

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

INCITS 547-201x  
T11/Project 547-D/Rev 1.05

# FIBRE CHANNEL

SWITCH FABRIC - 7  
(FC-SW-7)

REV 1.05

INCITS working draft proposed  
American National Standard  
for Information Technology

October 10, 2018

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. The T11 Technical Committee or anyone else 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.**

**Permission is granted to members of INCITS, its technical committees, and their associated task groups to reproduce this document for the purposes of INCITS standardization activities without further permission, provided this notice is included. All other rights are reserved. Any duplication of this document for commercial or for-profit use is strictly prohibited.**

Steven Wilson (T11 Chair)  
Broadcom Inc  
1320 Ridder Park Drive  
San Jose, CA 95131  
Voice: 408-433-8000  
steve.wilson@broadcom.com

Craig W. Carlson (T11 Vice Chair)  
Marvell Semiconductor  
12900 Whitewater Drive  
Minnetonka, MN 55343  
Voice: 952-687-2431  
craig.carlson@qlogic.com

Craig W. Carlson (T11.3 Chair)  
Marvell Semiconductor  
12900 Whitewater Drive  
Minnetonka, MN 55343  
Voice: 952-952-687-2431  
craig.carlson@qlogic.com

Craig W. Carlson (FC-SW-7 Chair)  
Marvell Semiconductor  
12900 Whitewater Drive  
Minnetonka, MN 55343  
Voice: 952-952-687-2431  
craig.carlson@qlogic.com

David Peterson (FC-SW-7 Editor)  
Broadcom Inc  
1230 Northland Drive  
Mendota Heights, MN 95120  
Voice: 408-433-8000  
david.peterson@broadcom.com

## **Revision History**

### **Rev 1.05**

2018-00094-v003 - Leafing Through the Fabric

### **Rev 1.04**

16-205v0 - VE Identification Server: FC-SW-7 Updates - missing value and editorial

### **Rev 1.03**

16-205v0 - VE Identification Server: FC-SW-7 Updates - fixup

### **Rev 1.02**

16-142v1 - EFCS text

16-205v0 - VE Identification Server: FC-SW-7 Updates

### **Rev 1.01**

16-016v1 - FC-SW-7: Application Services

16-057v0 - VE Identification Server: FC-SW-7 Additions

**Rev 1.00** - Initial draft standard based on ANSI/INCITS FC-SW-6 standard.

American National Standard  
for Information Technology

**Fibre Channel —  
Switch Fabric - 7 (FC-SW-7)**

Secretariat

**Information Technology Industry Council**

Approved (not yet approved)

**American National Standards Institute, Inc.**

**Abstract**

This standard describes the requirements for an interconnecting Fabric consisting of multiple Fabric Switch elements to support the ANSI/INCITS Fibre Channel - Framing and Signaling (FC-FS-5) and ANSI/INCITS Fibre Channel - Physical Interface (FC-PI-6) standards.

# 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 will in no circumstances give 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.

**CAUTION:** 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 and following calls for the identification of patents that may be required for the implementation of the standard, no such claims have been made. 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 licenses are not required to avoid infringement in the use of this standard.

Published by

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

Copyright © 201x by Information Technology Industry Council (ITI)  
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of ITI, 1250 Eye Street NW, Washington, DC 20005.

Printed in the United States of America

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

This standard describes the requirements for an interconnecting Fabric consisting of multiple Fabric Switch elements to support the INCITS Fibre Channel - Framing and Signaling - 5 (FC-FS-5) and INCITS Fibre Channel - Physical Interface - 6 (FC-PI-6) standards.

This standard was developed by Task Group T11 of Accredited Standards Committee INCITS during 2010-201x. The standards approval process started in 201x. 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:

Steven Wilson, Chair

Claudio DeSanti, Vice-Chair

Richard Johnson, Secretary

<b>Company</b>	<b>Name</b>
----------------	-------------

TBD.	
------	--

### **Introduction**

FC-SW-7 is one of the Fibre Channel family of standards. This family includes ANSI/INCITS FC-FS-5 and ANSI/INCITS FC-PI-6. ANSI/INCITS FC-GS-8, is a document related to Generic Fabric Services and is closely tied to FC-SW-7. ANSI/INCITS FC-BB-6 describes how Fabrics are extended over transports complementary to Fibre Channel. ANSI/INCITS FC-MI-3 and ANSI/INCITS FC-DA-2 describe interoperability profiles that assists in the interoperability of Switches. INCITS 332: 1999, FC-AL-2, specifies the arbitrated loop topology. ANSI/INCITS FC-SP-2 describes the Security requirements and protocols associated with Fibre Channel networks. ANSI/INCITS FC-IFR describes the requirements and protocols associated with Inter-Fabric Routing.

FC-SW-7 describes how Switches communicate and interact with one another to form a Fabric of Switches. Included are Fabric initialization and configuration, routing, Server communication, event distribution and repository exchanges (e.g., Zoning information).

Contents	Page
<b>1 Scope</b> .....	<b>1</b>
<b>2 Normative references</b> .....	<b>1</b>
2.1 Overview .....	1
2.2 Approved references .....	2
2.3 References under development .....	3
2.4 IETF references .....	3
<b>3 Definitions, conventions, abbreviations, acronyms, and symbols</b> .....	<b>4</b>
3.1 Definitions .....	4
3.2 Editorial conventions .....	11
3.3 State machine notation .....	12
3.4 Abbreviations, acronyms, and symbols .....	13
3.5 Definition of compliance terms .....	14
3.6 Keywords .....	14
3.7 T10 Vendor ID fields .....	15
3.8 Left-aligned ASCII data .....	15
<b>4 Structure and concepts</b> .....	<b>16</b>
4.1 Overview .....	16
4.2 E_Port operation .....	16
4.3 Fabric operation .....	16
4.4 Fabric definition .....	17
4.5 Switch .....	17
4.6 Switching characteristics .....	20
4.7 Distributed Switch and A_Port operations .....	20
4.7.1 Overview .....	20
4.7.2 A_Port operation .....	20
4.7.3 Controlling Switch functional model .....	21
4.7.4 FCDF functional model .....	22
4.8 Switch ports and Bridge ports .....	23
4.8.1 General characteristics .....	23
4.8.2 F_Port .....	24
4.8.3 FL_Port .....	24
4.8.4 E_Port .....	24
4.8.5 B_Port .....	24
4.8.6 A_Port .....	25
4.8.7 G_Ports and GL_Ports .....	25
4.8.8 PF_Port .....	25
4.8.9 VF_Port .....	25
4.8.10 PE_Port .....	25
4.8.11 VE_Port .....	25
4.8.12 PA_Port .....	25
4.8.13 VA_Port .....	25
4.9 Fabric addressing .....	25
4.10 Class F service .....	28
<b>5 Switch ports and Bridge ports</b> .....	<b>29</b>
5.1 Overview .....	29
5.2 Model elements .....	29
5.2.1 FC Transports .....	29
5.2.2 Switch Transport .....	29
5.2.3 Control Facilities .....	29
5.2.4 Link Services .....	29

5.3 F_Port operation	30
5.4 FL_Port operation	31
5.5 E_Port operation	33
5.6 B_Port operation	34
5.7 A_Port operation	35
5.8 Inter-Switch Link behavior	36
5.9 Class F service	38
5.9.1 Class F function	38
5.9.2 Class F rules	38
5.9.3 Class F frame format	39
5.9.4 Class F flow control	39
<b>6 Internal Link Services</b>	<b>41</b>
6.1 Switch Fabric Internal Link Services (SW_ILS)	41
6.2 Fabric SW_ILSs	43
6.2.1 Overview	43
6.2.2 Switch Fabric Internal Link Service Accept (SW_ACC)	43
6.2.3 Switch Fabric Internal Link Service Reject (SW_RJT)	43
6.2.4 Exchange Link Parameters (ELP)	47
6.2.4.1 ELP request	47
6.2.4.2 R_RDY flow control	52
6.2.4.3 VC_RDY flow control	53
6.2.4.4 ELP reply	54
6.2.5 Exchange Fabric Parameters (EFP)	55
6.2.6 Domain Identifier Assigned (DIA)	59
6.2.7 Request Domain_ID (RDI)	60
6.2.8 Hello (HLO)	63
6.2.8.1 HLO Overview	63
6.2.8.2 FSPF Header format	64
6.2.9 Link State Update (LSU)	65
6.2.9.1 LSU overview	65
6.2.9.2 Link State Record (LSR) format	66
6.2.9.3 Link State Header format	67
6.2.9.4 Link Descriptor format	69
6.2.10 Link State Acknowledgement (LSA)	69
6.2.11 Build Fabric (BF)	70
6.2.12 Reconfigure Fabric (RCF)	71
6.2.13 Inter-Switch Registered State Change Notification (SW_RSCN)	72
6.2.14 Distribute Registered Link Incident Records (DRLIR)	75
6.2.15 Merge Request (MR)	76
6.2.15.1 Merge Request payload	77
6.2.15.1.1 Merge Request payload in Basic Zoning	77
6.2.15.1.2 Merge Request payload in Enhanced Zoning	78
6.2.15.2 Merge Request reply	79
6.2.16 Acquire Change Authorization Request (ACA)	80
6.2.17 Release Change Authorization (RCA) request	82
6.2.18 Stage Fabric Configuration (SFC) request	83
6.2.18.1 SFC in Basic Zoning	85
6.2.18.2 SFC in Enhanced Zoning	85
6.2.18.2.1 Operation Request 'Activate Zone Set Enhanced'	86
6.2.18.2.2 Operation Request 'Deactivate Zone Set Enhanced'	86
6.2.18.2.3 Operation Request 'Distribute Zone Set Database'	86
6.2.18.2.4 Operation Request 'Activate Zone Set by Name'	87
6.2.18.2.5 Operation Request 'Set Zoning Policies'	87
6.2.19 Update Fabric Configuration (UFC) request	88



6.2.20 Check E_Port Connectivity (CEC)	89
6.2.21 Exchange Switch Capabilities	91
6.2.22 Exchange Switch Support (ESS)	94
6.2.22.1 ESS request payload	95
6.2.22.2 Interconnect Element Information Object	95
6.2.22.3 Capability Object	96
6.2.22.4 Service Specific Capability formats	96
6.2.22.4.1 Directory Server Capability	96
6.2.22.4.2 Fabric Controller Capability	97
6.2.22.4.3 ESS Fabric Configuration Server Capability Object	97
6.2.22.4.4 ESS Enhanced Zone Server Capability Object	98
6.2.22.4.5 Security Policy Server Capability Object	99
6.2.22.4.6 ESS Vendor Specific Capability Object	100
6.2.22.4.7 Domain Controller Capability Object	101
6.2.22.4.8 Event Server Capability Object	102
6.2.22.4.9 Switch Support Capability Object	102
6.2.22.4.10 Application Server Capability Object	102
6.2.22.4.11 Enhanced Fabric Configuration Server Capability Object	103
6.2.22.4.12 VE Identification Server Capability Object	104
6.2.22.5 ESS accept payload	104
6.2.23 Merge Request Resource Allocation (MRRA)	105
6.2.24 Switch Trace Route (STR)	106
6.2.24.1 Basic function	106
6.2.25 Exchange Virtual Fabrics Parameters (EVFP)	111
6.2.25.1 Basic function	111
6.2.25.2 EVFP_SYNC Message Payload	114
6.2.25.2.1 Overview	114
6.2.25.2.2 Tagging Administrative Status descriptor	116
6.2.25.2.3 Port VF_ID descriptor	117
6.2.25.2.4 Locally-Enabled VF_ID List descriptor	117
6.2.25.2.5 Vendor Specific descriptor	118
6.2.25.3 EVFP_COMMIT Message Payload	118
6.2.26 Enhanced Acquire Change Authorization request (EACA)	118
6.2.26.1 Commit Exchange Preamble	119
6.2.26.1.1 Transaction Identifier	119
6.2.26.1.2 Number of Switch Identifiers	120
6.2.26.1.3 Flags	120
6.2.26.1.4 ECS Switch List	120
6.2.27 Enhanced Stage Fabric Configuration (ESFC) request	122
6.2.28 Enhanced Update Fabric Configuration (EUFC) request	123
6.2.29 Enhanced Release Change Authorization (ERCA) request	123
6.2.30 Transfer Commit Ownership (TCO) request	124
6.3 Distributed Switch VA_Port SW_ILSs	125
6.3.1 Overview	125
6.3.2 VA_Port SW_ILS descriptors	127
6.3.2.1 Descriptor format	127
6.3.2.2 VN_Port Reachability descriptor	128
6.3.2.3 FLOGI/NPIV FDISC Parameters descriptor	128
6.3.2.4 VN_Port Unreachability descriptor	128
6.3.2.5 FCDF Reachability descriptor	129
6.3.2.6 Sequence Number descriptor	129
6.3.2.7 Controlling Switch Reachability descriptor	130
6.3.2.8 N_Port_IDs Reachability descriptor	130
6.3.2.9 Domain_IDs Reachability descriptor	132

6.3.2.10 Allocation Status descriptor . . . . .	133
6.3.2.11 Peering Status descriptor . . . . .	134
6.3.2.12 Membership Set descriptor . . . . .	135
6.3.2.13 Integrity descriptor . . . . .	136
6.3.2.14 FCDF Identification descriptor . . . . .	136
6.3.2.15 SW_ILS Request Information descriptor . . . . .	137
6.3.2.16 ELS Payload descriptor . . . . .	137
6.3.3 VA_Port SW_ILSs . . . . .	138
6.3.3.1 VN_Port Reachability Notification (VNRN) . . . . .	138
6.3.3.2 VN_Port Unreachability Notification (VNUN) . . . . .	139
6.3.3.3 FCDF Reachability Notification (FCRN) . . . . .	140
6.3.3.4 FCDF Unreachability Notification (FCUN) . . . . .	141
6.3.3.5 N_Port_ID Route Distribution (NPRD) . . . . .	143
6.3.3.6 N_Port_ID and Zoning ACL Distribution (NPZD) . . . . .	144
6.3.3.7 Active Zoning ACL Distribution (AZAD) . . . . .	146
6.3.3.8 Distributed Switch Membership Distribution (DSMD) . . . . .	148
6.3.3.9 Distributed ELS (DELS) . . . . .	149
6.3.4 VA_Port SW_ILS timeouts . . . . .	151
6.4 Controlling Switch redundancy protocol SW_ILSs . . . . .	151
6.4.1 Overview . . . . .	151
6.4.2 Controlling Switch redundancy protocol descriptors . . . . .	151
6.4.2.1 Descriptor format . . . . .	151
6.4.2.2 Controlling Switch State descriptor . . . . .	152
6.4.2.3 FCDF Topology descriptor . . . . .	152
6.4.2.4 FCDF N_Port_IDs descriptor . . . . .	154
6.4.2.5 RHello Interval descriptor . . . . .	155
6.4.2.6 Controlling Switch Parameters descriptor . . . . .	155
6.4.3 Controlling Switch redundancy protocol SW_ILSs . . . . .	156
6.4.3.1 Exchange Redundancy Parameters (ERP) . . . . .	156
6.4.3.2 Get FCDF Topology State (GFTS) . . . . .	157
6.4.3.3 Get FCDF N_Port_IDs State (GFNS) . . . . .	158
6.4.3.4 Secondary Synchronization Achieved (SSA) . . . . .	159
6.4.3.5 Redundancy Hello (RHello) . . . . .	160
6.4.3.6 Select Primary Controlling Switch (SPCS) . . . . .	160
6.4.3.7 Exchange Controlling Switch Parameters (ECSP) . . . . .	161
6.4.4 Controlling Switch redundancy protocol timeouts . . . . .	163
<b>7 Fabric Configuration . . . . .</b>	<b>164</b>
7.1 Fabric Configuration summary . . . . .	164
7.2 Switch port initialization . . . . .	165
7.2.1 Basic operation . . . . .	165
7.2.2 Switch_Name usage . . . . .	174
7.2.3 Exchange Switch Capabilities processing . . . . .	174
7.2.4 B_Port impact on ESC processing . . . . .	175
7.2.5 Extensions to support Virtual Fabrics . . . . .	176
7.3 Principal Switch Selection . . . . .	176
7.4 Address Distribution . . . . .	182
7.4.1 Address Distribution overview . . . . .	182
7.4.2 Domain_ID distribution by the Principal Switch . . . . .	184
7.4.3 Domain_ID requests by the Switches . . . . .	186
7.5 Principal ISL Recovery . . . . .	189
7.5.1 Overview . . . . .	189
7.5.2 Downstream Principal ISL discovery . . . . .	189
7.5.3 Upstream Principal ISL Recovery . . . . .	189

7.6 E_Port and Fabric isolation	189
7.7 B_Port operation	190
7.7.1 Differences between E_Ports and B_Ports	190
7.7.2 B_Port Internal Link Services	191
7.7.3 B_Port initialization	192
7.7.4 Example B_Port configuration	192
<b>8 Fabric Shortest Path First (FSPF)</b>	<b>193</b>
8.1 Overview	193
8.1.1 Basic components	193
8.1.2 Fabric connectivity	193
8.1.3 Addressing	193
8.1.4 Path selection and routing	193
8.1.5 FSPF path selection summary	194
8.2 FSPF message processing	194
8.2.1 Message transmission	194
8.2.2 Message reception and tests	194
8.3 Hello protocol	195
8.3.1 Basic functions	195
8.3.2 Hello message transmission	195
8.3.3 Hello message parameters	195
8.3.4 Hello message reception	196
8.4 The Link State Database	196
8.5 Usage of LSR fields	197
8.5.1 LSR Flags	197
8.5.2 LSR Age	197
8.5.3 LSR incarnation number	198
8.5.4 LSR instance rules	198
8.5.5 LSR checksum	199
8.5.6 Link Cost	201
8.6 Link State Database synchronization	202
8.6.1 Synchronization overview	202
8.6.2 Neighborhood and Adjacency	202
8.6.3 Continuous Link State Database synchronization	203
8.6.4 Reliable flooding	204
8.6.4.1 Basic operation	204
8.6.4.2 The flooding procedure	204
8.6.4.3 Generating a new LSR	205
8.6.4.4 Transmitting an LSR	205
8.6.4.5 Receiving an LSR	205
8.7 Neighbor finite state machine (FSM)	206
<b>9 Distributed Services</b>	<b>211</b>
9.1 Basic model	211
9.2 Distributed Services framework	211
9.2.1 Goals and characteristics of the Distributed Services framework	211
9.2.2 Distributed Service transport	211
9.2.2.1 Required FC-2 parameters	211
9.2.2.2 FC-CT Header usage	212
9.2.2.3 Frame distribution	212
9.2.2.4 Domain Controller Service Parameters	212
9.2.3 Common characteristics	212
9.2.4 Zoning considerations	213
9.2.5 Work categories	213

9.2.6	Frame formats	214
9.2.7	FC-CT command restrictions	214
9.3	Distributed Name Server	214
9.3.1	General behavior	214
9.3.2	FC-CT for distributed Name Servers	215
9.3.2.1	dNS command codes	215
9.3.2.2	FC-CT Header usage for dNS	219
9.3.3	Name Server Objects	220
9.3.4	FC-CT requests for dNS	222
9.3.4.1	Get Entry based on Port Identifier	222
9.3.4.2	Get Entry based on Port_Name	222
9.3.4.3	Get Entries based on Node_Name	223
9.3.4.4	Get Entries based on FC-4 TYPEs	224
9.3.4.5	Get Entries based on Port Type	224
9.3.4.6	Get Entries based on Zone Member	225
9.3.4.7	Get Entries based on Zone Name	226
9.3.4.8	Get Entries based on FC-4 Features	227
9.3.4.9	Get Entries based on Fabric Port_Name	228
9.4	Distributed Management Servers	228
9.4.1	General behavior	228
9.4.2	FC-CT Header	229
9.4.2.1	FC-CT Header parameters	229
9.4.2.2	FC-CT Header rule for Fabric internal requests	230
9.4.3	Fabric Configuration Service	230
9.4.4	Unzoned Name Service	233
9.4.5	Fabric Zone Service	233
9.4.6	Fabric-Device Management Service	233
9.4.6.1	Operational characteristics of the FDMI Server	233
9.4.6.2	Registration scenarios	234
9.4.6.2.1	HBA attached to a single Switch	234
9.4.6.2.2	HBA attached to multiple Switches	234
9.4.6.2.3	Resolution of the principal HBA manager	234
9.4.6.3	FDMI Inter-Switch messages	235
9.4.6.3.1	General format	235
9.4.6.3.2	FC-CT Header	235
9.4.6.3.3	FDMI Header	235
9.4.6.3.4	Payload	236
9.4.6.4	FDMI Inter-Switch requests	236
9.4.6.5	FDMI Inter-Switch responses	237
9.4.6.5.1	Reject response	237
9.4.6.5.2	Accept response	237
9.4.6.6	FDMI Inter-Switch operations	237
9.4.6.6.1	Registration Notification (FRN) operation	237
9.4.6.6.2	De-Register Notification (FDRN) operation	238
9.4.6.6.3	Update Notification (FUN) operation	238
9.4.6.6.4	Update Forward (FUF) operation	238
9.4.6.6.5	De-Register Forward (FDRF) operation	238
9.4.6.6.6	Fetch	238
9.4.6.7	GS client initiated FDMI requests	239
9.4.7	Other Fabric internal services	240
9.4.7.1	Fabric internal requests	240
9.4.7.2	Get Management Server Capabilities (GCAP) request	241
9.4.7.2.1	Overview	241
9.4.7.2.2	Capability Entry	241

9.4.7.2.3 Subtype Capability Bit Masks	242
9.4.8 Security Information Server	242
9.4.9 Application Server	243
9.4.10 Enhanced Fabric Configuration Service	244
9.5 Distributed Event Server	245
9.5.1 General behavior	245
9.5.2 FC-CT for distributed Event Server	245
9.5.2.1 FC-CT Header parameters	245
9.5.2.2 dES command codes	245
9.6 Distributed VE Identification Server	245
9.6.1 General behavior	245
9.6.2 FC-CT for distributed VE Identification Server	246
9.6.3 FC-CT Requests for dVEIS	246
9.6.3.1 Get VE Mappings - Domain_ID (G_VEM_D)	246
9.6.3.2 Update VE Mappings (U_VE_M)	248
<b>10 Switch Zone exchange &amp; merge</b>	<b>250</b>
10.1 Overview	250
10.2 Joining Switches	250
10.3 Enhanced Zoning support determination	250
10.4 Zoning framework and data structures	251
10.4.1 Basic Zoning framework	251
10.4.2 Basic Zoning data structures	254
10.4.2.1 Zoning Object List	254
10.4.2.2 Zoning Object format	255
10.4.2.3 General name format	256
10.4.2.4 Zone Member format	256
10.4.3 Enhanced Zoning framework	257
10.4.3.1 Introduction	257
10.4.3.2 Zone Set Database	257
10.4.3.3 Active Zone Set	260
10.4.4 Enhanced Zoning data structures	261
10.4.4.1 Zoning Object List	261
10.4.4.2 Zoning Object Types	261
10.4.4.3 Zone Set Object	262
10.4.4.3.1 Zone Set Object in the Zone Set Database	262
10.4.4.3.2 Zone Set Object in the Active Zone Set	263
10.4.4.4 Zone Reference Object	263
10.4.4.5 Zone Object in the Zone Set Database	264
10.4.4.6 Zone Object in the Active Zone Set	265
10.4.4.6.1 Zone Member format	266
10.4.4.7 Zone Alias Object	268
10.4.4.8 Zone Attribute Object	269
10.4.4.8.1 Zone Attribute Entry format	270
10.5 Merge Zone	273
10.5.1 Example merge operation	273
10.5.2 Merge Zone rules	276
10.5.2.1 Merge rules in Basic Zoning	276
10.5.2.2 Merge rules in Enhanced Zoning	276
10.6 Fabric Management Session protocol	277
10.6.1 Fabric Management Session protocol overview	277
10.6.2 Reserving Fabric Change Authorization	278
10.6.3 Staging the Fabric Configuration	278
10.6.4 Updating the Fabric Configuration	279

10.6.5 Releasing Fabric Change Authorization . . . . .	279
10.6.6 Mapping of a Server session to a Fabric Management Session . . . . .	279
10.6.7 Fabric behavior to handle the CT SFEZ request . . . . .	281
10.6.8 Fabric behavior to handle the CT AAPZ and RAPZ requests . . . . .	281
10.7 Switch behaviors during merge . . . . .	281
<b>11 Distributed broadcast . . . . .</b>	<b>282</b>
11.1 Overview . . . . .	282
11.2 Spanning tree . . . . .	282
11.2.1 Spanning tree example . . . . .	282
<b>12 Virtual Fabrics Switch support . . . . .</b>	<b>284</b>
12.1 Overview . . . . .	284
12.2 VF Capable Switch functional model . . . . .	285
12.3 Switch_Names usage . . . . .	286
12.4 Configuration information . . . . .	286
12.5 Enabling VFT tagging on Switch ports . . . . .	286
12.6 Exchange Virtual Fabrics parameters processing . . . . .	290
12.6.1 Overview . . . . .	290
12.6.2 Changing the VFT tagging mode . . . . .	292
12.6.3 Adding or removing Virtual Fabrics . . . . .	293
12.6.4 Changing the port VF_ID . . . . .	293
<b>13 Enhanced Commit Service . . . . .</b>	<b>295</b>
13.1 Overview . . . . .	295
13.2 Assisted mode protocol operations . . . . .	295
13.3 Autonomous mode protocol operations . . . . .	296
13.3.1 Protocol phases . . . . .	296
13.3.1.1 Overview . . . . .	296
13.3.1.2 Phase one . . . . .	296
13.3.1.3 Phase two . . . . .	296
13.3.1.4 Phase three . . . . .	296
13.3.1.5 Phase four . . . . .	296
13.3.2 Handling Fabric changes . . . . .	297
13.3.3 Error recovery . . . . .	297
13.3.3.1 Overview . . . . .	297
13.3.3.2 Managing Switch not functional . . . . .	297
13.3.3.2.1 Dead Man timer . . . . .	297
13.3.3.2.2 Basic procedure . . . . .	297
13.3.3.3 Resolution of multiple Managing Switches . . . . .	298
13.3.3.3.1 Two Managing Switches - same commit phase . . . . .	298
13.3.3.3.2 Two Managing Switches - different commit phases . . . . .	298
13.3.4 Ladder diagrams . . . . .	299
13.3.4.1 Normal case . . . . .	299
13.3.4.2 Unsuccessful case . . . . .	300
13.3.4.3 Transfer ownership case - recovery processing enabled . . . . .	301
13.3.5 State machines . . . . .	301
13.3.5.1 Overview . . . . .	301
13.3.5.2 States and transitions for the Managing Switch . . . . .	301
13.3.5.3 States and transitions for the managed Switch . . . . .	304
13.3.5.4 States and transitions for transfer commit ownership . . . . .	306
<b>14 Virtual Channels for Switched Fabric . . . . .</b>	<b>308</b>
14.1 Overview . . . . .	308
14.2 Assignment of Virtual Channels . . . . .	308

14.2.1 Overview of assignment	308
14.2.2 Simple	308
14.2.3 Fixed	308
14.2.4 Variable	309
14.3 VC parameter negotiation	310
14.3.1 Agreement of assignment schemes	310
14.3.2 Negotiation of number of VCs	310
14.4 Credit management	311
14.4.1 Overview	311
14.4.2 VC_RDY Primitive Signals	311
<b>15 Inter-Fabric Routing support</b>	<b>312</b>
15.1 F_RJT and F_BSY processing for Class 2/F	312
15.1.1 Overview	312
15.1.2 Encapsulated Class 2 F_RJT or Class 2 F_BSY frame format	312
15.1.2.1 Overview	312
15.1.2.2 Encapsulated Enc_Header field values	313
15.1.2.3 Encapsulated IFR_Header field values	314
15.1.2.4 Encapsulated Frame_Header field values	314
<b>16 Timers and constants</b>	<b>315</b>
16.1 General timers and constants	315
16.2 SW_ILS timeout values	316
<b>17 Distributed Switch</b>	<b>317</b>
17.1 Overview	317
17.2 FCDF handling of well-known addresses	323
17.3 A_Port to A_Port links (ASLs)	326
17.4 Distributed Switch operations	327
17.4.1 Overview	327
17.4.2 FCDF routing	328
17.4.3 N_Port_ID handling	329
17.4.4 Distribution tree	331
17.5 Controlling Switch redundancy protocol	331
17.5.1 Controlling Switch redundancy protocol overview	331
17.5.2 Controlling Switch redundancy protocol state machine	331
<b>18 Leaf Switch</b>	<b>340</b>
18.1 Overview	340
<b>Annex A</b>	<b>341</b>
A.1 Introduction	341
A.2 Example 1: two E/F/FL_Port-capable Switch ports	341
A.3 Example 2: two E/F/FL_Port-capable Switch ports and one PN_Port	342
A.4 Example 3: one E/F/FL_Port-capable Port and one E/F_Port-capable Port	343
<b>Annex B</b>	<b>344</b>
B.1 Introduction	344
B.2 ELP exchange protocol	344
B.2.1 ELP exchange without parameter negotiation	344
B.2.2 ELP exchange with parameter negotiation	346
<b>Annex C</b>	<b>348</b>
C.1 Introduction	348

C.2 Sample flows	348
C.2.1 HBA registration - single Switch	348
C.2.2 HBA registration - multiple Switches - caches updated	349
C.2.3 HBA registration - multiple Switches - caches not updated	349
C.2.4 HBA de-registration - primary HBA manager	350
C.2.5 HBA de-registration - non-primary HBA manager	351
<b>Annex D</b>	<b>353</b>
D.1 Overview	353
D.2 Background	353
D.3 Definitions	353
D.4 Characteristics of Avionics Fabrics	354
D.4.1 Overview	354
D.4.2 AE Switch port mode initialization	355
D.4.2.1 Overview	355
D.4.2.2 Switch port mode initialization state machine modifications	355
D.4.3 ELP payload requirements	359
D.4.4 AE Principal Switch	360
D.4.4.1 AE Principal Switch initialization process	360
D.4.4.2 Map update process	361
D.4.4.3 AE Principal Switch update process	361
D.4.5 FFI Domain Topology Map distribution	362
D.4.5.1 Overview	362
D.4.5.2 FFI Domain Topology Map distribution state machine diagram	362
D.4.5.3 FFI Domain Topology Map distribution state machine text	363
D.4.6 Fast Fabric Initialization (FFI) SW_ILS definition	370
D.4.6.1 Overview	370
D.4.6.2 Fast Fabric Initialization Link State Record (FFI LSR) format	375
D.4.6.3 Fast Fabric Initialization Link Descriptor format	376
D.5 FFI Domain Topology Map Distribution (Informative)	377
D.5.1 Sample configuration	377
D.5.2 Initialization procedure example	378
D.5.3 AE Principal Switch update example	380



<b>Figure</b>	<b>Page</b>
Figure 1 – Sample state machine . . . . .	12
Figure 2 – Basic Switch model . . . . .	17
Figure 3 – Enhanced Switch functional model . . . . .	18
Figure 4 – Multiple Switch Fabric example . . . . .	19
Figure 5 – Controlling Switch functional model . . . . .	21
Figure 6 – FCDF functional model . . . . .	22
Figure 7 – FCDF Switching Element . . . . .	23
Figure 8 – Domain, Area, and Port address partitioning . . . . .	25
Figure 9 – F_Port model . . . . .	30
Figure 10 – FL_Port model . . . . .	32
Figure 11 – E_Port model . . . . .	33
Figure 12 – B_Port model . . . . .	34
Figure 13 – A_Port model . . . . .	36
Figure 14 – Principal Inter-Switch Links . . . . .	37
Figure 15 – VA_Port SW_ILS relay example . . . . .	126
Figure 16 – Switch port mode initialization state machine . . . . .	166
Figure 17 – Switch port mode initialization state machine - continued . . . . .	167
Figure 18 – Simultaneous ELP processing- parameters acceptable to both Switches . . . . .	171
Figure 19 – ESC processing . . . . .	175
Figure 20 – Principal Switch selection state machine . . . . .	177
Figure 21 – Example Propagation of BF and RCF SW_ILS requests . . . . .	179
Figure 22 – Address Distribution state machines . . . . .	183
Figure 23 – RDI request processing by Principal Switch . . . . .	185
Figure 24 – RDI request Processing by non-Principal Switch . . . . .	188
Figure 25 – Example B_Port configuration - Virtual ISL . . . . .	192
Figure 26 – Neighbor finite state machine . . . . .	210
Figure 27 – Basic Zoning framework . . . . .	252
Figure 28 – Basic Zoning hierarchy . . . . .	253
Figure 29 – Basic Zoning Object structure . . . . .	254
Figure 30 – Logical structure of the Zone Set Database . . . . .	259
Figure 31 – Logical structure of the Active Zone Set . . . . .	260
Figure 32 – Merge operation between two Switches . . . . .	274
Figure 33 – Merge operation among several Switches . . . . .	275
Figure 34 – Broadcast path selection example . . . . .	283
Figure 35 – Virtual Fabrics . . . . .	284
Figure 36 – Functional model of a VF capable Switch . . . . .	285
Figure 37 – Switch port mode initialization state machine - Virtual Fabric support . . . . .	288
Figure 38 – A generic EVFP transaction . . . . .	290
Figure 39 – Normal commit ladder diagram . . . . .	299
Figure 40 – Unsuccessful commit ladder diagram . . . . .	300
Figure 41 – Transfer ownership ladder diagram . . . . .	301
Figure 42 – ECS Managing Switch state machine . . . . .	302
Figure 43 – ECS managed Switch state machine . . . . .	304
Figure 44 – ECS transfer commit ownership (TCO) state machine . . . . .	306
Figure 45 – VC_RDY Primitive Signal format . . . . .	311
Figure 46 – Encapsulated Class 2/F F_RJT and Class 2/F F_BSY frame format . . . . .	312
Figure 47 – Distributed Switch example - one Controlling Switch and two FCDFs . . . . .	317
Figure 48 – FSPF Topology of figure 47 . . . . .	318
Figure 49 – Distributed Switch example - two Controlling Switches - two FCDFs . . . . .	319
Figure 50 – Distributed Switch example - two Controlling Switches - four cascaded FCDFs . . . . .	320
Figure 51 – FSPF topology of figure 49 and figure 50 . . . . .	321
Figure 52 – Distributed Switch example - three Controlling Switches - two FCDFs . . . . .	322

Figure 53 – FSPF topology of figure 52 . . . . .	323
Figure 54 – FCDF with no link to Primary Controlling Switch WKA frame processing example . .	325
Figure 55 – Redundant Distributed Switch example . . . . .	331
Figure 56 – Controlling Switch redundancy protocol state machine . . . . .	333
Figure 57 – Leaf Switch example . . . . .	340
Figure A.1 – Initialization example 1 . . . . .	341
Figure A.2 – Initialization example 2 . . . . .	342
Figure A.3 – Initialization example 3 . . . . .	343
Figure B.1 – Reference ELP configuration . . . . .	344
Figure B.2 – A successful and complete ELP exchange . . . . .	345
Figure B.3 – An unsuccessful but complete ELP exchange . . . . .	345
Figure B.4 – A successful ELP exchange protocol parameter negotiation . . . . .	346
Figure B.5 – An unsuccessful ELP exchange protocol parameter negotiation . . . . .	347
Figure C.1 – Registration of HBA information - single Switch . . . . .	348
Figure C.2 – Registration of HBA information - multiple Switches caches updated . . . . .	349
Figure C.3 – Registration of HBA Information - multiple Switches caches not updated . . . . .	350
Figure C.4 – HBA de-registration - primary HBA manager . . . . .	351
Figure C.5 – HBA de-registration - non-primary HBA manager . . . . .	352
Figure D.1 – Modifications to Switch port mode initialization . . . . .	356
Figure D.2 – FFI Domain Topology Map distribution state machine, non-principal AE Switches .	362
Figure D.3 – FFI Domain Topology Map distribution state machine, AE Principal Switch . . . . .	363
Figure D.4 – Example Avionics Fabric . . . . .	378

<b>Table</b>	<b>Page</b>
Table 1 – Address identifier values . . . . .	27
Table 2 – SW_ILS command codes . . . . .	41
Table 3 – SW_RJT payload . . . . .	44
Table 4 – SW_RJT reason codes . . . . .	44
Table 5 – SW_RJT reason code explanations . . . . .	45
Table 6 – ELP request payload . . . . .	48
Table 7 – Fabric Controller Class F Service Parameters . . . . .	50
Table 8 – Class 2 Interconnect_Port Parameters . . . . .	51
Table 9 – Class 3 Interconnect_Port Parameters . . . . .	51
Table 10 – ISL Flow Control Mode values . . . . .	52
Table 11 – Flow Control Parameters. . . . .	52
Table 12 – VC_RDY flow control parameters . . . . .	53
Table 13 – Assignment schemes . . . . .	53
Table 15 – VC values - Fixed . . . . .	54
Table 16 – VC Values - Variable . . . . .	54
Table 14 – VC values - Simple . . . . .	54
Table 17 – ELP accept payload . . . . .	55
Table 18 – EFP request payload. . . . .	56
Table 19 – Switch_Priority field values . . . . .	57
Table 20 – Domain_ID_List record format . . . . .	57
Table 21 – Record_Type field values . . . . .	58
Table 22 – EFP accept payload . . . . .	58
Table 23 – DIA request payload . . . . .	59
Table 24 – DIA accept payload . . . . .	60
Table 25 – RDI request payload . . . . .	61
Table 26 – RDI accept payload . . . . .	62
Table 27 – HLO request payload . . . . .	63
Table 28 – FSPF Header . . . . .	64
Table 29 – FSPF command codes . . . . .	64
Table 30 – LSU request payload. . . . .	65
Table 31 – Flags field bit map . . . . .	66
Table 32 – Link State Record - Link Descriptor format . . . . .	66
Table 33 – Link State Header format . . . . .	67
Table 34 – Link State Record Type field values . . . . .	67
Table 35 – LSR Flags field bit descriptions. . . . .	68
Table 36 – Link Descriptor format . . . . .	69
Table 37 – Link Type values . . . . .	69
Table 38 – LSA request payload. . . . .	70
Table 39 – BF request payload . . . . .	71
Table 40 – BF accept payload. . . . .	71
Table 41 – RCF request payload . . . . .	72
Table 42 – RCF accept payload . . . . .	72
Table 43 – SW_RSCN request payload . . . . .	73
Table 44 – Device Entry format. . . . .	74
Table 45 – SW_RSCN accept payload . . . . .	75
Table 46 – DRLIR request payload. . . . .	75
Table 47 – DRLIR accept payload . . . . .	76
Table 48 – Merge Request request payload . . . . .	77
Table 49 – Protocol Version values. . . . .	77
Table 50 – Basic Zoning payload . . . . .	77
Table 51 – Enhanced Zoning payload. . . . .	78
Table 52 – Merge Request accept payload. . . . .	80

Table 53 – ACA request payload . . . . .	81
Table 54 – Acquire Change Authorization accept payload . . . . .	82
Table 55 – RCA request payload . . . . .	82
Table 56 – Release Change Authorization accept payload . . . . .	83
Table 57 – SFC request payload . . . . .	83
Table 58 – Operation Request values . . . . .	84
Table 59 – Stage Fabric Configuration accept payload . . . . .	85
Table 60 – Payload for Operation Request values 03 and 04 . . . . .	85
Table 61 – Payload for Operation Request ‘Activate Zone Set Enhanced’ . . . . .	86
Table 62 – Payload for Operation Request ‘Deactivate Zone Set Enhanced’ . . . . .	86
Table 63 – Payload for Operation Request ‘Distribute Zone Set Database’ . . . . .	86
Table 64 – Payload for Operation Request ‘Activate Zone Set by Name’ . . . . .	87
Table 65 – Payload for Operation Request ‘Set Zoning Policies’ . . . . .	87
Table 66 – Update Fabric Configuration request payload . . . . .	88
Table 67 – Update Fabric Configuration accept payload . . . . .	89
Table 68 – CEC request payload . . . . .	90
Table 69 – CEC accept payload . . . . .	91
Table 70 – ESC request payload . . . . .	92
Table 71 – Protocol Descriptor format . . . . .	92
Table 72 – Protocol ID values . . . . .	93
Table 73 – ESC accept payload . . . . .	93
Table 74 – ESS request payload . . . . .	95
Table 75 – Capability Object format . . . . .	96
Table 76 – Name Server Capability Flags . . . . .	96
Table 77 – Fabric Controller Capability flags . . . . .	97
Table 78 – Fabric Configuration Server Capability flags . . . . .	98
Table 79 – Enhanced Zone Server Capability flags . . . . .	98
Table 80 – Vendor Specific Capability Object . . . . .	100
Table 81 – Domain Controller Capability Object . . . . .	101
Table 82 – Event Server Capability Object . . . . .	102
Table 83 – Switch Support Capability Object . . . . .	102
Table 84 – Application Server Capability Object . . . . .	103
Table 85 – Enhanced Fabric Configuration Server Capability Object . . . . .	103
Table 86 – VE Identification Server Capability Flags . . . . .	104
Table 87 – ESS accept payload . . . . .	104
Table 88 – MRRA request payload . . . . .	105
Table 89 – Vendor Specific field . . . . .	105
Table 90 – MRRA response payload . . . . .	106
Table 91 – MRRA Response values . . . . .	106
Table 92 – STR request payload . . . . .	107
Table 93 – Nx_Port Tags . . . . .	108
Table 94 – Flags field values . . . . .	109
Table 95 – STR Reject Reason Code values . . . . .	109
Table 96 – Path Information . . . . .	110
Table 97 – STR accept payload . . . . .	111
Table 98 – EVFP request payload . . . . .	112
Table 99 – EVFP Message Codes . . . . .	112
Table 100 – EVFP accept payload . . . . .	113
Table 101 – SW_RJT reason codes . . . . .	114
Table 102 – EVFP_SYNC Message Payload format . . . . .	114
Table 103 – Descriptor format . . . . .	115
Table 104 – Descriptor Control codes . . . . .	115
Table 105 – Descriptor Types . . . . .	115
Table 106 – Tagging Administrative Status descriptor . . . . .	116

Table 107 – Administrative Tagging Modes . . . . .	116
Table 108 – Tagging Mode negotiation . . . . .	116
Table 109 – Port VF_ID descriptor . . . . .	117
Table 110 – Locally-Enabled VF_ID List descriptor . . . . .	117
Table 111 – Vendor Specific descriptor . . . . .	118
Table 112 – EACA request payload . . . . .	118
Table 113 – Commit Exchange Preamble . . . . .	119
Table 114 – Transaction Identifier . . . . .	119
Table 115 – Application ID values . . . . .	120
Table 116 – ECS Switch List . . . . .	121
Table 117 – Switch Identifier format . . . . .	121
Table 118 – ESFC request payload . . . . .	122
Table 119 – EUFC request payload . . . . .	123
Table 120 – ERCA request payload . . . . .	124
Table 121 – TCO request payload . . . . .	125
Table 122 – VA_Port SW_ILSs command codes . . . . .	126
Table 123 – Descriptor format . . . . .	127
Table 124 – Descriptor tags . . . . .	127
Table 125 – VN_Port Reachability descriptor format . . . . .	128
Table 126 – FLOGI/NPIV FDISC Parameters descriptor format . . . . .	128
Table 127 – VN_Port Unreachability descriptor format . . . . .	128
Table 128 – FCDF Reachability descriptor format . . . . .	129
Table 129 – Sequence Number descriptor format . . . . .	129
Table 130 – Controlling Switch Reachability descriptor format . . . . .	130
Table 131 – N_Port_IDs Reachability descriptor format . . . . .	130
Table 132 – N_Port_ID Reachability Entry format . . . . .	131
Table 133 – Domain_IDs Reachability descriptor format . . . . .	132
Table 134 – Reachable Domain_ID Entry format . . . . .	132
Table 135 – Allocation Status descriptor format . . . . .	133
Table 136 – Allocation / Deallocation Entry format . . . . .	133
Table 137 – Peering Status descriptor format . . . . .	134
Table 138 – Peering Entry format . . . . .	134
Table 139 – Membership Set descriptor format . . . . .	135
Table 140 – Integrity descriptor format . . . . .	136
Table 141 – Integrity Type field values . . . . .	136
Table 142 – FCDF Identification descriptor format . . . . .	136
Table 143 – SW_ILS Request Information descriptor format . . . . .	137
Table 144 – ELS Payload descriptor format . . . . .	137
Table 145 – VNRN request payload . . . . .	138
Table 146 – VNRN SW_ACC payload . . . . .	139
Table 147 – VNUN request payload . . . . .	139
Table 148 – VNUN SW_ACC payload . . . . .	140
Table 149 – FCRN request payload . . . . .	141
Table 150 – FCRN SW_ACC payload . . . . .	141
Table 151 – FCUN request payload . . . . .	142
Table 152 – FCUN SW_ACC payload . . . . .	142
Table 153 – NPRD request payload . . . . .	143
Table 154 – NPRD SW_ACC payload . . . . .	144
Table 155 – NPZD request payload . . . . .	145
Table 156 – NPZD SW_ACC payload . . . . .	146
Table 157 – AZAD request payload . . . . .	147
Table 158 – AZAD SW_ACC payload . . . . .	147
Table 159 – DSMD request payload . . . . .	148
Table 160 – DSMD SW_ACC payload . . . . .	149

