.  The RowStatus becomes
         active on successful creation of a VN_Port."
    ::= { t11FCoEVNPortEntry 12 }

--
-- FIP Snooping related objects used to configure FIP Snooping
-- on a (FIP-aware) Ethernet Bridge
--
t11FCoEFIPSnoopingEnable OBJECT-TYPE
    SYNTAX          INTEGER  {
                        enable(1),
                        disable(2)
                    }
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "This object is used to enable or disable FIP Snooping on an
         Ethernet Bridge."
    ::= { t11FCoEFIPSnoopingObjects 1 }

t11FCoEFIPSnoopingFcmap OBJECT-TYPE
    SYNTAX          OCTET STRING (SIZE (3))
```

```
    MAX-ACCESS      read-write
    STATUS          current
    DESCRIPTION
        "This object configures the FC-MAP value associated with the
         FIP snooping Ethernet Bridge."
    DEFVAL { '0EFC00'h }
    ::= { t11FCoEFIPSnoopingObjects 2 }

t11FCoEEnodeIfTable OBJECT-TYPE
    SYNTAX     SEQUENCE OF T11fcoeEnodeIfEntry
    MAX-ACCESS not-accessible
    STATUS     current
    DESCRIPTION
        "The conceptual table containing the interfaces on an
         Ethernet Bridge that are directly connected to ENodes.
         This table assumes that any necessary Access Control
         configuration on the Bridge is either pre-configured or
         or will get automatically configured based on FIP login
         requests and responses."
    ::= { t11FCoEFIPSnoopingObjects 3 }

t11FCoEEnodeIfEntry OBJECT-TYPE
    SYNTAX     T11fcoeEnodeIfEntry
    MAX-ACCESS not-accessible
    STATUS     current
    DESCRIPTION
        "An entry (conceptual row) in the t11FCoEEnodeIfTable,
         representing an interface directly connected to an ENode
         on the Ethernet Bridge."
    INDEX           { t11FCoEEnodeIfIfIndex }
    ::= { t11FCoEEnodeIfTable 1 }

T11fcoeEnodeIfEntry ::= SEQUENCE {
        t11FCoEEnodeIfIfIndex       InterfaceIndex,
        t11FCoEEnodeIfRowStatus     RowStatus
}

t11FCoEEnodeIfIfIndex OBJECT-TYPE
    SYNTAX          InterfaceIndex
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "The ifIndex of the interface on the Ethernet Bridge connected
         to an Enode."
    ::= { t11FCoEEnodeIfEntry 1 }

t11FCoEEnodeIfRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row."
    ::= { t11FCoEEnodeIfEntry 2 }
```

```
t11FCoEFIPSnoopingTable OBJECT-TYPE
    SYNTAX     SEQUENCE OF T11fcoeFIPSnoopingEntry
    MAX-ACCESS not-accessible
    STATUS     current
    DESCRIPTION
        "The conceptual table containing VLANs for which FIP snooping
         is individually enabled.  If t11FCoEFIPSnoopingEnable is
         set to enable, then FIP snooping is enabled for all VLANs and
         this table is ignored.  If t11FCoEFIPSnoopingEnable is set
         to disable, then FIP snooping is enabled for any VLAN that
         appears in this table.  If t11FCoEFIPSnoopingEnable is set
         to disable and this table is empty, then snooping is
         disabled."
    ::= { t11FCoEFIPSnoopingObjects 4 }

t11FCoEFIPSnoopingEntry OBJECT-TYPE
    SYNTAX     T11fcoeFIPSnoopingEntry
    MAX-ACCESS not-accessible
    STATUS     current
    DESCRIPTION
        "An entry (conceptual row) in the t11FCoEFIPSnoopingTable,
         representing a VLAN on which FIP Snooping is to be performed."
    INDEX           { t11FCoEFIPSnoopingVLANIndex }
    ::= { t11FCoEFIPSnoopingTable 1 }