Table 161 – DELS request payload . . . . .	149
Table 162 – DELS SW_ACC payload . . . . .	150
Table 163 – VA_Port SW_ILSs timeouts . . . . .	151
Table 164 – Controlling Switch redundancy protocol SW_ILSs command codes . . . . .	151
Table 165 – Controlling Switch State descriptor format . . . . .	152
Table 166 – FCDF Topology descriptor format . . . . .	152
Table 167 – FCDF Connectivity Record format . . . . .	153
Table 168 – ASL Record format . . . . .	153
Table 169 – FCDF N_Port_IDs descriptor format . . . . .	154
Table 170 – FCDF Allocation Record format . . . . .	154
Table 171 – RHello Interval descriptor format . . . . .	155
Table 172 – Controlling Switch Parameters descriptor format . . . . .	155
Table 173 – Controlling Switch Priority    Switch_Name Record format . . . . .	155
Table 174 – ERP request payload . . . . .	156
Table 175 – ERP SW_ACC payload . . . . .	157
Table 176 – GFTS request payload . . . . .	157
Table 177 – GFTS SW_ACC payload . . . . .	158
Table 178 – GFNS request payload . . . . .	158
Table 179 – GFNS SW_ACC payload . . . . .	159
Table 180 – SSA request payload . . . . .	159
Table 181 – SSA SW_ACC payload . . . . .	160
Table 182 – RHello request payload . . . . .	160
Table 183 – SPCS request payload . . . . .	161
Table 184 – SPCS SW_ACC payload . . . . .	161
Table 185 – ECSP request payload . . . . .	162
Table 186 – ECSP SW_ACC payload . . . . .	162
Table 187 – Controlling Switch redundancy protocol SW_ILSs timeouts . . . . .	163
Table 188 – Fabric Configuration summary . . . . .	164
Table 189 – Responses to ELP request for originating Interconnect_Port . . . . .	169
Table 190 – Recommended BF and RCF usage summary . . . . .	176
Table 191 – B_Port ILS support . . . . .	191
Table 192 – B_Port initialization summary . . . . .	192
Table 193 – Path selection (FSPF) operation summary . . . . .	194
Table 194 – Checksum byte order calculation . . . . .	200
Table 195 – DEFAULT_COST values . . . . .	201
Table 196 – Neighbor finite state machine . . . . .	207
Table 197 – Default Domain Controller Service Parameter values . . . . .	212
Table 198 – FC-CT command codes for dNS . . . . .	216
Table 199 – Name Server Entry Object . . . . .	220
Table 200 – Entry Object Format Indicator . . . . .	221
Table 201 – Name Server Entry Object description . . . . .	221
Table 202 – GE_ID request payload . . . . .	222
Table 203 – GE_ID accept payload . . . . .	222
Table 204 – GE_PN request payload . . . . .	222
Table 205 – GE_PN accept payload . . . . .	223
Table 206 – GE_NN request payload . . . . .	223
Table 207 – GE_NN accept payload . . . . .	223
Table 208 – GE_FT request payload . . . . .	224
Table 209 – GE_FT accept payload . . . . .	224
Table 210 – GE_PT request payload . . . . .	224
Table 211 – GE_PT accept payload . . . . .	225
Table 212 – GE_ZM request payload . . . . .	225
Table 213 – GE_ZM accept payload . . . . .	226
Table 214 – GE_ZN request payload . . . . .	226

Table 215 – GE_ZN accept payload . . . . .	226
Table 216 – GE_FF request payload . . . . .	227
Table 217 – GE_FF accept payload . . . . .	227
Table 218 – GE_FPN request payload . . . . .	228
Table 219 – GE_FPN accept payload . . . . .	228
Table 220 – Zoning effect on Servers of the distributed Management Service . . . . .	229
Table 221 – dMS FC-CT Header parameters . . . . .	229
Table 222 – Fabric Configuration Service command codes for dMS . . . . .	230
Table 223 – FDMI Inter-Switch message format . . . . .	235
Table 224 – FC-CT Header parameters . . . . .	235
Table 225 – FDMI Header format . . . . .	235
Table 226 – Vendor Specified field format . . . . .	236
Table 227 – FDMI Fabric Internal command codes . . . . .	236
Table 228 – Reason code explanation . . . . .	237
Table 229 – Registered HBA/Port list format . . . . .	238
Table 230 – HBA Entry format . . . . .	239
Table 231 – Port Entry format . . . . .	239
Table 232 – FDMI CT commands for the dMS . . . . .	240
Table 233 – Fabric internal Management Server requests . . . . .	240
Table 234 – GCAP request payload format . . . . .	241
Table 235 – GCAP CT_ACC payload format . . . . .	241
Table 236 – Capability Entry format . . . . .	241
Table 237 – Fabric Configuration Server (CT_Subtype 01h) . . . . .	242
Table 238 – Unzoned Name Server (CT_Subtype 02h) . . . . .	242
Table 239 – Security Information Server command codes for dMS . . . . .	242
Table 240 – Application Server command codes for dMS . . . . .	243
Table 241 – Enhanced Fabric Configuration Service command codes for dMS . . . . .	244
Table 242 – dES FC-CT Header parameters . . . . .	245
Table 243 – FC-CT command codes for dES . . . . .	245
Table 244 – dVEIS FC-CT Header parameters . . . . .	246
Table 245 – FC-CT Command Codes for dVEIS . . . . .	246
Table 246 – G_VEM_D Request payload . . . . .	246
Table 247 – G_VEM_D Accept payload . . . . .	247
Table 248 – VEM Mapping entry format . . . . .	247
Table 249 – VE Mapping record format . . . . .	247
Table 250 – Fabric VE ID format . . . . .	248
Table 251 – U_VE_M Request Payload . . . . .	248
Table 252 – VE Mapping entry format . . . . .	249
Table 253 – U_VE_M Accept Payload . . . . .	249
Table 254 – Zoning Object List format . . . . .	254
Table 255 – Zoning Object . . . . .	255
Table 256 – Zoning Object Type values . . . . .	255
Table 257 – Protocol Format values . . . . .	256
Table 258 – Zone Member format . . . . .	256
Table 259 – Zone Member Type and Identifier formats . . . . .	257
Table 260 – Zoning Object List format . . . . .	261
Table 261 – Zoning Object Type values . . . . .	261
Table 262 – Zone Set Object format in the Zone Set Database . . . . .	262
Table 263 – Zone Set Object format in the Active Zone Set . . . . .	263
Table 264 – Zone Reference Object format . . . . .	263
Table 265 – Zone Object format in the Zone Set Database . . . . .	264
Table 266 – Zone Object format in the Active Zone Set . . . . .	265
Table 267 – Zone Member format . . . . .	266
Table 268 – Zone Member Type and Identifier formats . . . . .	266

Table 269 – Wildcard Zone Member Identifier format . . . . .	267
Table 270 – Vendor Specified Zone Member Identifier format . . . . .	268
Table 271 – Zone Alias Object format . . . . .	268
Table 272 – Zone Attribute Object format . . . . .	269
Table 273 – Zone Attribute Block format . . . . .	269
Table 274 – Zone Attribute Entry format . . . . .	270
Table 275 – Zone Attribute Type values . . . . .	270
Table 276 – Protocol Attribute Value format . . . . .	271
Table 277 – Peer Zone Attribute Value format . . . . .	273
Table 278 – Vendor Specific Attribute Value format . . . . .	273
Table 279 – Basic Zoning merge rules . . . . .	276
Table 280 – Enhanced Zoning merge rules. . . . .	277
Table 281 – EACA phase - events and actions. . . . .	304
Table 282 – ESFC phase - events and actions. . . . .	305
Table 283 – EUFC phase - events and actions. . . . .	305
Table 284 – Simple assignment scheme. . . . .	308
Table 285 – Fixed assignment scheme. . . . .	308
Table 286 – VC assignments - Fixed. . . . .	309
Table 287 – Variable assignment scheme. . . . .	309
Table 288 – VC assignments - Variable . . . . .	310
Table 289 – VC_ID values for VC_RDY Primitive Signals. . . . .	311
Table 290 – Encapsulated Class 2/F F_RJT/F_BSY Enc_Header field values. . . . .	313
Table 291 – Encapsulated Class 2/F F_RJT/F_BSY IFR_Header field values . . . . .	314
Table 292 – FC-SW-7 timers and constants . . . . .	315
Table 293 – SW_ILS timeout values . . . . .	316
Table 294 – Forwarded Domain Controller and well-known address identifiers . . . . .	324
Table 295 – VA_Port ELP Flag bits . . . . .	326
Table 296 – Controlling Switch priority values. . . . .	332
Table D.1 – Responses to ELP request for originating Interconnect_Port . . . . .	357
Table D.2 – ELP required payload values for AE_Ports . . . . .	359
Table D.3 – Actions taken by non-Principal AE Switch for an FFI request Sequence . . . . .	365
Table D.4 – Action taken by AE Principal Switch for an FFI request Sequence . . . . .	369
Table D.5 – FFI request payload. . . . .	371
Table D.6 – FFI Type Flags definition . . . . .	372
Table D.7 – FFI Problem Detected Reason Codes. . . . .	373
Table D.8 – FFI accept payload . . . . .	374
Table D.9 – FFI Link State Record - Link Descriptor format . . . . .	375
Table D.10 – FFI LSR Flags definition . . . . .	375
Table D.11 – FFI Link Descriptor format . . . . .	376
Table D.12 – FFI Link Descriptor Flags definition . . . . .	376



draft proposed American National Standard  
for Information Technology—

# Fibre Channel — Switch Fabric - 7 (FC-SW-7)

## 1 Scope

This American National Standard for FC-SW-7 describes the operation and interaction of Fibre Channel Switches.

This standard includes:

- a) E\_Port operation and Fabric Configuration;
- b) path selection (FSPF);
- c) Bridge port (B\_Port) operation;
- d) distributed Server interaction and communication;
- e) exchange of information between Switches to support Zoning;
- f) distribution of Event Notifications between Switches;
- g) Virtual Fabrics Switch Support;
- h) Enhanced Commit Service;
- i) Virtual Channels; and
- j) Distributed Switch.

## 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 Varembe, 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](mailto: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](mailto:incits@itic.org)

Additional availability contact information is provided below as needed.

## 2.2 Approved references

- [1] INCITS 332-1999, *Fibre Channel - Second Generation Arbitrated Loop - 2 (FC-AL-2)*
- [2] INCITS TR-48-2012, *Fibre Channel - Methodologies for Interconnects - 3 (FC-MI-3)*
- [3] INCITS TR-49-2012, *Fibre Channel - Device Attach - 2 (FC-DA-2)*
- [4] INCITS 509-2014, *Fibre Channel - Backbone - 6 (FC-BB-6)*
- [5] INCITS 487-2018, *Fibre Channel - Link Services - 3 (FC-LS-3)*
- [6] INCITS 544-2018, *Fibre Channel - Single Byte Command Code Sets - 6 (FC-SB-6)*
- [7] INCITS 496-2012, *Fibre Channel - Security Protocols - 2 (FC-SP-2)*
- [8] INCITS 496-2012/AM 1-2015, *Fibre Channel - Security Protocols - 2 (FC-SP-2) - Amendment 1*
- [9] INCITS 475-2011, *Fibre Channel - Inter-Fabric Routing (FC-IFR)*

## 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 document, or regarding availability, contact the relevant Standards body or other organization as indicated.

NOTE 1 – For more information on the current status of a document, contact the INCITS Secretariat at the address listed in the front matter. To obtain copies of this document, contact Global Engineering at the address listed in the front matter, or the INCITS Secretariat.

[10] INCITS 553-201y, *Fibre Channel - Link Services - 4 (FC-LS-4)*

[11] INCITS 548-201y, *Fibre Channel - Generic Services - 8 (FC-GS-8)*

[12] INCITS 562-201y, *Fibre Channel - Framing and Signaling - 4 (FC-FS-4)*

[13] INCITS 543-201y, *Fibre Channel - Physical Interface - 7 (FC-PI-7)*

[14] INCITS 559-201y, *Fibre Channel - Physical Interface - 7p (FC-PI-7p)*

## 2.4 IETF references

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

RFC 905, *ISO Transport Protocol Specification, ISO DP 8073*, April 1984.

RFC 1008, *Implementation Guide for the ISO Transport Protocol*, June 1987.

RFC 4936, *Fibre Channel Zone Server MIB*, August 2007.

### **3 Definitions, conventions, abbreviations, acronyms, and symbols**

#### **3.1 Definitions**

##### **3.1.1 Active Zone Set**

Zone Set Definition currently in effect and enforced by the Fabric or other entity (e.g., the Name Server)

##### **3.1.2 address assignment**

process whereby addresses are dispensed to Switches and Switch ports

##### **3.1.3 address identifier**

unsigned 24-bit address value used to uniquely identify the source (S\_ID) and destination (D\_ID) of Fibre Channel frames (see FC-FS-5)

##### **3.1.4 Address Manager**

logical entity within a Switch that is responsible for address assignment

##### **3.1.5 Adjacent Switch**

remote Switch that does not require intermediate Switches in order to be reached

##### **3.1.6 Adjacency**

relationship between two Switches that have synchronized their topology databases

##### **3.1.7 Adjacent**

two Switches that have synchronized their topology databases are considered Adjacent

##### **3.1.8 AISL (Augmented ISL)**

E\_Port to E\_Port link between two Controlling Switches (see 17)

##### **3.1.9 AISL\_Set**

set of AISLs of a Controlling Switch

##### **3.1.10 Alternate Controlling Switch**

Controlling Switch operating as an Alternate (see 4.7.3)

##### **3.1.11 A\_Port (Adjacent port)**

combination of one PA\_Port and one VA\_Port operating together

##### **3.1.12 ASL (A\_Port Switch link)**

A\_Port to A\_Port link

##### **3.1.13 Area**

second level in the three-level address partitioning system specified by this standard (see 4.9)

##### **3.1.14 Area Identifier**

bits 15 through 8 of an address identifier

##### **3.1.15 Broadcast Address**

FFFFFFh value in the D\_ID field shall specify that the frame be broadcast to all Nx\_Ports

##### **3.1.16 Broadcast Zone**

Zone with the Broadcast Zone Attribute specified

**3.1.17 Broadcast Zoning Enforcement**

Zoning technique where the Fabric limits broadcast distribution among Zone Members using frame-by-frame filtering techniques

**3.1.18 B\_Port (Bridge port)**

Fabric inter-element port used to connect bridge devices with E\_Ports on a Switch

Note 1 to entry: The B\_Port provides a subset of the E\_Port functionality.

**3.1.19 Class F service**

service that multiplexes frames at frame boundaries and is used for control and coordination of the internal behavior of the Fabric

**3.1.20 Class N service**

any class of service other than Class F

**3.1.21 Controlling Switch**

Switch that is able to control a set of FCDFs in order to create a Distributed Switch

**3.1.22 Controlling Switch Set**

Switch\_Names of the Controlling Switches that are part of a Distributed Switch

**3.1.23 Core Switch**

set of entities with the same Core Switch\_Name that may host multiple Virtual Switches

**3.1.24 Core Switch\_Name**

Switch\_Name identifying the Core Switch (see 12.2) in a Virtual Fabric capable Switch

**3.1.25 Distributed Service**

implementation of a Generic Service operating at a well-known address (see FC-GS-8)

**3.1.26 Distributed Switch**

set of FCDFs associated with at least one Controlling Switch that controls the operations of the set of FCDFs

**3.1.27 Domain**

highest level in the three-level address partitioning system specified by this standard (see 4.9)

**3.1.28 Domain Address Manager**

Switch that is responsible for address assignment to other Switches outside of its Domain

**3.1.29 Domain Identifier (Domain\_ID)**

bits 23 through 16 of an address identifier

**3.1.30 Domain\_ID\_List**

list where each record contains a Domain\_ID value and the Switch\_Name of the Switch to which that Domain\_ID (see 6.2.5) is assigned

**3.1.31 Downstream Principal ISL**

Principal ISL to which frames may be sent from the Principal Switch to the destination Switch

Note 1 to entry: All Principal ISLs on the Principal Switch are downstream Principal ISLs.

### **3.1.32 Domain\_ID Overlap**

condition in which the Domain\_ID List of a Switch and the Domain\_ID List of a received EFP (see 7.3) are both non-null and have records that associate the same Domain\_ID to different Switch\_Names

### **3.1.33 Entry Switch**

role that a Switch assumes with respect to a Distributed Service request

### **3.1.34 E\_Port (Expansion port)**

Switch port that attaches to another Interconnect\_Port to create an Inter-Switch Link

### **3.1.35 E\_Port index**

index value associated with an E\_Port used by the Fabric Shortest Path First Protocol

### **3.1.36 Error\_Detect\_Timeout Value (E\_D\_TOV)**

time constant defined in FC-FS-5

### **3.1.37 F\_Port (Fabric port)**

combination of one PF\_Port and one VF\_Port operating together

### **3.1.38 Fabric**

entity that interconnects various Nx\_Ports attached to it, and is capable of routing frames using only the D\_ID information in an FC-2 Frame\_Header

### **3.1.39 Fabric Controller**

logical entity responsible for operation of the Fabric identified by the well-known address FFFFFFFh

### **3.1.40 Fabric VE\_ID**

{N\_Port\_ID, Local VE ID} pair of values used to identify a VE within a Fabric

### **3.1.41 FCDF (FC Data-Plane Forwarder)**

simplified FC switching entity that forwards FC frames among VA\_Ports and VF\_Ports through a FCDF Switching Element (see 4.7.4)

### **3.1.42 FCDF\_Set**

Switch\_Names of the FCDFs that are part of a Distributed Switch

### **3.1.43 F\_Port Controller**

logical entity at the well-known address FFFFFFFEh

### **3.1.44 flood**

process used to send information to all Switches within the Fabric

### **3.1.45 FL\_Port (Fabric Loop port)**

L\_Port that is able to perform the function of an F\_Port, attached via a link to one or more NL\_Ports in an Arbitrated Loop topology (see FC-AL-2)

Note 1 to entry: The AL\_PA of an FL\_Port is 00h.

Note 2 to entry: In this Standard, an FL\_Port is assumed to always refer to a port to which NL\_Ports are attached to a Fabric, and does not include F\_Ports.

### **3.1.46 Fx\_Port**

Switch port capable of operating as an F\_Port or FL\_Port

**3.1.47 Fabric\_Stability\_Timeout Value (F\_S\_TOV)**

time constant used to ensure that Fabric stability has been achieved during Fabric Configuration

**3.1.48 Fabric Shortest Path First (FSPF)**

link state protocol used for path selection

**3.1.49 G\_Port (Generic Fabric port)**

Switch port that may function either as an E\_Port, A\_Port, or as an F\_Port

**3.1.50 GL\_Port (Generic Fabric Loop port)**

Switch port that may function either as an E\_Port, A\_Port, or as an Fx\_Port

**3.1.51 Global VE ID**

identifier used to uniquely identify a Virtual Entity

Note 1 to entry: For example a 16-byte Universally Unique Identifier (UUID, see RFC 4122).

**3.1.52 Hard Zone**

Zone with the Hard Zone Attribute specified

**3.1.53 Hard Zoning**

Zoning enforcement in which the Fabric limits frame exchange by frame-by-frame filtering

**3.1.54 Interconnect\_Port**

generic reference to an E\_Port or a B\_Port

**3.1.55 Inter-Switch Link (ISL)**

link directly connecting the E\_Port of one Switch to the E\_Port of another Switch

**3.1.56 Isolated**

condition in which it has been determined that no Class N traffic may be transmitted across an ISL (see 7.6)

**3.1.57 L\_Port (Loop port)**

FC\_Port that contains Arbitrated Loop functions associated with the Arbitrated Loop topology (see FC-AL-2)

**3.1.58 Leaf Switch**

a Switch operating as a termination point for routes (i.e., does not allow routes through the Switch)

**3.1.59 Local VE\_ID**

value used to locally identify a VE within a VEM

Note 1 to entry: The local VE ID has a scope local to an N\_Port\_ID and is carried in the Priority field of the Frame\_Header.

**3.1.60 Locally-Enabled VF\_ID List**

configured list of VF\_IDs that an FC\_Port supporting Virtual Fabrics is able to enable on a link

**3.1.61 Multiplexer**

instance of the FC-2M sublevel, multiplexing and demultiplexing frames between physical and virtual ports based on the D\_ID/S\_ID and/or VF\_ID (see FC-FS-5)

**3.1.62 N\_Port (Node port)**

direct Fabric-attached PN\_Port (see FC-FS-5)

**3.1.63 N\_Port Identifier (N\_Port\_ID)**

address identifier assigned to an N\_Port

**3.1.64 Name\_Identifier**

64-bit identifier (see FC-FS-5)

**3.1.65 NL\_Port (Node Loop port)**

PN\_Port that is operating a Loop port state machine (see FC-AL-2)

Note 1 to entry: In this Standard, an NL\_Port is assumed to always refer to a loop-attached port, and does not include N\_Ports.

**3.1.66 Non-zero Domain\_ID\_List**

Domain\_ID\_List that contains at least one record (see 7.3)

**3.1.67 NPIV FDISC**

FDISC used to perform N\_Port\_ID Virtualization (see FC-LS-3)

**3.1.68 Nx\_Port**

end point for Fibre Channel frame communication (see FC-FS-5)

**3.1.69 PA\_Port (Physical A\_Port)**

Link\_Control\_Facility (see FC-FS-5) within the Fabric that attaches to another PA\_Port through a link

**3.1.70 path**

route through the Fabric between a source and a destination

**3.1.71 path selection**

process whereby paths are selected

**3.1.72 PE\_Port (Physical E\_Port)**

Link\_Control\_Facility (see FC-FS-5) within the Fabric that attaches to another PE\_Port or to a B\_Port through a link

**3.1.73 PF\_Port (Physical F\_Port)**

Link\_Control\_Facility (see FC-FS-5) within the Fabric that attaches to a PN\_Port (see FC-FS-5) through a link

**3.1.74 PN\_Port**

entity that includes a Link\_Control\_Facility and one or more Nx\_Ports (see FC-FS-5)

**3.1.75 port**

N\_Port, NL\_Port, F\_Port, FL\_Port, B\_Port, A\_Port, or E\_Port

**3.1.76 port VF\_ID**

configurable VF\_ID that is associated with any untagged frame received by a VF capable PE\_Port or PF\_Port

**3.1.77 point-to-point link**

Fibre Channel link connecting two ports



**3.1.78 Port Index**

three byte value used by FSPF to identify Switch ports

**3.1.79 port mode**

E\_Port, B\_Port, F\_Port, A\_Port, or FL\_Port operation

**3.1.80 Preferred Domain\_ID**

Domain\_ID previously granted to a Switch by the Domain Address Manager or through administrative means

**3.1.81 Principal ISL**

Inter-Switch Link that is used to communicate with the Principal Switch

**3.1.82 Principal Switch**

Switch that has been selected to perform as the Domain Address Manager for a Fabric (see 7)

**3.1.83 Primary Controlling Switch**

Controlling Switch operating as a Primary (see 4.7.3)

**3.1.84 remote Switch**

Switch that may be reached via one or more ISLs

**3.1.85 Resource\_Allocation\_Timeout Value (R\_A\_TOV)**

time constant defined in FC-FS-5

**3.1.86 Router**

entity within a Switch responsible for the routing of connectionless frames (see 4.5)

**3.1.87 routing**

process to identify the appropriate Switch port(s) to deliver a connectionless frame towards its destination

**3.1.88 Secondary Controlling Switch**

Controlling Switch operating as a Secondary (see 4.7.3)

**3.1.89 Soft Zoning**

Zoning enforcement in which the Fabric enforces membership through Name Server visibility

**3.1.90 Switch**

smallest entity that may function as a Fibre Channel Fabric (see 4.5)

**3.1.91 Switch Construct**

entity within a Switch responsible for transporting frames between Switch ports (see 4.5)

**3.1.92 Switching Element**

set of functions performed by the Path Selector, Router, Switch Construct, Address Manager, and Fabric Controller (see 4.5)

**3.1.93 Switch\_Name**

Name\_Identifier that identifies a Switch or a Bridge device

Note 1 to entry: The format of the name is specified in FC-FS-5.

### **3.1.94 Switch port**

E\_Port, A\_Port, F\_Port, or FL\_Port

### **3.1.95 Switch\_Priority**

value used during Principal Switch selection to cause one Switch to be favored over another

### **3.1.96 T10 Vendor ID**

character string that uniquely identifies a vendor

### **3.1.97 topology**

communication infrastructure that provides Fibre Channel communication among a set of PN\_Ports (e.g., a Fabric, an Arbitrated Loop, or a combination of the two)

### **3.1.98 Upstream Principal ISL**

Principal ISL to which frames may be sent from the local Switch to the Principal Switch

Note 1 to entry: A Switch that is not the Principal Switch always has exactly one upstream Principal ISL.

Note 2 to entry: The Principal Switch does not have an upstream Principal ISL.

### **3.1.99 VA\_Port (Virtual A\_Port)**

instance of the FC-2V sublevel of Fibre Channel that connects to another VA\_Port

### **3.1.100 VEM ID**

identifier used to uniquely identify a VEM

Note 1 to entry: For example a 16-byte Universally Unique Identifier (UUID, see RFC 4122).

### **3.1.101 Virtual Entities Manager (VEM)**

entity managing Virtual Entities

Note 1 to entry: For example a Hypervisor.

### **3.1.102 Virtual Entity (VE)**

virtualized resource

Note 1 to entry: For example a Virtual Machine (VM).

### **3.1.103 Virtual Fabric**

interconnected set of Virtual Switches and/or Switches identified by a Virtual Fabric ID (VF\_ID)

### **3.1.104 Virtual Fabric Tagging Header (VFT\_Header)**

Extended\_Header (see FC-FS-5) used to support Virtual Fabrics

### **3.1.105 VE\_Port (Virtual E\_Port)**

instance of the FC-2V sublevel that connects to another VE\_Port or to a B\_Port to create an Inter-Switch Link

Note 1 to entry: A VE\_Port is addressable by the VE\_Port or B\_Port connected to it through the Fabric Controller well-known address identifier (i.e., FFFFFFFDh).

### **3.1.106 VF\_Port (Virtual F\_Port)**

instance of the FC-2V sublevel that connects to one or more VN\_Ports (see FC-FS-5)

Note 1 to entry: A VF\_Port is addressable by a VN\_Port connected to it through the F\_Port Controller well-known address identifier (i.e., FFFFh).

### **3.1.107 Zero Domain\_ID\_List**

Domain\_ID\_List that is empty (see 7.3)

### **3.1.108 Zone**

group of Zone Members

### **3.1.109 Zone Definition**

parameters that define a Zone

### **3.1.110 Zone Member**

device included in a Zone

### **3.1.111 Zone Member Definition**

parameters that define a Zone Member

### **3.1.112 Zone Name**

name assigned to a Zone

### **3.1.113 Zone Set**

set of Zones that are used simultaneously

### **3.1.114 Zone Set Database**

database that contains the Zone Sets not enforced by the Fabric

### **3.1.115 Zone Set Name**

name assigned to a Zone Set

### **3.1.116 Zone Set State**

state of a Zone Set (i.e., Activated or Deactivated)

### **3.1.117 Zoning ACLs**

Access Control Lists used by an FCDF to enforce Zoning

### **3.1.118 Zoning Database**

term used to indicate both the Active Zone Set and the Zone Set Database

## **3.2 Editorial conventions**

In FC-SW-7, 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.

### 3.3 State machine notation

State machines in this standard use the style shown in figure 1.

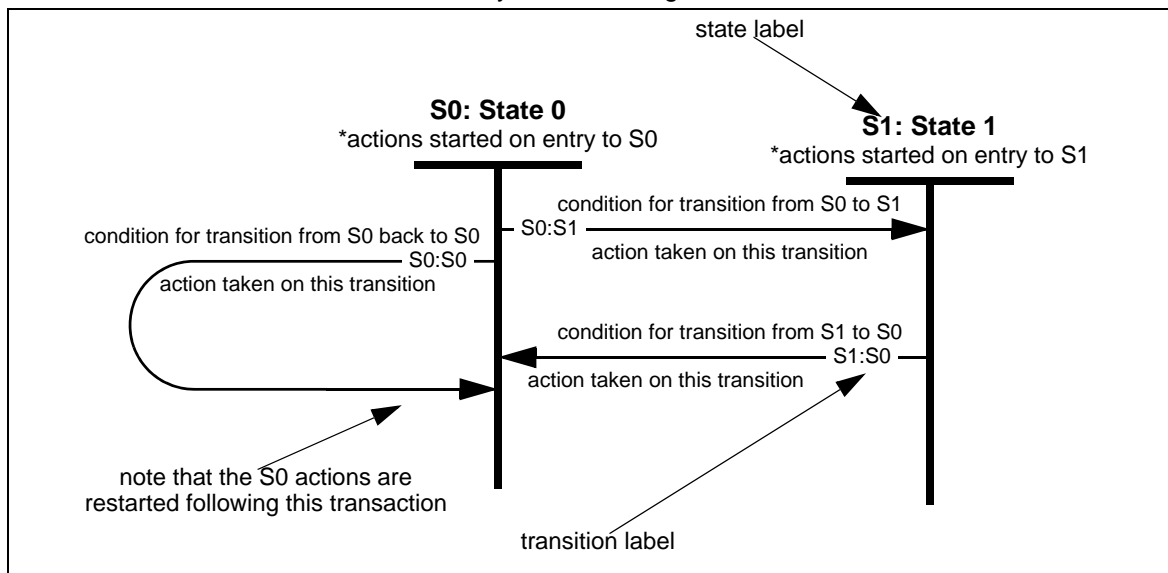


Figure 1 – Sample state machine

These state machines make three assumptions:

- a) time elapses only within discrete states;
- b) state transitions are logically instantaneous, so the only actions taken during a transition are setting flags and variables and sending signals. These actions complete before the next state is entered; and
- c) every time a state is entered, the actions of that state are started. Note that this means that a transition that points back to the same state repeats the actions from the beginning. All the actions started upon entry complete before any tests are made to exit the state.

### 3.4 Abbreviations, acronyms, and symbols

Abbreviations and acronyms applicable to this International Standard are listed below. Definitions of several of these items are included in 3.1. Abbreviations used that are not listed below are defined in FC-FS-5.

<b>Area_ID</b>	Area Identifier
<b>CT</b>	Common Transport
<b>Domain_ID</b>	Domain Identifier
<b>D_S_TOV</b>	Distributed_Services_Timeout Value
<b>E_D_TOV</b>	Error_Detect_Timeout value
<b>ELS</b>	Extended Link Service
<b>FAN</b>	Fabric Address Notification Extended Link Service
<b>FSM</b>	finite state machine
<b>FC-AL-2</b>	Fibre Channel Arbitrated Loop - 2, reference [1]
<b>FC-BB-6</b>	Fibre Channel Backbone - 6, reference [4]
<b>FC-DA-2</b>	Fibre Channel- Device Attach - 2, reference [3]
<b>FC-FS-4</b>	Fibre Channel - Framing and Signaling - 4, reference [12]
<b>FC-GS-7</b>	Fibre Channel - Generic Services - 7, reference [11]
<b>FC-IFR</b>	Fibre Channel - Inter-Fabric Routing, reference [9]
<b>FC-LS-2</b>	Fibre Channel - Link Services - 2, reference [5]
<b>FC-LS-3</b>	Fibre Channel - Link Services - 3, reference [10]
<b>FC-MI-3</b>	Fibre Channel - Methodologies for Interconnects - 3, reference [2]
<b>FC-PI-6</b>	Fibre Channel -Physical Interface - 6, reference [13]
<b>FC-SB-5</b>	Fibre Channel - Single Byte Command Sets - 5 reference [6]
<b>FC-SP-2</b>	Fibre Channel - Security Protocols - 2 reference [7]
<b>F_S_TOV</b>	Fabric_Stability_Timeout value
<b>FDMI</b>	Fabric Device Management Interface
<b>FSPF</b>	Fabric Shortest Path First
<b>ISL</b>	Inter-Switch Link
<b>IU</b>	Information Unit
<b>LCF</b>	Link Control Facility
<b>LFA</b>	Loop Fabric Address
<b>LSR</b>	Link State Record
<b>R</b>	Reserved
<b>R_A_TOV</b>	Resource_Allocation_Timeout value
<b>RFC</b>	Request For Comment
<b>SM</b>	State Machine
<b>SW_ACC</b>	Switch Fabric Link Service Accept
<b>SW_ILS</b>	Switch Internal Link Service
<b>SWN</b>	Switch Name
<b>SWP</b>	Switch Priority
<b>SW_RJT</b>	Switch Fabric Link Service Reject
<b>VF_ID</b>	Virtual Fabric Identifier
<b>WKA</b>	well-known address
<b>WWN</b>	World Wide Name
<b>1xAL_TIME</b>	One times the AL_TIME
<b>1xF_S_TOV</b>	One times the F_S_TOV
<b>2xAL_TIME</b>	Two times the AL_TIME
<b>2xF_S_TOV</b>	Two times the F_S_TOV
<b>3xAL_TIME</b>	Three times the AL_TIME
<b>=</b>	Is equal to

### 3.5 Definition of compliance terms

The usual definitions of the following terms do not apply in this standard and therefore they are defined below:

**Prohibited:** If a feature or parameter value is Prohibited, it means that it shall not be used between compliant implementations.

**Required:** If a feature or parameter value is Required, it means that it shall be used between compliant implementations.

**Allowed:** If a feature or parameter value is Allowed, it means that it may be used between compliant implementations.

### 3.6 Keywords

**3.6.1 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.

**3.6.2 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.

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

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

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

**3.6.6 obsolete:** A keyword indicating that an item was defined in prior Fibre Channel standards but has been removed from this standard.

**3.6.7 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.

**3.6.8 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 error.

**3.6.9 restricted:** A keyword referring to bits, bytes, words, and fields that are set aside for use in other Fibre Channel standards. A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word or field for the purposes of the requirements defined in this standard.

**3.6.10 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.

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

### 3.7 T10 Vendor ID fields

A T10 Vendor ID shall be a string of one to eight characters that is recorded in an informal list of Vendor IDs maintained by INCITS Technical Committee T10 (see <http://www.t10.org>).

A field described as containing a T10 Vendor ID shall contain the first character of the T10 Vendor ID in the highest order byte of the field, and successive characters of the T10 Vendor ID in successively lower order bytes of the field. Any bytes of the field not filled by characters of the T10 Vendor ID shall be filled with ASCII space characters (20h).

### 3.8 Left-aligned ASCII data

ASCII data fields described as being left-aligned shall contain only ASCII printable characters (i.e., code values 20h to 7Eh), shall contain the first character of the data in the highest order byte of the field, and successive characters of the data in successively lower order bytes of the field. Any bytes of the field not filled by characters of the data shall be filled with ASCII space characters (i.e., 20h).

## 4 Structure and concepts

### 4.1 Overview

This standard describes the operation and interaction of Fibre Channel Switches. This includes E\_Port operation, Fabric operation, A\_Port operation, and Distributed Switch operation.

### 4.2 E\_Port operation

E\_Port operation specifies the tools and algorithms for interconnection and initialization of Fibre Channel Switches to create a multi-Switch Fabric. Fabric operation defines an E\_Port (i.e., Expansion port) that operates in a manner similar to a PN\_Port and F\_Port, as defined in FC-FS-5, with additional functionality provided for interconnecting Switches.

E\_Port operation defines credit models and management between E\_Ports for the various classes of service other than Class F. E\_Ports conforming to this standard support Class F, Class 2 and/or Class 3. Support for other classes of service is not defined by Fabric operation.

E\_Port operation defines how ports that are capable of being an E\_Port, F\_Port, and/or FL\_Port discover and self-configure for their appropriate operating mode. Once a port establishes that it is connected to another Switch and is operating as an E\_Port, an address assignment algorithm is executed to allocate port addresses throughout the Fabric.

### 4.3 Fabric operation

Fabric operation includes the following:

- a) Fabric Configuration - describes how a Principal Switch is selected and describes the address assignment algorithm;
- b) Exchange Switch Capabilities - allows Switches to exchange certain operational capabilities such as which path selection protocols are supported;
- c) B\_Port - a simplified E\_Port that allows Bridge type devices to participate in Fabric operation;
- d) Path selection - path selection is the process by which a Switch determines the best path from a source domain to a destination domain. These paths may then be used in any appropriate manner by the Switch to move frames to their destinations. This path selection process does not require nor preclude the use of static or dynamic load-balancing. This standard defines the Fabric Shortest Path First (FSPF) protocol;
- e) Distributed Server communication - the Distributed Server model allows the Generic Services operating at well-known addresses (see FC-GS-7) to be distributed among Switches that comprise the Fabric. Both the distributed Name Server and the distributed Management Server are described. In addition, the Inter-Switch FDML protocol has been defined;
- f) Exchange of Zoning information - defines how Zoning information is communicated between Switches in the Fabric. Zoning information is exchanged between Switches when two Fabrics are merged, and when changes to Zoning information are propagated between Switches;
- g) Distributed Event Notification - defines how Registered State Change Notifications (RSCNs) and Distribute Registered Link Incident Records (DRLIR) are distributed between Switches in the Fabric;
- h) Virtual Fabrics Switch Support - defines the operation of Virtual Fabrics from a frame tagging perspective;
- i) Enhanced Commit Service - defines a commit service that allows serialization based on application and advanced error recovery; and
- j) Virtual Channel Architecture - defines virtual channels between E\_Ports.



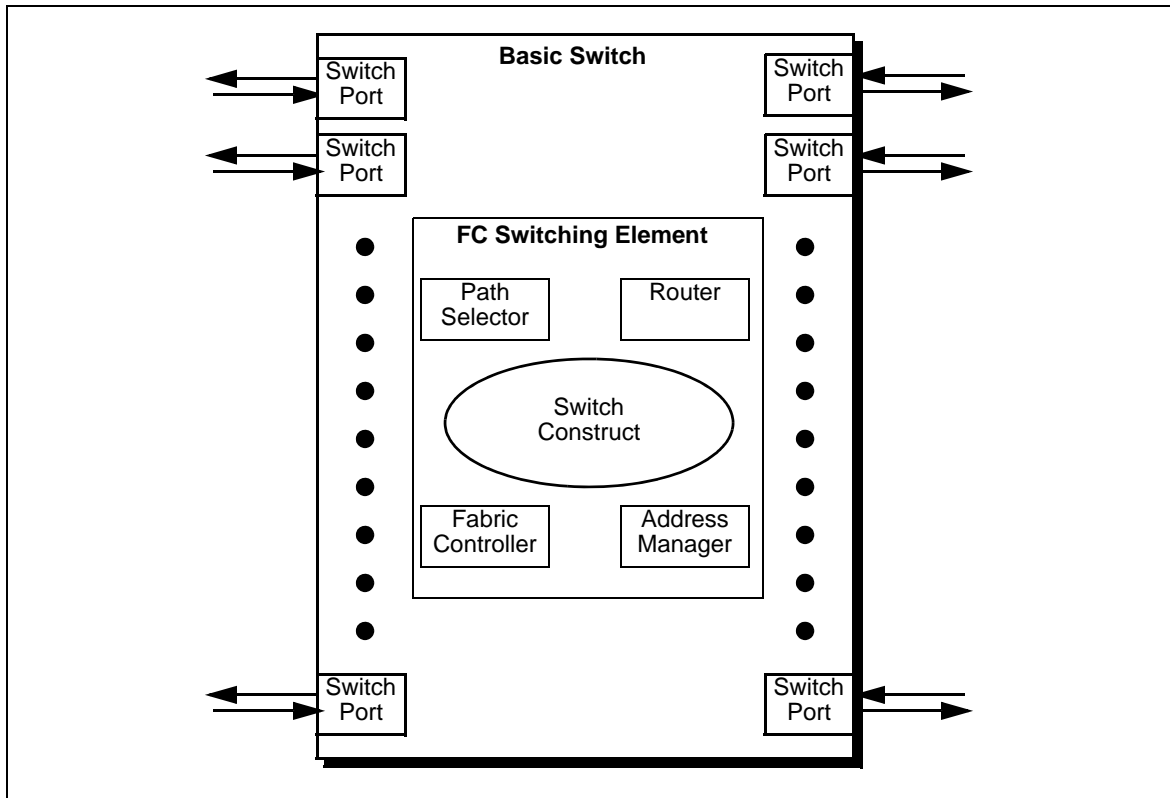
In addition to normal Fabric operation and topologies the Distributed Switch operation is described in 17. A\_Port operation (see 4.8.2) is limited to the Distributed Switch operation.

#### 4.4 Fabric definition

The Fabric serves as a transport that provides a switched interconnect between Nx\_Ports.

#### 4.5 Switch

A Switch is the smallest entity that may function as a Switch-based Fibre Channel Fabric. Figure 2 illustrates the conceptual model of a Basic Switch.



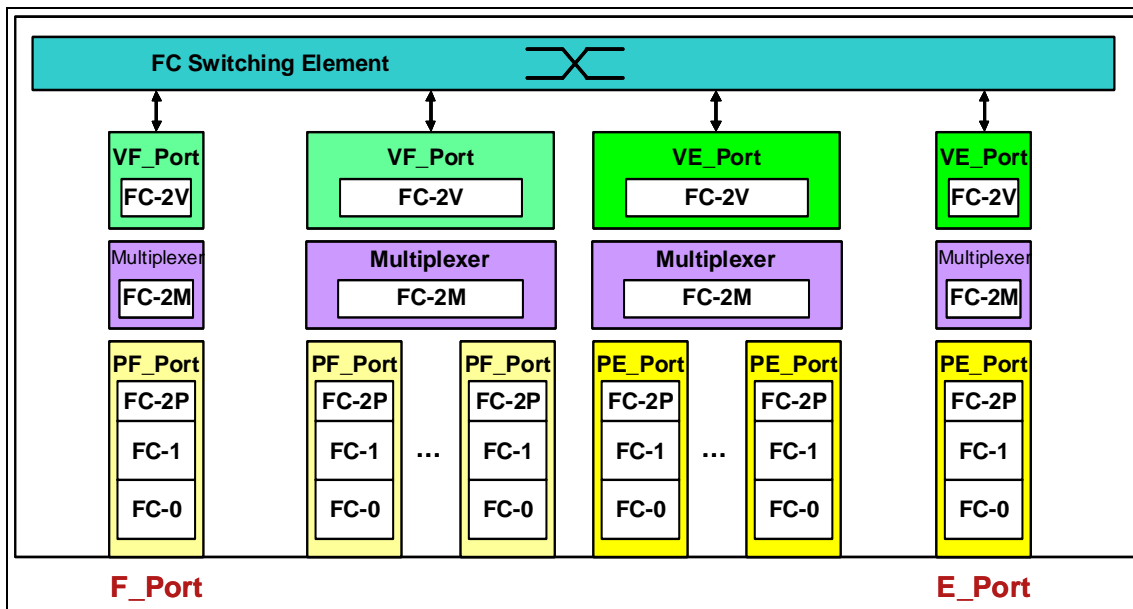
**Figure 2 – Basic Switch model**

A Switch is composed of the following major components:

- a) two or more Switch ports;
- b) a Switch Construct, capable of either multiplexed frame switching or circuit switching, or both;
- c) an Address Manager;
- d) a Path Selector, which performs path selection;
- e) a Router; and
- f) a Fabric Controller.

The set of functions performed by the Path Selector, The Router, The Switch Construct, the Address Manager and the Fabric Controller is referred to as a FC Switching Element.

The FC Switching Element abstraction allows the definition of an Enhanced Switch functional model, based on the functional layering of Fibre Channel (see FC-FS-5). Figure 3 shows the functional model of an Enhanced Switch.




**Figure 3 – Enhanced Switch functional model**

An Enhanced Switch is able to aggregate its physical ports in sets that behave as virtual ports, providing higher bandwidth than the one available to a single physical port. A physical port is an LCF (see FC-FS-5), that may behave as a Physical F\_Port (PF\_Port) or as a Physical E\_Port (PE\_Port). A virtual port is an instance of the FC-2V sublevel of Fibre Channel (see FC-FS-5), that may behave as a Virtual F\_Port (VF\_Port) or as a Virtual E\_Port (VE\_Port).

As defined, a Switch port may be either an E\_Port, an F\_Port, or an FL\_Port. A Switch port that is capable of assuming more than one of these roles is called a G\_Port or GL\_Port. Once a Switch port assumes a role, via the Switch port initialization procedure, it shall remain in that role until an event occurs that causes re-initialization.

The link joining a pair of E\_Ports is called an Inter-Switch Link (ISL). ISLs carry frames originating from Nx\_Ports and those frames generated within the Fabric. The frames generated within the Fabric serve as control, management and support for the Fabric.

Switches may be joined freely or in a structured fashion to form a larger Fabric, as illustrated in figure 4.



**Figure 4 – Multiple Switch Fabric example**

The structure of the Switch Construct in the Switch, as seen in figure 2, is undefined and beyond the scope of this standard. It may support either or both circuit switching and multiplexed frame switching. It may be non-blocking, allowing concurrent operation of all possible combinations or it may be blocking, restricting operations. The Switch Construct may also contain redundancy, as may be required for high availability configurations.

The Address Manager is responsible for the assignment of addresses within some portion of the Fabric. Within the Switch, the Address Manager is responsible for Domain\_ID(s) for the Switch, and allocating Area\_IDs and Port\_IDs within the acquired Domain(s).

The Path Selector is a logical entity that establishes frame routing paths.

The Router is a logical entity that performs the routing of Class F, Class 2, and Class 3 frames to their final destination.

The Fabric Controller is a logical entity that performs the management of the Switch. The Fabric Controller is addressed as a VN\_Port, though it may or may not be attached to the Fabric via a link.

#### **4.6 Switching characteristics**

Path, circuit switching, and frame routing within a Switch occurs synchronously to the current word alignment of the outbound fibre.

Synchronous switching guarantees retention of the established word alignment on the outbound fibre of the Switch port.

A switching event occurs every time a connectionless frame is transmitted and when a connection based service is established, suspended or terminated.

Synchronous switching associated with connectionless frame routing and connection oriented Dedicated Connections or virtual connection Services shall guarantee the word alignment on the outbound fibre. Switches shall ensure that synchronous switching only occurs between frames.

#### **4.7 Distributed Switch and A\_Port operations**

##### **4.7.1 Overview**

A Distributed Switch is a set of FCDFs associated with at least one Controlling Switch that controls the operations of the set of FCDFs. Distributed Switch operation is specified in 17.

##### **4.7.2 A\_Port operation**

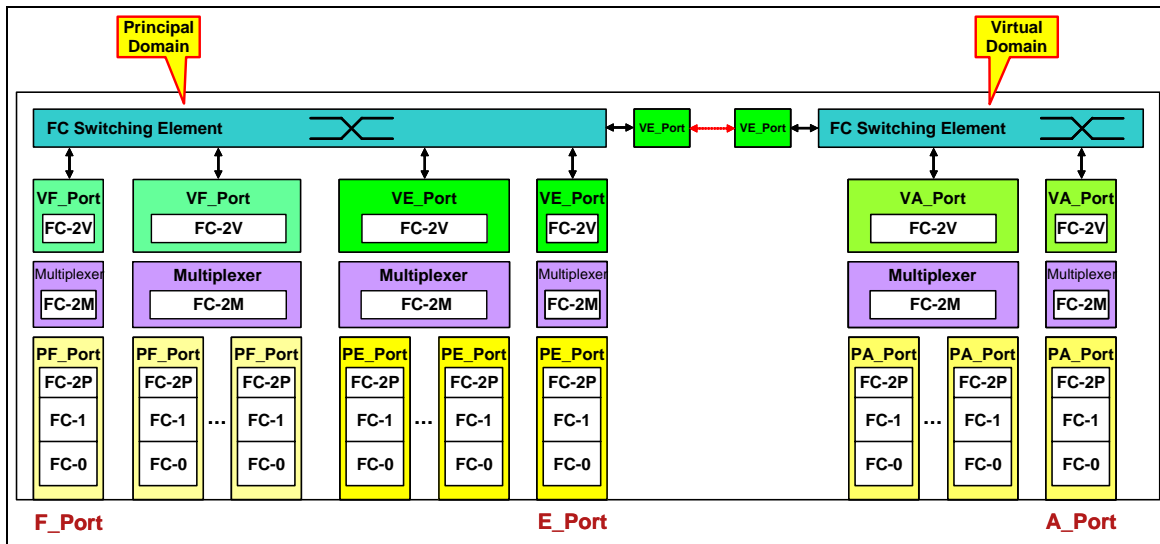
A\_Port operation specifies the tools and algorithms for interconnection and initialization of Controlling Switches and FCDFs to create a Distributed Switch.

A\_Port operation defines credit models and management between A\_Ports for the various classes of service other than Class F. A\_Ports conforming to this standard support Class F, Class 2, and/or Class 3. A\_Port support for these classes of service is the same as for E\_Ports.

A\_Port operation defines how ports that are capable of being an A\_Port, E\_Port, F\_Port, and/or FL\_Port discover and self-configure for their appropriate operating mode.

### 4.7.3 Controlling Switch functional model

Figure 5 shows the functional model of a Controlling Switch.



**Figure 5 – Controlling Switch functional model**

A Controlling Switch is an FC Switch that supports the instantiation of VA\_Ports, VF\_Ports, and VE\_Ports. As any Enhanced FC Switch, a Controlling Switch is able to aggregate its physical ports in sets that behave as virtual ports, providing higher bandwidth than the one available to a single physical port.

For a Controlling Switch, a physical port is an LCF (see FC-FS-5), that may behave as a Physical F\_Port (PF\_Port), as a Physical E\_Port (PE\_Port), or as a Physical A\_Port (PA\_Port). A virtual port is an instance of the FC-2V sublevel of Fibre Channel (see FC-FS-5), that may behave as a Virtual F\_Port (VF\_Port), as a Virtual E\_Port (VE\_Port), or as a Virtual A\_Port (VA\_Port).

As shown in figure 5, a Controlling Switch is functionally modeled as having two FC Switching Elements, one for the Principal Domain and one for the Virtual Domain, connected by an internal VE\_Port to VE\_Port link. The Switching Element associated with the Principal Domain supports the instantiation of VF\_Ports and VE\_Ports, the Switching Element associated with the Virtual Domain supports the instantiation of VA\_Ports.

If operating as a Primary, a Controlling Switch:

- controls the allocation of N\_Port\_IDs to N\_Ports connected to the FCDFs of the Distributed Switch;
- maintains a topology map of the Distributed Switch;
- computes and distributes Virtual Domain routes for the FCDFs and Alternate Controlling Switches of the Distributed Switch; and
- distributes Zoning information to the FCDFs of the Distributed Switch.

If operating as a Secondary, a Controlling Switch:

- receives N\_Port\_ID allocation information from the Primary Controlling Switch;
- maintains a topology map of the Distributed Switch; and

- c) keeps its state synchronized with the Primary Controlling Switch and is able to take its place in case of failure (see 17.5).

If operating as an Alternate, a Controlling Switch:

- a) receives N\_Port\_ID allocation information from the Primary Controlling Switch;
- b) receives Virtual Domain routes for the Distributed Switch from the Primary Controlling Switch; and
- c) is able to take the place of the Secondary Controlling Switch in case of failure of the Secondary Controlling Switch or if the Secondary Controlling Switch becomes a Primary Controlling Switch (i.e., the Primary Controlling Switch has failed). See 17.5.

#### 4.7.4 FCDF functional model

Figure 6 shows the functional model of an FCDF.

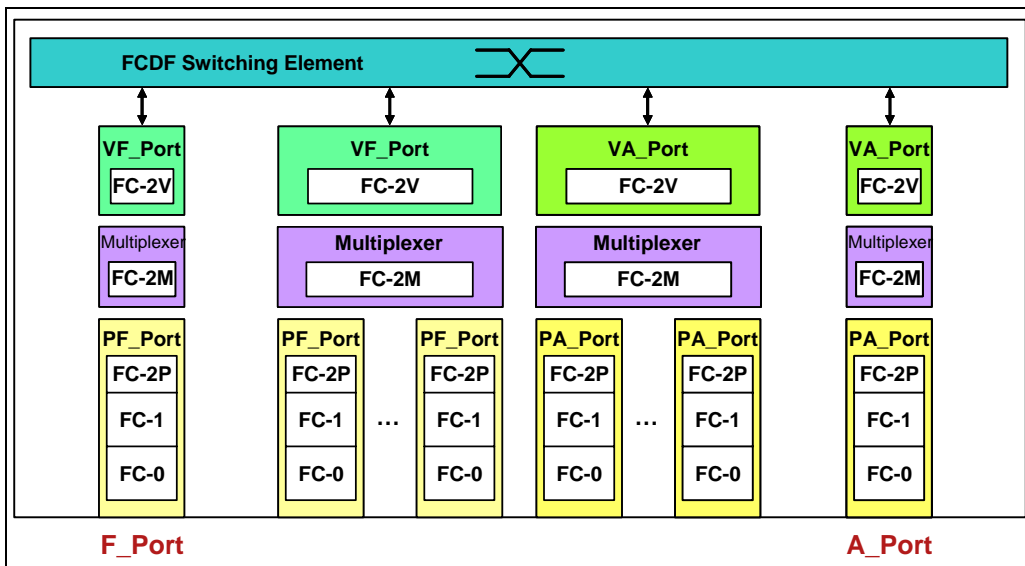
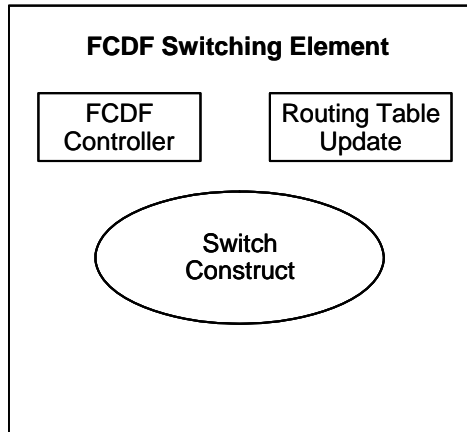


Figure 6 – FCDF functional model

An FCDF is a simplified FC switching entity that forwards FC frames among VA\_Ports and VF\_Ports through a FCDF Switching Element. As any FC Switch, an FCDF is able to aggregate its physical ports in sets that behave as virtual ports, providing higher bandwidth than the one available to a single physical port. An FCDF shall support at least one VA\_Port operating together with a PA\_Port (i.e., an A\_Port) and may support one or more F\_Ports.

For an FCDF, a physical port is an LCF (see FC-FS-5), that may behave as a Physical F\_Port (PF\_Port) or as a Physical A\_Port (PA\_Port). A virtual port is an instance of the FC-2V sublevel of Fibre Channel (see FC-FS-5), that may behave as a Virtual F\_Port (VF\_Port) or as a Virtual A\_Port (VA\_Port).

Figure 7 shows the model of the FCDF Switching Element, composed of a Switch Construct, a Routing Table Update function, and an FCDF Controller function.



**Figure 7 – FCDF Switching Element**

The Switch Construct is the entity performing FC frames forwarding based on the FC frame's D\_ID field according to a routing table. The structure of the Switch Construct is undefined and beyond the scope of this standard.

The Routing Table Update is a logical entity that updates the Switch Construct's routing table through the VA\_Port protocols.

The FCDF Controller is a logical entity that performs the management of the FCDF through the VA\_Port protocols.

An FCDF also:

- a) receives N\_Port\_ID allocation information from the Primary Controlling Switch;
- b) receives routing information from the Primary Controlling Switch; and
- c) receives Zoning information for the Distributed Switch from the Primary Controlling Switch.

## **4.8 Switch ports and Bridge ports**

### **4.8.1 General characteristics**

A Switch shall have two or more Switch ports. A Switch equipped only with F\_Ports or FL\_Ports forms a non-expandable Fabric. To be part of an expandable Fabric, a Switch shall incorporate at least one Switch port capable of E\_Port operation.

A Switch port supports one or more of the following port modes:

- a) F\_Port (see 4.8.2);
- b) FL\_Port (see 4.8.3);
- c) E\_Port (see 4.8.4);
- d) B\_Port (see 4.8.5); or
- e) A\_Port (see 4.8.6).

Switch ports that assume either the E\_Port or B\_Port mode are generally referred to as Interconnect\_Ports. A Switch port that is capable of supporting more than one port mode attempts to

configure itself first as an FL\_Port (see FC-AL-2), then as an E\_Port or a B\_Port or an A\_Port, and finally as an F\_Port (see FC-FS-5), depending on which of the five port modes are supported by the Switch port. A Bridge port shall only support B\_Port operation.

NOTE 2 – The characteristics of a Bridge device are not described in this standard. However, one type of Bridge device is described in the FC-BB-6 standard.

The detailed procedure for port mode selection is described in 7.2.

#### **4.8.2 F\_Port**

An F\_Port is the point at which all frames originated by an N\_Port enter the Fabric, and all frames destined for an N\_Port exit the Fabric. An F\_Port may also be the Fabric entry point for frames originated by an N\_Port destined for an internal Fabric destination, such as the Fabric Controller. Similarly, an F\_Port may also be the Fabric exit point for frames originated internal to the Fabric and destined for an N\_Port. Frames shall not be communicated across a link between an F\_Port and anything other than an N\_Port. Functionally, an F\_Port is the combination of one PF\_Port and one VF\_Port operating together.

F\_Ports are described in detail in 5.3.

#### **4.8.3 FL\_Port**

An FL\_Port is the point at which all frames originated by an NL\_Port enter the Fabric, and all frames destined for an NL\_Port exit the Fabric. An FL\_Port may also be the Fabric entry point for frames originated by an NL\_Port destined for an internal Fabric destination, such as the Fabric Controller. Similarly, an FL\_Port may also be the Fabric exit point for frames originated internal to the Fabric and destined for an NL\_Port. Frames shall not be communicated across a link between an FL\_Port and anything other than an NL\_Port.

FL\_Ports are described in detail in 5.4.

#### **4.8.4 E\_Port**

An E\_Port is the point at which frames pass between the Switches within the Fabric. Frames with a destination other than the local Switch or any N\_Port or NL\_Port attached to the local Switch exit the local Switch through an E\_Port. Frames that enter a Switch via an E\_Port are forwarded to a local destination, or are forwarded towards their ultimate destination via another E\_Port. Frames shall not be communicated across a link between an E\_Port and anything other than an E\_Port or a B\_Port. Functionally, an E\_Port is the combination of one PE\_Port and one VE\_Port operating together.

E\_Ports are described in detail in 5.5.

#### **4.8.5 B\_Port**

A Bridge port (B\_Port) is a port on a Bridge device. It normally functions as a conduit between the Switch and the Bridge for frames destined for or through a Bridge device. A B\_Port is also used to carry frames between a Switch and the Bridge device for purposes of configuring the Bridge device.

B\_Ports are described in detail in 5.6.



**4.8.6 A\_Port**

An A\_Port (Adjacent port) is the point at which frames pass between FCDFs and between FCDFs and Controlling Switches within a Distributed Switch. Functionally, an A\_Port is the combination of one PA\_Port and one VA\_Port operating together.

A\_Ports are described in detail in 5.7.

**4.8.7 G\_Ports and GL\_Ports**

A G\_Port is a Switch port that is capable of either operating as an E\_Port, F\_Port, or A\_Port. A G\_Port determines through port initialization whether it operates as an E\_Port, as an F\_Port, or as an A\_Port. A GL\_Port is a G\_Port that is also capable of operating as an FL\_Port.

**4.8.8 PF\_Port**

A PF\_Port is the LCF within the Fabric that attaches to a PN\_Port (see FC-FS-5) through a link.

**4.8.9 VF\_Port**

A VF\_Port is an instance of the FC-2V sublevel that connects to one or more VN\_Ports (see FC-FS-5). A VF\_Port is addressable by a VN\_Port connected to it through the F\_Port Controller well-known address identifier (i.e., FF FF FEh).

**4.8.10 PE\_Port**

A PE\_Port is the LCF within the Fabric that attaches to another PE\_Port or to a B\_Port through a link.

**4.8.11 VE\_Port**

A VE\_Port is an instance of the FC-2V sublevel that connects to another VE\_Port or to a B\_Port to create an Inter-Switch Link. A VE\_Port is addressable by the VE\_Port or B\_Port connected to it through the Fabric Controller well-known address identifier (i.e., FF FF FDh).

**4.8.12 PA\_Port**

A PA\_Port is the LCF within the Fabric that attaches to another PA\_Port through a link.

**4.8.13 VA\_Port**

A VA\_Port is an instance of the FC-2V sublevel that connects to another VA\_Port to create an A\_Port Switch Link (ASL). A VA\_Port is addressable by the VA\_Port connected to it through the Fabric Controller well-known address identifier (i.e., FF FF F9h).

**4.9 Fabric addressing**

Switches use the address partitioning model as shown in figure 8.

2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	
Domain_ID								Area_ID								Port_ID								
Address identifier																								

Figure 8 – Domain, Area, and Port address partitioning

A Domain is one or more Switches that have the same Domain\_ID for all N\_Ports and NL\_Ports within or attached to those Switches, except for well-known addresses. If there is more than one Switch in the Domain, each Switch within the Domain shall be directly connected via an ISL to at least one other Switch in the same Domain.

An Area\_ID shall apply to either of the following:

- a) one or more N\_Ports within and attached to a single Switch, except for well-known addresses;  
or
- b) an Arbitrated Loop of NL\_Ports attached to a single FL\_Port.

A single Arbitrated Loop shall have exactly one Area\_ID.

A Port\_ID shall apply to either of the following:

- a) a single N\_Port within a Domain/Area, except for well-known addresses; or
- b) the valid AL\_PA of a single NL\_Port or FL\_Port on an Arbitrated Loop.

Address identifier values for this standard are listed in table 1. Any value listed as reserved is not meaningful within this standard.

**Table 1 – Address identifier values (Part 1 of 2)**

Address Identifier (hex)			Description
Domain_ID	Area_ID	Port_ID	
00	00	00	Undefined <sup>a</sup>
00	00	AL_PA	E_Port: Reserved F_Port: Reserved FL_Port: Private Loop NL_Port <sup>b</sup> and <sup>g</sup>
00	00	non-AL_PA	Reserved
00	01 - FF	00 - FF	Reserved
01 - EF	00 - FF	00	F_Port: N_Port Identifier <sup>h</sup> FL_Port: Loop Fabric Address <sup>c</sup>
01 - EF	00 - FF	AL_PA	F_Port: N_Port Identifier <sup>h</sup> FL_Port: N_Port Identifier for Public Loop NL_Port <sup>c</sup>
01 - EF	00 - FF	non-AL_PA	F_Port: N_Port Identifier <sup>h</sup> FL_Port: Reserved
F0 - FE	00 - FF	00 - FF	Reserved
FF	00 - F9	00 - FF	Reserved
FF	FA	00-0F	Reserved for Internal Loopback Addresses
FF	FA	10-1F	Reserved for External Loopback Addresses
FF	FA	20-FF	Reserved
FF	FB	00 - FF	Obsolete
FF	FC	00	Reserved
FF	FC	01 - EF	N_Port Identifier for Domain Controller <sup>d</sup>
FF	FC	F0 - FF	Reserved
FF	FD - FE	00 - FF	Reserved
FF	FF	00 - EF	Reserved

<sup>a</sup> This value is used by an N\_Port requesting an address identifier during FLOGI.

<sup>b</sup> See FC-AL-2 for a definition of AL\_PA and FC-DA for a definition of Private Loop and FL\_Port operation with Private Loop devices.

<sup>c</sup> See FC-AL-2 for the definition and use of Loop Fabric Address, and for a definition of Public Loop.

<sup>d</sup> A Domain Controller Identifier may be used to address the Fabric Controller of a remote Switch that may or may not be connected via an ISL to the originating Switch. The Port\_ID field is set to the Domain\_ID of the remote Switch.

<sup>e</sup> The usage of well-known addresses FFFFF0h through FFFFFCh, are not defined by this standard. FC-FS-5 defines or reserves these values for well-known addresses.

<sup>f</sup> This address identifier has special usage depending on the originator. If the originator is an attached external N\_Port or NL\_Port (i.e., attached via an F\_Port or FL\_Port) then the destination of a frame sent to FFFFFDh is the Fabric Controller of the local Switch. If the originator is the Fabric Controller of the local Switch, then the destination of a frame sent to FFFFFDh via an ISL is the Fabric Controller of the remote Switch at the other end of the ISL.

<sup>g</sup> This value is used by a public loop NL\_Port requesting an address identifier during FLOGI.

<sup>h</sup> A Switch may use the same Domain\_ID and Area\_ID for the N\_Port Identifier of N\_Ports attached to different F\_Ports.

**Table 1 – Address identifier values (Part 2 of 2)**

Address Identifier (hex)			Description
Domain_ID	Area_ID	Port_ID	
FF	FF	F0 - FC	well-known address <sup>e</sup>
FF	FF	FD	N_Port Identifier for Fabric Controller <sup>f</sup>
FF	FF	FE	N_Port Identifier for F_Port Controller
FF	FF	FF	Broadcast Address
<p><sup>a</sup> This value is used by an N_Port requesting an address identifier during FLOGI.</p> <p><sup>b</sup> See FC-AL-2 for a definition of AL_PA and FC-DA for a definition of Private Loop and FL_Port operation with Private Loop devices.</p> <p><sup>c</sup> See FC-AL-2 for the definition and use of Loop Fabric Address, and for a definition of Public Loop.</p> <p><sup>d</sup> A Domain Controller Identifier may be used to address the Fabric Controller of a remote Switch that may or may not be connected via an ISL to the originating Switch. The Port_ID field is set to the Domain_ID of the remote Switch.</p> <p><sup>e</sup> The usage of well-known addresses FFFFF0h through FFFFFCh, are not defined by this standard. FC-FS-5 defines or reserves these values for well-known addresses.</p> <p><sup>f</sup> This address identifier has special usage depending on the originator. If the originator is an attached external N_Port or NL_Port (i.e., attached via an F_Port or FL_Port) then the destination of a frame sent to FFFFFDh is the Fabric Controller of the local Switch. If the originator is the Fabric Controller of the local Switch, then the destination of a frame sent to FFFFFDh via an ISL is the Fabric Controller of the remote Switch at the other end of the ISL.</p> <p><sup>g</sup> This value is used by a public loop NL_Port requesting an address identifier during FLOGI.</p> <p><sup>h</sup> A Switch may use the same Domain_ID and Area_ID for the N_Port Identifier of N_Ports attached to different F_Ports.</p>			

#### 4.10 Class F service

Class F service is a connectionless service similar to Class 2 that is used for internal control of the Fabric. Class F service as used by this standard is defined in 5.9.

## 5 Switch ports and Bridge ports

### 5.1 Overview

This clause defines the specific behaviors for all modes of a Switch port and a Bridge port. Note that the models described below are defined for purposes of describing behavior. No implication is made as to whether the actual implementation of an element is in hardware or software. An element may be implemented on a per-port basis, or may be a logical entity that is embodied in a single physical implementation shared by multiple ports.

A Switch port may be able to operate in more than one mode, and configure itself to the appropriate mode during the initialization process (see 7.2). During initialization, the Switch port may assume a mode for purposes of determining if that mode is appropriate. For example, a Switch port operates in FL\_Port mode to determine if it is attached to a loop of NL\_Ports. If that is not successful, it then tries operating as an E\_Port to see if another E\_Port or B\_Port is attached. The Switch port continues until it finds a mode in which to operate.

A Bridge device may contain one or more Bridge ports (B\_Ports).

Ports that operate in the E\_Port or B\_Port mode are generically referred to as Interconnect\_Ports. A single Inter-Switch Link (ISL) connects two Interconnect\_Ports together. Valid combinations of Interconnect\_Ports are:

- a) E\_Port to E\_Port; and
- b) E\_Port to B\_Port.

B\_Port to B\_Port ISLs are not allowed.

### 5.2 Model elements

#### 5.2.1 FC Transports

The FC-FS-5 Transport includes all of the functionality described in FC-FS-5 to construct and deconstruct a frame, to encode and decode the words that make up the frame, and to transmit and receive the frame on the physical media. The FC-AL-2 Transport contains additional functionality to support the Arbitrated Loop protocols. See FC-BB-6 for additional Fibre Channel transports.

#### 5.2.2 Switch Transport

The Switch Transport is an abstraction to show the “back end” of the Switch port as it interacts with the Switch Construct and/or other Switch ports within the Switch. The Switch Transport exists to move frames between the Switch port and the rest of the Switch. No other implementation details are implied by this element.

#### 5.2.3 Control Facilities

The Control Facilities are internal logical ports that receive and perform requests, and generate responses. Each Control Facility has associated with it an address identifier, and support for classes of service. The Control Facilities also manage the various Transport elements.

#### 5.2.4 Link Services

The Link Services represent the various Link Services that are supported by the corresponding Control Facility.

### 5.3 F\_Port operation

An F\_Port is the point at which an external PN\_Port is attached to the Fabric. It normally functions as a conduit to the Fabric for frames transmitted by the PN\_Port, and as a conduit from the Fabric for frames destined for the PN\_Port.

An F\_Port shall support one or more of the following classes of service:

- a) Class 2 service;
- b) Class 3 service.

An F\_Port shall not intentionally transmit Class F frames on its outbound fibre. An PN\_Port that receives a Class F frame shall discard it, as required by FC-FS-5.

An F\_Port shall not admit to the Fabric any Class F frames, any Primitive Sequences, or any Primitive Signals other than Idle, that the F\_Port receives on its inbound fibre.

The model of an F\_Port on an FC-FS-5 Transport is shown in figure 9.

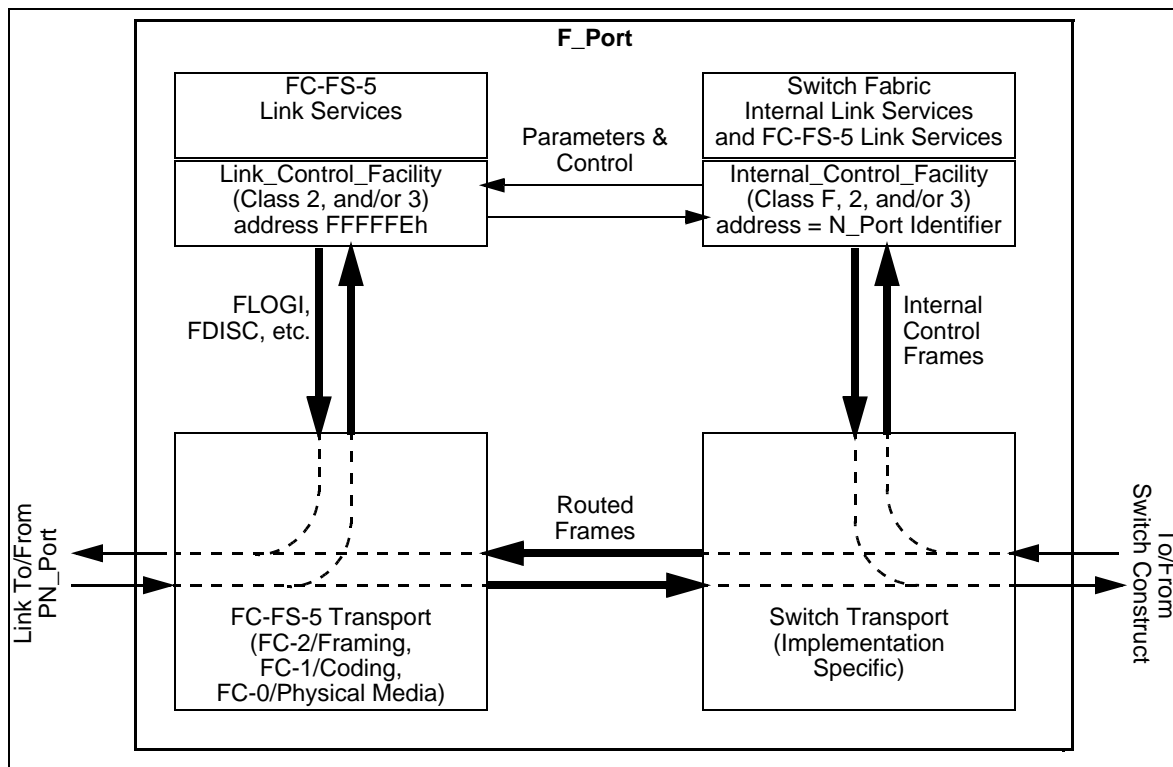


Figure 9 – F\_Port model

An F\_Port contains an FC-FS-5 Transport element through which passes all frames and Primitives transferred across the link to and from the PN\_Port. Frames received from the PN\_Port are either directed to the Switch Construct via the Switch Transport element, or directed to the Link\_Control\_Facility. The Link\_Control\_Facility receives frames related to Link Services such as FLOGI, and transmits responses to those Link Service frames.

Frames received from the FC-FS-5 Transport element that are destined for other ports are directed by the Switch Transport to the Switch Construct for further routing. Frames received from the Switch Construct by the Switch Transport are directed either to the FC-FS-5 Transport for transmission to the PN\_Port, or to the Internal\_Control\_Facility. The Internal\_Control\_Facility receives frames related to Switch Fabric Internal Link Services, and transmits responses to those Internal Link Services frames. Information is passed between the Internal\_Control\_Facility and the Link\_Control\_Facility to effect the control and configuration of the Transport elements.

The F\_Port is used by Switches to transmit and receive frames with a single PN\_Port. A link to an F\_Port always connects to exactly one PN\_Port.

An F\_Port link follows the FC-0, FC-1, and FC-2 protocols defined for point-to-point Links as defined in FC-FS-5.

See FC-BB-6 for additional Fibre Channel transports.

#### **5.4 FL\_Port operation**

An FL\_Port is the point at which one or more external NL\_Ports are attached to the Fabric. It normally functions as a conduit to the Fabric for frames transmitted by the attached NL\_Ports, and as a conduit from the Fabric for frames destined for the attached NL\_Ports.

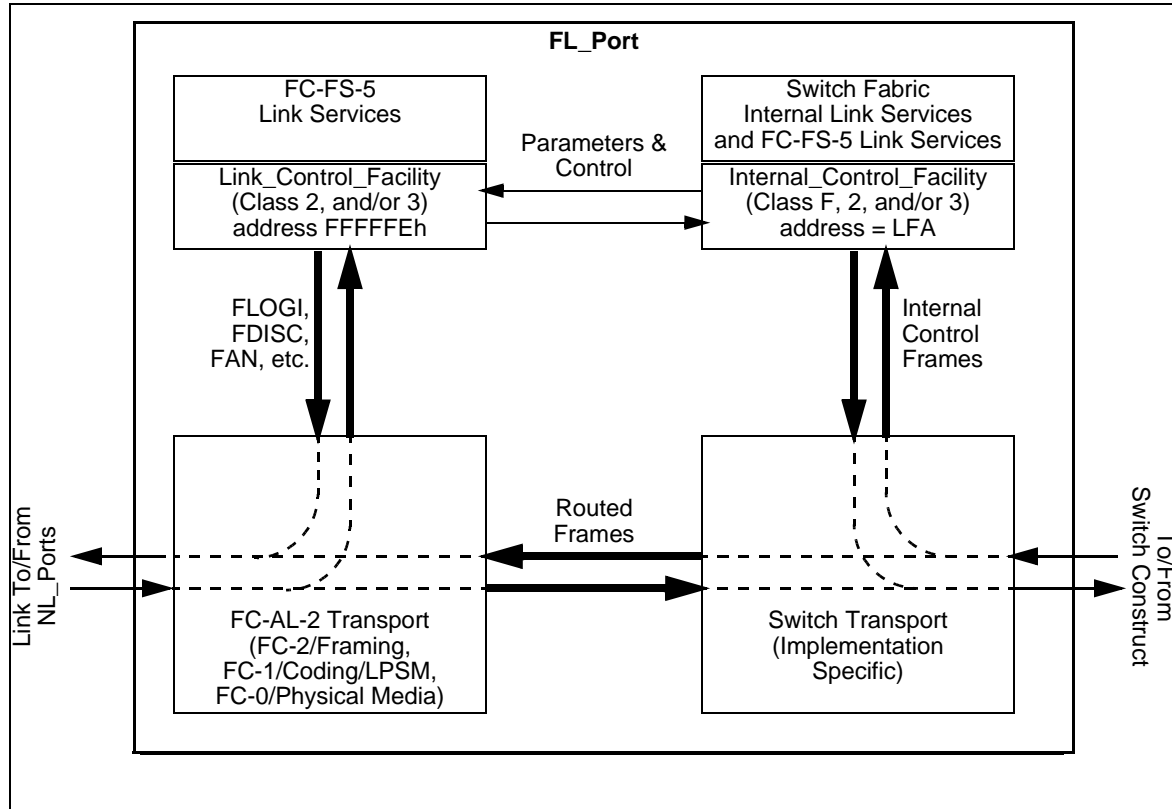
An FL\_Port shall support one or more of the following classes of service:

- a) Class 2 service; or
- b) Class 3 service.

An FL\_Port shall not intentionally transmit Class F frames on its outbound fibre. An FL\_Port shall not admit to the Fabric any Class F frames, any Primitive Sequences, or any Primitive Signals other than Idle, that the FL\_Port receives on its inbound fibre.

An FL\_Port that conforms to this standard should also conform to the FL\_Port requirements specified in FC-DA-2 and FC-MI-3.

The model of an FL\_Port is shown in figure 10.



**Figure 10 – FL\_Port model**

An FL\_Port contains an FC-AL-2 Transport element that passes all frames and Primitives transferred across the link to and from the multiple NL\_Ports. Frames received from the NL\_Ports are either directed to the Switch Construct via the Switch Transport element, or directed to the Link\_Control\_Facility. The Link\_Control\_Facility receives frames related to Link Services such as FLOGI, and transmits responses to those Link Service frames. The Link\_Control\_Facility also transmits and receives Loop Initialization Sequences and transmits the FAN ELS.

Frames received from the FC-AL-2 Transport element that are destined for other ports are directed by the Switch Transport to the Switch Construct for further routing. Frames received from the Switch Construct by the Switch Transport are directed either to the FC-AL-2 Transport for transmission to the destination NL\_Port, or to the Internal\_Control\_Facility. The Internal\_Control\_Facility receives frames related to Switch Fabric Internal Link Services and Loop management Extended Link Services (see FC-LS), and transmits responses to those Link Services frames. Information is passed between the Internal\_Control\_Facility and the Link\_Control\_Facility to effect the control and configuration of the Transport elements.

The FL\_Port is used by Switches to transmit and receive frames with one or more attached NL\_Ports. A link to an FL\_Port connects to one or more NL\_Ports.

An FL\_Port link follows the FC-0, FC-1, and FC-2 protocols defined in FC-FS-5, with the additional Arbitrated Loop protocols defined in FC-AL-2.



### 5.5 E\_Port operation

An E\_Port is the point at which a Switch is connected to another Switch to create a multi-Switch Fabric. Also, an E\_Port is the point at which a Switch is connected to a Bridge device. It normally functions as a conduit between the Switches for frames destined for remote N\_Ports and NL\_Ports. An E\_Port is also used to carry frames between Switches for purposes of configuring and maintaining the Fabric.

An E\_Port shall support the Class F service. An E\_Port shall also be capable of routing one or more of the following classes of service:

- a) Class 2 service; or
- b) Class 3 service.

An E\_Port shall not admit to the Fabric any Primitive Sequences, or any Primitive Signals other than Idle, that the E\_Port receives on its inbound fibre.

The model of an E\_Port on an FC-FS-5 Transport is shown in figure 11.

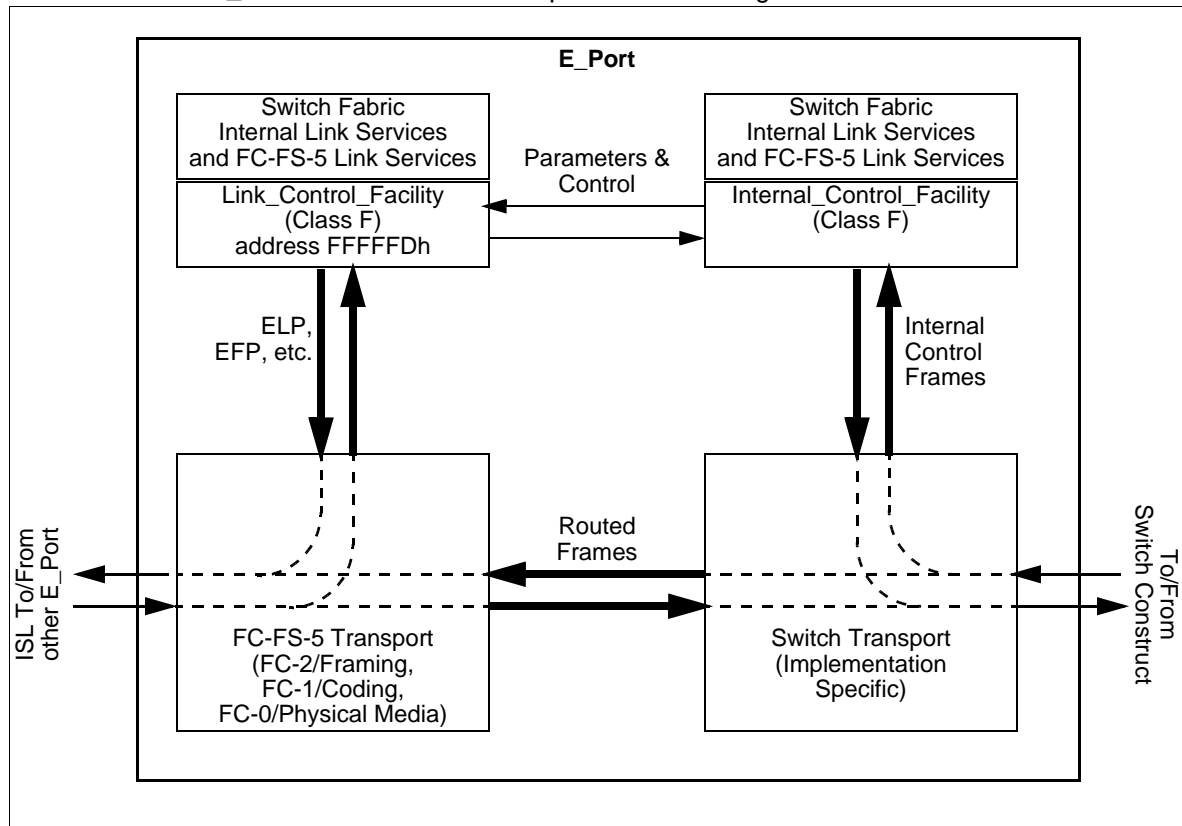


Figure 11 – E\_Port model

An E\_Port contains an FC-FS-5 Transport element through which all frames are passed, and Primitives are transferred across the link to and from the other E\_Port. Frames received from the other E\_Port are either directed to the Switch Construct via the Switch Transport element, or directed to the Link\_Control\_Facility. The Link\_Control\_Facility receives frames related to Switch Fabric Internal Link Services such as ELP, and transmits responses to those Link Service frames.

Frames received from the FC-FS-5 Transport element that are destined for other ports are directed by the Switch Transport to the Switch Construct for further routing. Frames received from the Switch Construct by the Switch Transport are directed either to the FC-FS-5 Transport for transmission to the other E\_Port, or to the Internal\_Control\_Facility. The Internal\_Control\_Facility receives frames related to Switch Fabric Internal Link Services, and transmits responses to those Internal Link Services frames. Information is passed between the Internal\_Control\_Facility and the Link\_Control\_Facility to effect the control and configuration of the Transport elements.

See FC-BB-6 for additional Fibre Channel transports.

### 5.6 B\_Port operation

A Bridge port (B\_Port) is a port that is used to connect a Switch to Bridge device. It normally functions as a conduit between the Switch and the Bridge for frames destined for or through a Bridge device. A B\_Port is also used to carry frames between a Switch and the Bridge device for purposes of configuring the Bridge device.

A B\_Port shall support Class F service. A B\_Port shall also be capable of forwarding one or more of the following classes of service:

- a) Class 2 service; or
- b) Class 3 service.

A B\_Port shall not admit to the Fabric any Primitive Sequences, or any Primitive Signals other than Idle, that the B\_Port receives on its inbound fibre.

The model of a B\_Port is shown in figure 12.

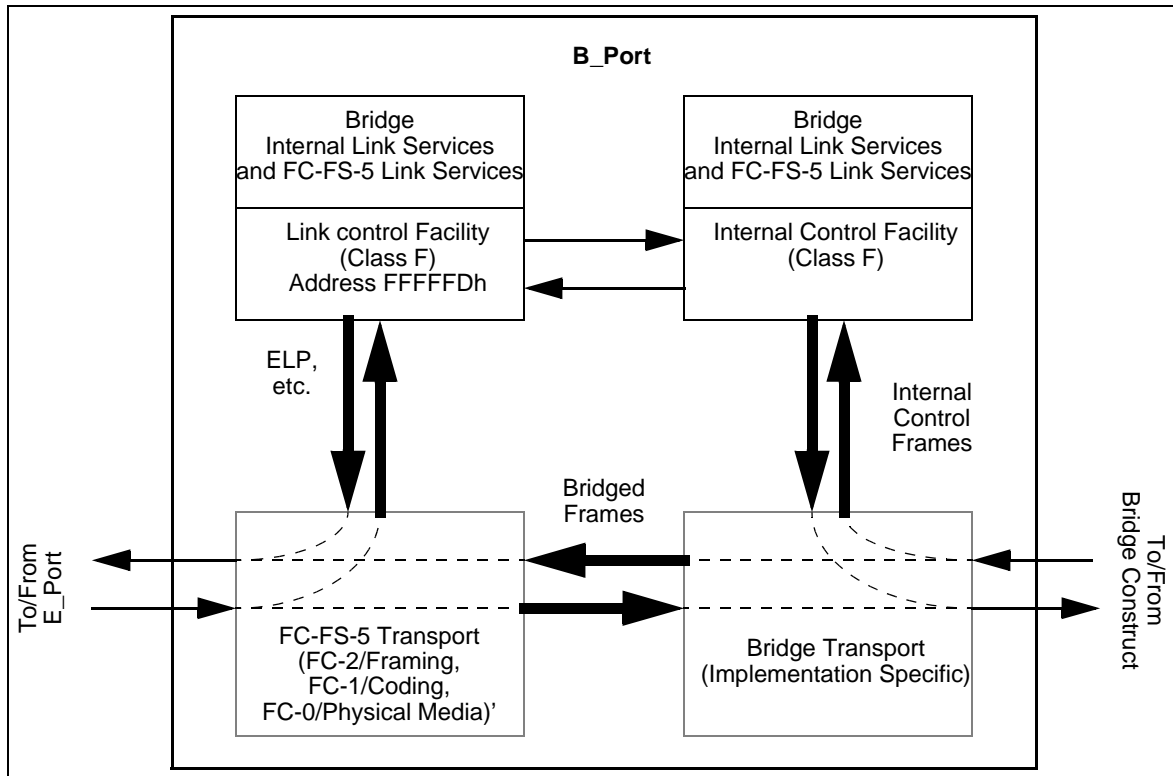


Figure 12 – B\_Port model

A B\_Port contains an FC-FS-5 Transport element through which pass all frames and Primitives transferred across the link to and from the E\_Port. Frames received from the attached E\_Port are either directed to the Bridge Construct via the Bridge Transport element, or directed to the Link\_Control\_Facility. The Link\_Control\_Facility receives frames related to Bridge Fabric Internal Link Services such as ELP, and transmits responses to those Link Service frames.