T11fcoeFIPSnoopingEntry ::= SEQUENCE {
        t11FCoEFIPSnoopingVLANIndex VlanIndex,
        t11FCoEFIPSnoopingRowStatus RowStatus
}

t11FCoEFIPSnoopingVLANIndex OBJECT-TYPE
    SYNTAX          VlanIndex
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        "This object identifies the VLAN-ID that the FCoE FIP Snooping
         function is being enabled for."
    ::= { t11FCoEFIPSnoopingEntry 1 }

t11FCoEFIPSnoopingRowStatus OBJECT-TYPE
    SYNTAX          RowStatus
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        "The status of this conceptual row."
    ::= { t11FCoEFIPSnoopingEntry 2 }

-- Conformance

t11FCoEMIBCompliances  OBJECT IDENTIFIER
    ::= { t11FCoEMIBConformance 1 }

t11FCoEMIBGroups  OBJECT IDENTIFIER
```

```
    ::= { t11FCoEMIBConformance 2 }

t11FCoEMIBCompliance MODULE-COMPLIANCE
    STATUS          current
    DESCRIPTION
        "The compliance statement for entities which implement the
         T11-FCOE-MIB module."
    MODULE          -- this module

    GROUP           t11FCoECfgConformanceObjects
    DESCRIPTION
       "This group is mandatory for FCFs"

    GROUP           t11FCoEVLANConformanceObjects
    DESCRIPTION
       "This group is mandatory for FCFs"

    GROUP           t11FCoEStaticVfcConformanceObjects
    DESCRIPTION
       "This group is mandatory for FCFs"

    GROUP           t11FCoEDynamicVfcConformanceObjects
    DESCRIPTION
       "This group is mandatory for FCFs"

    GROUP           t11FCoEFIPSnoopingConformanceObjects
    DESCRIPTION
        "This group is mandatory for Ethernet Bridges which support
         FIP Snooping."

    GROUP           t11FCoEFIPSnoopingVLANConformanceObjects
    DESCRIPTION
        "This group is mandatory for Ethernet Bridges which support
         FIP Snooping."

    GROUP           t11FCoEVNPortConformanceObjects
    DESCRIPTION
       "This group is mandatory for VN_Ports"

    GROUP           t11FCoEEnodeIfObjects
    DESCRIPTION
        "This group is mandatory for Ethernet Bridges which need
         configuration information indicating which of their
         interfaces are directly attached to ENodes."

    OBJECT          t11FCoECfgFcmap
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not required."

    OBJECT          t11FCoECfgDynamicVfcCreation
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not required."
```

```
OBJECT          t11FCoECfgDefaultFCFPriority
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoECfgDATov
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoECfgAddressingMode
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEVLANRowStatus
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcFCFPriority
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcBindType
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcBindPortWWN
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcBindIfIndex
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcBindMACAddress
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcBindVLANIndex
MIN-ACCESS      read-only
DESCRIPTION
    "Support for write access is not required."

OBJECT          t11FCoEStaticVfcRowStatus
MIN-ACCESS      read-only
DESCRIPTION
```

```
        "Support for write access is not required."


    OBJECT          t11FCoEDynamicVfcBindPortWWN
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not permitted."

    OBJECT          t11FCoEDynamicVfcBindIfIndex
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not permitted."

    OBJECT          t11FCoEDynamicVfcBindMACAddress
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not permitted."

    OBJECT          t11FCoEDynamicVfcBindVLANIndex
    MIN-ACCESS  read-only
    DESCRIPTION
        "Support for write access is not permitted."

    OBJECT          t11FCoEDynamicVfcRowStatus
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not permitted."

    OBJECT          t11FCoEFIPSnoopingEnable
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not required."

    OBJECT          t11FCoEFIPSnoopingFcmap
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not required."

    OBJECT          t11FCoEEnodeIfRowStatus
    MIN-ACCESS      read-only
    DESCRIPTION
        "Support for write access is not required."