Frames received from the FC-FS-5 Transport element that are destined for other ports are directed by the Bridge Transport to the Bridge Construct for further forwarding. Frames received from the Bridge Construct by the Bridge Transport are directed either to the FC-FS-5 Transport for transmission to the other E\_Port, or to the Internal\_Control\_Facility. The Internal\_Control\_Facility receives frames related to Bridge Fabric Internal Link Services, and transmits responses to those Internal Link Services frames. Information is passed between the Internal\_Control\_Facility and the Link\_Control\_Facility to effect the control and configuration of the Transport elements.

The Bridge port utilizes Class F service as a connectionless service exchanging frames between E\_Ports and B\_Ports. The definition of Class F function and Class F rules apply to both E\_Ports and B\_Ports as defined in clause 5.9.

### 5.7 A\_Port operation

An A\_Port is the combination of one PA\_Port and one VA\_Port operating together. A PA\_Port is the LCF within the Fabric that attaches to another PA\_Port through a link. A VA\_Port is an instance of the FC-2V sublevel of Fibre Channel that connects to another VA\_Port. A VA\_Port is uniquely identified by an A\_Port\_Name Name\_Identifier and is addressable by the VA\_Port connected to it through the A\_Port Controller address identifier (i.e., FFFFF9h).

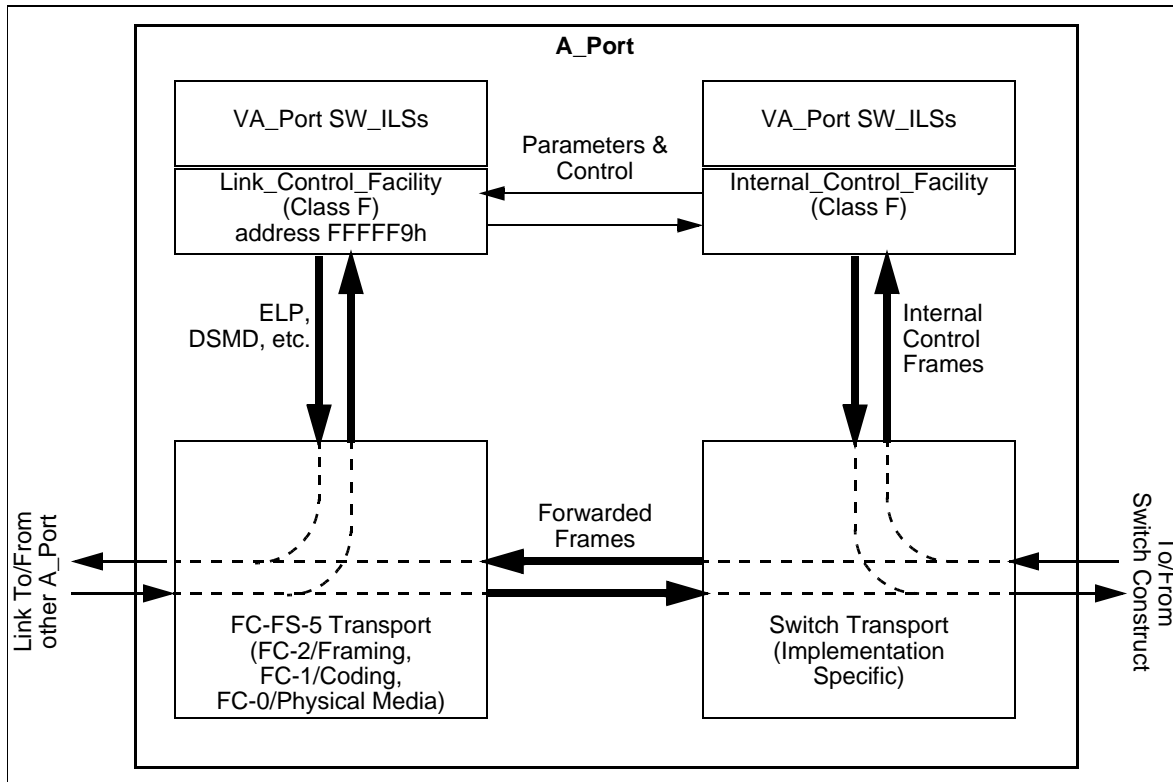
An A\_Port is the point at which a Controlling Switch is connected to an FCDF to create a Distributed Switch. Also, an A\_Port is the point at which an FCDF is connected to another FCDF. It normally functions as a conduit among FCDFs and between FCDFs and Controlling Switches for frames destined for remote N\_Ports and NL\_Ports. An A\_Port is also used to carry frames between Controlling Switch and FCDFs for purposes of configuring and maintaining the Distributed Switch.

An A\_Port shall support the Class F service. An A\_Port shall also be capable of forwarding one or more of the following classes of service:

- a) Class 2 service; or
- b) Class 3 service.

An A\_Port shall not admit to its FCDF or Controlling Switch any Primitive Sequences, or any Primitive Signals other than Idle, that the A\_Port receives on its inbound fibre.

The model of an A\_Port on an FC-FS-5 Transport is shown in figure 13..



**Figure 13 – A\_Port model**

An A\_Port contains an FC-FS-5 Transport element through which all frames are passed, and Primitives are transferred across the link to and from the other A\_Port. Frames received from the other A\_Port are either directed to the Switch Construct via the Switch Transport element, or directed to the Link\_Control\_Facility. The Link\_Control\_Facility receives frames related to the ELP SW\_ILS and the VA\_Port SW\_ILSs, and transmits responses to those frames.

Frames received from the FC-FS-5 Transport element that are destined for other ports are directed by the Switch Transport to the Switch Construct for further forwarding. Frames received from the Switch Construct by the Switch Transport are directed either to the FC-FS-5 Transport for transmission to the other A\_Port, or to the Internal\_Control\_Facility. The Internal\_Control\_Facility receives frames related to VA\_Port SW\_ILSs, and transmits responses to those frames.

Information is passed between the Internal\_Control\_Facility and the Link\_Control\_Facility to effect the control and configuration of the Transport elements.

### 5.8 Inter-Switch Link behavior

Inter-Switch Links (ISLs) are used by Switches to transmit and receive frames with other Switches or Bridge devices. An ISL always connects exactly one E\_Port on a Switch to exactly one E\_Port on another Switch or exactly one B\_Port on a Bridge device.

An ISL on an FC-FS-5 Transport follows the FC-0, FC-1, and FC-2 protocols defined for point-to-point Links as defined in FC-FS-5, with the exception that Class F frames are allowed to transit the link. R\_RDY shall be used for the management of buffer-to-buffer flow control of Class F frames on the ISL prior to the completion of the exchange of link parameters (see 6.2.4 and 7.2). An alternate

method of buffer-to-buffer flow control may be defined in that process. See FC-BB-6 for additional Fibre Channel transports.

For purposes of defining and maintaining the Fabric Configuration, an ISL may be designated as a Principal ISL. The Principal ISL is a path that is used during configuration and address assignment to route Class F configuration frames, and is therefore a known path between two Switches. If a Principal ISL is lost, there may be no other available paths between the two affected Switches, so as a result the Fabric Configuration is possibly broken and shall be rebuilt (i.e., by issuing the BF SW\_ILS, see 6.2.11). If a non-Principal ISL is lost, at least one other path is known to be available between the Switches (i.e., the Principal ISL), therefore the lost ISL may be resolved via a routing change.

A Switch discovers the Principal ISL(s) during the process of Principal Switch Selection (see 7.3) and Address Distribution (see 7.4). During this process, the Switch identifies two kinds of Principal ISLs. The Principal ISL that leads towards the Principal Switch is called the upstream Principal ISL. All frames from the Switch to the Principal Switch are sent via the upstream Principal ISL. The Principal Switch has no upstream Principal ISL. All other Switches have exactly one upstream Principal ISL.

A Principal ISL that leads away from the Principal Switch is called the downstream Principal ISL. Any frame sent by the Switch to another Switch as a result of a frame received on the upstream Principal ISL is sent via the downstream Principal ISL that leads towards that Switch. The Principal Switch may have one or more downstream Principal ISLs. All other Switches may have zero or more downstream Principal ISLs.

Principal ISLs are further illustrated in figure 14.

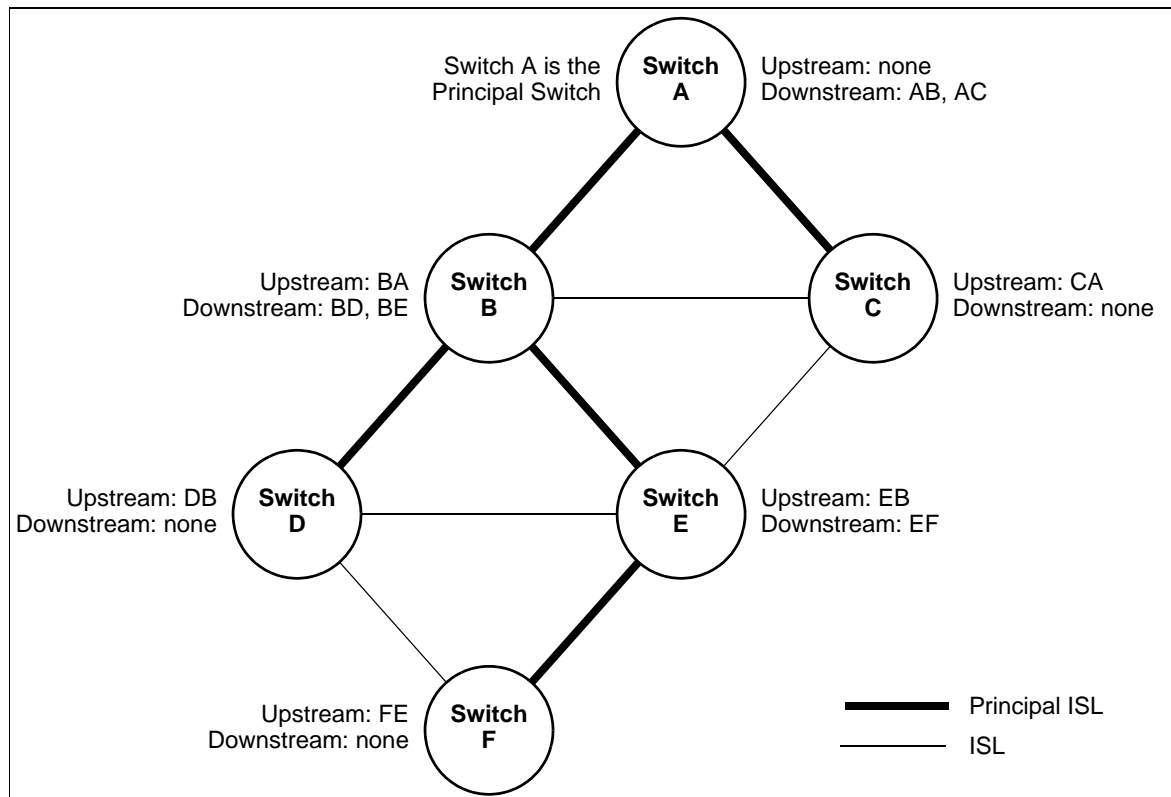


Figure 14 – Principal Inter-Switch Links

## 5.9 Class F service

### 5.9.1 Class F function

Class F service is a connectionless service with notification of non-delivery between Interconnect\_Ports. Class F service is used for control, coordination, and configuration of the Fabric. Class F service is defined by this standard for use by Switches communicating across Inter-Switch Links.

A Class F service is requested by an Interconnect\_Port on a frame by frame basis. The Fabric routes the frame to the destination Interconnect\_Port. If an Interconnect\_Port transmits consecutive frames to multiple destinations, the Fabric demultiplexes them to the requested destinations. Class F delimiters are used to indicate the requested service and to initiate and terminate one or more Sequences as described in FC-FS-5. Class F shall follow the rules for Class 2 except where otherwise stated in this standard.

### 5.9.2 Class F rules

To provide Class F service, the transmitting and receiving Interconnect\_Ports and the Fabric shall obey the following rules:

- a) except for some Switch Fabric Internal Link Service protocols, an Interconnect\_Port is required to have exchanged link parameters (see 6.2.4 and 7.2) with the associated destination with which it intends to communicate;
- b) the Fabric routes the frames without establishing a Dedicated Connection between communicating Interconnect\_Ports. To obtain Class F service, the Interconnect\_Port shall use Class F delimiters as defined in 5.9.3;
- c) an Interconnect\_Port is allowed to send consecutive frames to one or more destinations. This enables an Interconnect\_Port to demultiplex multiple Sequences to a single or multiple destinations concurrently;
- d) a given Interconnect\_Port may receive consecutive frames from different sources. Each source is allowed to send consecutive frames for one or more Sequences;
- e) an Interconnect\_Port addressed by a Class F frame shall provide an acknowledgment to the source for each valid Data frame received. The destination Interconnect\_Port shall use ACK\_1 for the acknowledgment. If a Switch is unable to deliver the ACK\_1 frame, the Switch shall return an F\_BSY or F\_RJT;
- f) the Sequence Initiator shall increment the SEQ\_CNT field of each successive frame transmitted within a Sequence. However, the Switches may not guarantee delivery to the destination in the same order of transmission;
- g) since the SOFf delimiter does not indicate whether a frame is the first frame of a Sequence, the starting SEQ\_CNT of every Sequence shall be zero;
- h) an Interconnect\_Port may originate multiple Exchanges and initiate multiple Sequences with one or more Interconnect\_Ports. The Interconnect\_Port originating an Exchange shall assign an X\_ID unique to the Originator called OX\_ID and the Responder of the Exchange shall assign an X\_ID unique to the responder called RX\_ID. The value of OX\_ID or RX\_ID is unique to a given Interconnect\_Port. The Sequence Initiator shall assign a SEQ\_ID, for each Sequence it initiates, that is unique to the Sequence Initiator and the respective Sequence Recipient pair while the Sequence is Open;
- i) each Interconnect\_Port exercises buffer-to-buffer flow control with the Interconnect\_Port to which it is directly attached. End-to-end flow control is performed by communicating Interconnect\_Ports. ACK\_1 frames are used to perform end-to-end flow control and R\_RDY is used for buffer-to-buffer flow control. However, some other agreed upon methods outside the scope of this standard may be used for buffer-to-buffer flow control (e.g., see FC-BB-6);

- j) if a Switch is unable to deliver the frame to the destination Interconnect\_Port, then the source is notified of each frame not delivered by an F\_BSY or F\_RJT frame with corresponding D\_ID, S\_ID, OX\_ID, RX\_ID, SEQ\_ID, and SEQ\_CNT from the Switch. The source is also notified of valid frames busied or rejected by the destination Interconnect\_Port by P\_BSY or P\_RJT;
- k) a busy or reject may be issued by an intermediate Interconnect\_Port or the destination Interconnect\_Port with a valid reason code;
- l) if a Class F Data frame is busied, the sender shall retransmit the busied frame up to the ability of the sender to retry, including zero;
- m) the Credit established during the ELP protocol by interchanging link parameters shall be honored. Class F may share Credit with other classes of service;
- n) effective transfer rate between any given Interconnect\_Port pair is dependent upon the number of Interconnect\_Ports a given Interconnect\_Port is multiplexing and demultiplexing;
- o) frames within a Sequence are tracked on a Sequence\_Qualifier and SEQ\_CNT basis;
- p) an Interconnect\_Port shall be able to recognize SOF delimiters for Class N service, whether or not all classes of service are supported by the port;
- q) an Interconnect\_Port addressed by a Vendor Specific Class F frame, shall send an LS\_RJT if it does not understand the frame. A Vendor Specific Class F frame is indicated by an R\_CTL field value of F0h;
- r) an Interconnect\_Port shall use R\_RDY and FC-FS-5 buffer-to-buffer flow control with the Interconnect\_Port to which it is directly attached, until after the exchange of link parameters (see 6.2.4 and 7.2). The BB\_Credit prior to the exchange of link parameters shall be 1. An Interconnect\_Port may agree to use an alternate buffer-to-buffer credit model for Class F following the successful exchange of link parameters. See FC-BB-6 for additional Fibre Channel transports; and
- s) a Class F frame shall be forwarded to its destination without checking by an intermediate entity. A Class F frame not destined for the receiving E\_Port (or E\_Port's domain) shall always be forwarded regardless of whether or not the receiving E\_Port recognizes the frame.

### 5.9.3 Class F frame format

Class F frames shall use the Frame Content format defined in FC-FS-5. The Start\_of\_Frame Fabric (SOFF) delimiter shall precede the frame content of all Class F frames. The Data Field size of all Class F frames shall be less than or equal to 256 bytes prior to the successful completion of Exchange Link Parameters (see 6.2.4; Exchange Link Parameters establishes the maximum receive frame size for Class F frames). All Class F frames shall include the CRC defined in FC-FS-5. The End\_of\_Frame Normal (EOFn) delimiter shall immediately follow the CRC of all normally completed Class F Data frames and all normally completed Class F Link\_Control frames except the last frame of a Sequence. The End\_of\_Frame Terminate (EOFT) delimiter shall immediately follow the CRC of all Class F Link\_Control frames that indicate the last frame of a normally terminated Sequence.

An Interconnect\_Port or Switch may invalidate or discard without notification any incorrectly formed Class F frame, or any Class F frame with a code violation or CRC error.

### 5.9.4 Class F flow control

Class F service uses end-to-end flow control in all Transports. ACK\_1 frames are used to perform end-to-end flow control. ACK\_1 frames shall be formatted as described in 5.9.3. The ACK\_0 Link Control frame shall not be used for Class F service.

In the FC-FS-5 Transport, R\_RDY is used for buffer-to-buffer flow control. R\_RDY is transmitted by the Interconnect\_Port at one end of the ISL, to the Interconnect\_Port at the other end of the ISL, to indicate that a buffer is available for further frame reception by the first Interconnect\_Port. This process operates in both directions on the ISL. After the successful exchange of link parameters, an alternate method of buffer-to-buffer flow control may be established on an ISL (see 7.2). This

alternate method of buffer-to-buffer flow control remains in effect until a Link Offline or Link Failure occurs, or a new set of link parameters is successfully exchanged between the Interconnect\_Ports.

See FC-BB-6 for additional Fibre Channel transports.



## 6 Internal Link Services

### 6.1 Switch Fabric Internal Link Services (SW\_ILS)

SW\_ILSs operate internal to the Fabric between Switches and internal to the Distributed Switch between FCDFs and Controlling Switches. All SW\_ILS frames shall be transmitted using the Class F service. The following defines the header fields of all SW\_ILS frames:

**R\_CTL:** This field shall be set to 02h for all request frames, and to 03h for all reply frames.

**CS\_CTL:** This field shall be set to 00h.

**D\_ID and S\_ID:** Set as indicated for the specific SW\_ILS.

**TYPE:** This field shall be set to 22h, indicating Fibre Channel Fabric Switch Services.

All other fields shall be set as appropriate according to the rules defined in FC-FS-5. The first word in the payload specifies the command code. The command codes are summarized in table 2.

**Table 2 – SW\_ILS command codes (Part 1 of 3)**

Encoded Value (hex)	Description	Abbr.
01000000	Switch Fabric Internal Link Service Reject	SW_RJT
02xxxxxx	Switch Fabric Internal Link Service Accept	SW_ACC
10000000	Exchange Link Parameters	ELP
11xxxxxx	Exchange Fabric Parameters	EFP
12000000	Domain Identifier Assigned	DIA
1300xxxx	Request Domain_ID	RDI
14000000	Hello	HLO
15000000	Link State Update	LSU
16000000	Link State Acknowledgement	LSA
17000000	Build Fabric	BF
18000000	Reconfigure Fabric	RCF
1B000000	Inter-Switch Registered State Change Notification	SW_RSCN
1E000000	Distribute Registered Link Incident Records	DRLIR
20000000	Obsolete	DSCN
<sup>a</sup> ASF and EBP are currently defined in FC-BB-6.		

**Table 2 – SW\_ILS command codes (Part 2 of 3)**

Encoded Value (hex)	Description	Abbr.
21000000	Obsolete	LOOPD
22xxxxxx	Merge Request	MR
2300xxxx	Acquire Change Authorization	ACA
24000000	Release Change Authorization	RCA
25xxxxxx	Stage Fabric Configuration	SFC
26000000	Update Fabric Configuration	UFC
28xxxxxx	Reserved for FC-BB-6 use <sup>a</sup>	
29000000	Check E_Port Connectivity	CEC
2A010000	Enhanced Acquire Change Authorization	EACA
2A020000	Enhanced Stage Fabric Configuration	ESFC
2A030000	Enhanced Update Fabric Configuration	EUFC
2A040000	Enhanced Release Change Authorization	EACA
2A050000	Transfer Commit Ownership	TCO
3000xxxx	Exchange Switch Capabilities	ESC
31000000	Exchange Switch Support	ESS
32000000	Reserved for Legacy Implementations (see FC-SP-2)	
33000000	Reserved for Legacy Implementations (see FC-SP-2)	
34000000	Merge Request Resource Allocation	MRRA
35010000	Switch Trace Route	STR
36000000	Exchange Virtual Fabrics Parameters	EVFP
40xxxxxx	Reserved for FC-SP-2 use	
41xxxxxx	Reserved for FC-SP-2 use	
42xxxxxx	Reserved for FC-SP-2 use	
50000000	Fast Fabric Initialization for the Avionics Environment (See annex D)	FFI

<sup>a</sup> ASF and EBP are currently defined in FC-BB-6.

**Table 2 – SW\_ILS command codes (Part 3 of 3)**

Encoded Value (hex)	Description	Abbr.
70000000 to 7FFFFFFF	Vendor Specific	
90000000 to 9FFFFFFF	Vendor Specific	
A0xxxxxx	Distributed Switch VA_Port (see 6.3)	
A1xxxxxx	Controlling Switch redundancy protocol (see 6.4)	
others	Reserved	
<sup>a</sup> ASF and EBP are currently defined in FC-BB-6.		

Unless otherwise specified, the rules regarding the following aspects of Switch Fabric Internal Link Services are as defined for the Extended Link Services in FC-FS-5 (e.g., Sequence and Exchange Management, error detection and recovery).

## 6.2 Fabric SW\_ILSs

### 6.2.1 6.2.1 Overview

This subclause describes Link Services that operate internal to the Fabric between Switches. In the case of Exchange Link Parameters (ELP), Link Services also operate internal to the Fabric between Switches and Bridge devices. Timeout values for specific SW\_ILS's and the actions following a timeout expiration are specified in 16.2.

### 6.2.2 Switch Fabric Internal Link Service Accept (SW\_ACC)

The Switch Fabric Internal Link Service Accept reply Sequence shall notify the transmitter of an SW\_ILS request that the SW\_ILS request Sequence has been accepted and the requested operation has been either completed or initiated, depending on the request. The first word of the payload shall contain 02 xx xx xxh. The remainder of the payload is unique to the specific SW\_ILS request.

**Protocol:** SW\_ACC may be sent as a reply Sequence to an SW\_ILS request. An SW\_ACC shall not be sent for HLO, LSU, and LSA request Sequences.

**Addressing:** The S\_ID field shall be set to the value of the D\_ID field in the SW\_ILS request. The D\_ID field shall be set to the value of the S\_ID field in the SW\_ILS request.

**Payload:** The payload content following the first word is defined within individual SW\_ILS requests.

### 6.2.3 Switch Fabric Internal Link Service Reject (SW\_RJT)

The Switch Fabric Internal Link Service Reject shall notify the transmitter of an SW\_ILS request that the SW\_ILS request Sequence has been rejected. A four-byte reason code shall be contained in the

Data\_Field. SW\_RJT may be transmitted for a variety of conditions that may be unique to a specific SW\_ILS request.

**Protocol:** SW\_RJT may be sent as a reply Sequence to an SW\_ILS request. An SW\_RJT shall not be sent for HLO, LSU, and LSA request Sequences.

**Addressing:** The S\_ID field shall be set to the value of the D\_ID field in the SW\_ILS request. The D\_ID field shall be set to the value of the S\_ID field in the SW\_ILS request.

**Payload:** The format of the SW\_RJT reply payload is specified in table 3.

**Table 3 – SW\_RJT payload**

Item	Size (bytes)
01 00 00 00h	4
Reserved	1
Reason Code	1
Reason Code Explanation	1
Vendor Specific	1

**Reason Code:** The reason codes are summarized in table 4.

**Table 4 – SW\_RJT reason codes**

Encoded value (hex)	Description
01	Invalid SW_ILS command code
02	Invalid revision level
03	Logical error
04	Invalid payload size
05	Logical busy
07	Protocol error
09	Unable to perform command request
0B	Command not supported
0C	Invalid Attachment
FF	Vendor Specific error
others	Reserved

**Invalid SW\_ILS command code:** The command code is not recognized by the recipient.

**Invalid revision level:** The recipient does not support the specified revision level.

**Logical error:** The request identified by the command code and the payload content is invalid or logically inconsistent for the conditions present.

**Invalid payload size:** The size of the payload is inconsistent with the command code and/or any Length fields in the payload.

**Logical busy:** The recipient is busy and is unable to process the request at this time.

**Protocol error:** An error has been detected that violates the protocol.

**Unable to perform command request:** The recipient is unable to perform the request.

**Command not supported:** The command code is not supported by the recipient.

**Invalid Attachment:** The recipient is in the Invalid Attachment state.

**Vendor Specific Error:** The Vendor Specific field indicates the error condition.

**Reason Code Explanation:** The reason code explanations are specified in table 5.

**Table 5 – SW\_RJT reason code explanations (Part 1 of 3)**

Encoded value (hex)	Description
00	No additional explanation
01	Class F Service Parameter error
03	Class N Service Parameter error
04	Unknown Flow Control code
05	Invalid Flow Control Parameters
0D	Invalid Port_Name
0E	Invalid Switch_Name
0F	R_A_TOV or E_D_TOV mismatch
10	Invalid Domain_ID_List
19	Command already in progress
29	Insufficient resources available
<p><sup>a</sup> The range of values 30h-3Fh are used to indicate Security reason code explanations.</p> <p><sup>b</sup> The range of values 40h-4Fh are used to indicate Zoning reason code explanations.</p>	

**Table 5 – SW\_RJT reason code explanations (Part 2 of 3)**

Encoded value (hex)	Description
2A	Domain_ID not available
2B	Invalid Domain_ID
2C	Request not supported
2D	Link Parameters not yet established
2E	Requested Domain_IDs not available
2F	E_Port is Isolated
31	Authorization Failed <sup>a</sup>
32	Authentication Failed
33	Incompatible Security Attribute
34	Security Checks in Progress
35	Policy Summary Not Equal
36	FC-SP Zoning Summary Not Equal
41	Invalid Data Length <sup>b</sup>
42	Unsupported Command
44	Not Authorized
45	Invalid Request
46	Fabric Changing
47	Update Not Staged
48	Invalid Zone Set Format
49	Invalid Data
4A	Unable to Merge
4B	Zone Set Size Not Supported
50	Unable to Verify Connection
<p><sup>a</sup> The range of values 30h-3Fh are used to indicate Security reason code explanations.</p> <p><sup>b</sup> The range of values 40h-4Fh are used to indicate Zoning reason code explanations.</p>	

**Table 5 – SW\_RJT reason code explanations (Part 3 of 3)**

Encoded value (hex)	Description
58	Requested Application Not Supported
59	Transaction Specified by Request Does Not Exist
5A	Invalid Phase Transition in Transaction
5B	In Advanced Phase
5C	Switch Not Authorized
5D	Out of Order
others	Reserved
<p><sup>a</sup> The range of values 30h-3Fh are used to indicate Security reason code explanations.</p> <p><sup>b</sup> The range of values 40h-4Fh are used to indicate Zoning reason code explanations.</p>	

**Vendor Specific:** This field is valid when the reason code indicates a Vendor Specific error, otherwise this field is reserved.

## 6.2.4 Exchange Link Parameters (ELP)

### 6.2.4.1 ELP request

The Exchange Link Parameters Switch Fabric Internal Link Service requests the exchange of link parameters between two Interconnect\_Ports connected via an ISL. The exchange of link parameters establishes the operating environment between the two Interconnect\_Ports, and the capabilities of the Switches or Bridge devices that are connected by the Interconnect\_Ports. When an ELP is received by an Interconnect\_Port, any Active or Open Class F Sequences between the two Interconnect\_Ports, and any Dedicated Connections, shall be abnormally terminated prior to transmission of the SW\_ACC reply Sequence.

Use of the ELP SW\_ILS for Switch port initialization is described in 7.2.

#### Protocol:

Exchange Link Parameters (ELP) request Sequence  
Accept (SW\_ACC) reply Sequence

**Error Detection and Recovery:** See table 293.

**Addressing:** For use in Switch port initialization, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch, and the D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the ELP request payload is specified in table 6.

**Table 6 – ELP request payload**

Item	Size (bytes)
10000000h	4
Revision	1
Flags	2
BB_SC_N	1
R_A_TOV	4
E_D_TOV	4
Requester Interconnect_Port_Name	8
Requester Switch_Name	8
Fabric Controller Class F Service Parameters	16
Obsolete	4
Class 2 Interconnect_Port Parameters	4
Class 3 Interconnect_Port Parameters	4
Reserved	20
ISL Flow Control Mode	2
Flow Control Parameter Length (N)	2
Flow Control Parameters	N

**Revision:** This field denotes the revision of the ELP payload format and the Switch port initialization protocol.

The Revision field value is increased only if there is a substantial change in the ELP payload format, or if there are procedural changes in the Switch port initialization protocol that are not compatible with the previous version. Examples of substantial changes in the ELP payload format are:

- a) changes in the lengths of existing ELP payload fields; or
- b) the addition of new fields in the ELP payload. The use of bits that were previously reserved does not constitute a substantial change to the ELP payload.

For this standard, the Revision field value shall be 03h.

If a Fabric Controller receives an ELP request containing a Revision field value that is higher than its supported value, the Fabric Controller shall respond with its highest supported Revision field value. If



a Fabric Controller receives an ELP request containing a Revision field value that is equal to or lower than its supported value, the Fabric Controller shall respond with the Revision field value received in the ELP request.

**Flags:** This field contains flag bits that provide additional information about the ELP. The following flag bits are defined.

Bit 15, the Bridge port bit, shall indicate whether the sending port is a B\_Port. If bit 15 is zero, the sending port is an E\_Port and not a B\_Port. If bit 15 is one, the sending port is a B\_Port.

Bit 14, the Bridge Virtual Fabrics bit, is meaningful only for a B\_Port and shall indicate whether the sending B\_Port supports Virtual Fabric Tagging. If bit 14 is zero, the sending B\_Port does not support the passing through of VFT tagged frames (see FC-FS-5). If bit 14 is one, the sending B\_port supports the passing through of VFT tagged frames.

Bit 13, is the Controlling FCF/Switch bit. This bit set to one indicates that the originator of the ELP request or SW\_ACC is a Controlling Switch. This bit set to zero indicates that the originator of the ELP request or SW\_ACC is not a Controlling Switch.

Bit 12, is the FDF/FCDF bit. This bit set to one indicates that the originator of the ELP request or SW\_ACC is an FCDF. This bit set to zero indicates that the originator of the ELP request or SW\_ACC is not an FCDF.

Bit 11-10, the Energy Efficient Operation Support bits. If set to 00b, the port does not support Energy Efficient operation. If set to 01b, the port supports Energy Efficient Fast Wake mode only (see FC-FS-5). If set to 10b, the port supports both Energy Efficient modes of Fast Wake, and Quiet (see FC-FS-5). The value 11b is reserved.

Bits 9-0 shall be reserved.

**BB\_SC\_N:** This field indicates the Buffer-to-Buffer State Change number. The BB\_SC\_N field is valid only if the R\_RDY flow control mode is specified in the ISL Flow Control Mode field. A value between 0 and 15 indicates that the sender of the ELP frame is requesting a  $2^{BB\_SC\_N}$  number of frames be sent between two consecutive BB\_SCs Primitive Signals, and a  $2^{BB\_SC\_N}$  number of R\_RDY Primitive Signals be sent between two consecutive BB\_SCr Primitive Signals. When the two ports exchanging link parameters specify different non-zero values of BB\_SC\_N, the larger value shall be used. If either port specifies a BB\_SC\_N value of zero, then the BB\_Credit recovery process shall not be performed and no BB\_SCx Primitive Signals shall be sent. If a port specifies a non-zero BB\_SC\_N value it shall support the BB\_SCs and BB\_SCr Primitive Signals. See FC-FS-5 for a description of the BB\_Credit recovery process. The following BB\_SC\_N bits are defined:

Bits 7-4, Reserved

Bits 3-0, Buffer-to-buffer State Change Number (BB\_SN\_N).

If all frames or R\_RDY Primitive Signals sent between two BB\_SCx Primitive Signals are lost, then  $2^{BB\_SC\_N}$  number of BB\_Credits are lost, and are unable to be recovered by the scheme outlined in FC-FS-5. Therefore BB\_SC\_N should be chosen so that the probability of losing  $2^{BB\_SC\_N}$  number of consecutive frames or R\_RDY Primitive Signals is deemed negligible. Therefore the recommended value of BB\_SC\_N is 8.

**R\_A\_TOV:** This field shall be set to the value of R\_A\_TOV required by the Switch.

**E\_D\_TOV:** This field shall be set to the value of E\_D\_TOV required by the Switch.

The values of R\_A\_TOV and E\_D\_TOV may be established by a Profile(s) or other means.

**Interconnect\_Port\_Name:** The Interconnect\_Port\_Name is an eight-byte field that identifies an Interconnect\_Port. The format of the name is specified in FC-FS-5. Each Interconnect\_Port shall provide a unique Interconnect\_Port\_Name within the Fabric.

**Switch\_Name:** The Switch\_Name is an eight-byte field that identifies a Switch or Bridge device. The format of the name is specified in FC-FS-5. Each Switch\_Name shall be unique within the Fabric.

**Fabric Controller Class F Service Parameters:** This field contains the E\_Port Class F Service Parameters. The format of the Parameters is specified in table 7.

**Table 7 – Fabric Controller Class F Service Parameters**

Word	3 1	3 0	2 9	2 8	2 7	2 6	2 5	2 4	2 3	2 2	2 1	1 0	1 9	1 8	1 7	1 6	1 5	1 4	1 3	1 2	1 1	1 0	9	8	7	6	5	4	3	2	1	0
0	VAL		Reserved														Reserved															
1	R	X	Reserved														Receive Data Field Size															
2	Concurrent Sequences														End-to-End Credit																	
3	Open Sequences per Exchange														Reserved																	

**VAL (Class Valid):** This bit shall be set to one.

**XII (X\_ID Interlock):** This bit when one indicates that the Fabric Controller supplying this parameter requires that an interlock be used during X\_ID assignment in Class F. In X\_ID assignment, the Sequence Initiator shall set the Recipient X\_ID value to FFFFh in the first Data frame of a Sequence, and the Recipient shall supply its X\_ID in the ACK frame corresponding to the first Data frame of a Sequence. The Sequence Initiator shall not transmit additional frames until the corresponding ACK is received. Following reception of the ACK, the Sequence Initiator continues transmission of the Sequence using both assigned X\_ID values.

**Receive Data Field Size:** This field shall specify the largest Data Field size in bytes for a frame that may be received by the Fabric Controller supplying the Parameters as a Sequence Recipient for a Class F frame. Values less than 256 or greater than 2112 are invalid. Values shall be a multiple of four bytes.

**Concurrent Sequences:** This field shall specify the number of Sequence Status Blocks provided by the Fabric Controller supplying the Parameters for tracking the progress of a Sequence as a Sequence Recipient. The maximum number of Concurrent Sequences that may be specified is 255. A value of zero in this field is reserved. In Class F, the value of SEQ\_ID shall range from 0 to 255, independent of the value in this field. A Fabric Controller is allowed to respond with P\_BSY to a frame initiating a new Sequence if Interconnect\_Port resources are not available.

**End-to-End Credit:** End-to-end credit is the maximum number of Class F Data frames that may be transmitted by a Fabric Controller without receipt of accompanying ACK or Link\_Response frames. The minimum value of end-to-end credit is one. The end-to-end credit field specified is associated

with the number of buffers available for holding the Data\_Field of a Class F frame and processing the contents of that Data\_Field by the Interconnect\_Port supplying the Parameters. Bit 15 of this field shall be set to zero. A value of zero for this field is reserved.

**Open Sequences per Exchange:** The value of the Open Sequences per Exchange shall specify the maximum number of Sequences that may be Open at one time at the Recipient between a pair of Fabric Controllers for one Exchange. This value plus two shall specify the number of instances of Sequence Status that shall be maintained by the Recipient for a single Exchange in the Exchange Status Block. This value is used for Exchange and Sequence tracking. The value in this field limits the link facility resources required for error detection and recovery.

Interconnect\_Port Parameters indicate that the Interconnect\_Port is capable of transporting the indicated Class of Service, and the conditions under which it may transport the Class. One word of the ELP payload is allocated for each Class.

**Class 2 Interconnect\_Port Parameters:** This field contains the Class 2 Interconnect\_Port Parameters. The format of the Parameters is specified in table 8.

**Table 8 – Class 2 Interconnect\_Port Parameters**

Word	3 1	3 0	2 9	2 8	2 7	2 6	2 5	2 4	2 3	2 2	2 1	1 0	1 9	1 8	1 7	1 6	1 5	1 4	1 3	1 2	1 1	1 0	9	8	7	6	5	4	3	2	1	0
0	V A L	R	R	R	R	S	Reserved										Receive Data Field Size															

**VAL (Class Valid):** This bit shall be set to one if the Interconnect\_Port supports Class 2. If this bit is zero, all other Class 2 Interconnect\_Port Parameters shall be invalid.

**SEQ (Sequential Delivery):** If this bit is set to one by an Interconnect\_Port, it is indicating that the Switch is able to guarantee sequential delivery (as defined in FC-FS-5) of Class 2 frames. Sequential Delivery shall be functional only if both Interconnect\_Ports indicate support for this feature.

**Receive Data Field Size:** This field shall specify the largest Data Field size in bytes for a frame that may be received by the Interconnect\_Port supplying the Parameters for a Class 2 frame. Values less than 256 or greater than 2112 are invalid. Values shall be a multiple of four bytes.

**Class 3 Interconnect\_Port Parameters:** This field contains the Class 3 Interconnect\_Port Parameters. The format of the Parameters is specified in table 9.

**Table 9 – Class 3 Interconnect\_Port Parameters**

Word	3 1	3 0	2 9	2 8	2 7	2 6	2 5	2 4	2 3	2 2	2 1	1 0	1 9	1 8	1 7	1 6	1 5	1 4	1 3	1 2	1 1	1 0	9	8	7	6	5	4	3	2	1	0
0	V A L	R	R	R	R	S	Reserved										Receive Data Field Size															

**VAL (Class Valid):** This bit shall be set to one if the Interconnect\_Port supports Class 3. If this bit is zero, all other Class 3 Interconnect\_Port Parameters shall be invalid.

**SEQ (Sequential Delivery):** If this bit is set to one by an Interconnect\_Port, it is indicating that the Switch is able to guarantee sequential delivery (as defined in FC-FS-5) of Class 3 frames. Sequential Delivery shall be functional only if both Interconnect\_Ports indicate support for this feature.

**Receive Data Field Size:** This field shall specify the largest Data Field size in bytes for a frame that may be received by the Interconnect\_Port supplying the Parameters for a Class 3 frame. Values less than 256 or greater than 2112 are invalid. Values shall be a multiple of four bytes.

**ISL Flow Control Mode:** This field contains a code that specifies the flow control method supported by the Interconnect\_Port. Table 10 shows the allowed values for this field.

**Table 10 – ISL Flow Control Mode values**

Value (hex)	Usage
0001	Vendor Specific
0002	R_RDY flow control
0003 - 1FFF	Vendor Specific
2000	VC_RDY flow control
AE02	Reserved for AE Use
Other Values	Reserved

**Flow Control Parameter Length:** This field specifies the length in bytes of the Flow Control Parameters that follow. Values shall be a multiple of four. A value of zero indicates no parameters follow.

**Flow Control Parameters:** These parameters contain information used to configure flow control for the ISL. Flow control parameters are specific to a given flow control mode.

**6.2.4.2 R\_RDY flow control**

A value of 0002h in the ISL Flow Control Mode field indicates that R\_RDY flow control, as defined in FC-FS-5, shall be used. When R\_RDY flow control mode is used, the Flow Control Parameter Length field shall be set to 20h and the Flow Control Parameters field shall contain the fields as specified in table 11. Values other than 0002h for the ISL Flow Control Mode field are not required to follow the Flow Control Parameter field format specified in table 11.

**Table 11 – Flow Control Parameters**

Item	Size
BB_Credit	4
Compatibility Parameters	16

**Buffer-to-buffer Credit:** The BB\_Credit field specified shall be associated with the number of buffers available for holding Class 2, Class 3 or Class F frames received from the Interconnect\_Port. The Buffer-to-buffer Credit shall be a single value that represents the total buffer-to-buffer Credit available for all Class 2 frames, and all Class 3 frames. The buffer-to-buffer credit value may also be applied to Class F frames.

**Compatibility Parameters:** This field contains associated compatibility parameters to assist in assuring backward compatibility with existing implementations.

Compatibility Parameter values that should be used for R\_RDY Flow Control mode are described in FC-MI-3 (reference [2]).

**6.2.4.3 VC\_RDY flow control**

A value of 2000h in the ISL Flow Control Mode field indicates that VC\_RDY flow control shall be used. When VC\_RDY flow control mode is used, the Flow Control Parameter Length field shall be based on the number of VCs supported and the Flow Control Parameters field shall contain fields as specified in table 12.

**Table 12 – VC\_RDY flow control parameters**

Item	Size
BB_Credit	4
Assignment Scheme	2
VC Value	2
VC Credit 0	4
...	
VC Credit n-1	4

If VC\_RDY flow control mode is specified, and the recipient of the ELP does not support the VC\_RDY flow control mode, then the ELP is rejected with a reason code explanation of "Invalid Flow Control Code".

**Buffer-to-buffer Credit:** The BB\_Credit field specified shall be associated with the number of buffers available for holding Class 2, Class 3 or Class F frames received from the Interconnect\_Port independent of VC\_Credit. The Buffer-to-buffer Credit shall be a single value that represents the total buffer-to-buffer Credit available for all Class 2 frames, and all Class 3 frames. The buffer-to-buffer credit value may also be applied to Class F frames.

**Assignment Scheme:** The assignment scheme identifies how frames are assigned to virtual channels. The assignment schemes and their values are provided in table 13.

**Table 13 – Assignment schemes**

Value	Scheme
0001h	Simple
0002h	Fixed
0003h	Variable
FF00h-FFFFh	Vendor Specific
other values	Reserved

**VC Value:** This field contains an integer that maps to the number of VCs defined for a given assignment scheme.

Valid VC values for the Simple assignment scheme are specified in table 14.

**Table 14 – VC values - Simple**

VC value	Number of VCs (n)
0001h	2
other values	reserved

Valid VC values for the Fixed assignment scheme are specified in table 15.

**Table 15 – VC values - Fixed**

VC value	Number of VCs (n)
0001h	8
0002h	12
0003h	20
0004h	36
0005h	68
0006h	132
other values	Reserved

Valid VC values for the Variable assignment scheme are specified in table 16.

**Table 16 – VC Values - Variable**

VC value	Number of VCs (n)
0002h	4
0003h	8
0004h	16
0005h	32
0006h	64
0007h	128
0008h	256
other values	Reserved

Clause 14 contains additional information regarding Virtual Channels and their associated assignment schemes.

#### 6.2.4.4 ELP reply

##### Reply Switch Fabric Internal Link Service Sequence:

- Service Reject (SW\_RJT)
  - Signifies the rejection of the ELP request
- Accept (SW\_ACC)
  - Signifies acceptance of the ELP request
  - Accept payload

**Payload:** The format of the ELP accept payload is specified in table 17.

**Table 17 – ELP accept payload**

Item	Size (bytes)
02000000h	4
Revision	1
Flags	2
BB_SC_N	1
R_A_TOV	4
E_D_TOV	4
Responder Interconnect_Port_Name	8
Responder Switch_Name	8
Fabric Controller Class F Service Parameters	16
Obsolete	4
Class 2 Interconnect_Port Parameters	4
Class 3 Interconnect_Port Parameters	4
Reserved	20
ISL Flow Control Mode	2
Flow Control Parameter Length (N)	2
Flow Control Parameters	N

The fields in table 17 are the same as defined in table 6.

### 6.2.5 Exchange Fabric Parameters (EFP)

The Exchange Fabric Parameters Switch Fabric Internal Link Service requests the exchange of Fabric Parameters between two E\_Ports connected via an ISL. The exchange of Fabric Parameters is used to establish the address allocation within the Fabric. When an E\_Port receives EFP from another E\_Port, all Active or Open Class F Sequences and Dedicated Connections shall be unaffected.

Use of the EFP SW\_ILS for Fabric Configuration is described in 7.3 and 7.4.

#### Protocol:

Exchange Fabric Parameters (EFP) request Sequence

Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the EFP request payload is specified in table 18.

**Table 18 – EFP request payload**

Item	Size (bytes)
Command code = 11h	1
Record length = 10h	1
Payload length	2
Reserved	3
Principal Switch_Priority	1
Principal Switch_Name	8
Domain_ID_List	N
Obsolete	N

**Record Length:** This field contains an 8-bit unsigned binary integer that specifies the total length of each record in the payload (see below). The value shall be 10h.

**Payload Length:** This field contains a 16-bit unsigned binary integer that specifies the total length in bytes of the Payload. The value specified shall be greater than or equal to 16, and less than or equal to 65532.



**Principal Switch\_Priority:** This field shall specify the priority level of the Switch that the transmitting Switch believes is the Principal Switch. Values for this field are summarized in table 19.

**Table 19 – Switch\_Priority field values**

Value (hex)	Description
00	Reserved
01	Highest priority value <sup>a</sup>
02	The Switch was the Principal Switch prior to sending or receiving BF <sup>b</sup>
03 to FE	Higher to lower priority values <sup>c</sup>
FF	The Switch is not capable of acting as a Principal Switch
<sup>a</sup> This value allows the system administrator to establish which Switch becomes the Principal Switch. <sup>b</sup> This allows the same Switch to become Principal Switch if it is still part of the Fabric after sending and/or receiving the Build Fabric SW_ILS. <sup>c</sup> The Switch_Priority value for a given Switch is established by means not defined by this standard.	

**Principal Switch\_Name:** This field shall specify the Switch\_Name of the Switch that the transmitting Switch believes is the Principal Switch.

**Domain\_ID\_List:** This field shall contain a list of records that specify the Domain\_ID and corresponding Switch\_Name of the Switch that has been granted the Domain\_ID by the Principal Switch. The Domain\_ID\_List shall contain a record for each value of Domain\_ID that has been assigned. If no Switch has been assigned a Domain\_ID, the Domain\_ID\_List shall contain no records. The format of a Domain\_ID\_List record is specified in table 20.

**Table 20 – Domain\_ID\_List record format**

Item	Size (bytes)
Record_Type	1
Domain_ID	1
Reserved	2
Reserved	4
Switch_Name for Domain_ID	8

**Record\_Type:** This field shall specify the type of record. Values for this field are summarized in table 21.

**Table 21 – Record\_Type field values**

Value (hex)	Description
00	Reserved
01	Domain_ID_List record
02	Obsolete
all others	Reserved

**Domain\_ID:** This field shall specify the Domain\_ID assigned by the Principal Switch.

**Switch\_Name for Domain\_ID:** This field shall specify the Switch\_Name of the Switch that has been assigned the Domain\_ID by the Principal Switch.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)  
Signifies the rejection of the EFP request
- Accept (SW\_ACC)  
Signifies acceptance of the EFP request
- Accept payload

**Payload:** The format of the EFP accept payload is specified in table 22.

**Table 22 – EFP accept payload**

Item	Size (bytes)
Command code = 02h	1
Page length = 10h	1
Payload length	2
Reserved	3
Principal Switch_Priority	1
Principal Switch_Name	8
Domain_ID_List	N

The fields in table 22 are the same as defined for table 18 with the following exception. The Domain\_ID\_List in the EFP request payload specifies the current Domain\_ID\_List of the originating Switch. The Domain\_ID\_List in the EFP accept payload specifies the Domain\_ID\_List of the responding Switch prior to the merging of the received Domain\_ID\_List from the originating Switch

with the Domain\_ID\_List of the responding Switch. This ensures that both the sending Switch and the responding Switch each have the same Domain\_ID\_List following the EFP exchange.

**6.2.6 Domain Identifier Assigned (DIA)**

The Domain Identifier Assigned Switch Fabric Internal Link Service indicates that a Principal Switch has been selected, and that the upstream neighbor Switch has been assigned a Domain Identifier. This communication signals that the Recipient may request a Domain Identifier from the Originating E\_Port.

Use of the DIA SW\_ILS for Fabric Configuration is described in 7.4.

**Protocol:**

- Domain Identifier Assigned (DIA) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the DIA request payload is specified in table 23.

**Table 23 – DIA request payload**

Item	Size (bytes)
12000000h	4
Originating Switch_Name	8
Not Meaningful	4

**Originating Switch\_Name:** This field shall contain the Switch\_Name of the Switch that originated the DIA request.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)
  - Signifies the rejection of the DIA request
- Accept (SW\_ACC)
  - Signifies acceptance of the DIA request
  - Accept payload

**Payload:** The format of the DIA accept payload is specified in table 24.

**Table 24 – DIA accept payload**

Item	Size (bytes)
02000000h	4
Responding Switch_Name	8
Not Meaningful	4

**Responding Switch\_Name:** This field shall contain the Switch\_Name of the Switch that responds to the DIA request.

### 6.2.7 Request Domain\_ID (RDI)

The Request Domain\_ID Switch Fabric Internal Link Service is sent by a Switch to request a Domain\_ID from the Domain Address Manager. RDI shall not be sent by a Switch unless the Switch has received a DIA SW\_ILS since the last reconfiguration event.

Use of the RDI SW\_ILS for Fabric Configuration is described in 7.4.

#### Protocol:

- Request Domain\_ID (RDI) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the RDI request payload is specified in table 25.

**Table 25 – RDI request payload**

Item	Size (bytes)
13h	1
Reserved	1
Payload Length	2
Requesting Switch_Name	8
Reserved	3
Requested Domain_ID #1	1
Reserved	3
Requested Domain_ID #2	1
...	
Reserved	3
Requested Domain_ID #n	1

**Payload Length:** This field contains a 16-bit unsigned binary integer that specifies the total length in bytes of the payload. The value specified shall be greater than or equal to 16, and less than or equal to 964.

**Requesting Switch\_Name:** This field specifies the Switch\_Name of the Switch requesting a Domain\_ID.

**Requested Domain\_ID:** This field shall contain one or more requested Domain\_IDs for the requesting Switch. If there is a Preferred Domain\_ID, this field is set to the Preferred Domain\_ID, otherwise it is set to zero. If more than one Domain\_ID is requested then none of the requested Domain\_IDs shall be zero.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)
  - Signifies the rejection of the RDI request
- Accept (SW\_ACC)
  - Signifies acceptance of the RDI request
  - Accept payload

**Payload:** The format of the RDI accept payload is specified in table 26.

**Table 26 – RDI accept payload**

Item	Size (bytes)
02h	1
Reserved	1
Payload Length	2
Requesting Switch_Name	8
Reserved	3
Granted Domain_ID #1	1
Reserved	3
Granted Domain_ID #2	1
...	
Reserved	3
Granted Domain_ID #n	1

**Payload Length:** This field contains a 16-bit unsigned binary integer that specifies the total length of the payload. The least significant two bits shall be zero. The value specified shall be equal to the value specified in the request payload.

**Requesting Switch\_Name:** This field specifies the Switch\_Name of the Switch requesting a Domain\_ID.

**Granted Domain\_ID:** This field shall contain the Domain\_ID granted by the Domain Address Manager to the requesting Switch. The Granted Domain\_ID field is set to:

- a) the Preferred Domain\_ID specified in the request if it is available;
- b) another Domain\_ID if the Preferred Domain\_ID specified in the request is not available; or
- c) another Domain\_ID if the Requested Domain\_ID value specified in the request is zero.

If no Domain\_ID is available then an SW\_RJT shall be returned. An SW\_RJT may be returned if the requested Domain\_ID is not available, or for other reasons the Principal Switch is unable to grant a Domain\_ID to the requesting Switch. If more than one Requested Domain\_ID was specified in the request, the response shall contain a number of Granted Domain\_IDs equal to the number requested. If the Domain Address Manager is unable to grant the full set of Domain\_IDs, it shall reject the request.

## 6.2.8 Hello (HLO)

### 6.2.8.1 HLO Overview

The Hello Switch Fabric Internal Link Service is used to determine when two way communication is established with a neighbor Switch. The exchange of Domain\_IDs is also used to determine the health of the ISL. When an E\_Port receives HLO from another E\_Port, all Active or Open Class F Sequences and Dedicated Connections shall be unaffected.

The HLO SW\_ILS shall be sent as a unidirectional Exchange.

Use of the HLO SW\_ILS for path selection is described in clause 8. Other uses of HLO are not defined by this standard.

#### Protocol:

Hello (HLO) request Sequence

**Addressing:** For use in path selection, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the HLO request payload is specified in table 27.

**Table 27 – HLO request payload**

Item	Size (bytes)
FSPF Header	20
Reserved	4
Hello_Interval	4
Dead_Interval	4
Recipient Domain_ID	4
Reserved	1
Originating Port Index	3

**FSPF Header:** The format of the FSPF Header is described in 6.2.8.2.

**Hello\_Interval:** This field shall specify in seconds, the interval between two consecutive HLO messages generated by the Switch during the life of the Adjacency with the neighbor Switch (see table 292).

**Dead\_Interval:** This field shall specify in seconds, the maximum interval the requesting Switch shall wait for reception of a Hello from its neighbor. If the interval expires and no Hello has been received, then the detecting Switch shall bring down the Adjacency (see table 292). Switches may also reset this timer on reception of an FSPF LSA or LSU.

The Hello\_Interval and Dead\_Interval values are configured separately for each port. Two E\_Ports connected by an ISL shall share the same two values for proper operation of the FSPF protocol.

**Recipient Domain\_ID:** This field shall specify the Domain\_ID of the neighbor Switch. If the neighbor Domain\_ID is known, then the Recipient Domain\_ID value shall be set to 000000hIIIDomain\_ID'. If the neighbor Domain\_ID is unknown, then the Recipient Domain\_ID value shall be FFFFFFFFh.

Valid values for the Domain\_ID are: 01h-EFh.

**Originating Port Index:** This field shall specify the source E\_Port Index.

### 6.2.8.2 FSPF Header format

The format of the FSPF Header is specified in table 28.

**Table 28 – FSPF Header**

Item	Size (bytes)
Command	4
FSPF Version	1
Obsolete	1
Authentication Type	1
Reserved	1
Originating Domain_ID	4
Authentication	8

**Command:** This field indicates the command code for the FSPF ILS. FSPF command code values are specified in table 29.

**Table 29 – FSPF command codes**

Value (hex)	Description
14000000	Hello
15000000	Link State Update
16000000	Link State Acknowledgement

**FSPF Version:** This field contains a code that indicates the FSPF protocol version. The value shall be 02h.

**Authentication Type:** This field shall specify the usage of the Authentication field. This value shall be set to 00h.



**Originating Domain\_ID:** This field contains the Domain\_ID of the Switch sending this request. The Domain\_ID value shall be set to 000000h||Domain\_ID. If multiple Domain\_IDs are in use by the Switch, then the Switch shall use the lowest value Domain\_ID as the Originating Domain\_ID. Valid values for the Domain\_ID are: 01h-EFh.

**Authentication:** This field shall specify the Authentication information appropriate for the specified Authentication Type. This field shall contain 0000000000000000h.

**6.2.9 Link State Update (LSU)**

**6.2.9.1 LSU overview**

The Link State Update Switch Fabric Internal Link Service requests the transfer of one or more Link State Records from one Switch to another Switch. The transfer may be of updated Link State Records, or may be a transfer of an entire Link State Database. When an E\_Port receives an LSU from another E\_Port, all Active or Open Class F Sequences and Dedicated Connections shall be unaffected.

The LSU SW\_ILS shall be sent as a unidirectional Exchange.

Use of the LSU SW\_ILS for path selection is described in clause 8. Other uses of LSU are not defined by this standard.

**Protocol:**

Link State Update (LSU) request Sequence

**Addressing:** For use in path selection, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the LSU request payload is specified in table 30.

**Table 30 – LSU request payload**

Item	Size (bytes)
FSPF Header	20
Reserved	3
Flags	1
Number of Link State Records	4
Link State Records	n

**FSPF Header:** The format of the FSPF header is described in 6.2.8.2.

**Flags:** This field shall contain information used to synchronize the Link State Database. The bit map values are listed in table 31.

**Table 31 – Flags field bit map**

Bit	Description
0	Data Base Exchange - Value b'1' - LSU is used for initial database synchronization Value b'0' - LSU is used for a topology update
1	Database Complete Value b'1' - Last sequence of data base synchronization. LSU contains no LSRs. Value b'0' - Not the last sequence of data base synchronization
2-7	Reserved

**Number of Link State Records:** This field shall specify the number of Link State Records that follow this field.

### 6.2.9.2 Link State Record (LSR) format

**Link State Record:** There is one format for the LSR, the Link Descriptor format. The Link Descriptor format is specified in table 32. One or more descriptors may be contained in a single LSR.

**Table 32 – Link State Record - Link Descriptor format**

Item	Size (bytes)
Link State Record Header (LSR Type 01h)	24
Reserved	2
Number of Links	2
Link Descriptor #1	16
...	16
...	16
Link Descriptor #n	16

**Link State Header:** The format of the Link State Header is described in 6.2.9.3.

**Number of Links:** This field specifies the number of Link Descriptors contained in the Link State Record. A Switch keeps a list of all its ISLs, but only ISLs that are in the full state shall be advertised in the LSR.

**Link Descriptor:** The format of the Link Descriptor is described in 6.2.9.4.

### 6.2.9.3 Link State Header format

The format of the Link State Header is described in table 33.

**Table 33 – Link State Header format**

Item	Size (bytes)
LSR Type	1
LSR Flags	1
LSR Age	2
Reserved	4
Link State Identifier	4
Advertising Domain_ID	4
Link State Incarnation Number	4
Checksum	2
LSR Length	2

**LSR Type:** The LSR types are specified in table 34.

**Table 34 – Link State Record Type field values**

Value (hex)	Description
01	Switch Link Record
02	Obsolete
F0-FF	Vendor Specific
all others	Reserved

**LSR Flags:** This field contains flag bits that provide additional information about the LSR. The

following LSR flag bits are specified in table 35.

**Table 35 – LSR Flags field bit descriptions**

Bit	Description
0	Leaf Switch: If this bit is set to one, then the Switch shall be a route termination point and other Switches in the Fabric shall not route through this Switch. If this bit is set to zero, then the Switch shall allow other Switches in the Fabric to establish routes through this Switch.
1 to 7	Reserved

**LSR Age:** This field contains a value that indicates the time in seconds since the record has been generated. LSR Age may be used to flush old records from the database.

**Link State Identifier:** This field contains the Domain\_ID of the Switch that owns the LSR. The format of Link State Identifier shall be set to '000000hllDomain\_ID'.

**Advertising Domain\_ID:** This field contains the Domain\_ID of the Switch that is advertising the LSR on behalf of the owning Switch.

**Incarnation Number:** This field contains the current incarnation of the LSR.

**Checksum:** This field contains the checksum value of the Link State Record. This value shall be calculated on all bytes of the LSR except for the LSR Age field. A complete description of how the checksum is calculated is given in 8.5.5.

NOTE 3 – Not calculating the checksum on the Age value allows the Age value to advance without requiring the recalculation of the checksum.

**LSR Length:** This field contains the length of the LSR in bytes. The LSR length includes the LSR Age field.

#### 6.2.9.4 Link Descriptor format

This field contains a descriptor that defines the state of the ISL, as specified in table 36.

**Table 36 – Link Descriptor format**

Item	Size (bytes)
Link ID	4
Reserved	1
Output Port Index	3
Reserved	1
Neighbor Port Index	3
Link Type	1
Reserved	1
Link Cost	2

**Link Identifier:** This field identifies the link and contains the Domain\_ID of the neighbor Switch at the other end of the ISL, relative to the owning Switch.

**Output Port Index:** This field shall specify the source E\_Port Index.

**Neighbor Port Index:** The field shall specify the destination E\_Port Index.

**Link Type:** This field shall specify the type of ISL. Values are specified in table 37.

**Table 37 – Link Type values**

Value (hex)	Description
01	Point to Point Link
F0-FF	Vendor Specific
all others	Reserved

**Link Cost:** This field contains a value that describes the cost of transmitting a frame over the ISL. See 8.5.6 for a complete description of Link Cost calculation.

#### 6.2.10 Link State Acknowledgement (LSA)

The Link State Acknowledgement Switch Fabric Internal Link Service is used to acknowledge the receipt of an LSR. When an E\_Port receives LSA from another E\_Port, all Active or Open Class F Sequences and Dedicated Connections shall be unaffected.

The LSA SW\_ILS shall be sent as a unidirectional Exchange.

Use of the LSA SW\_ILS for path selection is described in clause 8. Other uses of LSA are not defined by this standard.

**Protocol:**

Link State Acknowledgement (LSA) request Sequence

**Addressing:** For use in path selection, the S\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the LSA request payload is specified in table 38.

**Table 38 – LSA request payload**

Item	Size (bytes)
FSPF Header	20
Reserved	3
Flags	1
Number of Link State Record Headers	4
Link State Record Headers	n

**FSPF Header:** The format of the FSPF Header is described in 6.2.8.2.

**Flags:** The bit settings shall match the bit settings specified in the Flags field of the corresponding LSU.

**Number of Link State Record Headers:** This field shall specify the number of Link State Record Headers that follow this field.

**Link State Record Header:** The format of the Link State Record header is described in 6.2.9.3.

**6.2.11 Build Fabric (BF)**

The Build Fabric Switch Fabric Internal Link Service requests a non-disruptive reconfiguration of the entire Fabric. Fabric Configuration is performed as described in clause 7.

NOTE 4 – The BF SW\_ILS allows the Fabric to attempt reconfiguration without loss of or change of address. Examples of situations in which BF is appropriate include certain losses of a Principal ISL (Link Failure or Offline), or when two Fabrics are joined.

A BF shall cause the Domain\_ID\_List to be cleared.

The transmission or reception of BF shall not of itself cause the loss of Class N frames, or cause a busy response to any Class N frames. Active or Open Class F Sequences between the two E\_Ports, and any Dedicated Connections, shall not be abnormally terminated.

Use of the BF SW\_ILS for Fabric Configuration is described in 7.3 and 7.4.

**Protocol:**

- Build Fabric (BF) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the BF request payload is specified in table 39.

**Table 39 – BF request payload**

Item	Size (bytes)
17000000h	4

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)
  - Signifies the rejection of the BF command
- Accept (SW\_ACC)
  - Signifies acceptance of the BF request
  - Accept payload

**Payload:** The format of the BF accept payload is specified in table 40.

**Table 40 – BF accept payload**

Item	Size (bytes)
02000000h	4

**6.2.12 Reconfigure Fabric (RCF)**

The Reconfigure Fabric Switch Fabric Internal Link Service requests a disruptive reconfiguration of the entire Fabric. Fabric Configuration is performed as described in clause 7.

NOTE 5 – Since the RCF causes a complete reconfiguration of the Fabric, and may cause addresses allocated to a Switch to change, this SW\_ILS is recommended to be used with caution. The BF SW\_ILS allows the Fabric to attempt reconfiguration without loss of or change of address and therefore should be attempted before an RCF. Examples of situations in which RCF may be appropriate include resolution of overlapped Domains, or the failure of a Fabric Reconfiguration initiated by a BF.

An RCF shall cause the Domain\_ID\_List to be cleared.

When an RCF is transmitted by an E\_Port, any Active or Open Class F Sequences between the two E\_Ports, and any Dedicated Connections, shall be abnormally terminated. Also, all Class N frames shall be discarded, and all Dedicated Connections shall be abnormally terminated.

When an RCF is received and accepted by an E\_Port, any Active or Open Class F Sequences between the two E\_Ports, and any Dedicated Connections, shall be abnormally terminated prior to transmission of the SW\_ACC reply Sequence. Also, all Class N frames shall be discarded, and all Dedicated Connections shall be abnormally terminated prior to transmission of the SW\_ACC reply Sequence. If an E\_Port rejects the RCF, the Switch to which it belongs shall not propagate the RCF over its other E\_Ports, nor send an ELP over its Isolated Interconnect\_Ports. The rejecting E\_Port shall go in Isolated state and send an SW\_RJT reply Sequence with reason code explanation "E\_Port is Isolated".

Use of the RCF SW\_ILS for Fabric Configuration is described in 7.3 and 7.4.

**Protocol:**

Reconfigure Fabric (RCF) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the RCF request payload is specified in table 41.

**Table 41 – RCF request payload**

Item	Size (bytes)
18000000h	4

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the RCF command  
 Accept (SW\_ACC)  
 Signifies acceptance of the RCF request  
 – Accept payload

**Payload:** The format of the RCF accept payload is specified in table 42.

**Table 42 – RCF accept payload**

Item	Size (bytes)
02000000h	4

**6.2.13 Inter-Switch Registered State Change Notification (SW\_RSCN)**

The Fabric shall distribute RSCNs between Switches using the Inter-Switch RSCN payload.

The Inter-Switch RSCN format is similar to the format used for Nx\_Ports, except:

- a) the affected N\_Port field is as defined in the payload description in this subclause. The upper nibble is masked before delivery to an Nx\_Port;



- b) a "detection function" code is contained in the payload;
- c) an SW\_ILS is used as the transport; and
- d) if a Switch has any directly attached Nx\_Ports registered to receive RSCNs, it shall convert a received SW\_RSCN SW\_ILS to an appropriate RSCN ELS.

**Protocol:**

Inter-Switch Registered State Change Notification (SW\_RSCN) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh designating the Domain Controller ID of the Switch that generates the SW\_RSCN. The D\_ID shall be set to FFFCyh to designate the Domain Controller ID of the recipient Switch.

**Payload:** The format of the SW\_RSCN request payload is specified in table 43.

**Table 43 – SW\_RSCN request payload**

Item	Size (bytes)
1B00000h	4
Affected N_Port	4
Detection Function	4
Number of Device Entries (m)	4
Device Entry 1	20
Device Entry 2	20
.....	
Device Entry m	20

**Affected N\_Port**

This field specifies the address of the affected N\_Port.

For Fabric events (see above). The first nibble in the high order byte shall be:

- 0xh = no additional info;
  - 1xh = port is online;
  - 2xh = port is offline;
- where (x indicates a valid hexadecimal value).

The second nibble in the high order byte shall be:

- x0h = port address format
  - x1h = area address format
  - x2h = domain address format
  - x3h = Fabric address format
- where (x indicates a valid hexadecimal value).

The remaining three bytes contain the 24 bit address.

**Detection function**

The value used by SCR (see FC-FS-5) to describe the detector of the change:

00000001h = Fabric detected

00000002h = N\_Port detected

### Number of Device Entries

This field contains the number of device entries in the payload.

### Device Entry

The format of the device entry is specified in table 44.

**Table 44 – Device Entry format**

Item	Size (bytes)
Port State	1
N_Port_ID	3
N_Port_Name	8
Node_Name	8

For an N\_Port device the number of devices would be 1 and the N\_Port\_ID entry in the only device entry would be identical to the value in the N\_Port\_ID portion of the affected N\_Port field in the payload. In case of a Loop port and where the SW\_RSCN format is a AREA wide format, the number of devices would be the total number of devices in the loop port that is either coming online or going offline. Also note that if there are 126 devices in a loop port then the SW\_RSCN itself may become a multi-frame sequence. An AREA format SW\_RSCN should be converted to an RSCN ELS with only one Affected N\_Port\_ID page.

### Port State

This byte may contain the Port State. The state values are the same values as defined in the Affected Port description.

### N\_Port\_ID

This field contains the 24 bit Fibre Channel Address of the device.

### N\_Port\_Name

This field contains the Name\_Identifier of the port associated with the device.

### Node\_Name

This field contains the Name\_Identifier of the Node associated with the device.

### Reply Switch Fabric Internal Link Service Sequence:

Service Reject (SW\_RJT)

Signifies the rejection of the SW\_RSCN request

Accept (SW\_ACC)

Signifies acceptance of the SW\_RSCN request and its RSCN information

– Accept payload

**Payload:** The format of the SW\_RSCN accept payload is specified in table 45.

**Table 45 – SW\_RSCN accept payload**

Item	Size (bytes)
02000000h	4

**6.2.14 Distribute Registered Link Incident Records (DRLIR)**

Distribute Registered Link Incident Records (DRLIR) Switch Fabric Internal Link Service provides a method for a Fabric built RLIR to be distributed to every Switch in the Fabric. The normal response to a DRLIR SW\_ILS sequence shall be an SW\_ACC. If the recipient Switch does not support the DRLIR SW\_ILS, the recipient Switch shall reply with an SW\_RJT with a reason code of “command not supported”. If the recipient Switch does not support the RLIR format contained in the DRLIR, the recipient Switch shall reply with an SW\_RJT with a reason code of “unable to perform command request”.

When a Switch creates an RLIR, the Switch shall generate the corresponding DRLIRs. A DRLIR shall be created for every Established Registration List that the originating Switch supports, even if that Switch has no registrants in the Established Registration List. The Switch shall distribute the DRLIRs to every Switch in the Fabric via the Domain Controller Identifier.

When a Switch receives a DRLIR, the Switch shall extract the RLIR. The RLIR shall then be sent to the local registrants of the given RLIR format as if the RLIR was generated in the local Switch.

**Protocol:**

- Distribute Registered Link Incident Record (DRLIR) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh designating the Domain Controller ID of the Switch that generates the DRLIR. The D\_ID shall be set to FFFCyyh to designate the Domain Controller ID of the recipient Switch.

**Payload:** The format of the DRLIR request payload is specified in table 46.

**Table 46 – DRLIR request payload**

Item	Size (bytes)
1E000000h	4
Embedded RLIR	28 to 328

**Embedded RLIR:** The format of the embedded RLIR is defined in FC-LS-3.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)
  - Signifies the rejection of the DRLIR request
- Accept (SW\_ACC)
  - Signifies acceptance of the DRLIR request and its Link Incident Record

- Accept payload

**Payload:** The format of the DRLIR accept payload is specified in table 47.

**Table 47 – DRLIR accept payload**

Item	Size (bytes)
02000000h	4

### 6.2.15 Merge Request (MR)

The Merge Request SW\_ILS requests that the recipient merge any Zoning data with the Zoning data supplied in the MR payload according to the rules specified in table 279. The Merge Request provides a mechanism to distribute Zoning information between Adjacent Switches. Use of the Merge Request is described in 10.2.

To distinguish between Enhanced and Basic Zoning, a Protocol Version field is used.

If Protocol Version is 00h:

- a) the payload contains Basic Zoning structures; and
- b) the Fabric is working in Basic Zoning mode.

If Protocol Version is 01h:

- a) the payload contains Enhanced Zoning structures; and
- b) the Fabric is working in Enhanced Zoning mode.

The Zone Merge may be successful only between Switches working in the same Zoning mode (i.e., both in Basic Zoning mode or both in Enhanced Zoning mode). This means that the value of the received Protocol Version field shall match the current Zoning operational mode of the Switch, otherwise the link is Isolated. In particular, if a Switch working in Enhanced Zoning mode receives over a certain link a MR with Protocol Version field equal to 00h, then that link shall be Isolated.

**Protocol:**

- Merge Request (MR) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the MR request payload is specified in 6.2.15.1.

### 6.2.15.1 Merge Request payload

The format of the MR request payload is specified in table 48.

**Table 48 – Merge Request request payload**

Item	Size
Merge Request - 22h	1
Protocol Version	1
Version Specific Payload	x

#### Protocol Version

The Protocol Version field contains a number that identifies the Zoning operational mode of the Fabric (i.e., Basic or Enhanced) and the format of the Zoning structures conveyed in the payload. Table 49 specifies the defined values.

**Table 49 – Protocol Version values**

Value	Meaning
00	Basic Zoning
01	Enhanced Zoning
others	Reserved

#### 6.2.15.1.1 Merge Request payload in Basic Zoning

The format of the Version Specific payload for Protocol Version = 00h is specified in table 50.

**Table 50 – Basic Zoning payload**

Item	Size
Active Zone Set Length	2
Active Zone Set Name	a
Active Zone Set Object List	x
Zone Set Database Object Length	4
Zone Set Database Object List	y

The Zone Set Database Object list may not be present, and if so the Zone Set Database Object Length shall be set to zero.

#### Active Zone Set Length

The Active Zone Set Length field contains the length of the Active Zone Set Name and Active Zone Set Object List, x+y (in bytes).

**Active Zone Set Name**

The Active Zone Set name field contains the name of the Active Zone Set. The format of the Active Zone Set name follows the structure and rules for the Name Entry described in 10.4.2.3.

**Active Zone Set Object List**

The Active Zone Set may only contain Zone Objects (type 2) in the Active Zone Set Object List.

In the Basic Zoning Framework each of the Zone Object members may be of member type N\_Port\_Name (type 1), Domain\_ID and physical port (type 2), or N\_Port\_ID (type 3). All other Zone Member types are not allowed.

**Zone Set Database Object List**

The Zone Set Database Object list contains information regarding all Zone configurations plus all objects that comprise the Zone Sets. The Active Zone Set, name and object list, shall not be included in the Zone Set Database Object list. Support of the Zone Set Database Object list is optional. A Zone Set Database Object length of 0 is required if the Zone Set Database is not supported.

In the Basic Zoning Framework the Zone Set Database does not use all Zoning Object types in the Zone Set Database Object List. Zone Set type objects shall have members that are only Zone Objects (type 2). Each of the Zone Object members may be of member type N\_Port\_Name (type 1), Domain\_ID and physical port (type 2), N\_Port\_ID (type 3), or Alias Name (type 4). Each Zone Alias Object member may be of member type N\_Port\_Name (type 1), Domain\_ID and physical port (type 2), or N\_Port\_ID (type 3). All other combinations are not allowed.

**6.2.15.1.2 Merge Request payload in Enhanced Zoning**

The format of the Version Specific payload for Protocol Version = 01h is specified in table 51.

**Table 51 – Enhanced Zoning payload**

Item	Size
Reserved	2
Enhanced Zoning Flags	4
Active Zone Set Length	4
Active Zone Set Object List	x
Zone Set Database Object Length	4
Zone Set Database Object List	y

**Enhanced Zoning Flags**

The format of the Enhanced Zoning Flags field is as follows:

Bit 0- reserved.

Bit 1- reserved.

Bit 2- Indicates the Merge Control Setting. When this bit is one, this Switch is working in Restrict mode, so it may join a Fabric only if the Fabric's Zoning Database is equal to its Zoning Database. When this bit is zero, this Switch is working in Allow mode, so it may join a Fabric only if the Fabric's Zoning Database is mergeable with its Zoning Database.

Bit 3- Indicates the Default Zone Setting. When this bit is one this Switch denies traffic between members of the Default Zone. When this bit is zero this Switch permit traffic between members of the Default Zone.

Bit 4- Indicates that the Zone Set Database is supported. When this bit is one, the Zone Server on this Switch is able to maintain a Zone Set Database. When this bit is zero, the Zone Server on this Switch is not able to maintain a Zone Set Database.

Bit 5- Indicates that the Zone Set Database is enabled. When this bit is one, the Zone Server on this Switch is maintaining a Zone Set Database. When this bit is zero, the Zone Server on this Switch is not maintaining a Zone Set Database.

Bit 6-31 reserved.

### **Active Zone Set Length**

The Active Zone Set Length field is extended to 4 bytes and contains the length of the Active Zone Set Object List, in bytes.

### **Active Zone Set Object List**

The Active Zone Set may only contain Zone Objects (type '02') in the Active Zone Set Object List.

Any Zone Member Identifier type may be used as Zone Member in the Active Zone Set's Zone Objects, with the exception of the Alias Name identifier (type '04').

### **Zone Set Database Object List**

The Zone Set Database Object list contains information regarding all Zone configurations plus all objects that comprise the Zone Sets. The Active Zone Set, name and object list, shall not be included in the Zone Set Database Object list. Support of the Zone Set Database Object list is optional. A Zone Set Database Object length of 0 is required if the Zone Set Database is not supported.

In the Enhanced Zoning Framework the Zone Set Database may use all Zoning Object types in the Zone Set Database Object List. Zone Set type objects shall have members that are only Zone Reference Objects (type '04'). Any Zone Member Identifier type may be used as Zone Member in the Zone Set Database's Zone Objects. Any Alias Member Identifier type may be used as Zone Alias Member in the Zone Set Database's Zone Alias Objects, with the exception of the Alias Name identifier (type '04').

#### **6.2.15.2 Merge Request reply**

##### **Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)

- Signifies the rejection of the MR request
- Accept (SW\_ACC)
- Signifies acceptance of the MR request

Successful completion of the Merge Request is indicated by an SW\_ACC. If the recipient is unable to complete the Merge Request, an SW\_RJT with reason code “Unable to Complete Command Requested” and reason code explanation indicating why the Merge Request was not completed shall be returned, and the E\_Port shall enter the Isolated State.

The format of the Merge Request accept payload is specified in table 52.

**Table 52 – Merge Request accept payload**

Item	Size (bytes)
02000000h	4
Reserved	1
Obsolete	1
Obsolete	1
Obsolete	1

**6.2.16 Acquire Change Authorization Request (ACA)**

Acquire Change Authorization requests are SW\_ILSs addressed from the Domain Controller of the Managing Switch to the Domain Controller of a Managed Switch. Acquire Change Authorization request messages are sent by a Managing Switch to Managed Switches to reserve local resources in each Switch.

The Acquire Change Authorization (ACA) request Switch Fabric Internal Link Service requests that the recipient reserve local resources for the purposes of changing Switch or Switch service resources. The Acquire Change Authorization request provides a mechanism to lock a Fabric to distribute information (e.g., Zoning) amongst Switches. Use of the Acquire Change Authorization is described in 10.6.2.

**Protocol:**

- Acquire Change Authorization (ACA) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh, indicating the Domain Controller of the originating Switch. The D\_ID shall be set to FFFCyh, indicating the Domain Controller of the destination Switch.

**Payload:** The format of the ACA request payload is specified in table 53.



**Table 53 – ACA request payload**

Item	Size (bytes)
23h	1
Reserved	1
Domain_ID List Length	2
Reserved	3
Domain_ID #1	1
Reserved	3
Domain_ID #2	1
...	
Reserved	3
Domain_ID #n	1

**Domain\_ID\_List Length:** This field specifies the length of the Domain\_ID List in bytes.

**Domain\_ID List:** The payload contains a list of Domain\_ID's known to the Managing Switch. The Domain\_ID List begins with the Reserved field immediately following the Domain\_ID List Length field. The recipient checks the list of Domain\_ID's against those it knows to be active within the Fabric. If the list differs from the Domain\_ID's known to the Managed Switch, the request is rejected with an SW\_RJT with a reason code "Unable to Perform Command Requested", and a reason code explanation of "Fabric Changing".

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)

Signifies the rejection of the ACA request

Accept (SW\_ACC)

Signifies acceptance of the ACA request

An SW\_ACC indicates that the operation completed successfully.

If the Managed Switch is unable to accept the ACA due to another pending ACA, an SW\_RJT with reason code "Logical Busy" shall be returned.

The format of the Acquire Change Authorization accept payload is specified in table 54.

**Table 54 – Acquire Change Authorization accept payload**

Item	Size (bytes)
02000000h	4
Reserved	1
Obsolete	1
Obsolete	1
Obsolete	1

After an unsuccessful attempt to acquire change authorization, a Switch should release any acquired change authorization, and wait a random time before attempting ACA again.

**6.2.17 Release Change Authorization (RCA) request**

Release Change Authorization requests are SW\_ILSs addressed from the Domain Controller of the Managing Switch to the Domain Controller of a Managed Switch. Release Change Authorization request messages are sent by a Managing Switch to Managed Switches to release local resources in each Switch.

**Protocol:**

Release Change Authorization (RCA) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh, indicating the Domain Controller of the originating Switch. The D\_ID shall be set to FFFCyh, indicating the Domain Controller of the destination Switch.

**Payload:** The format of the RCA request payload is specified in table 55.

**Table 55 – RCA request payload**

Item	Size (bytes)
24000000h'	4

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the RCA request  
 Accept (SW\_ACC)  
 Signifies acceptance of the RCA request

The format of the Release Change Authorization accept payload is specified in table 56.

**Table 56 – Release Change Authorization accept payload**

Item	Size (bytes)
02000000h	4
Reserved	1
Obsolete	1
Obsolete	1
Obsolete	1

**6.2.18 Stage Fabric Configuration (SFC) request**

Stage Fabric Configuration requests are SW\_ILSs addressed from the Domain Controller of the Managing Switch to the Domain Controller of a Managed Switch. Stage Fabric Configuration request messages are sent by a Managing Switch to Managed Switches to stage changes to local resources in each Switch.

The Stage Fabric Configuration request provides a mechanism to distribute information to other Switches in the Fabric. Use of the Stage Fabric Configuration is described in 10.6.3.

**Protocol:**

Stage Fabric Configuration (SFC) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh, indicating the Domain Controller of the originating Switch. The D\_ID shall be set to FFFCyh, indicating the Domain Controller of the destination Switch.

**Payload:** The format of the SFC request payload is specified in table 57.

**Table 57 – SFC request payload**

Item	Size (bytes)
25h	1
Operation Request (see table 58)	1
Operation Specific Payload	x

**Operation Request:** The operation request value further specifies the operation to be attempted by the recipient.

**Operation Specific Payload:** The remaining part of the SFC payload is dependent on the operation requested. Table 58 depicts the currently defined Operation Request values.

**Table 58 – Operation Request values**

<b>Value (hex)</b>	<b>Description</b>
00-02	Reserved
03	Activate Zone Set
04	Deactivate Zone Set
05-07	Reserved for FC-SP-2 use (See reference [7])
08	Activate Zone Set Enhanced
09	Deactivate Zone Set Enhanced
0A	Distribute Zone Set Database
0B	Activate Zone Set by Name
0C	Set Zoning Policies
0D-1F	Reserved
20-3F	Reserved for FC-SP-2 use (See reference [7])
40 thru DF	Reserved
E0 thru FF	Vendor Specific

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)  
Signifies the rejection of the SFC request
- Accept (SW\_ACC)  
Signifies acceptance of the SFC request

Stage Fabric Configuration responses are Class F frames addressed from the Domain Controller of a Managed Switch to the Domain Controller of the Managing Switch. A Stage Fabric Configuration accept is sent by a Managed Switch to a Managing Switch when a Stage Fabric Configuration request has been received.

The format of the Stage Fabric Configuration accept payload is specified in table 59.

**Table 59 – Stage Fabric Configuration accept payload**

Item	Size (bytes)
02000000h	4
Reserved	1
Obsolete	1
Obsolete	1
Obsolete	1

### 6.2.18.1 SFC in Basic Zoning

Operation Requests values 03 and 04 are used in the context of Basic Zoning. Only the Basic Zoning Data structures defined in 10.4.2 shall be used with them. Enhanced Zoning Data Structures shall not be used with them. Table 60 specifies the payload structure for them.

**Table 60 – Payload for Operation Request values 03 and 04**

Item	Size (bytes)
Zone Set Length	2
Zone Set Name	a
Zoning Object List	x
Zone Set Database Object Length	4
Zone Set Database Object List	y

The Zone Set Length field specifies the length in bytes of the following:

- a) Zone Set Name; and
- b) Zoning Object List.

The Zone Set Database Object Length field specifies the length in bytes of the Zone Set Database Object List. Refer to clause 6.2.15.1 for implementation notes.

If the request value is 03h, then the remainder of the SFC payload contains the Zone Set configuration utilized by the recipient to determine if a Activate Zone Set operation may be attempted.

If the request value is 04h, the remainder of the SFC payload is ignored.

### 6.2.18.2 SFC in Enhanced Zoning

The following Operation Requests are used in the context of Enhanced Zoning. Only the Enhanced Zoning Data structures defined in 10.4.4 shall be used with them. Basic Zoning Data Structures shall not be used with them.

### 6.2.18.2.1 Operation Request 'Activate Zone Set Enhanced'

Operation Request 'Activate Zone Set Enhanced' is used in Enhanced Zoning to activate a Zone Set distributing its definition across the Fabric. Together with the Zone Set to be activated, also the entire Zone Set Database may be distributed. Table 61 depicts the payload format.

**Table 61 – Payload for Operation Request 'Activate Zone Set Enhanced'**

Item	Size (bytes)
Reserved	2
Active Zone Set Length	4
Active Zone Set Object List	x
Zone Set Database Object Length	4
Zone Set Database Object List	y

#### 6.2.18.2.1.1 Length fields

The Active Zone Set Length field contains the length of the Active Zone Set Object List, in bytes. The Active Zone Set Length field is extended to 4 bytes.

The Zone Set Database Object length field specifies the length of the Zone Set Database Object List, in bytes. If set to zero, then the Zone Set Database is not included in the payload.

#### 6.2.18.2.1.2 Object Lists

The Object Lists shall use the appropriate Enhanced Zoning payloads for the Zone Set to be activated and the Zone Set Database, as described in 10.4.4.

### 6.2.18.2.2 Operation Request 'Deactivate Zone Set Enhanced'

Operation Request 'Deactivate Zone Set Enhanced' is used in Enhanced Zoning to deactivate the current Active Zone Set. Table 62 depicts the payload format.

**Table 62 – Payload for Operation Request 'Deactivate Zone Set Enhanced'**

Item	Size (bytes)
Reserved	2

### 6.2.18.2.3 Operation Request 'Distribute Zone Set Database'

Operation Request 'Distribute Zone Set Database' applies to the Zone Set Database. Its purpose is to distribute in the Fabric a new definition of the Zone Set Database, without affecting the Active Zone Set. Table 63 defines its payload.

**Table 63 – Payload for Operation Request 'Distribute Zone Set Database'**

Item	Size (bytes)
Reserved	2

**Table 63 – Payload for Operation Request ‘Distribute Zone Set Database’**

Item	Size (bytes)
Zone Set Database Object Length	4
Zone Set Database Object List	y

**6.2.18.2.3.1 Zone Set Database Object Length**

The Zone Set Database Object Length field specifies the length of the Zone Set Database Object List. If the Zone Set Database Object Length is zero, the Zone Set Database Object List is not present, and this operation clears the entire Zone Set Database.

**6.2.18.2.3.2 Zone Set Database Object Lists**

The Object List shall use the appropriate Enhanced Zoning payloads for the Zone Set Database, as described in 10.4.4.

**6.2.18.2.4 Operation Request ‘Activate Zone Set by Name’**

Operation Request ‘Activate Zone Set by Name’ applies to both Active Zone Set and Zone Set Database. Its purpose is to activate a Zone Set defined in the Zone Set Database without having to transmit over the Fabric its definition. Table 64 depicts the payload format.

**Table 64 – Payload for Operation Request ‘Activate Zone Set by Name’**

Item	Size (bytes)
Reserved	2
Zone Set Name	a

**6.2.18.2.4.1 Zone Set Name**

This field contains the Name of the Zone Set to be activated. It shall be defined in the Zone Set Database.

**6.2.18.2.5 Operation Request ‘Set Zoning Policies’**

Operation Request ‘Set Zoning Policies’ is used in Enhanced Zoning to establish the Fabric Zoning Policies. Table 65 depicts the payload format.

**Table 65 – Payload for Operation Request ‘Set Zoning Policies’**

Item	Size (bytes)
Reserved	2
Enhanced Zoning Flags	4

**6.2.18.2.5.1 Enhanced Zoning Flags**

The format of the Enhanced Zoning Flags field is as follows:

Bit 0-1 reserved.