    ::= { t11FCoEMIBCompliances 1 }

-- Units of Conformance

t11FCoECfgConformanceObjects OBJECT-GROUP
    OBJECTS         {
                        t11FCoECfgFcmap,
                        t11FCoECfgDynamicVfcCreation,
                        t11FCoECfgDefaultFCFPriority,
                        t11FCoECfgDATov,
                        t11FCoECfgAddressingMode
```

```
                    }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to all implementations
        of FCoE."
    ::= { t11FCoEMIBGroups 1 }

t11FCoEVLANConformanceObjects OBJECT-GROUP
    OBJECTS         {
                        t11FCoEVLANOperState,
                        t11FCoEVLANRowStatus
                    }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        FCoE VLANs."
    ::= { t11FCoEMIBGroups 2 }

t11FCoEStaticVfcConformanceObjects OBJECT-GROUP
    OBJECTS         {
                        t11FCoEStaticVfcFCFPriority,
                        t11FCoEStaticVfcBindType,
                        t11FCoEStaticVfcBindPortWWN,
                        t11FCoEStaticVfcBindIfIndex,
                        t11FCoEStaticVfcBindMACAddress,
                        t11FCoEStaticVfcBindVLANIndex,
                        t11FCoEStaticVfcIfIndex,
                        t11FCoEStaticVfcCreationTime,
                        t11FCoEStaticVfcFailureCause,
                        t11FCoEStaticVfcRowStatus
                    }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        static VFC interfaces."
    ::= { t11FCoEMIBGroups 3 }

t11FCoEDynamicVfcConformanceObjects OBJECT-GROUP
    OBJECTS         {
                        t11FCoEDynamicVfcBindPortWWN,
                        t11FCoEDynamicVfcBindIfIndex,
                        t11FCoEDynamicVfcBindMACAddress,
                        t11FCoEDynamicVfcBindVLANIndex,
                        t11FCoEDynamicVfcIfIndex,
                        t11FCoEDynamicVfcCreationTime,
                        t11FCoEDynamicVfcRowStatus
                    }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        virtual VFC interfaces."
    ::= { t11FCoEMIBGroups 4 }

t11FCoEFIPSnoopingConformanceObjects OBJECT-GROUP
```

```
    OBJECTS         {
                        t11FCoEFIPSnoopingEnable,
                        t11FCoEFIPSnoopingFcmap
                    }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        FIP Snooping."
    ::= { t11FCoEMIBGroups 5 }

t11FCoEEnodeIfObjects OBJECT-GROUP
    OBJECTS         { t11FCoEEnodeIfRowStatus }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        ENode interfaces on a FIP-snooping Ethernet Bridge."
    ::= { t11FCoEMIBGroups 6 }

t11FCoEFIPSnoopingVLANConformanceObjects OBJECT-GROUP
    OBJECTS         { t11FCoEFIPSnoopingRowStatus }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        FIP-snooping per VLAN on an Ethernet Bridge."
    ::= { t11FCoEMIBGroups 7 }

t11FCoEVNPortConformanceObjects OBJECT-GROUP
    OBJECTS         {
                t11FCoEVNPortAddressingMode,
                t11FCoEVNPortFcmap,
                t11FCoEVNPortLocalMACAddress,
                t11FCoEVNPortLocalPortWWN,
                t11FCoEVNPortBindMACAddress,
                t11FCoEVNPortBindPortWWN,
                t11FCoEVNPortBindVLANIndex,
                t11FCoEVNPortIfIndex,
                t11FCoEVNPortDATov,
                t11FCoEVNPortCreationTime,
                t11FCoEVNPortRowStatus
            }
    STATUS          current
    DESCRIPTION
        "A collection of objects related to implementation of
        VN_Ports."
    ::= { t11FCoEMIBGroups 8 }

END
```

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

F.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.

F.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.

F.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.

F.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).

F.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).

F.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.

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

When transmitting a FLOGI ELS or FDISC ELS request, the ENode sets the DA to 0EFC00FFFFFEh 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.