Bit 2- **Merge Control Setting.** If this bit is one the Fabric shall work in Restrict mode, so a Switch may join the Fabric only if its Zoning Database is equal to the Fabric's Zoning Database. If this bit is zero the Fabric shall work in Allow mode, so a Switch may join the Fabric only if its Zoning Database is mergeable with the Fabric's Zoning Database.

Bit 3- **Default Zone Setting.** If this bit is one the Fabric shall deny traffic between members of the Default Zone. If this bit is zero the Fabric shall permit traffic between members of the Default Zone.

Bit 4-31 reserved.

### 6.2.19 Update Fabric Configuration (UFC) request

Update Fabric Configuration requests are SW\_ILSs addressed from the Domain Controller of the Managing Switch to the Domain Controller of a Managed Switch. Update Fabric Configuration request messages are sent by a Managing Switch to Managed Switches to effect the changes to local resources in each Switch. There is no data included in this message.

**Protocol:**

Update Fabric Configuration (UFC) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh, indicating the Domain Controller of the originating Switch. The D\_ID shall be set to FFFCyh, indicating the Domain Controller of the destination Switch.

**Payload:** The format of the UFC request payload is specified in table 66.

**Table 66 – Update Fabric Configuration request payload**

Item	Size (bytes)
26000000h	4

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the UFC request  
 Accept (SW\_ACC)  
 Signifies acceptance of the UFC request



The format of the Update Fabric Configuration accept payload is specified in table 67.

**Table 67 – Update Fabric Configuration accept payload**

Item	Size (bytes)
02000000h	4
Reserved	1
Obsolete	1
Obsolete	1
Obsolete	1

### 6.2.20 Check E\_Port Connectivity (CEC)

The Check E\_Port Connectivity (CEC) SW\_ILS requests the exchange of link parameters between two E\_Ports connected through B\_Ports. The exchange of link parameters establishes the operating environment between the two E\_Ports, and the capabilities of the Switches that are connected by the E\_Ports. The CEC SW\_ILS is transparent to B\_Ports. Use of the CEC SW\_ILS for Switch port initialization is described in 7.2.

#### Protocol:

Check E\_Port Connectivity (CEC) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Switch port initialization, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch; the D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the CEC request payload is specified in table 68.

**Table 68 – CEC request payload**

Item	Size (bytes)
2900 0000h	4
Revision	1
Flags	2
Reserved	1
R_A_TOV	4
E_D_TOV	4
Requester E_Port_Name	8
Requester Switch_Name	8
Fabric Controller Class F Service Parameters	16
Obsolete	4
Class 2 E_Port Parameters	4
Class 3 E_Port Parameters	4
Reserved	20

The descriptions of the fields in the CEC request payload are as defined in the ELP request payload (see 6.2.4). In this case the Interconnect\_Port is functioning as an E\_Port.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)
  - Signifies the rejection of the CEC request.
- Accept (SW\_ACC)
  - Signifies acceptance of the CEC request
  - Accept payload

**Payload:** The format of the CEC accept payload is specified in table 69.

**Table 69 – CEC accept payload**

Item	Size (bytes)
0200 0000h	4
Revision	1
Flags	2
Reserved	1
R_A_TOV	4
E_D_TOV	4
Responder E_Port_Name	8
Responder Switch_Name	8
Fabric Controller Class F Service Parameters	16
Obsolete	4
Class 2 E_Port Parameters	4
Class 3 E_Port Parameters	4
Reserved	20

The descriptions of the fields in the CEC accept payload are as defined in the ELP accept payload (see 6.2.4). In this case the Interconnect\_Port is functioning as an E\_Port.

### 6.2.21 Exchange Switch Capabilities

The Exchange Switch Capabilities SW\_ILS defines a mechanism for two Switches to exchange vendor and protocol information.

A Switch is not required to support the ESC SW\_ILS. If the receiving Switch does not support the ESC SW\_ILS, it shall respond with an SW\_RJT with a reason code of “Command Not Supported”.

#### Protocol:

Exchange Switch Capabilities (ESC) request Sequence

Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the ESC request payload is specified in table 70.

**Table 70 – ESC request payload**

Item	Size (bytes)
Command code = 30h	1
Flags	1
Payload Length	2
Vendor ID String	8
Protocol Descriptor #1	12
...	
Protocol Descriptor #n	12

**Flags:** This field contains flag bits (7:0) that provide additional information about the ESC request. The following flag bits are defined.

**Bit 0, Multiple Protocol Descriptors.** This bit shall indicate whether the ESC accept may contain multiple Accepted Protocol Descriptors. When the bit is set to one, the ESC accept may contain multiple Accepted Protocol Descriptors. When the bit is set to zero, the ESC accept shall contain only one Accepted Protocol Descriptor.

Bits 1-7 shall be reserved.

**Payload Length:** This field contains a 16-bit unsigned binary integer that specifies the total length in bytes of the payload. The least significant two bits shall be zero. The value specified shall be greater than or equal to 24, and less than or equal to 65532.

**Vendor ID String:** This field shall contain a T10 Vendor ID of either the manufacturer of the requesting Switch, or an OEM of the requesting Switch.

**Protocol Descriptor:** This field allows the requesting Switch to identify which Switch-to-Switch protocols it supports. There may be more than one Protocol Descriptors specified in the ESC SW\_ILS frame. This list of Protocol Descriptors allows a single port on the requesting Switch to specify that it supports more than one Switch-to-Switch protocol. The format of the Protocol Descriptor is specified in table 71.

**Table 71 – Protocol Descriptor format**

Item	Size (bytes)
Vendor ID String	8
Reserved	2
Protocol ID	2

**Vendor ID String:** For non-vendor specific protocols, this field shall be zero filled. For vendor specific protocols, this field shall contain the T10 Vendor ID associated with the Protocol ID field. It is the intention that this field contain the T10 Vendor ID of the original Switch manufacturer that designed the protocol being described.

**Protocol ID:** This field shall contain a value identifying the protocol. If the value of this field is in the range 8000h to FFFFh, then this field combined with the Vendor ID String field specifies a vendor specific protocol. If the value of this field is in the range 0000h - 7FFFh then a non-vendor specific protocol is specified. Values for this field are summarized in table 72.

**Table 72 – Protocol ID values**

Value	Use
0000h	Reserved
0001h	Obsolete
0002h	FSPF Protocol
0003h	Virtual Fabrics Supported
0004h - 7FFFh	Reserved
8000h - FFFFh	Vendor Specific (see note)
Note: Vendor Specific values are only meaningful when combined with the Vendor ID String field.	

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)  
Signifies the rejection of the ESC request
- Accept (SW\_ACC)  
Signifies acceptance of the ESC request
- Accept payload

**Payload:** The format of the ESC accept payload is specified in table 73.

**Table 73 – ESC accept payload (Part 1 of 2)**

Item	Size (bytes)
Command code = 02h	1
Reserved	1
Payload Length	2

**Table 73 – ESC accept payload (Part 2 of 2)**

Item	Size (bytes)
Vendor ID String	8
Accepted Protocol Descriptor #1	12
...	
Accepted Protocol Descriptor #n	12

**Payload Length:** If the ESC accept contains only one Accepted Protocol Descriptor, this field shall be set to zero. If the Multiple Protocol Descriptors bit is set to one in the ESC request and the ESC Accept contains two or more Accepted Protocol Descriptors, then this field specifies the total length in bytes of the payload. The least significant two bits shall be zero. The value specified shall be greater than or equal to 36, and less than or equal to 65532. If the Multiple Protocol Descriptors bit is set to zero in the ESC request, the ESC Accept shall contain only one Accepted Protocol Descriptor.

**Vendor ID String:** This field shall contain a T10 Vendor ID of the responding Switch. This field shall contain either an identifier for the Switch manufacturer, or an OEM identifier.

**Accepted Protocol Descriptor:** This field shall contain a Protocol Descriptor chosen by the responding Switch. This Protocol Descriptor shall be chosen from the list presented in the ESC request. The format of this field is as specified in table 71.

### 6.2.22 Exchange Switch Support (ESS)

The Exchange Switch Support (ESS) SW\_ILS defines a mechanism for two Switches to exchange vendor and support information relative to various supported features within the Fabric services and Switch link services payloads.

Exchange Switch Support requests are addressed from the Domain Controller of a requesting Switch to the Domain Controller of a responding Switch.

Path selection shall complete before an ESS request may be issued to a destination Switch.

**Protocol:**

Exchange Switch Support (ESS) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** For use in determining Switch support of Fabric services and SW\_ILS support, the S\_ID shall be set to FFFCxxh, indicating the Domain Controller of the originating Switch. The D\_ID shall be set to FFFCyyh indicating the Domain Controller of the destination Switch.

### 6.2.22.1 ESS request payload

The format of the ESS request payload is specified in table 74.

**Table 74 – ESS request payload**

Item	Size (bytes)
31000000h	4
Revision	4
Payload Length	4
Interconnect Element Information Object	256
Number of Capability Objects	2
Reserved	2
Capability Object #1	Length of Object
Capability Object #2	Length of Object
...	
Capability Object #n	Length of Object

**Revision:** The revision field shall contain a value of 01h.

**Payload Length:** The length shall specify the number of bytes in the ESS request payload. This value does not include the request code, revision and payload length bytes. This value shall be a multiple of 4.

**Number of Capability Objects:** Number of Capability Objects contained in the payload.

### 6.2.22.2 Interconnect Element Information Object

The Interconnect Element Information object contains vendor name, model name and number, release code and vendor specific information related to the Switch. The Interconnect Element object shall be supported.

The format of the Interconnect Element Information is described in FC-GS-7.

### 6.2.22.3 Capability Object

The Capability Object is used to convey levels of support for FC-SW-7 and FC-GS-8 functionality. The format of the Capability Object is specified in table 75.

**Table 75 – Capability Object format**

Item	Size (bytes)
Well-Known Address Type	1
Well-Known Address Subtype	1
reserved	1
Number of Capability Entries	1
Capability Entry #1	8
Capability Entry #2	8
...	
Capability Entry #n	8

**Well-Known Address Type:** The Well-Known Address (WKA) type represents the type of service that the Capability Object represents. Allowed values are specified in FC-GS-8.

**Well-Known Address Subtype:** This field specifies a sub-service type for the specific service.

**Number of Capability Entries:** Number of Capability Entries within this Capability Object.

**Capability Entry:** Each Capability Entry shall be eight bytes in length and contain information specific to a particular service.

### 6.2.22.4 Service Specific Capability formats

#### 6.2.22.4.1 Directory Server Capability

The Well-Known Address Type shall be set to FCh and the Well-Known Address Subtype shall be set to 02h

Table 76 defines the bit definition for identifying specific support for the Name Server subtype of the Directory Server.

**Table 76 – Name Server Capability Flags**

Item	Size (bytes)
Name Server Support Flags	4
NS Vendor Specific Support Flags	4

The format of the Name Server Support Flags field is as follows:



Bit 0- Name Server Entry Object 00h Support - When set indicates that the Name Server instance may accept large Name Server objects.

Bit 1- Name Server Entry Object 01h Support - When set indicates that the Name Server instance may accept small Name Server objects.

Bit 2- Name Server Entry Object 02h Support - When set indicates that the Name Server instance may accept Large + FC-4 Features Name Server objects.

Bit 3- Name Server Entry Object 03h Support - When set indicates that the Name Server instance may accept Small + FC-4 Features Name Server objects (see 9.3.3).

Bit 4 - GE\_PT Zero Length Accept - When set indicates that the Name Server may support receipt of a 0 length accept payload from an interswitch GE\_PT (or other GE\_\*) query.

Bits 5-31 reserved.

The format of the Name Server Vendor Specific Support Flags field is vendor specific and dependent on the Vendor Name.

**6.2.22.4.2 Fabric Controller Capability**

The Well-Known Address Type shall be set to FDh and the Well-Known Address Subtype shall be set to 00h.

Table 77 defines the bit definition for identifying specific support for the Fabric Controller.

**Table 77 – Fabric Controller Capability f lags**

Item	Size (bytes)
Fabric Controller Support Flags	4
Fabric Controller Vendor Specific Support Flags	4

The format of the Fabric Controller Support Flags field is as follows:

Bit 0- SW\_RSCN Support - When set indicates that the transmitting Fabric Controller supports receiving the SW\_RSCN request.

Bits 1-31 Reserved.

The format of the Fabric Controller Vendor Specific Support Flags field is vendor specific and dependent on the Vendor Name.

**6.2.22.4.3 ESS Fabric Configuration Server Capability Object**

The WKA Type shall be set to FAh and the WKA Subtype shall be set to 01h.

Table 78 specifies the bit definitions for identifying specific support for the Fabric Configuration Server subtype of the Management Server.

**Table 78 – Fabric Configuration Server Capability flags**

Item	Size (bytes)
Fabric Configuration Server support flags	4
Reserved	4

The format of the Fabric Configuration Server support flags field is as follows:

Bit 0- Basic Configuration Services - When this bit is one, the Switch supports commands that are members of the Basic Configuration Service class (see table 222). When this bit is zero, the Switch does not support commands that are members of the Basic Configuration Service class.

Bit 1- Platform Configuration Services - When this bit is one, the Switch supports commands that are members of the Platform Configuration Service class (see table 222). When this bit is zero, the Switch does not support commands that are members of the Platform Configuration Service class.

Bit 2- Topology Discovery Configuration Services - When this bit is one, the Switch supports commands that are members of the Topology Discovery Configuration Service class (see table 222). When this bit is zero, the Switch does not support commands that are members of the Topology Discovery Configuration Service class.

Bit 3- Enhanced Configuration Services - When this bit is one, the Switch supports commands that are members of the Enhanced Configuration Service class (see table 222). When this bit is zero, the Switch does not support commands that are members of the Enhanced Configuration Service class.

Bits 4-31 reserved.

#### 6.2.22.4.4 ESS Enhanced Zone Server Capability Object

The WKA Type shall be set to FAh and the WKA Subtype shall be set to 03h.

Table 79 specifies the bit definitions for identifying specific support for the Zone Server subtype of the Management Server.

**Table 79 – Enhanced Zone Server Capability flags**

Item	Size (bytes)
Switch Enhanced Zoning support flags	4
Reserved	4

The format of the Switch Enhanced Zoning Server Support Flags field is as follows:

Bit 0- Enhanced Zoning supported - When this bit is one, the Switch is able to work in Enhanced Zoning mode. When this bit is zero, the Switch is not able to work in Enhanced Zoning mode.

Bit 1- Enhanced Zoning enabled - When this bit is one, the Switch is working in Enhanced Zoning mode. When this bit is zero, the Switch is working in Basic Zoning mode.

Bit 2- Merge Control Setting - When this bit is one, this Switch is working in Restrict mode, so it may join a Fabric only if the Fabric's Zoning Database is equal to its Zoning Database. When this bit is zero, this Switch is working in Allow mode, so it may join a Fabric only if the Fabric's Zoning Database is mergeable with its Zoning Database.

Bit 3- Default Zone Setting - When this bit is one this Switch denies traffic between members of the Default Zone. When this bit is zero this Switch permit traffic between members of the Default Zone.

Bit 4- Zone Set Database supported - When this bit is one, the Zone Server on this Switch is able to maintain a Zone Set Database. When this bit is zero, the Zone Server on this Switch is not able to maintain a Zone Set Database.

Bit 5- Zone Set Database enabled - When this bit is one, the Zone Server on this Switch is maintaining a Zone Set Database. When this bit is zero, the Zone Server on this Switch is not maintaining a Zone Set Database.

Bit 6- Activate Direct command supported - When this bit is one, this Switch supports the Activate Direct command. When this bit is zero, this Switch does not support the Activate Direct command.

Bit 7- Hard Zoning supported - When this bit is one, this Switch supports Hard Zoning. When this bit is zero, this Switch does not support Hard Zoning.

Bit 8- FC-SP Zoning supported (see FC-SP-2).

Bit 9- FC-SP Zoning enabled (see FC-SP-2).

Bits 10-31 reserved.

#### **6.2.22.4.5 Security Policy Server Capability Object**

The Security Policy Server Capability Object allows a Switch to discover the level of support for individual Policy Object types provided by other Switches of a Fabric. See FC-SP-2 for the definition of the Security Policy Server Capability Object.

#### 6.2.22.4.6 ESS Vendor Specific Capability Object

The general format of the Vendor Specific Capability Object closely follows the format of the Capability Object currently defined for ESS. The format of the Vendor Specific Capability Object is specified in table 80.

**Table 80 – Vendor Specific Capability Object**

Item	Size (bytes)
Well-Known Address Type	1
Well-Known Address Subtype	1
Reserved	1
Length (n)	1
T10 Vendor ID	8
Vendor Specific Information	n*8

**Well-Known Address Type:** The Well-Known Address Type field shall be set to E0h to indicate that the Capability Object is a Vendor Specific Type Capability Object.

**Well-Known Address Subtype:** The Well-Known Address Subtype field shall contain a value specified by the vendor.

**Length:** The Length field shall contain a value between 01h and FFh to indicate the number of doublewords of vendor specific information contained in the Capability Object. The T10 Vendor ID field is included in the doubleword count.

**T10 Vendor ID:** The T10 Vendor ID field shall contain the vendor's eight byte T10 Vendor ID.

**Vendor Specific Information:** The Vendor Specific Information field contains the vendor's information. The format of the information is defined by the vendor and not by this standard. When the vendor specific information does not align on a doubleword boundary the information is padded with nulls (00h) to the right to complete the final doubleword.

#### 6.2.22.4.7 Domain Controller Capability Object

The Well-Known Address Type shall be set to ECh and the Well-Known Address Subtype shall be set to 00h. The Domain Controller Capability Object is specified in table 81.

**Table 81 – Domain Controller Capability Object**

Item	Size (bytes)
Receive Data Field Size	2
End-to-End Credit	2
Concurrent Sequences	2
Open Sequences per Exchange	2

**Receive Data Field Size:** This field shall specify the largest Data Field size in bytes for a frame that may be received by the Domain Controller supplying the Parameters as a Sequence Recipient for a Class F frame. Values less than 256 or greater than 2112 are invalid. Values shall be a multiple of four bytes.

**End-to-End Credit:** End-to-end credit is the maximum number of Class F Data frames that may be transmitted by a Domain Controller without receipt of accompanying ACK or Link\_Response frames. The minimum value of end-to-end credit is one. The End-to-End Credit field specified is associated with the number of buffers available for holding the Data\_Field of a Class F frame and processing the contents of that Data\_Field by the Domain Controller supplying the Parameters. Bit 15 of this field shall be set to zero. A value of zero for this field is reserved.

**Concurrent Sequences:** This field shall specify the number of Sequence Status Blocks provided by the Domain Controller supplying the Parameters for tracking the progress of a Sequence as a Sequence Recipient. The maximum number of Concurrent Sequences that may be specified is 255. A value of zero in this field is reserved. In Class F, the value of SEQ\_ID shall range from 0 to 255, independent of the value in this field. A Domain Controller is allowed to respond with P\_BSY to a frame initiating a new Sequence if Domain Controller resources are not available.

**Open Sequences per Exchange:** The value of the Open Sequences per Exchange shall specify the maximum number of Sequences that may be Open at one time at the Recipient between a pair of Domain Controllers for one Exchange. This value plus two shall specify the number of instances of Sequence Status that shall be maintained by the Recipient for a single Exchange in the Exchange Status Block. This value is used for Exchange and Sequence tracking. The value in this field limits the link facility resources required for error detection and recovery.

#### 6.2.22.4.8 Event Server Capability Object

The Well-Known Address Type shall be set to F4h and the Well-Known Address Subtype shall be set to 01h. The bit definitions for identifying specific support for the Event Server are defined in table 82.

**Table 82 – Event Server Capability Object**

Item	Size (bytes)
Event Server Flags	4
Event Server Vendor Specific Support Flags	4

The format of the Event Server Flags field is as follows:

Bit 0- When set indicates that the Event Server instance is supported.

Bits 1-31 - reserved.

The format of the Event Server Vendor Specific Flags field is vendor specific and dependent on the Vendor Name.

#### 6.2.22.4.9 Switch Support Capability Object

The Well-Known Address Type shall be set to 20h and the Well-Known Address Subtype shall be set to 01h. The Switch Support Capability Object is specified in table 83.

**Table 83 – Switch Support Capability Object**

Item	Size (bytes)
Switch Support Flags	4
Reserved	4

The format of the Switch Support Flags fields is as follows:

Bit 0- Encapsulated F\_RJT/F\_BSY - When this bit is set to one, the Switch supports the encapsulated Class 2/F F\_RJT and F\_BSY frame format (see 15).

Bits 1-31 - reserved.

#### 6.2.22.4.10 Application Server Capability Object

The WKA Type shall be set to FAh (i.e., Management Service, see FC-GS-8) and the WKA Subtype shall be set to 20h (i.e., Application Server, see FC-GS-8).

Table 84 specifies the bit definitions for identifying specific support for the Application Server subtype of the Management Server.

**Table 84 – Application Server Capability Object**

Item	Size (bytes)
Application Server Support Flags	4
Reserved	4

The format of the Application Server Support Flags field is as follows:

Bit 0- Application Server supported - If this bit is set to one, the Switch supports the Application Server. If this bit is set to zero, the Switch does not support the Application Server.

Bits 1-31 - reserved.

#### 6.2.22.4.11 Enhanced Fabric Configuration Server Capability Object

The WKA Type shall be set to FAh (i.e., Management Service, see FC-GS-8) and the WKA Subtype shall be set to 08h (i.e., Enhanced Fabric Configuration Server, see FC-GS-8).

Table 85 specifies the bit definitions for identifying specific support for the Enhanced Fabric Configuration Server subtype of the Management Server.

**Table 85 – Enhanced Fabric Configuration Server Capability Object**

Item	Size (bytes)
Enhanced Fabric Configuration Server Support Flags	4
Reserved	4

The format of the Enhanced Fabric Configuration Server Support Flags field is as follows:

Bit 0- Basic Configuration Services - If this bit is set to one, the Switch supports commands that are members of the Basic Configuration Service class (see table 241). If this bit is set to zero, the Switch does not support commands that are members of the Basic Configuration Service class.

Bit 1- Platform Configuration Services - If this bit is set to one, the Switch supports commands that are members of the Platform Configuration Service class (see table 241). If this bit is set to zero, the Switch does not support commands that are members of the Platform Configuration Service class.

Bit 2- Topology Discovery Configuration Services - If this bit is set to one, the Switch supports commands that are members of the Topology Discovery Configuration Service class (see table 241). If this bit is set to zero, the Switch does not support commands that are members of the Topology Discovery Configuration Service class.

Bit 3- Enhanced Configuration Services - If this bit is set to one, the Switch supports commands that are members of the Enhanced Configuration Service class (see table 241). If this bit is set to zero,

the Switch does not support commands that are members of the Enhanced Configuration Service class.

Bits 4-31 reserved.

#### 6.2.22.4.12 VE Identification Server Capability Object

The Well-Known Address Type shall be set to FCh and the Well-Known Address Subtype shall be set to 04h.

Table 86 shows the bit definition for identifying specific support for the VE Identification Server subtype of the Directory Service.

**Table 86 – VE Identification Server Capability Flags**

Item	Size (Bytes)
VE Identification Server Support Flags	4
Reserved	4

The format of the VE Identification Server Support Flags field is as follows:

Bit 0- If set to one indicates that the VE Identification Server is supported. If set to zero indicates that the VE Identification Server is not supported.

Bits 1-31 - reserved.

#### 6.2.22.5 ESS accept payload

The format of the accept payload is specified in table 87.

**Table 87 – ESS accept payload**

Item	Size (bytes)
02000000h	4
Revision	4
Payload Length	4
Interconnect Element Information Object	256
Number of Capability Objects	2
Reserved	2
Capability Object #1	Length of Object
Capability Object #2	Length of Object
...	
Capability Object #n	Length of Object



**6.2.23 Merge Request Resource Allocation (MRRA)**

The Merge Request Resource Allocation (MRRA) SW\_ILS defines a mechanism for Switches to request resources to be allocated for the transfer of a Merge Request SW\_ILS. MRRA enables buffer management in the Fabric Controller.

MRRA SW\_ILSs are addressed from the Fabric Controller of a requesting Switch to the Fabric Controller of a responding Switch.

**Protocol:**

Merge Request Resource Allocation (MRRA) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** For use in determining resource availability, the S\_ID shall be set to FFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID shall be set to FFFFDh, indicating the Fabric Controller of the destination Switch.

The format of the MRRA request payload is specified in table 88.

**Table 88 – MRRA request payload**

Item	Size (bytes)
34000000h	4
Revision	4
Merge Request Size	4
Vendor Specific	16

**Revision:** Shall be set to 00000001h.

**Merge Request Size:** The Merge Request Size is the number of words in the entire Merge Request SW\_ILS that is subsequently sent to the Adjacent Switch.

**Vendor Specific:** The format of the Vendor Specific field is specified in table 89.

**Table 89 – Vendor Specific field**

Item	Size (bytes)
Vendor ID	8
Vendor Specific Information	8

**Vendor ID:** This field contains the T10 Vendor ID of the vendor that defines the content of the Vendor Specific Information field.

**Vendor Specific Information:** The format of this field is specific to the vendor.

**Reply Merge Request Resource Allocation Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the MRRA request

**Accept (SW\_ACC)**

Signifies the acceptance of the MRRA request

- Accept payload

**Payload:** The format of the MRRA accept payload is specified in table 90.

**Table 90 – MRRA response payload**

Item	Size (bytes)
02000000h	4
Vendor ID	8
MRRA Response	4
Maximum Resources Available	4
Waiting Period	4

**Vendor ID:** This field contains the T10 Vendor ID of the vendor that defines the content of the vendor specific fields in the Merge Response field.

**MRRA Response:** This field shall specify the response to the MRRA request. The values are defined in table 91.

**Table 91 – MRRA Response values**

Value (Hex)	Description
0	Reserved
1	Requested resources available
2	Requested resources are not available
E0-FF	Vendor Specific
Others	Reserved

**Maximum Resources Available:** This field specifies the maximum size in words in the entire Merge Request SW\_ILS that the Fabric Controller is able to accept.

**Waiting Period:** This field specifies the time in seconds that the requesting port should wait before retrying the MRRA request. The Waiting Period should be less than R\_A\_TOV. This value is only meaningful when the Merge Response value is set to 2.

**6.2.24 Switch Trace Route (STR)**

**6.2.24.1 Basic function**

The STR request initiates the STR operation to find the route between two Nx\_Ports (e.g., source port and destination port) in a common Zone. After receiving a FC Trace Route (FTR) request (see

FC-GS-8), the Managing Switch shall send a STR request to the Domain Controller of the Switch to which the source port is attached.

The source port's Switch shall append its Path Information to the request, increment the number or Path Information Entries, and forward the STR request to the next Switch in the Fabric ingress path (i.e., from source port to destination port). Each Switch in the Fabric ingress path shall append its Path Information to the STR request, increment the Number of Path Information Entries, and forward the STR request to the next Switch in the path.

The destination port's Switch shall append the Path Information for the Fabric ingress path and the Fabric egress path (i.e., from the destination port to the source port) and change the Fabric Egress Path bit in the Flags field. The Fabric Egress Path bit shall inform Switches as to the direction that the STR request is traveling. When the source port's Switch receives the STR request for the second time, it shall append its Path Information to source port, increment the Number of Path Information Entries, and forward the frame to the Fabric Configuration Server of the Entry Switch.

If the two N\_Ports are not in a common Zone, the STR request shall be rejected with the SW\_RJT reason code of "Unable to Perform Command Request" and a SW\_RJT reason code explanation of "Invalid Operation".

If a Switch in the path rejects or does not accept the STR request, the last Switch to append its information shall set the STR Reject Reason Code in the payload and send the STR request to the Domain Controller ID of the Managing Switch. The Switch shall make a best effort to return the STR request to the Managing Switch.

**Protocol:**

Switch Trace Route (STR) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to FFFCxxh which designates the Domain Controller ID of the Switch that generates the STR. The D\_ID shall be set to FFFCyyh to designate the Domain Controller ID of the recipient Switch. The Switch shall forward the STR request on the egress port that was entered in the Path Information Entry to ensure that the data path is operational.

**Payload:** The format of the STR request payload is specified in table 92.

**Table 92 – STR request payload (Part 1 of 2)**

Item	Size (bytes)
35010000h	4
Revision	4
Source Port Tag	2
Source Port Length	2
Source Port Value	n
Destination Port Tag	2
Destination Port Length	2
Destination Port Value	n
Token	4

**Table 92 – STR request payload (Part 2 of 2)**

Item	Size (bytes)
T10 Vendor ID	8
Vendor Specific Information	8
Flags	4
Remaining Hop Count	4
STR Reject Reason Code	4
Managing Switch’s Domain Controller Address	1
Requesting Port’s N_Port_ID	3
Number or Path Information Entries	1
Reserved	3
Source Port’s Fabric Ingress Path Information	36
Intermediate Switch’s Path Information	36
...	36
Destination Port’s Fabric Ingress Path Information	36
Destination Port’s Fabric Egress Path Information	36
...	36
Source Port’s Fabric Egress Path Information	36

**Revision:** The revision shall contain a value of 01h.

**Source Port Tag:** The tag used to identify the source port as specified in table 93.

**Table 93 – Nx\_Port Tags**

Tag (hex)	Item
01	N_Port_ID
02	Nx_Port Name_Identifier

**Source Port Length:** The length of the Source Port Value in bytes. The length shall be a multiple of four.

**Source Port Value:** The value of source port. Fill bytes are added as necessary to the end of the actual value in order to ensure that the length of the value field is a multiple of four. Fill bytes shall be nulls (00h). The number of fill bytes (f) is zero, one, two or three depending on the length of the actual value (m). The total length of the value field is (n= f + m).

**Destination Port Tag:** The tag used to identify the destination port as specified in table 93.

**Destination Port Length:** The length of the Destination Port Value in bytes. The length shall be a multiple of four.

**Destination Port Value:** The value of destination port. Fill bytes are added as necessary to the end of the actual value in order to ensure that the length of the value field is a multiple of four. Fill bytes

shall be nulls (00h). The number of fill bytes (f) is zero, one, two or three depending on the length of the actual value (m). The total length of the value field is (n= f + m).

**Token:** An identifier for the FTR request that is provided by the requesting Nx\_Port.

**T10 Vendor ID:** Contains the T10 Vendor ID of the vendor that defines the content of the Vendor Specific Information field.

**Vendor Specific Information:** The Vendor Specific Information field shall contain the vendor's information. The format of the information is defined by the vendor and not by this standard.

**Flags:** The Flags field has 32 bits that are specified in table 94.

**Table 94 – Flags field values**

Bit	Description
0	Fabric Egress Path - Value b'0' - The STR request is in the Fabric ingress path. Value b'1' - The STR request is in the Fabric egress path.
1	Trace Complete bit Value b'0' - The STR request is still tracing the route. Value b'1' - The STR request has completed tracing the route.
28-32	Vendor Specific Information
Others	Reserved

**Remaining Hop Count:** The Remaining Hop Count is the Maximum Hop Count specified by the user minus the number of hops that the STR request has traveled between the source and destination port and back to the source port. Starting with the source port's Switch, each Switch in the path decrements the Remaining Hop Count by 1. If the Remaining Hop Count is decremented to 0, then the Switch shall fill in its path routing information and forward the STR request to the Entry Switch's Domain Controller and include the STR Reject Reason Code of "Reached Maximum Hop Count".

**STR Reject Reason Code:** If the STR request is completed successfully, the STR Reject Reason Code is set to 00h to signify a successful completion. If the STR request is not acceptable by the next Switch in the path, then the Switch with the STR payload shall send the STR request to the Entry Switch's Domain Controller. The STR payload shall contain one of the following STR Reject Reason Code values.

**Table 95 – STR Reject Reason Code values (Part 1 of 2)**

Value (hex)	Item
00	Command Completed Successfully
01	Command Not Supported in Next Switch
02	No Response from Next Switch
03	Maximum Hop Count Reached
04	Source Port not in Fabric
05	Destination Port not in Fabric
06	Devices not in Common Zone

**Table 95 – STR Reject Reason Code values (Part 2 of 2)**

Value (hex)	Item
07	No Route Between Designated Ports
08	No Additional Explanation
09	Fabric Busy
0A	Fabric Build in Progress
F0-FF	Vendor Specific Error Codes
Others	Reserved

**Managing Switch’s Domain Controller Address:** The Domain\_ID for the Domain Controller of the Managing Switch that received the FTR request. This address may be the N\_Port\_ID for an Nx\_Port that is the Fabric Configuration Server of the Fabric.

**Requesting Port’s N\_Port\_ID:** The N\_Port\_ID of the requesting port is needed so that the Token will not be mistaken for another N\_Port\_ID’s duplicate token.

**Number of Path Information Entries:** The number of Path Information Entries in the request. As each Switch appends its path routing information onto the STR payload, it shall increment the number of path information entries by one. The destination port’s Switch shall increment the count by two since it shall append two entries.

**Source Port’s Fabric Ingress Path Information:** The Path Information for source port’s Switch in the Fabric ingress path. The format for the Path Information is specified in table 96.

**Table 96 – Path Information**

Item	Size (bytes)
Switch Name	8
Domain_ID	4
Ingress Port_Name	8
Ingress Physical Port Number	4
Egress Port_Name	8
Egress Physical Port Number	4

**Switch Name:** The Switch Name of the Switch in the path that is appending the path information.

**Domain\_ID:** The Domain\_ID of the Switch reporting the Path Information. The format of Domain\_ID shall be set to 000000h||Domain\_ID’.

**Ingress Port\_Name:** The Port\_Name of the F\_Port or E\_Port on the Switch that the frame enters.

**Ingress Physical\_Port\_Number:** The Physical\_Port\_Number of the F\_Port or E\_Port on the Switch that the frame enters.

**Egress Port\_Name:** The Port\_Name of the F\_Port or E\_Port on the Switch that the frame exits.

**Egress Physical\_Port\_Number:** The Physical\_Port\_Number of the F\_Port or E\_Port on the Switch that the frame exits.

**Intermediate Switch's Path Information:** The Path Information for the second Switch in the Fabric ingress path, if any.

**Destination Port's Fabric Ingress Route Information:** The Path Information for destination port's Switch in the Fabric ingress path.

**Destination Port's Fabric Egress Route Information:** The Path Information for destination port's Switch in the Fabric egress path.

**Source Port's Fabric Egress Route Information:** The Path Information for source port's Switch in the Fabric egress path.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)  
Signifies rejection of the STR request
- Accept (SW\_ACC)  
Signifies acceptance of the STR request
- Accept payload

**Payload:** The format of the STR accept payload is specified in table 97.

**Table 97 – STR accept payload**

Item	Size (bytes)
02000000h	4

**6.2.25 Exchange Virtual Fabrics Parameters (EVFP)**

**6.2.25.1 Basic function**

The Exchange Virtual Fabrics Parameters (EVFP) SW\_ILS provides support for Virtual Fabrics (see 12).

**Protocol:** Exchange Virtual Fabrics Parameters (EVFP) request Sequence

**Addressing:** The S\_ID field shall be set to FFFFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** Two types of EVFP messages are defined. All EVFP request messages share the same message structure, specified in table 98.

**Table 98 – EVFP request payload**

Item	Size (bytes)
3600 0000h	4
Protocol Version	1
EVFP Message Code	1
Transaction Identifier	2
Core Switch_Name	8
Reserved	2
Message Payload Length	2
Message Payload	variable

**Protocol Version:** This field denotes the version of the EVFP protocol. A value of 01h shall be used to indicate the version specified in this standard. If a Fabric Controller receives an EVFP request containing a Version field value that is higher than its supported value, the Fabric Controller shall respond with its highest supported Revision field value. If a Fabric Controller receives an EVFP request containing a Version field value that is equal to or lower than its supported value, the Fabric Controller shall respond with the Version field value received in the EVFP request.

**EVFP Message Code:** Specifies the EVFP message that is to be transmitted from the source to the destination. The defined EVFP message codes are specified in table 99.

**Table 99 – EVFP Message Codes**

Value	Description	Reference
01h	EVFP_SYNC	6.2.25.2
02h	EVFP_COMMIT	6.2.25.3
all others	Reserved	

**Transaction Identifier:** Uniquely identifies an EVFP transaction between two entities. The Transaction Identifier shall be set by the EVFP initiator, and each subsequent EVFP message shall contain the same value, until the EVFP transaction is completed.

NOTE 6 – The usage of the Transaction Identifier is very similar to the usage of an OX\_ID when an Exchange Originator is enforcing uniqueness via the OX\_ID mechanism (see FC-FS-5), but it is not related in any way to the OX\_ID present in the Fibre Channel frames carrying the EVFP messages.

**Core Switch\_Name:** Core Switch\_Name of the originating Switch.



**Payload Length:** Shall be set to the total length in bytes of the EVFP payload (i.e., 20 + the Message Payload length).

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)

Signifies rejection of the EVFP request

Accept (SW\_ACC)

Signifies acceptance of the EVFP request

-Accept payload

**Accept payload:** All EVFP accept messages share the same message structure, specified in table 100. With the exception of the first four bytes, the fields in table 100 are the same as the fields in table 98.

**Table 100 – EVFP accept payload**

Item	Size (bytes)
0200 0000h	4
Protocol Version	1
EVFP Message Code	1
Transaction_ID	2
Core Switch_Name	8
Reserved	2
Message Payload Length	2
Message Payload	variable

Table 101 shows the use of reason codes and reason code explanations under some error conditions.

**Table 101 – SW\_RJT reason codes**

Error condition	Reason code	Reason code explanation
EVFP SW_ILS not supported	Command Not Supported	No Additional Explanation
EVFP collision	Unable to perform command request	Command Already in Progress
EVFP Protocol Version not supported	Invalid Revision Level	No Additional Explanation
EVFP_COMMIT before EVFP_SYNC	Logical Error	No Additional Explanation
Insufficient Resources	Unable to Perform Command Request	Insufficient Resources Available
Invalid Payload Message	Protocol Error	No Additional Explanation

### 6.2.25.2 EVFP\_SYNC Message Payload

#### 6.2.25.2.1 Overview

The EVFP\_SYNC Message Payload carries a list of descriptors. The format of the EVFP\_SYNC Message Payload is specified in table 102. This Message Payload is used in both EVFP\_SYNC request and EVFP\_SYNC accept.

**Table 102 – EVFP\_SYNC Message Payload format**

Item	Size (bytes)	Reference
Descriptor #1 = Tagging Administrative Status	8	6.2.25.2.2
Descriptor #2 = Port VF_ID	8	6.2.25.2.3
Descriptor #3 = Locally-Enabled VF_ID List	516	6.2.25.2.4
...		
Descriptor #m	variable	

All descriptors share the same format, as specified in table 103.

**Table 103 – Descriptor format**

Item	Size (bytes)
Descriptor Control	1
Descriptor Type	1
Descriptor Length	2
Descriptor Value	variable

**Descriptor Control:** Specifies the behavior of the receiving entity if the descriptor is unsupported. The defined codes are specified in table 104.

**Table 104 – Descriptor Control codes**

Value	Description
01h	Critical. Abort the EVFP transaction if the descriptor is unsupported. <sup>a</sup>
02h	Non critical. Skip the descriptor if unsupported and continue the EVFP transaction. <sup>a</sup>
all others	Reserved
<sup>a</sup> The Descriptor Control provides extensibility to the protocol. An implementation supporting a subset of the descriptors is able to process the unknown ones as specified by the Descriptor Control value.	

**Descriptor Type:** Specifies the type of the descriptor. The defined descriptors are summarized in table 105.

**Table 105 – Descriptor Types**

Value	Description	Reference
01h	Tagging Administrative Status Descriptor	6.2.25.2.2
02h	Port VF_ID Descriptor	6.2.25.2.3
03h	Allowed VF_ID List Descriptor	6.2.25.2.4
F0h .. FEh	Vendor Specific Descriptor	6.2.25.2.5
all others	Reserved	

**Descriptor Length:** Specifies the length in bytes of the Descriptor Value.

### 6.2.25.2.2 Tagging Administrative Status descriptor

The format of the Tagging Administrative Status descriptor is specified in table 106.

**Table 106 – Tagging Administrative Status descriptor**

Item	Size (bytes)
Descriptor Control = 01h	1
Descriptor Type = 01h	1
Descriptor Length = 0004h	2
Administrative Tagging Mode	4

The defined Administrative Tagging Modes are specified in table 107.

**Table 107 – Administrative Tagging Modes**

Value	Notation	Description
0000 0001h	OFF	The Interconnect_Port shall not perform VFT Tagging
0000 0002h	ON	The Interconnect_Port may perform VFT Tagging if the peer does not prohibit it
0000 0003h	AUTO	The Interconnect_Port may perform VFT Tagging if the peer request it

In absence of any explicit configuration, the default Administrative Tagging Mode of a Switch port of a VF capable Switch should be AUTO.

Table 108 specifies how VFT tagging is negotiated between peer Interconnect\_Ports.

**Table 108 – Tagging Mode negotiation**

		Peer Tagging Mode		
		OFF	ON	AUTO
Local Tagging Mode	OFF	Non Tagging	Non Tagging	Non Tagging
	ON	Non Tagging	Tagging	Tagging
	AUTO	Non Tagging	Tagging	Non Tagging

### 6.2.25.2.3 Port VF\_ID descriptor

The format of the Port VF\_ID descriptor is specified in table 109.

**Table 109 – Port VF\_ID descriptor**

Item	Size (bytes)
Descriptor Control = 01h	1
Descriptor Type = 02h	1
Descriptor Length = 0004h	2
Port Flags	2
Port VF_ID	2

**Port Flags:** Reserved. Shall be set to zero.

**Port VF\_ID:** The 12 least significant bit of this field shall be set to the port VF\_ID. The four most significant bits shall be set to zero. In absence of any explicit configuration, the value 001h should be used as the port VF\_ID.

### 6.2.25.2.4 Locally-Enabled VF\_ID List descriptor

The format of the Locally-Enabled VF\_ID List descriptor is specified in table 110.

**Table 110 – Locally-Enabled VF\_ID List descriptor**

Item	Size (bytes)
Descriptor Control = 01h	1
Descriptor Type = 03h	1
Descriptor Length = 0200h	2
VF_ID Bitmap	512

**VF\_ID Bitmap:** Each Virtual Fabric is identified by a bit in the VF\_ID Bitmap. The high-order bit represents VF\_ID zero, each successive bit represents the successive VF\_ID, and the low-order bit represents VF\_ID 4095. Virtual Fabric K is allowed on the Interconnect\_Port if the Kth bit of the VF\_ID Bitmap is set to one and is disallowed if the Kth bit of the VF\_ID Bitmap is set to zero. The bit representing the Control VF\_ID (see FC-FS-5) shall be set to zero.

The list of Virtual Fabrics operational over a link is computed by performing a bit-wise 'AND' between the received VF\_ID Bitmap and the locally configured VF\_ID Bitmap.

### 6.2.25.2.5 Vendor Specific descriptor

The format of the Vendor Specific descriptor is specified in table 111.

**Table 111 – Vendor Specific descriptor**

Item	Size (bytes)
Descriptor Control	1
Descriptor Type	1
Descriptor Length	2
T10 Vendor ID	8
Vendor Specific	n

**T10 Vendor ID:** Shall be set to the Vendor's T10 Vendor ID.

### 6.2.25.3 EVFP\_COMMIT Message Payload

Both EVFP\_COMMIT request and EVFP\_COMMIT accept have no Message Payload.

### 6.2.26 Enhanced Acquire Change Authorization request (EACA)

The EACA request is sent from the Domain Controller of the Managing Switch to the Domain Controller of each Managed Switch contained in the ECS Switch List (see 6.2.26.1.4). The EACA request instructs the Managed Switch to reserve local resources associated with the designated application for the purpose of ensuring the consistency of the application's data.

The use if this ILS is detailed in 13.

**Protocol:**

- Enhanced Acquire Change Authorization (EACA) request Sequence
- Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to hex'FFFCxx', indicating the Domain Controller of the originating Switch. The D\_ID shall be set to hex'FFFCyy', indicating the Domain Controller of the destination Switch.

**Payload:** The format of the EACA request payload is specified in table 112.

**Table 112 – EACA request payload**

Item	Size (bytes)
2A01h	2
Reserved	2
Commit Exchange Preamble Length	4
Commit Exchange Preamble	n

**Commit Exchange Preamble Length:** This field specifies the length of the Commit Exchange Preamble in bytes.

#### 6.2.26.1 Commit Exchange Preamble

The format of the Commit Exchange Preamble is described in table 113 below.

**Table 113 – Commit Exchange Preamble**

Item	Size (bytes)
Transaction_Identifier	12
Number of Switch Identifiers (m)	1
Flags	1
Reserved	2
ECS Switch List	m*12

##### 6.2.26.1.1 Transaction Identifier

The Transaction Identifier is used to uniquely identify a transaction in the Fabric. The scope of a Transaction Identifier is from the EACA request that begins the transaction to the ERCA request that ends the transaction. This allows multiple applications to use the Enhanced Commit Service simultaneously. The format of the Transaction Identifier is specified in table 114 below.

**Table 114 – Transaction Identifier**

Item	Size (bytes)
Generation Count	4
Application_ID	1
Domain_ID	1
Reserved	2
Timestamp	4

**Generation Count:** The generation count contains a monotonically increasing value that is used to distinguish transactions related to the same application. When a Switch originates an EACA, it shall increment the previous generation count by 1. The generation count shall be 32-bit unsigned integers that shall wrap to zero on exceeding FFFF FFFFh.

**Application\_ID:** The Application\_ID contains a value that represents an application, service, or function in the Fabric. The Application ID values are specified in table 115.

**Table 115 – Application ID values**

Value (hex)	Description
00	Vendor Specific
01	Fabric Policies (see FC-SP-2)
E0-EF	Vendor Specific
FF	Reserved for RFC 4936
other values	Reserved

**Domain\_ID:** This field contains the Domain\_ID of the originating Switch.

**Timestamp:** This field contains a timestamp that specifies the time that the request was sent from the originating Switch. The timestamp indicates in milli-seconds the elapsed time since the last system boot on the originating Switch.

Together the Generation Count, Application\_ID, Domain\_ID, and the Timestamp provide the means for a Fabric to distinguish all ECS requests. All fields of the Transaction Identifier are established by the Managing Switch when a transaction is initiated with an EACA. All ESFC, EUFC, and TCO requests within the transaction and the ERCA ending the transaction are identified by the Transaction Identifier established by the EACA request.

#### 6.2.26.1.2 Number of Switch Identifiers

Specifies the number of Switch Identifiers in the ECS Switch List including the Managing Switch.

#### 6.2.26.1.3 Flags

This field contains flag bits (7:0) that provide additional information about the specified transaction. The following flag bits are defined.

Bit 7, the Assisted/Autonomous Mode bit, shall indicate whether ECS is operating in assisted mode or in autonomous mode for the specified transaction. When the bit is set to one, ECS is operating in assisted mode for the specified transaction (see 13.2). When the bit is set to zero, ECS is operating in autonomous mode for the specified transaction (see 13.3).

Bits 6-0 shall be reserved.

#### 6.2.26.1.4 ECS Switch List

The ECS Switch List contains a list of all Switches participating in the specified transaction. This list represents a subset of all Switches in the Fabric. The ECS Switch List contains authorized Switches and non-authorized Switches. Only the authorized Switches shall participate in the ECS error recovery processing. It is the responsibility of the initial Managing Switch to construct the ECS Switch List according to the requirements of the specified application.



The Managing Switch shall send ECS requests to Managed Switches in the order indicated in the ECS Switch list. The first Switch in the list shall be the initial Managing Switch, followed by Switches that are authorized, and then by Switches that are not authorized. The format of the ECS Switch List is specified in table 116.

**Table 116 – ECS Switch List**

Item	Size (bytes)
Managing Switch Identifier	12
Switch Identifier 1	12
Switch Identifier 2	12
Switch Identifier...	12
Switch Identifier...	12
Switch Identifier m-1	12

**Switch Identifier:** The Switch Identifier contains the Domain\_ID and the Switch\_Name for a Switch in the Fabric. The format of the Switch Identifier is specified in table 117.

**Table 117 – Switch Identifier format**

Item	Size (bytes)
Domain_ID	1
Flags	1
Reserved	2
Switch_Name	8

**Flags:** This field contains flag bits (7:0) that provide additional information about the designated Switch.

Bit 7, the Switch Authorized bit, shall indicate whether the Switch shall participate in ECS recovery processing if the ECS protocol is operating in autonomous mode. When the bit is set to one, the Switch is authorized and shall participate in ECS recovery processing. When the bit is set to zero, the Switch is not authorized and shall not participate in ECS recovery processing.

Bits 6-0 shall be reserved.

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)

Signifies the rejection of the EACA request

Accept (SW\_ACC)

Signifies acceptance of the EACA request

An SW\_ACC indicates that the operation completed successfully.

SW\_RJT shall be returned as a reply to signify the rejection of the EACA\_ILS request Sequence.

### 6.2.27 Enhanced Stage Fabric Configuration (ESFC) request

An ESFC request is sent from the Domain Controller of the Managing Switch to the Domain Controller of each Managed Switch contained in the ECS Switch List (see 6.2.26.1.4). The ESFC request allows application specific data to be transported to all Managed Switches and signals each Switch to validate and stage the data locally. During the staging process the Managed Switch performs the necessary consistency checks to ensure that the operation will complete successfully with respect to that Switch.

The use of the ESFC SW\_ILS is specified in 13.

**Protocol:**

Enhanced Stage Fabric Configuration (ESFC) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to hex'FFFCxx', indicating the Domain Controller of the originating Switch. The D\_ID shall be set to hex'FFFCyy', indicating the Domain Controller of the destination Switch.

**Payload:** The format of the ESFC request payload is specified in table 118.

**Table 118 – ESFC request payload**

Item	Size (bytes)
2A02h	2
Reserved	2
Commit Exchange Preamble Length	4
Commit Exchange Preamble (see 6.2.26.1)	n
Application Data Length	4
Application Data	m

**Commit Exchange Preamble Length:** This field specifies the length of the Commit Exchange Preamble in bytes.

**Application Data Length:** The field contains the length of the application specific data in bytes.

**Application Data:** This field contains any application specific operations and the actual application data that is to be operated upon. An application shall define the format of its Application Data.

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the ESFC request  
 Accept (SW\_ACC)  
 Signifies acceptance of the ESFC request

An SW\_ACC indicates that the operation completed successfully.  
 An SW\_RJT shall be returned as a reply to signify the rejection of the ESFC request Sequence.

**6.2.28 Enhanced Update Fabric Configuration (EUFC) request**

An EUFC request is sent from the Domain Controller of the Managing Switch to the Domain Controller of each Managed Switch contained in the ECS Switch List (see 6.2.26.1.4). The EUFC request is sent to each Managed Switch to request that the Managed Switch commit the changes specified by the application data contained in the proceeding ESFC request.

The use of the EUFC SW\_ILS is specified in 13.

**Protocol:**

Enhanced Stage Fabric Configuration (EUFC) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to hex'FFFCxx', indicating the Domain Controller of the originating Switch. The D\_ID shall be set to hex'FFFCyy', indicating the Domain Controller of the destination Switch.

**Payload:** The format of the EUFC request payload is specified in table 119.

**Table 119 – EUFC request payload**

Item	Size (bytes)
2A03h	2
Reserved	2
Commit Exchange Preamble Length	4
Commit Exchange Preamble (see 6.2.26.1)	n

**Commit Exchange Preamble Length:** This field specifies the length of the Commit Exchange Preamble in bytes.

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the EUFC request  
 Accept (SW\_ACC)  
 Signifies acceptance of the EUFC request

An SW\_ACC indicates that the operation completed successfully.  
 An SW\_RJT shall be returned as a reply to signify the rejection of the EUFC request Sequence.

**6.2.29 Enhanced Release Change Authorization (ERCA) request**

The ERCA request is sent from the Domain Controller of a Switch in the ECS Switch List (see 6.2.26.1.4) to all other Switches contained in the ECS Switch List. Typically, ERCA requests are sent by the Managing Switch to all Managed Switches. ERCA requests are sent to Managed Switches in

order to free the resources reserved by the previous EACA request and completes the specified transaction with respect to each Managed Switch.

To facilitate recovery and minimize the window for denial of service attacks, an ERCA may be sent by any Switch that is a member of the ECS Switch List.

The use of the ERCA SW\_ILS is specified in 13.

**Protocol:**

Enhanced Release Change Authorization (ERCA) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to hex'FFFCxx', indicating the Domain Controller of the originating Switch. The D\_ID shall be set to hex'FFFCyy', indicating the Domain Controller of the destination Switch.

**Payload:** The format of the ERCA request payload is specified in table 120.

**Table 120 – ERCA request payload**

Item	Size (bytes)
2A04h	2
Reserved	2
Commit Exchange Preamble Length	4
Commit Exchange Preamble (see 6.2.26.1)	n

**Commit Exchange Preamble Length:** This field specifies the length of the Commit Exchange Preamble in bytes.

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the ERCA request  
 Accept (SW\_ACC)  
 Signifies acceptance of the ERCA request

An SW\_ACC indicates that the operation completed successfully.  
 SW\_RJT shall be returned as a reply to signify the rejection of the ERCA request Sequence.

**6.2.30 Transfer Commit Ownership (TCO) request**

The TCO request is sent from the Domain Controller of one Managing Switch to the Domain Controller of another Managing Switch. The TCO request transfers the role of Managing Switch to another Switch in the ECS Switch List (see 6.2.26.1.4). This provides a mechanism to ensure that only one Managing Switch is chosen to complete the transaction.

The use of the TCO SW\_ILS is further specified in 13.

**Protocol:**

Transfer Commit Ownership (TCO) request Sequence  
 Accept (SW\_ACC) reply Sequence

**Addressing:** The S\_ID shall be set to hex'FFFCxx', indicating the Domain Controller of the originating Switch. The D\_ID shall be set to hex'FFFCyy', indicating the Domain Controller of the destination Switch.

**Payload:** The format of the TCO request payload is specified in table 121.

**Table 121 – TCO request payload**

Item	Size (bytes)
2A05h	2
Reserved	2
Commit Exchange Preamble Length	4
Commit Exchange Preamble (see 6.2.26.1)	n

**Commit Exchange Preamble Length:** This field specifies the length of the Commit Exchange Preamble in bytes.

**Reply Switch Fabric Internal Link Service Sequence:**

Service Reject (SW\_RJT)  
 Signifies the rejection of the TCO request  
 Accept (SW\_ACC)  
 Signifies acceptance of the TCO request

An SW\_ACC indicates that the operation completed successfully.  
 SW\_RJT shall be returned as a reply to signify the rejection of the TCO request Sequence.

**6.3 Distributed Switch VA\_Port SW\_ILSs****6.3.1 Overview**

VA\_Port SW\_ILSs are used for Distributed Switch operations (see 17.4).

VA\_Port SW\_ILSs are used to exchange information between Controlling Switches and FCDFs (i.e., they are not used to exchange information between FCDFs).

If a Distributed Switch includes cascaded FCDFs or multiple Controlling Switches, the intermediate FCDFs or Controlling Switches relay the SW\_ILSs through a chain of Exchanges, as shown in figure 15. If one Exchange of this chain of Exchanges is abnormally terminated, then the other Exchanges in the chain shall be terminated as well.

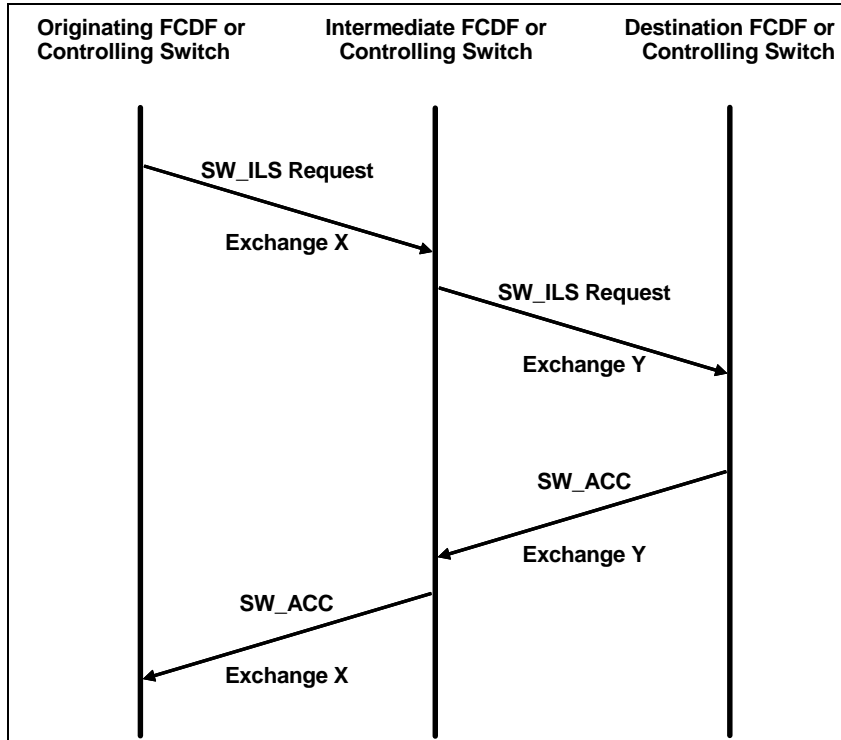


Figure 15 – VA\_Port SW\_ILS relay example

To enable the establishment of a chain of Exchanges, all VA\_Port SW\_ILS requests include the originating and destination FCDF or Controlling Switch Switch\_Names in the first two fields of their payload. The subsequent part of a VA\_Port SW\_ILS request is a list of self-identifying descriptors, as defined in 6.3.2. The response to a VA\_Port SW\_ILS request (i.e., an SW\_ACC or an SW\_RJT) shall be processed by intermediate FCDFs to follow the chain of Exchanges established when relaying the request to its destination in the opposite direction (i.e., the originating and destination FCDF or Controlling Switch Switch\_Names in a VA\_Port SW\_ACC shall not be used to relay the SW\_ACC to its destination).

Table 122 shows the VA\_Port SW\_ILSs command codes.

Table 122 – VA\_Port SW\_ILSs command codes

Encoded value	Description	Abbreviation
A000 0001h	VN_Port Reachability Notification	VNRN
A000 0002h	VN_Port Unreachability Notification	VNUN
A000 0003h	FCDF Reachability Notification	FERN
A000 0004h	FCDF Unreachability Notification	FCUN
A000 0005h	N_Port_ID Route Distribution	NPRD
A000 0006h	N_Port_ID and Zoning ACL Distribution	NPZD
A000 0007h	Active Zoning ACL Distribution	AZAD
A000 0008h	Distributed Switch Membership Distribution	DSMD
A000 0009h	Distributed ELS	DELS

## 6.3.2 VA\_Port SW\_ILS descriptors

### 6.3.2.1 Descriptor format

Each VA\_Port SW\_ILS descriptor has the format specified in table 123. This format applies also to the descriptors for the Controlling Switch redundancy protocol SW\_ILSs (see 6.4).

**Table 123 – Descriptor format**

Item	Size (bytes)
Descriptor Tag	4
Descriptor Length	4
Descriptor Value	variable

**Descriptor Tag:** the two most significant bytes of this field are reserved, the two least significant bytes contain the tag value. The descriptor tag values are specified in table 124.

**Table 124 – Descriptor tags**

Tag value	Descriptor	Reference
0001h	VN_Port Reachability	6.3.2.2
0002h	FLOGI/NPIV FDISC Parameters	6.3.2.3
0003h	VN_Port Unreachability	6.3.2.4
0004h	FCDF Reachability	6.3.2.5
0005h	Sequence Number	6.3.2.6
0006h	Controlling Switch Reachability	6.3.2.7
0007h	N_Port_IDs Reachability	6.3.2.8
0008h	Domain_IDs Reachability	6.3.2.9
0009h	Allocation Status	6.3.2.10
000Ah	Peering Status	6.3.2.11
000Bh	Membership Set	6.3.2.12
000Ch	Integrity	6.3.2.13
000Dh	FCDF Identification	6.3.2.14
000Eh	SW_ILS Request Information	6.3.2.15
000Fh	ELS Payload	6.3.2.16
0011h	Controlling Switch State	6.4.2.2
0012h	FCDF Topology	6.4.2.3
0013h	FCDF N_Port_IDs	6.4.2.4
0014h	RHello Interval	6.4.2.5
0015h	Controlling Switch Parameters	6.4.2.6
all others	Reserved	

**Descriptor Length:** contains the length in bytes of the Descriptor Value.

**Descriptor Value:** contains the specific information carried in the descriptor.

### 6.3.2.2 VN\_Port Reachability descriptor

The format of the VN\_Port Reachability descriptor is specified in table 125.

**Table 125 – VN\_Port Reachability descriptor format**

Item	Size (bytes)
Tag Value = 0001h	4
Length = 12	4
F_Port_Name	8
Physical Port Number	4

**F\_Port\_Name:** contains the F\_Port\_Name of the VF\_Port to which the newly reachable VN\_Port is being associated.

**Physical Port Number:** contains the physical port number where an FLOGI or NPIV FDISC request has been received.

### 6.3.2.3 FLOGI/NPIV FDISC Parameters descriptor

The format of the FLOGI/NPIV FDISC Parameters descriptor is specified in table 126.

**Table 126 – FLOGI/NPIV FDISC Parameters descriptor format**

Item	Size (bytes)
Tag Value = 0002h	4
Length = variable	4
FLOGI/NPIV FDISC Parameters	variable

**FLOGI/NPIV FDISC Parameters:** contains the payload of an FLOGI or NPIV FDISC request, LS\_ACC, or LS\_RJT (see FC-LS-3).

### 6.3.2.4 VN\_Port Unreachability descriptor

The format of the VN\_Port Unreachability descriptor is specified in table 127.

**Table 127 – VN\_Port Unreachability descriptor format**

Item	Size (bytes)
Tag Value = 0003h	4
Length = 20	4
Flags	1
Unreachable N_Port_ID	3
Unreachable N_Port_Name	8
F_Port_Name	8

**Flags:** 8 flag bits. The following flag bits are defined:

Bit 7.. 1: reserved.



Bit 0: indicates if only one VN\_Port is unreachable or if all the VN\_Ports associated to a VF\_Port are unreachable. This bit is set to zero to indicate that only one VN\_Port is unreachable and to one to indicate that all the VN\_Ports associated to a VF\_Port are unreachable.

**Unreachable N\_Port\_ID:** when bit 0 of the Flags field is set to zero contains the N\_Port\_ID of the unreachable VN\_Port. When bit 0 of the Flags field is set to one contains 000000h.

**Unreachable N\_Port\_Name:** when bit 0 of the Flags field is set to zero contains the N\_Port\_Name of the unreachable VN\_Port. When bit 0 of the Flags field is set to one contains 0000 0000 0000 0000h.

**F\_Port\_Name:** contains the F\_Port\_Name of the involved VF\_Port.

### 6.3.2.5 FCDF Reachability descriptor

The format of the FCDF Reachability descriptor is specified in table 128.

**Table 128 – FCDF Reachability descriptor format**

Item	Size (bytes)
Tag Value = 0004h	4
Length = 28	4
FCDF or Controlling Switch Switch_Name	8
Local A_Port_Name	8
Adjacent A_Port_Name	8
Reserved	2
Link Cost	2

**FCDF or Controlling Switch Switch\_Name:** contains the Switch\_Name of the Adjacent entity with which an ASL has been instantiated or deinstantiated.

**Local A\_Port\_Name:** contains the local A\_Port\_Name of the instantiated or deinstantiated ASL.

**Adjacent A\_Port\_Name:** contains the Adjacent A\_Port\_Name of the instantiated or deinstantiated ASL.

**Link Cost:** contains the cost of the instantiated or deinstantiated ASL.

### 6.3.2.6 Sequence Number descriptor

The format of the Sequence Number descriptor is specified in table 129.

**Table 129 – Sequence Number descriptor format**

Item	Size (bytes)
Tag Value = 0005h	4
Length = 8	4
Sequence Number	8

**Sequence Number:** contains a monotonically increasing sequence number. When the sequence number reaches the value FFFFFFFF FFFFFFFFh it wraps to 00000000 00000000h.

### 6.3.2.7 Controlling Switch Reachability descriptor

The format of the Controlling Switch Reachability descriptor is specified in table 130.

**Table 130 – Controlling Switch Reachability descriptor format**

Item	Size (bytes)
Tag Value = 0006h	4
Length = variable	4
Controlling Switch Switch_Name	8
Number of Paths to the Controlling Switch (j)	4
Next-hop Switch_Name #1	8
Local A_Port_Name #1	8
Path #1 cost	4
Next-hop Switch_Name #2	8
Local A_Port_Name #2	8
Path #2 cost	4
...	
Next-hop Switch_Name #j	8
Local A_Port_Name #j	8
Path #j cost	4

**Controlling Switch Switch\_Name:** contains the Switch\_Name of the Controlling Switch.

**Number of Paths to the Controlling Switch:** contains the number of paths toward the Controlling Switch. Each path that follows is expressed by the Switch\_Name of the next-hop FCDF or Controlling Switch followed by the local A\_Port\_Name of the involved ASL and by the path cost.

### 6.3.2.8 N\_Port\_IDs Reachability descriptor

The format of the N\_Port\_IDs Reachability descriptor is specified in table 131.

**Table 131 – N\_Port\_IDs Reachability descriptor format**

Item	Size (bytes)
Tag Value = 0007h	4
Length = variable	4
Number of N_Port_ID Reachability Entries (p)	4
N_Port_ID Reachability Entry #1	see table 132
N_Port_ID Reachability Entry #2	see table 132
...	
N_Port_ID Reachability Entry #p	see table 132

**Number of N\_Port\_ID Reachability Entries:** contains the number of N\_Port\_ID Reachability Entries that follow. There shall be an N\_Port\_ID Reachability Entry for each FCDF currently belonging to the Distributed Switch. The N\_Port\_ID Reachability Entry format is specified in table 132.

**Table 132 – N\_Port\_ID Reachability Entry format**

Item	Size (bytes)
Reachable FCDF Switch_Name	8
Number of Equal Cost Paths to the Reachable FCDF (w)	4
Next-hop Switch_Name #1	8
Local A_Port_Name #1	8
Next-hop Switch_Name #2	8
Local A_Port_Name #2	8
...	
Next-hop Switch_Name #w	8
Local A_Port_Name #w	8
Number of N_Port_ID Ranges (q)	4
N_Port_ID Range #1	4
N_Port_ID Range #2	4
...	
N_Port_ID Range #q	4

**Reachable FCDF Switch\_Name:** contains the Switch\_Name of the FCDF to which the subsequent next-hops and N\_Port\_ID Ranges refer.

**Number of Equal Cost Paths to the Reachable FCDF:** contains the number of equal cost paths having the lowest cost toward the destination FCDF. Each path that follows is expressed as the Switch\_Name of the next-hop FCDF or Controlling Switch followed by the local A\_Port\_Name of the involved ASL. Only the first path listed shall be used to relay VA\_Port SW\_ILSs to the reachable FCDF.

**Number of N\_Port\_ID Ranges:** contains the number of N\_Port\_ID Range Entries that follow. The N\_Port\_ID Range is defined by an N\_Port\_ID in the least significant three bytes, and by the number of bits defining the range in the most significant byte (e.g., the range 020200h to 02027Fh is expressed as '07h || 020200h'). The set of N\_Port\_ID Range Entries encodes in a compact form all the N\_Port\_IDs currently allocated to VN\_Ports logged into the reachable FCDF.

### 6.3.2.9 Domain\_IDs Reachability descriptor

The format of the Domain\_IDs Reachability descriptor is specified in table 133.

**Table 133 – Domain\_IDs Reachability descriptor format**

Item	Size (bytes)
Tag Value = 0008h	4
Length = variable	4
Number of Reachable Domain_ID Entries (r)	4
Reachable Domain_ID Entry #1	see table 134
Reachable Domain_ID Entry #2	see table 134
...	
Reachable Domain_ID Entry #r	see table 134

**Number of Reachable Domain\_ID Entries:** contains the number of Reachable Domain\_ID Entries that follow. The Reachable Domain\_ID Entry format is specified in table 134.

**Table 134 – Reachable Domain\_ID Entry format**

Item	Size (bytes)
Reachable Domain_ID	4
Number of Equal Cost Paths to the Reachable Domain_ID (y)	4
Next-hop Switch_Name #1	8
Local A_Port_Name #1	8
Next-hop Switch_Name #2	8
Local A_Port_Name #2	8
...	
Next-hop Switch_Name #y	8
Local A_Port_Name #y	8

**Reachable Domain\_ID:** contains the reachable Domain\_ID. The three most significant bytes of this field are reserved.

**Number of Equal Cost Paths to the Reachable Domain\_ID:** contains the number of equal cost paths having the lowest cost toward the destination Domain\_ID. Each path that follows is expressed as the Switch\_Name of the next-hop FCDF or Controlling Switch followed by the local A\_Port\_Name of the involved ASL.

### 6.3.2.10 Allocation Status descriptor

The format of the Allocation Status descriptor is specified in table 135.

**Table 135 – Allocation Status descriptor format**

Item	Size (bytes)
Tag Value = 0009h	4
Length = variable	4
Number of Allocation / Deallocation Entries (z)	4
Allocation / Deallocation Entry #1	see table 136
Deallocation Entry #2	see table 136
...	
Deallocation Entry #z	see table 136

**Number of Allocation / Deallocation Entries:** contains the number of Allocation / Deallocation Entries that follow. Only one Allocation Entry may be present, multiple Deallocation Entries may be present. The Allocation / Deallocation Entry format is specified in table 136.

**Table 136 – Allocation / Deallocation Entry format**

Item	Size (bytes)
Flags	4
Allocated / Deallocated N_Port_ID	4
N_Port_Name associated with the Allocated/Deallocated N_Port_ID	8
Switch_Name of the FCDF associated with the Allocated/Deallocated N_Port_ID	8
FLOGI / NPIV FDISC LS_ACC Parameters	116

**Flags:** 32 flag bits. The following Flags field bits are defined:

Bit 31 .. 2: reserved.

Bit 1: indicates if the FLOGI / NPIV FDISC LS\_ACC Parameters field is present in the payload. The field is present when this bit is set to one and not present when this bit is set to zero. This bit shall not be set to one when bit 0 indicates deallocation (i.e., the FLOGI / NPIV FDISC LS\_ACC Parameters field may be present only when an N\_Port\_ID allocation is performed).

Bit 0: indicates if the operation is an allocation or a deallocation. This bit is set to zero to indicate allocation and to one to indicate deallocation.

**Allocated / Deallocated N\_Port\_ID:** contains the N\_Port\_ID that the Primary Controlling Switch allocated or deallocated in the least significant three bytes. The most significant byte is reserved.

**N\_Port\_Name associated with the Allocated/Deallocated N\_Port\_ID:** contains the N\_Port\_Name of the VN\_Port for which an N\_Port\_ID is allocated or deallocated.

**Switch\_Name of the FCDF associated with the Allocated/Deallocated N\_Port\_ID:** contains the Switch\_Name of the FCDF associated with the VN\_Port for which an N\_Port\_ID is allocated or deallocated.

**FLOGI / NPIV FDISC LS\_ACC Parameters:** this field is present when bit 1 of the Flags field is set to one. It contains the payload of the LS\_ACC generated by the Primary Controlling Switch in response to the FLOGI or NPIV FDISC payload provided in the VNRN request Sequence.

### 6.3.2.11 Peering Status descriptor

The format of the Peering Status descriptor is specified in table 137.

**Table 137 – Peering Status descriptor format**

Item	Size (bytes)
Tag Value = 000Ah	4
Length = variable	4
Number of Peering Entries (h)	4
Peering Entry #1	see table 138
Peering Entry #2	see table 138
...	
Peering Entry #h	see table 138

**Number of Peering Entries:** contains the number of Peering Entries that follow.

Each Peering Entry contains a complete list of the Peer N\_Port\_IDs with which the Principal N\_Port\_ID is allowed to communicate according to the current Fabric Zoning configuration. The Peering Entry format is specified in table 138.

**Table 138 – Peering Entry format**

Item	Size (bytes)
Principal N_Port_ID	4
Number of Allowed Peers (k)	4
Peer N_Port_ID #1	4
Peer N_Port_ID #2	4
...	
Peer N_Port_ID #k	4

**Principal N\_Port\_ID:** contains in the least significant three bytes the N\_Port\_ID to which the subsequent Peer N\_Port\_IDs refer. The most significant byte is reserved.

**Number of Allowed Peers:** contains the number of N\_Port\_IDs that follow to which the Principal N\_Port\_ID is allowed to communicate.

**Peer N\_Port\_ID:** contains an N\_Port\_ID in the least significant three bytes and the most significant byte is reserved.

### 6.3.2.12 Membership Set descriptor

The format of the Membership Set descriptor is specified in table 139.

**Table 139 – Membership Set descriptor format**

Item	Size (bytes)
Tag Value = 000Bh	4
Length = variable	4
Fabric_Name	8
Virtual Domain Switch_Name	8
Primary Controlling Switch Switch_Name	8
Secondary Controlling Switch Switch_Name	8
Number of Controlling Switches in the Controlling Switch Set (k)	4
Controlling Switch Switch_Name #1	8
Controlling Switch Switch_Name #2	8
...	
Controlling Switch Switch_Name #k	8
Number of FCDFs in the FCDF_Set (n)	4
FCDF Switch_Name #1	8
FCDF Switch_Name #2	8
...	
FCDF Switch_Name #n	8

**Fabric\_Name:** contains the Fabric\_Name of the Distributed Switch's associated Fabric.

**Virtual Domain Switch\_Name:** contains the Switch\_Name of the Virtual Domain.

**Primary Controlling Switch Switch\_Name:** contains the Switch\_Name of the Primary Controlling Switch.

**Secondary Controlling Switch Switch\_Name:** contains the Switch\_Name of the Secondary Controlling Switch. This field shall be set to 00000000 00000000h when there is no Secondary Controlling Switch.

**Number of Controlling Switches in the Controlling Switch Set:** contains the number of Controlling Switch Switch\_Names that follow. This list of Controlling Switch Switch\_Names is the Controlling Switch Set of the Distributed Switch.

**Number of FCDFs in the FCDF\_Set:** contains the number of FCDF Switch\_Names that follow. This list of FCDF Switch\_Names is the FCDF\_Set of the Distributed Switch. If the number of FCDF Switch\_Names is zero, then any FCDF is allowed in the Distributed Switch.

### 6.3.2.13 Integrity descriptor

The format of the Integrity descriptor is specified in table 140.

**Table 140 – Integrity descriptor format**

Item	Size (bytes)
Tag Value = 000Ch	4
Length = variable	4
Integrity Type	4
Integrity Check Value Length	4
Integrity Check Value	variable

**Integrity Type:** indicates, in the least significant byte, the type of cryptographic integrity that protects the payload. Integrity Type field values are specified in table 141.

**Table 141 – Integrity Type field values**

Value (hex)	Description
00	No integrity
01	HMAC-SHA-256-128 integrity
02 to FF	Reserved

**Integrity Check Value Length:** contains the length expressed in bytes of the Integrity Check Value.

**Integrity Check Value:** contains the cryptographic hash of the payload computed using the shared key according to the specified Integrity Type.

### 6.3.2.14 FCDF Identification descriptor

The format of the FCDF Identification descriptor is specified in table 142.

**Table 142 – FCDF Identification descriptor format**

Item	Size (bytes)
Tag Value = 000Dh	4
Length = 32	4
Number of Physical Ports	4
T10 Vendor ID	8
Product Identification	16
Product Revision Level	4

**Number of Physical Ports:** contains the number of physical ports that the FCDF has.

**T10 Vendor ID:** contains eight bytes of left-aligned ASCII data identifying the vendor of the FCDF. The T10 vendor identification is assigned by INCITS.



**Product Identification:** contains sixteen bytes of left-aligned ASCII data defined by the vendor.

**Product Revision Level:** contains four bytes of left-aligned ASCII data defined by the vendor.

### 6.3.2.15 SW\_ILS Request Information descriptor

The format of the SW\_ILS Request Information descriptor is specified in table 143.

**Table 143 – SW\_ILS Request Information descriptor format**

Item	Size (bytes)
Tag Value = 000Eh	4
Length = 4	4
SW_ILS Request Opcode	4

**SW\_ILS Request Opcode:** contains the opcode of the SW\_ILS request to which the SW\_ACC containing this descriptor is replying.

### 6.3.2.16 ELS Payload descriptor

The format of the ELS Payload descriptor is specified in table 144.

**Table 144 – ELS Payload descriptor format**

Item	Size (bytes)
Tag Value = 000Fh	4
Length = variable	4
ELS Control	4
N_Port_ID	4
ELS Payload	variable

**ELS Control:** Contains control information for use in a Distributed ELS SW\_ILS. The following bits are defined:

Bit 31 .. 1: Reserved.

Bit 0: Send ELS. The Send ELS bit is valid if the ELS Payload descriptor is used in a DELS SW\_ILS. This bit is used if a Controlling Switch requires that an F\_Port Controller in an FCDF transmit an ELS request to the destination N\_Port\_ID provided. If set to one, it indicates that the recipient of the DELS SW\_ILS is to transmit the ELS Payload as an ELS Exchange with the intended recipient. If set to zero, it indicates that the recipient should perform the ELS function provided in the ELS payload (see 6.3.3.9).

**N\_Port\_ID:** The N\_Port ID field is meaningful only if ELS Control Flags bit 0 (Send ELS) is set to one. If ELS Control bit 0 is set to one, it shall contain the N\_Port\_ID of the VN\_Port to which the encapsulated ELS payload is to be sent. If meaningful, this field contains an N\_Port\_ID in the least significant three bytes and the most significant byte is reserved.

**ELS Payload:** The ELS Payload field contains the payload of an Extended Link Service (see FC-LS-3).

### 6.3.3 VA\_Port SW\_ILSs

#### 6.3.3.1 VN\_Port Reachability Notification (VNRN)

The VN\_Port Reachability Notification SW\_ILS is used by an FCDF to communicate to the Primary Controlling Switch that a VN\_Port is attempting Fabric login through an FLOGI request or a NPIV FDISC request. If the FCDF does not have an ASL with the Primary Controlling Switch, the VNRN SW\_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs or Controlling Switches.

#### VNRN request Sequence

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFF9h, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFF9h, indicating the destination VE\_Port.

**Payload:** the format of the VNRN request Sequence Payload is specified in table 145.

**Table 145 – VNRN request payload**

Item	Size (bytes)
SW_ILS Code = A000 0001h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
VN_Port Reachability descriptor	see 6.3.2.2
FLOGI/NPIV FDISC Parameters descriptor	see 6.3.2.3

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating FCDF Switch\_Name:** contains the Switch\_Name of the originating FCDF.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**VN\_Port Reachability descriptor:** describes the VF\_Port associated with the newly reachable VN\_Port. See 6.3.2.2.

**FLOGI/NPIV FDISC Parameters descriptor:** contains the payload of the received FLOGI or NPIV FDISC request (see FC-LS-3).

#### VNRN reply Sequence

**SW\_RJT:** indicates the rejection of the VNRN request Sequence. As a result, a FLOGI LS\_RJT or a NPIV FDISC LS\_RJT with reason code set to 'Logical busy' and reason code explanation set to "No additional explanation" is sent as response to the FLOGI request or NPIV FDISC request that caused the issuance of the VNRN request (see 17.4.3).

**SW\_ACC:** indicates the acceptance of the VNRN request Sequence (see 17.4.3). The format of the VNRN SW\_ACC payload is specified in table 146.

**Table 146 – VNRN SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12
FLOGI / NPIV FDISC Parameters descriptor	see 6.3.2.3

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

**FLOGI / NPIV FDISC Parameters descriptor:** this descriptor contains the payload of the LS\_ACC or LS\_RJT generated by the Primary Controlling Switch in response to the FLOGI or NPIV FDISC payload provided in the VNRN request Sequence.

### 6.3.3.2 VN\_Port Unreachability Notification (VNUN)

The VN\_Port Unreachability Notification SW\_ILS is used by an FCDF to communicate to the Primary Controlling Switch that one or more of its VN\_Ports have been logged out. If the FCDF does not have an ASL with the Primary Controlling Switch, the VNUN SW\_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs or Controlling Switches.

#### VNUN request Sequence

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the VNUN request Sequence payload is specified in table 147.

**Table 147 – VNUN request payload**

Item	Size (bytes)
SW_ILS Code = A000 0002h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
VN_Port Unreachability descriptor	see 6.3.2.4

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating FCDF Switch\_Name:** contains the Switch\_Name of the requesting FCDF.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**VN\_Port Unreachability descriptor:** describes the unreachable VN\_Port(s) and the associated VF\_Port. See 6.3.2.4.

**VNUN reply Sequence**

**SW\_RJT:** indicates the rejection of the VNUN request Sequence.

**SW\_ACC:** indicates the acceptance of the VNUN request Sequence. The format of the VNUN SW\_ACC payload is specified in table 148.

**Table 148 – VNUN SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

**6.3.3.3 FCDF Reachability Notification (FCRN)**

The FCDF Reachability Notification SW\_ILS is used by an FCDF to communicate to the Primary Controlling Switch that it has instantiated an ASL with another FCDF or with a Controlling Switch. If the FCDF does not have an ASL with the Primary Controlling Switch, the FCRN SW\_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs or Controlling Switches.

The FCRN SW\_ILS is also used between Primary and Secondary Controlling Switch to keep their state synchronized.

**FCRN request Sequence**

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the FCRN request Sequence payload is specified in table 149.

**Table 149 – FCRN request payload**

Item	Size (bytes)
SW_ILS Code = A000 0003h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
FCDF Reachability descriptor	see 6.3.2.5

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating FCDF Switch\_Name:** contains the Switch\_Name of the requesting FCDF.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**FCDF Reachability descriptor:** describes the instantiated ASL (see 6.3.2.5).

#### FCRN reply Sequence

**SW\_RJT:** indicates the rejection of the FCRN request Sequence.

**SW\_ACC:** indicates the acceptance of the FCRN request Sequence. The format of the FCRN SW\_ACC payload is specified in table 150.

**Table 150 – FCRN SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

#### 6.3.3.4 FCDF Unreachability Notification (FCUN)

The FCDF Unreachability Notification SW\_ILS is used by an FCDF to communicate to the Primary Controlling Switch that it has deinstantiated an ASL with another FCDF or with a Controlling Switch. If the FCDF does not have an ASL with the Primary Controlling Switch, the FCUN SW\_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs or Controlling Switches.

The FCUN SW\_ILS is also used between Primary and Secondary Controlling Switch to keep their state synchronized.

**FCUN request Sequence**

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the FCUN request Sequence payload is specified in table 151.

**Table 151 – FCUN request payload**

Item	Size (bytes)
SW_ILS Code = A000 0004h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
FCDF Reachability descriptor	see 6.3.2.5

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating FCDF Switch\_Name:** contains the Switch\_Name of the requesting FCDF.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**FCDF Reachability descriptor:** describes the deinstantiated ASL (see 6.3.2.5)..

**FCUN reply Sequence**

**SW\_RJT:** indicates the rejection of the FCUN request Sequence.

**SW\_ACC:** indicates the acceptance of the FCUN request Sequence. The format of the FCUN SW\_ACC payload is specified in table 152.

**Table 152 – FCUN SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

### 6.3.3.5 N\_Port\_ID Route Distribution (NPRD)

The N\_Port\_ID Route Distribution SW\_ILS is used by the Primary Controlling Switch to communicate to FCDFs or to Alternate Controlling Switches, if any, the N\_Port\_ID routing information for the Distributed Switch. If the Primary Controlling Switch does not have an ASL with the destination FCDF or an AISL with the destination Alternate Controlling Switch, the NPRD SW\_ILS is relayed to the destination FCDF by the intermediate FCDFs or Controlling Switches.

#### NPRD request Sequence

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the NPRD request Sequence payload is specified in table 153.

**Table 153 – NPRD request payload**

Item	Size (bytes)
SW_ILS Code = A000 0005h	4
Destination FCDF or Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Sequence Number descriptor	see 6.3.2.6
Primary Controlling Switch Reachability descriptor	see 6.3.2.7
Secondary Controlling Switch Reachability descriptor	see 6.3.2.7
N_Port_IDs Reachability descriptor	see 6.3.2.8
Domain_IDs Reachability descriptor	see 6.3.2.9

**Destination FCDF or Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination FCDF or Alternate Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the requesting Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Sequence Number descriptor:** contains a sequence number for this instance of an NPRD request. See 6.3.2.6.

**Primary Controlling Switch Reachability descriptor:** contains the reachability information toward the Primary Controlling Switch.

NOTE 7 – Paths toward the Primary Controlling Switch are fundamental for the operation of an FCDF. Specifying higher cost paths enables more redundancy, because if the lowest cost path toward the Primary Controlling Switch fails, a higher cost path may be used.

**Secondary Controlling Switch Reachability descriptor:** contains the reachability information toward the Secondary Controlling Switch.

**N\_Port\_IDs Reachability descriptor:** contains reachability information for the N\_Port\_IDs currently allocated from the Virtual Domain. See 6.3.2.8.

**Domain\_IDs Reachability descriptor:** contains reachability information for other Domain\_IDs. See 6.3.2.9.

NPRD requests sent to the Alternate Controlling Switches, if any, shall not include the Domain\_ID Reachability descriptor. NPRD requests sent to the FCDFs shall include the Domain\_ID Reachability descriptor.

**NPRD reply Sequence**

**SW\_RJT:** indicates the rejection of the NPRD request Sequence.

**SW\_ACC:** indicates the acceptance of the NPRD request Sequence. The format of the NPRD SW\_ACC payload is specified in table 154.

**Table 154 – NPRD SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF or Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

**6.3.3.6 N\_Port\_ID and Zoning ACL Distribution (NPZD)**

The N\_Port\_ID and Zoning ACL Distribution SW\_ILS is used by the Primary Controlling Switch to communicate to the Secondary Controlling Switch (if present) and to Alternate Controlling Switches, if any, the allocation of an N\_Port\_ID and/or the deallocation of one or more N\_Port\_IDs, and to communicate to FCDFs the allocation of an N\_Port\_ID and its associated Zoning ACL information and/or the deallocation of one or more N\_Port\_IDs and their associated Zoning ACL information. Upon receiving an NPZD request, an FCDF shall update its Zoning enforcement for the listed Principal N\_Port\_IDs according to the received Zoning ACLs. If the Primary Controlling Switch does not have an ASL with the destination FCDF or an AISL with the destination Alternate Controlling Switch, the NPZD SW\_ILS is relayed to the destination by the intermediate FCDFs or Controlling Switches.

**NPZD request Sequence**

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.



**Payload:** the format of the NPZD request Sequence payload is specified in table 155.

**Table 155 – NPZD request payload**

Item	Size (bytes)
SW_ILS Code = A000 0006h	4
Destination FCDF or Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Sequence Number descriptor	see 6.3.2.6
Allocation Status descriptor	see 6.3.2.10
Peering Status descriptor	see 6.3.2.11

**Destination FCDF or Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination FCDF or Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the requesting Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Sequence Number:** contains a sequence number for this instance of an NPZD request. See 6.3.2.6.

**Allocation Status descriptor:** describes the allocated or deallocated N\_Port\_IDs. See 6.3.2.10.

When an N\_Port\_ID is deallocated and allocated at the same time (e.g., as a result of a re-login), that N\_Port\_ID is listed only in the allocation entry (i.e., there is no deallocation entry for that N\_Port\_ID in the Allocation Status descriptor).

NPZD requests sent to the Secondary Controlling Switch, if present, shall include the FLOGI / NPIV FDISC LS\_ACC Parameters in the Allocation Status descriptor. NPZD requests sent to the FCDFs or to the Alternate Controlling Switches, if any, shall not include the FLOGI / NPIV FDISC LS\_ACC Parameters in the Allocation Status descriptor.

**Peering Status descriptor:** contains Peering entries per each VN\_Port currently logged into the destination FCDF. See 6.3.2.11.

NPZD requests sent to the Secondary Controlling Switch, if present, or to the Alternate Controlling Switches, if any, shall not include the Peering Status descriptor. NPZD requests sent to the FCDFs may include the Peering Status descriptor.

When present, the Peering Status descriptor contains Peering entries per each VN\_Port currently logged into the destination FCDF and with which the allocated N\_Port\_ID is allowed to communicate or with which the deallocated N\_Port\_IDs were allowed to communicate, according to the current Fabric Zoning configuration. In case of allocation, the Peering Status descriptor for the FCDF to which the N\_Port\_ID is allocated also contains a Peering Entry with a Principal N\_Port\_ID equal to the allocated N\_Port\_ID, if Zoning is enforced by the Fabric. In case of deallocation, the Zoning ACLs for the deallocated N\_Port\_IDs are implicitly removed (i.e., are no longer enforced) and the Peering Status descriptor for the FCDF that had the deallocated N\_Port\_IDs does not contain Peering Entries with a Principal N\_Port\_ID equal to any of the deallocated N\_Port\_IDs. If there is no Peering Status

descriptor for the FCDF to which the N\_Port\_ID is allocated, that N\_Port\_ID is allowed to communicate with any other VN\_Port.

If Zoning is enforced in the Fabric, an NPZD request resulting from a Fabric Login within the Virtual Domain or a Fabric Logout within the Virtual Domain shall contain an Allocation Status descriptor and:

- a) shall contain a Peering Status descriptor if such a Fabric Login or Fabric Logout results in a change in the Zoning ACLs that the receiving FCDF has to enforce; or
- b) shall not contain a Peering Status descriptor if such a Fabric Login or Fabric Logout does not result in a change in the Zoning ACLs that the receiving FCDF has to enforce.

If Zoning is enforced in the Fabric and a Fabric Login outside the Virtual Domain or a Fabric Logout outside the Virtual Domain results in a change in some Zoning ACLs that an FCDF has to enforce, then the resulting NPZD request shall not contain an Allocation Status descriptor and shall contain a Peering Status descriptor and is sent only to the affected FCDFs (see 17.9.3).

If Zoning is not enforced in the Fabric:

- a) an NPZD request resulting from a Fabric Login within the Virtual Domain or a Fabric Logout within the Virtual Domain shall contain an Allocation Status descriptor and shall not contain a Peering Status descriptor; and
- b) an NPZD request shall not be generated as a result of a Fabric Login outside the Virtual Domain or a Fabric Logout outside the Virtual Domain.

### NPZD reply Sequence

**SW\_RJT:** indicates the rejection of the NPZD request Sequence.

**SW\_ACC:** indicates the acceptance of the NPZD request Sequence. The format of the NPZD SW\_ACC payload is specified in table 156.

**Table 156 – NPZD SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF or Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

### 6.3.3.7 Active Zoning ACL Distribution (AZAD)

The Active Zoning ACL Distribution SW\_ILS is used by the Primary Controlling Switch to communicate to an FCDF new Zoning ACL information when a new Zone Set is activated in the fabric. Upon receiving an AZAD request, an FCDF shall completely replace its Zoning enforcement according to the received Zoning ACLs. If the Primary Controlling Switch does not have an ASL with the destination FCDF, the AZAD SW\_ILS is relayed to the destination FCDF by the intermediate FCDFs or Controlling Switches.

**AZAD request Sequence**

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the AZAD request Sequence payload is specified in table 157.

**Table 157 – AZAD request payload**

Item	Size (bytes)
SW_ILS Code = A000 0007h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Sequence Number descriptor	see 6.3.2.6
Peering Status descriptor	see 6.3.2.11

**Destination FCDF Switch\_Name:** contains the Switch\_Name of the destination FCDF.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the requesting Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Sequence Number:** contains a sequence number for this instance of an AZAD request. See 6.3.2.6.

**Peering Status descriptor:** contains a Peering entry per each VN\_Port currently logged into the destination FCDF specifying the N\_Port\_IDs with which that VN\_Port is allowed to communicate, according to the current Fabric Zoning configuration. See 6.3.2.11.

**AZAD reply Sequence**

**SW\_RJT:** indicates the rejection of the AZAD request Sequence.

**SW\_ACC:** indicates the acceptance of the AZAD request Sequence. The format of the AZAD SW\_ACC payload is specified in table 158.

**Table 158 – AZAD SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

### 6.3.3.8 Distributed Switch Membership Distribution (DSMD)

The Distributed Switch Membership Distribution SW\_ILS is used by the Primary Controlling Switch to communicate to an FCDF or to an Alternate Controlling Switch the identities of the Primary and Secondary Controlling Switches and of all the Controlling Switches and FCDFs that compose the Distributed Switch. The DSMD payload may be integrity protected by a cryptographic hash; in this case the involved entities shall be provided with a shared key. If the Primary Controlling Switch does not have an ASL with the destination FCDF or an AISL with the destination Alternate Controlling Switch, the DSMD SW\_ILS is relayed to the destination FCDF by the intermediate FCDFs or Controlling Switches.

#### DSMD request Sequence

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the DSMD request Sequence payload is specified in table 159.

**Table 159 – DSMD request payload**

Item	Size (bytes)
SW_ILS Code = A000 0008h	4
Destination FCDF or Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Membership Set descriptor	see 6.3.2.12
Integrity descriptor	see 6.3.2.13

**Destination FCDF or Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination FCDF or Alternate Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Membership Set descriptor:** contains the Distributed Switch membership information. See 6.3.2.12.

**Integrity descriptor:** contains the cryptographic hash protecting the DSMD payload. See 6.3.2.13.

#### DSMD reply Sequence

**SW\_RJT:** indicates the rejection of the DSMD request Sequence.

**SW\_ACC:** indicates the acceptance of the DSMD request Sequence. The format of the DSMD SW\_ACC payload is specified in table 160.

**Table 160 – DSMD SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF or Controlling Switch Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12
FCDF Identification descriptor	see 6.3.2.14

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

**FCDF Identification descriptor:** contains identification information for the replying FCDF. See 6.3.2.14.

### 6.3.3.9 Distributed ELS (DELS)

The Distributed ELS SW\_ILS is used by the Primary Controlling Switch to:

- a) request an F\_Port Controller in an FCDF to perform an ELS function; or
- b) request an F\_Port Controller in an FCDF to perform an ELS Exchange with an attached N\_Port.

#### DELS request Sequence

**Addressing:** when used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port. When used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** the format of the DELS request Sequence payload is specified in table 161.

**Table 161 – DELS request payload**

Item	Size (bytes)
SW_ILS Code = A000 0009h	4
Destination FCDF Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
ELS Payload descriptor	see 6.3.2.16

**Destination FCDF Switch\_Name:** contains the Switch\_Name of the destination FCDF.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**ELS Payload descriptor:** contains the payload of the ELS request carried in the DELS SW\_ILS. See 6.3.2.16.

The ELS requests that are allowed in a DELS SW\_ILS request are:

- a) RNID forwarded to an FCDF when received by the Fabric Controller;
- b) RLS forwarded to an FCDF when received by the Domain Controller;
- c) RDP forwarded to an FCDF when received by the Domain Controller;
- d) LCB request from a Controlling Switch to an FCDF; and
- e) LOGO request from a Controlling Switch to an FCDF.

**DELS reply Sequence**

**SW\_RJT:** indicates the rejection of the DELS request Sequence. As a result, the Controlling Switch shall either answer a received ELS request that caused the issuance of the DELS request, if possible, or reject the received ELS request with a LS\_RJT.

**SW\_ACC:** indicates the acceptance of the DELS request Sequence. The format of the DELS SW\_ACC payload is specified in table 162.

**Table 162 – DELS SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating FCDF Switch_Name	8
Descriptor List Length	4
SW_ILS Request Information descriptor	12
ELS Payload descriptor	see 6.3.2.16

**SW\_ILS Request Information descriptor:** see 6.3.2.15.

**ELS Payload descriptor:** see 6.3.2.16. This descriptor contains the payload of the LS\_RJT or LS\_ACC generated by the FCDF in response to the ELS request provided in the DELS request.

### 6.3.4 VA\_Port SW\_ILS timeouts

Table 163 shows the timeouts associated to each VA\_Port SW\_ILS.

**Table 163 – VA\_Port SW\_ILSs timeouts**

Description	Abbreviation	Timeout
VN_Port Reachability Notification	VNRN	2000 ms
VN_Port Unreachability Notification	VNUN	500 ms
FCDF Reachability Notification	FCRN	500 ms
FCDF Unreachability Notification	FCUN	500 ms
N_Port_ID Route Distribution	NPRD	1000 ms
N_Port_ID and Zoning ACL Distribution	NPZD	1000 ms
Active Zoning ACL Distribution	AZAD	1000 ms
Distributed Switch Membership Distribution	DSMD	1000 ms

## 6.4 Controlling Switch redundancy protocol SW\_ILSs

### 6.4.1 Overview

The Controlling Switch redundancy protocol SW\_ILSs are used to exchange redundancy information between Controlling Switches. Controlling Switch redundancy protocol SW\_ILSs include the originating and destination Controlling Switch Switch\_Names in the first two fields of their payload. The subsequent part of a Controlling Switch redundancy protocol SW\_ILS is a list of self-identifying descriptors, as defined in 6.4.2. The descriptor list may be null.

Table 164 shows the Controlling Switch redundancy protocol SW\_ILSs command codes.

**Table 164 – Controlling Switch redundancy protocol SW\_ILSs command codes**

Encoded Value	Description	Abbreviation
A100 0001h	Exchange Redundancy Parameters	ERP
A100 0002h	Get FCDF Topology State	GFTS
A100 0003h	Get FDCF N_Port_IDs State	GFNS
A100 0004h	Secondary Synchronization Achieved	SSA
A100 0005h	Redundancy Hello	RHello
A100 0006h	Select Primary Controlling Switch	SPCS
A100 0007h	Exchange Controlling Switch Parameters	ECSP

### 6.4.2 Controlling Switch redundancy protocol descriptors

#### 6.4.2.1 Descriptor format

The Controlling Switch redundancy protocol descriptors have the same format of the VA\_Port SW\_ILS descriptors (see 6.3.2.1). Descriptor tags are specified in table 124.

### 6.4.2.2 Controlling Switch State descriptor

The format of the Controlling Switch State descriptor is specified in table 165.

**Table 165 – Controlling Switch State descriptor format**

Item	Size (bytes)
Tag Value = 0011h	4
Length = variable	4
Originating Controlling Switch Priority	4
Virtual Domain Switch_Name	8
Number of Allocated N_Port_ID Ranges (q)	4
Allocated N_Port_ID Range #1	4
Allocated N_Port_ID Range #2	4
...	
Allocated N_Port_ID Range #q	4

**Originating Controlling Switch Priority:** contains the operational Priority of the originating Controlling Switch in the least significant byte and three reserved bytes in the three most significant bytes.

**Virtual Domain Switch\_Name:** contains the Switch\_Name of the Virtual Domain. This field shall be set to 00000000 00000000h if the Switch\_Name of the Virtual Domain has not yet been assigned.

**Number of Allocated N\_Port\_ID Ranges:** contains the number of Allocated N\_Port\_ID Range Entries that follow. This list of Allocated N\_Port\_ID Ranges identifies the N\_Port\_IDs allocated by the originating Controlling Switch. The N\_Port\_ID Range is defined by an N\_Port\_ID in the least significant three bytes, and by the number of bits defining the range in the most significant byte (e.g., the range 020200h .. 02027Fh is expressed as '7 || 020200h').

### 6.4.2.3 FCDF Topology descriptor

The format of the FCDF Topology descriptor is specified in table 166.

**Table 166 – FCDF Topology descriptor format**

Item	Size (bytes)
Tag Value = 0012h	4
Length = variable	4
Number of FCDF Connectivity Records (n)	4
FCDF Connectivity Record #1	see table 167
FCDF Connectivity Record #2	see table 167
...	
FCDF Connectivity Record #n	see table 167



**Number of FCDF Connectivity Records:** contains the number of FCDF Connectivity Records that follow. The format of the FCDF Connectivity Record is specified in table 167.

**Table 167 – FCDF Connectivity Record format**

Item	Size (bytes)
FCDF Switch_Name	8
Number of ASL Records (m)	4
ASL Record #1	28
ASL Record #2	28
...	
ASL Record #m	8

**FCDF Switch\_Name:** contains the Switch\_Name of the FCDF whose ASLs are being described.

**Number of ASL Records:** contains the number of ASL Records that follow. The format of the ASL Record is specified in table 168.

**Table 168 – ASL Record format**

Item	Size (bytes)
Switch_Name of Neighbor	8
Local A_Port_Name	8
Adjacent A_Port_Name	8
Link Cost	4

**Switch\_Name of Neighbor:** contains the Switch\_Name of the FCDF or Controlling Switch at the other end of the described ASL.

**Local A\_Port\_Name:** contains the local A\_Port\_Name of the described ASL.

**Adjacent A\_Port\_Name:** contains the Adjacent A\_Port\_Name of the described ASL.

**Link Cost:** contains the link cost of the described ASL in the two least significant bytes and two reserved bytes in the two most significant bytes.

#### 6.4.2.4 FCDF N\_Port\_IDs descriptor

The format of the FCDF N\_Port\_IDs descriptor is specified in table 169.

**Table 169 – FCDF N\_Port\_IDs descriptor format**

Item	Size (bytes)
Tag Value = 0013h	4
Length = variable	4
Distributed Switch Switch_Name	8
Virtual Domain_ID Value	4
Number of FCDF Allocation Records (n)	4
FCDF Allocation Record #1	see table 170
FCDF Allocation Record #2	see table 170
...	
FCDF Allocation Record #n	see table 170

**Distributed Switch Switch\_Name:** contains the Switch\_Name for the Distributed Switch, Switch\_Name associated with the Virtual Domain\_ID value.

**Virtual Domain\_ID Value:** contains the Virtual Domain\_ID for the Distributed Switch in the least significant byte and three reserved bytes in the three most significant bytes.

**Number of FCDF Allocation Records:** contains the number of FCDF Allocation Records that follow. The format of the FCDF Allocation Record is specified in table 170.

**Table 170 – FCDF Allocation Record format**

Item	Size (bytes)
FCDF Switch_Name	8
Number of Allocated N_Port_ID Ranges (s)	4
Allocated N_Port_ID Range #1	4
Allocated N_Port_ID Range #2	4
...	
Allocated N_Port_ID Range #s	4

**FCDF Switch\_Name:** contains the Switch\_Name of the FCDF whose N\_Port\_IDs allocation is provided.

**Number of Allocated N\_Port\_ID Ranges:** contains the number of Allocated N\_Port\_ID Range Entries that follow. This list of Allocated N\_Port\_ID Ranges identifies the N\_Port\_IDs allocated to the described FCDF. The N\_Port\_ID Range is defined by an N\_Port\_ID in the least significant three bytes, and by the number of bits defining the range in the most significant byte (e.g., the range 020200h .. 02027Fh is expressed as '07h || 020200h').

**6.4.2.5 RHello Interval descriptor**

The format of the RHello Interval descriptor is specified in table 171.

**Table 171 – RHello Interval descriptor format**

Item	Size (bytes)
Tag Value = 0014h	4
Length = 4	4
RHello_Interval	4

**RHello\_Interval:** contains the RHello\_Interval value expressed in ms.

**6.4.2.6 Controlling Switch Parameters descriptor**

The format of the Controlling Switch Parameters descriptor is specified in table 172.

**Table 172 – Controlling Switch Parameters descriptor format**

Item	Size (bytes)
Tag Value = 0015h	4
Length = variable	4
Primary Controlling Switch Priority    Switch_Name Record	12
Number of Controlling Switch Priority    Switch_Name Records (t)	4
Controlling Switch Priority    Switch_Name Record #1	12
Controlling Switch Priority    Switch_Name Record #2	12
...	
Controlling Switch Priority    Switch_Name Record #t	12

**Primary Controlling Switch Priority || Switch\_Name Record:** contains the Controlling Switch Priority || Switch\_Name Record of the Controlling Switch that the originating Controlling Switch believes is the Primary Controlling Switch.

**Number of Controlling Switch Priority || Switch\_Name Records:** contains the number of Controlling Switch Priority || Switch\_Name Records that follow.

The format of the Controlling Switch Priority || Switch\_Name Record is specified in table 173.

**Table 173 – Controlling Switch Priority || Switch\_Name Record format**

Item	Size (bytes)
Controlling Switch Priority	4
Controlling Switch Switch_Name	8

**Controlling Switch Priority:** contains the Priority of the Controlling Switch (see table 296) in the least significant byte and the three most significant bytes are reserved.

**Controlling Switch Switch\_Name:** contains the Switch\_Name of the Controlling Switch.

### 6.4.3 Controlling Switch redundancy protocol SW\_ILSs

#### 6.4.3.1 Exchange Redundancy Parameters (ERP)

In a Distributed Switch with a Controlling Switch Set containing two Switch\_Names the Exchange Redundancy Parameter (ERP) SW\_ILS is used to determine which Controlling Switch behaves as Primary Controlling Switch and which one behaves as Secondary Controlling Switch. In a Distributed Switch with a Controlling Switch Set containing more than two Switch\_Names the Exchange Redundancy Parameter (ERP) SW\_ILS is used by the Primary Controlling Switch to select the Secondary Controlling Switch.

#### ERP request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the ERP request Sequence payload is specified in table 174.

**Table 174 – ERP request payload**

Item	Size (bytes)
SW_ILS Code = A100 0001h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Controlling Switch State descriptor	see 6.4.2.2

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the requesting Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Controlling Switch State descriptor:** contains information about the Controlling Switch and the N\_Port\_IDs currently allocated by the Controlling Switch. See 6.4.2.2.

#### ERP reply Sequence

**SW\_RJT:** indicates the rejection of the ERP request Sequence.

**SW\_ACC:** indicates the acceptance of the ERP request Sequence. The format of the ERP SW\_ACC payload is specified in table 175.

**Table 175 – ERP SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Controlling Switch State descriptor	see 6.4.2.2

#### 6.4.3.2 Get FCDF Topology State (GFTS)

The Get FCDF Topology State (GFTS) SW\_ILS is used by the Secondary Controlling Switch to request to the Primary the current FCDF topology, in order to synchronize its state with the one of the Primary.

##### GFTS request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the GFTS request Sequence payload is specified in table 176.

**Table 176 – GFTS request payload**

Item	Size (bytes)
SW_ILS Code = A100 0002h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

##### GFTS reply Sequence

**SW\_RJT:** indicates the rejection of the GFTS request Sequence.

**SW\_ACC:** indicates the acceptance of the GFTS request Sequence. The format of the GFTS SW\_ACC payload is specified in table 177.

**Table 177 – GFTS SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
FCDF Topology descriptor	see 6.4.2.3

**FCDF Topology descriptor:** describes the current FCDF topology information. See 6.4.2.3.

### 6.4.3.3 Get FCDF N\_Port\_IDs State (GFNS)

The Get FDCF N\_Port\_IDs State (GFNS) SW\_ILS is used by the Secondary Controlling Switch to request to the Primary the Virtual Domain\_ID value and the current allocation of N\_Port\_IDs to each FCDF of the Distributed Switch, in order to synchronize its state with the one of the Primary.

#### GFNS request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the GFNS request Sequence payload is specified in table 178.

**Table 178 – GFNS request payload**

Item	Size (bytes)
SW_ILS Code = A100 0003h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

#### GFNS reply Sequence

**SW\_RJT:** indicates the rejection of the GFNS request Sequence.

**SW\_ACC:** indicates the acceptance of the GFNS request Sequence. The format of the GFNS SW\_ACC payload is specified in table 179.

**Table 179 – GFNS SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
FCDF N_Port_IDs descriptor	see 6.4.2.4

**FCDF N\_Port\_IDs descriptor:** contains information about the Virtual Domain and its allocated N\_Port\_IDs. See 6.4.2.4.

#### 6.4.3.4 Secondary Synchronization Achieved (SSA)

The Secondary Synchronization Achieved (SSA) SW\_ILS is used by the Secondary Controlling Switch to communicate to the Primary that it achieved state synchronization.

##### SSA request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the SSA request Sequence payload is specified in table 180.

**Table 180 – SSA request payload**

Item	Size (bytes)
SW_ILS Code = A100 0004h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

##### SSA reply Sequence

**SW\_RJT:** indicates the rejection of the SSA request Sequence.

**SW\_ACC:** indicates the acceptance of the SSA request Sequence. The format of the SSA SW\_ACC payload is specified in table 181.

**Table 181 – SSA SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

#### 6.4.3.5 Redundancy Hello (RHello)

The Redundancy Hello (RHello) SW\_ILS is used by the Controlling Switch redundancy protocol (see 17.5). The RHello SW\_ILS is transmitted in a unidirectional Exchange (i.e., it does not have a reply Sequence).

##### RHello request Sequence

**Addressing:** when used over an AISL, the S\_ID field shall be set to FFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFDh, indicating the destination VE\_Port. When used over an ASL, the S\_ID field shall be set to FFFFF9h, indicating the originating VA\_Port, and the D\_ID field shall be set to FFFFF9h, indicating the destination VA\_Port.

**Payload:** The format of the RHello request Sequence payload is specified in table 182.

**Table 182 – RHello request payload**

Item	Size (bytes)
SW_ILS Code = A100 0005h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
RHello Interval descriptor	see 6.4.2.5

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**RHello\_Interval descriptor:** contains the RHello interval value. See 6.4.2.5.

#### 6.4.3.6 Select Primary Controlling Switch (SPCS)

The Select Primary Controlling Switch (SPCS) SW\_ILS is used to initiate the Primary Controlling Switch selection process when the Controlling Switch Set contains more than two Switch\_Names.



### SPCS request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the SPCS request Sequence payload is specified in table 183.

**Table 183 – SPCS request payload**

Item	Size (bytes)
SW_ILS Code = A100 0006h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

### SPCS reply Sequence

**SW\_RJT:** indicates the rejection of the SPCS request Sequence.

**SW\_ACC:** indicates the acceptance of the SPCS request Sequence. The format of the SPCS SW\_ACC payload is specified in table 184.

**Table 184 – SPCS SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length = 0000 0000h	4

#### 6.4.3.7 Exchange Controlling Switch Parameters (ECSP)

The Exchange Controlling Switch Parameters (ECSP) SW\_ILS is used to select the Primary Controlling Switch if the Controlling Switch Set contains more than two Switch\_Names.

### ECSP request Sequence

**Addressing:** the S\_ID field shall be set to FFFFFFFDh, indicating the originating VE\_Port, and the D\_ID field shall be set to FFFFFFFDh, indicating the destination VE\_Port.

**Payload:** The format of the ECSP request Sequence payload is specified in table 185.

**Table 185 – ECSP request payload**

Item	Size (bytes)
SW_ILS Code = A100 0007h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Controlling Switch Parameters descriptor	see 6.4.2.6

**Destination Controlling Switch Switch\_Name:** contains the Switch\_Name of the destination Controlling Switch.

**Originating Controlling Switch Switch\_Name:** contains the Switch\_Name of the originating Controlling Switch.

**Descriptor List Length:** contains the length in bytes of the subsequent list of descriptors.

**Controlling Switch Parameters descriptor:** contains the Controlling Switch Priority || Switch\_Name Records known to the originator of the ECSP request. See 6.4.2.6.

**ECSP reply Sequence**

**SW\_RJT:** indicates the rejection of the ECSP request Sequence.

**SW\_ACC:** indicates the acceptance of the ECSP request Sequence. The format of the ECSP SW\_ACC payload is specified in table 186.

**Table 186 – ECSP SW\_ACC payload**

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Destination Controlling Switch Switch_Name	8
Originating Controlling Switch Switch_Name	8
Descriptor List Length	4
Controlling Switch Parameters descriptor	see 6.4.2.6

**Controlling Switch Parameters descriptor:** see 6.4.2.6

#### 6.4.4 Controlling Switch redundancy protocol timeouts

Timeouts associated to each Controlling Switch redundancy protocol SW\_ILS are specified in table 187.

**Table 187 – Controlling Switch redundancy protocol SW\_ILSs timeouts**

Description	Abbreviation	Timeout
Exchange Redundancy Parameter	ERP	1000 ms
Get FCDF Topology State	GFTS	1000 ms
Get FDCF N_Port_IDs State	GFNS	1000 ms
Secondary Synchronization Achieved	SSA	1000 ms
Select Primary Controlling Switch	SPCS	1000 ms
Exchange Controlling Switch Parameters	ECSP	1000 ms

## 7 Fabric Configuration

### 7.1 Fabric Configuration summary

The Fabric Configuration process enables a Switch port to determine its operating mode, exchange operating parameters, and provides for distribution of addresses. This process is summarized in table 188. See 17 for Distributed Switch operations.

**Table 188 – Fabric Configuration summary**

Operation	Starting Condition	Process	Ending Condition
Establish link parameters and Switch port operating mode	Switch port has achieved word synchronization.	The Switch port attempts to discover whether it is an FL_Port, an E_Port, an A_Port, or an F_Port.	Switch port mode is known. If a port is an E_Port, link parameters have been exchanged and Credit has been initialized.
Select Principal Switch	BF or RCF SW_ILS transmitted or received.	Switch_Names are exchanged over all ISLs to select a Principal Switch, the Principal Switch becomes the Domain Address Manager.	The Principal Switch is selected.
Domain_ID Acquisition	Domain Address Manager has been selected.	Switch requests a Domain_ID from the Domain Address Manager.	Switch has a Domain_ID.
Zoning Merge	Switch has a Domain_ID.	Zoning data are exchanged over the E_Ports, following the merge protocol defined in clause 10.	The Zoning definitions are consistent across the E_Ports.
Path selection	Switch has a Domain_ID.	Path selection (FSPF) is defined in clause 8.	Switch is operational with routes established.

Domain\_IDs may be assigned statically or dynamically. When Domain\_IDs are assigned statically, the administrator shall configure a Domain\_ID and a Fabric\_Name on each Switch of the Fabric, and the operations described in 7.3 and 7.4 shall not be performed by a Switch. A configured Fabric\_Name shall conform to the rules regarding Name\_Identifiers specified in FC-FS-5. When Domain\_IDs are assigned dynamically, the operations described in 7.3 and 7.4 shall be performed by a Switch.

NOTE 8 – An erroneous condition in which two Switches have been assigned the same Domain\_ID may be detected when FSPF begins its operations.

Once path selection has completed, routes for Class N Frames are established and Class N Frames shall traverse the Fabric using established routes. Class N Frames shall continue to traverse the Fabric until an RCF clears the routes or a previously determined route is invalidated (e.g., Max\_Age expires for an LSR, physical link is removed).

## **7.2 Switch port initialization**

### **7.2.1 Basic operation**

Switch ports shall initialize as described below. Figure 16 and figure 17 show the state machine of the process. If the state machine is different than the text, the state machine shall apply. A Switch port that is running this state machine shall be capable of either E/F/FL\_Port, E/A/F/FL\_Port, E/F\_Port, E/A/F\_Port, E/FL\_Port, E/A/FL\_Port, E/A\_Port, E\_Port, or A\_Port operation. Initialization of Switch ports that are F/FL\_Port, FL\_Port, or F\_Port operation is specified in FC-FS-5 and FC-AL-2. This state machine is also applicable to B\_Port operation.

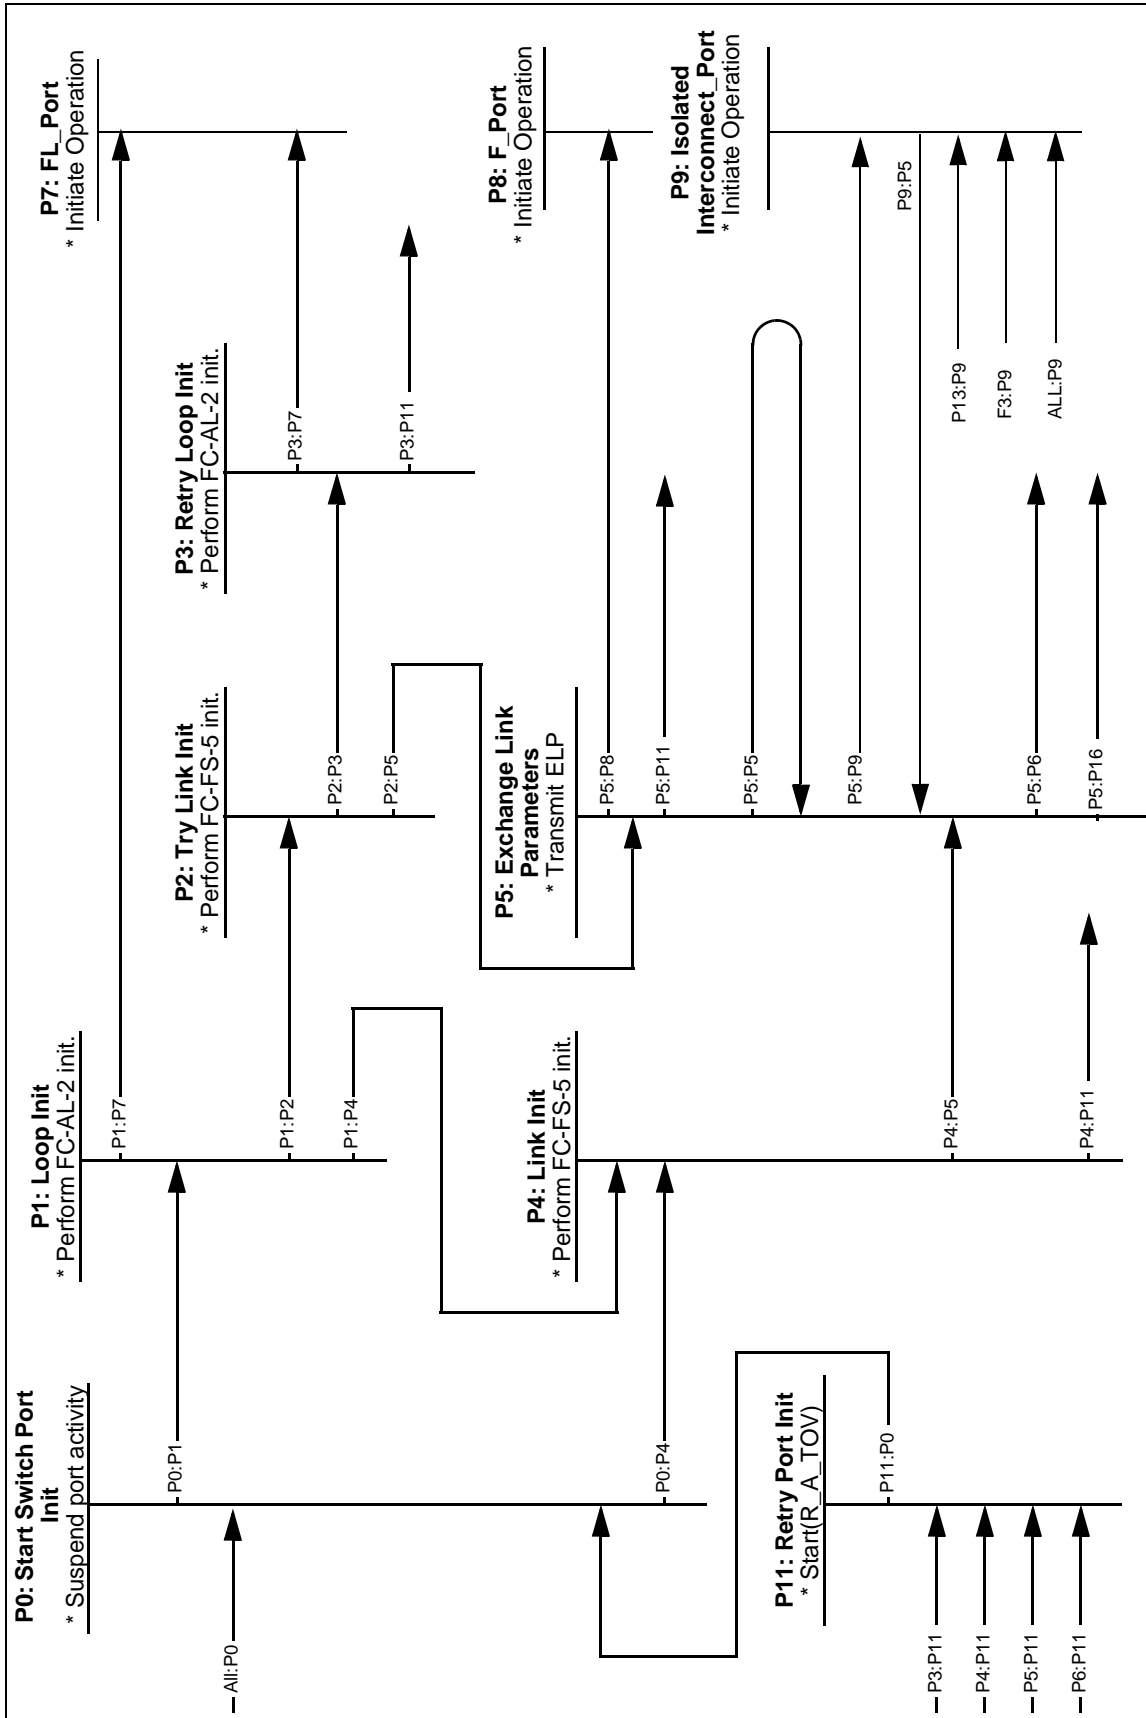


Figure 16 – Switch port mode initialization state machine

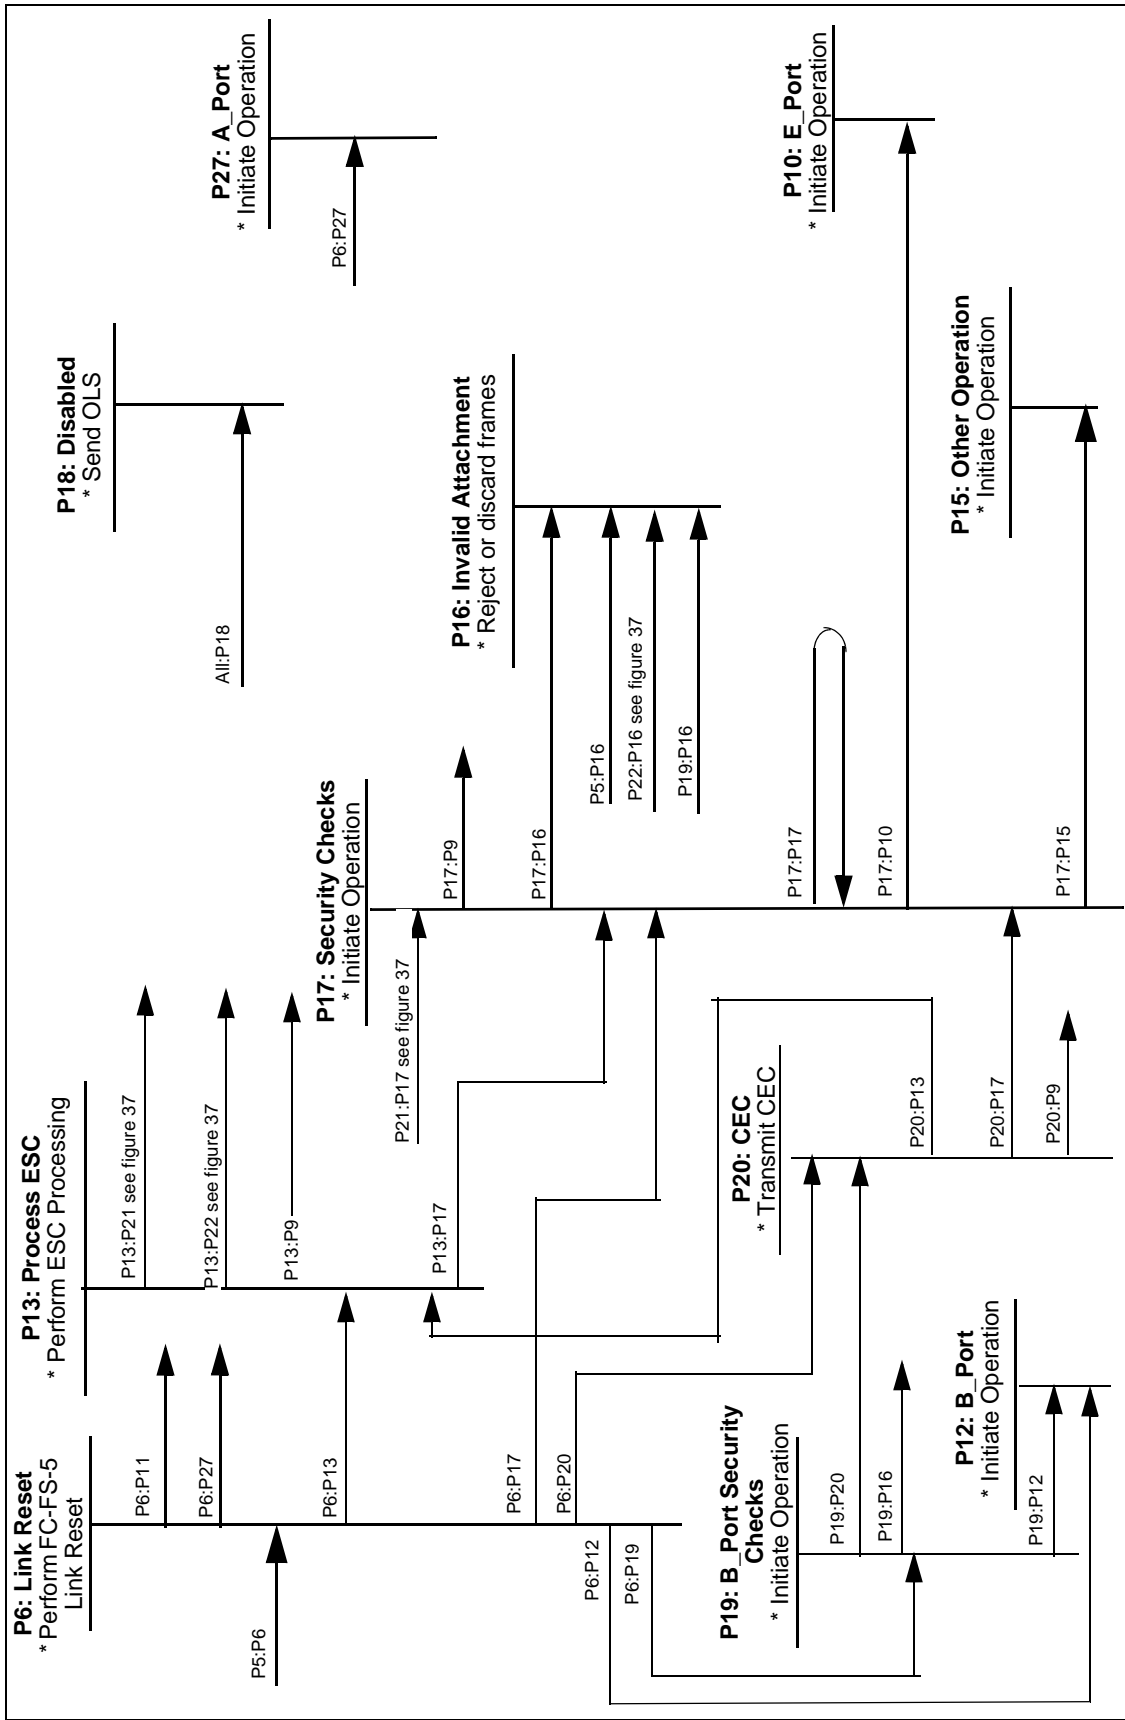


Figure 17 – Switch port mode initialization state machine - continued

**Transition All:P0.** This transition occurs whenever 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) receiving an initialization Primitive Sequence, such as OLS, NOS, LIP;
- c) outside intervention requesting an initialization;
- d) a transition to Link Offline, as defined in FC-FS-5;
- e) a loss of word synchronization for greater than R\_T\_TOV; or
- f) a failure to successfully complete a prior initialization attempt, and the timeout period has expired.

LR is not considered an Initialization Event, but shall operate as specified in FC-FS-5.

**State P0: Start Switch Port Initialization.** This state marks the beginning of Switch port initialization. All activity on the Switch port is suspended until the Initialization is complete.

**Transition P0:P1.** The Switch port is capable of becoming an FL\_Port. Attempt Loop Initialization first (as defined in FC-AL-2).

**Transition P0:P4.** The Switch port is not capable of becoming an FL\_Port. Attempt Link Initialization.

**State P1: Loop Initialization.** An FL\_Port-capable Switch port attempts Loop Initialization (as defined in FC-AL-2).

**Transition P1:P7.** This transition occurs if the FL\_Port transitions from the OPEN\_INIT state to the MONITORING state, is in participating mode, and the resulting AL\_PA bitmap generated during the LISA Loop Initialization Sequence indicates that one or more L\_Port (other than the Switch port) is attached. This transition also occurs if Switch port is in non-participating mode.

**Transition P1:P2.** This transition occurs if the FL\_Port transitions from the OPEN\_INIT state to the MONITORING state, is in participating mode, and the resulting AL\_PA bitmap generated during the LISA Loop Initialization Sequence indicates zero or one L\_Port (other than the Switch port) is attached; or, if the Loop Initialization procedure did not complete and OLS or NOS is received (see annex A).

**Transition P1:P4.** This transition occurs if the Loop Initialization does not complete successfully. This may occur if the Switch port is attached to a non-L\_Port capable port, so the next thing to try is a Link Initialization.

**State P2: Try Link Initialization.** The Switch port is FL\_Port-capable, is in participating mode, and has detected zero attached NL\_Ports, then there is a possibility that the Switch port is point-to-point attached to another FL\_Port-capable Switch port. In this case the Switch port shall attempt Link Initialization by transmitting LIPs and, when receiving LIPs, OLSs for up to 2xAL\_TIME until receiving NOS or LR (see annex A), and then complete Link Initialization as defined in FC-FS-5. Otherwise the Switch port shall complete the Link Initialization protocol initiated by the other Switch port.

**Transition P2:P3.** This transition occurs if the Link Initialization does not complete successfully.

**Transition P2:P5.** This transition occurs if the Link Initialization completes successfully.

**State P3: Retry Loop Initialization.** The Switch port had detected that it may be able to operate point-to-point with another loop device, but the attempt to do so failed. In this case, the Switch port shall then attempt to go back to loop operation by retrying Loop Initialization (as defined in FC-AL-2).



**Transition P3:P7.** This transition occurs if the Loop Initialization succeeds (the FL\_Port transitions from the OPEN\_INIT state to the MONITORING state and participating).

**Transition P3:P11.** This transition occurs if the Loop Initialization fails following a re-attempt of Loop Initialization.

**State P4: Link Initialization.** The Switch port shall attempt Link Initialization as defined in FC-FS-5.

**Transition P4:P5.** This transition occurs if the Link Initialization procedure succeeds.

**Transition P4:P11.** This transition occurs if the Link Initialization procedure fails.

**State P5: Exchange Link Parameters.** The Switch port shall originate an ELP SW\_ILS request Sequence (see 6.2.4). Table 189 below defines the responses and actions to an ELP request for the originating Interconnect\_Port.

**Table 189 – Responses to ELP request for originating Interconnect\_Port (Part 1 of 2)**

Response to ELP	Indication	Originating Interconnect_Port Action
1. R_RDY	Request received at destination	Wait E_D_TOV+4 for response frame
2. ACK_1	Request received at destination	Wait E_D_TOV+4 for response frame
3. SW_ACC	Destination Interconnect_Port received and processed request	Send ACK_1, Transition (P5:P6)
4. F_BSY or P_BSY	Destination is busy	Retry <sup>a</sup> , Transition (P5:P11)
5. F_RJT or P_RJT	The frame is not acceptable	Respond accordingly <sup>c</sup> , Transition (P5:P11) if appropriate
6. ELP (rcvd Switch_Name > own Switch_Name)	Both Interconnect_Ports sent ELP at the same time	Send SW_ACC or SW_RJT based on the values of the ELP parameters, Transition (P5:P6) (see figure 18 for an example of this response)
<p><sup>a</sup> The retry is performed following a timeout period, as defined in P11 below.</p> <p><sup>b</sup> The reason code shall be “Unable to perform command request” with a reason code explanation of “Command already in progress”.</p> <p><sup>c</sup> Response is defined in FC-FS-5.</p> <p><sup>d</sup> An SW_ACC is sent for the other ELP Exchange in progress, as described in Response #6, if the received ELP parameters are acceptable. If the received ELP parameters are not acceptable, then an SW_RJT is sent. See figure 18.</p>		

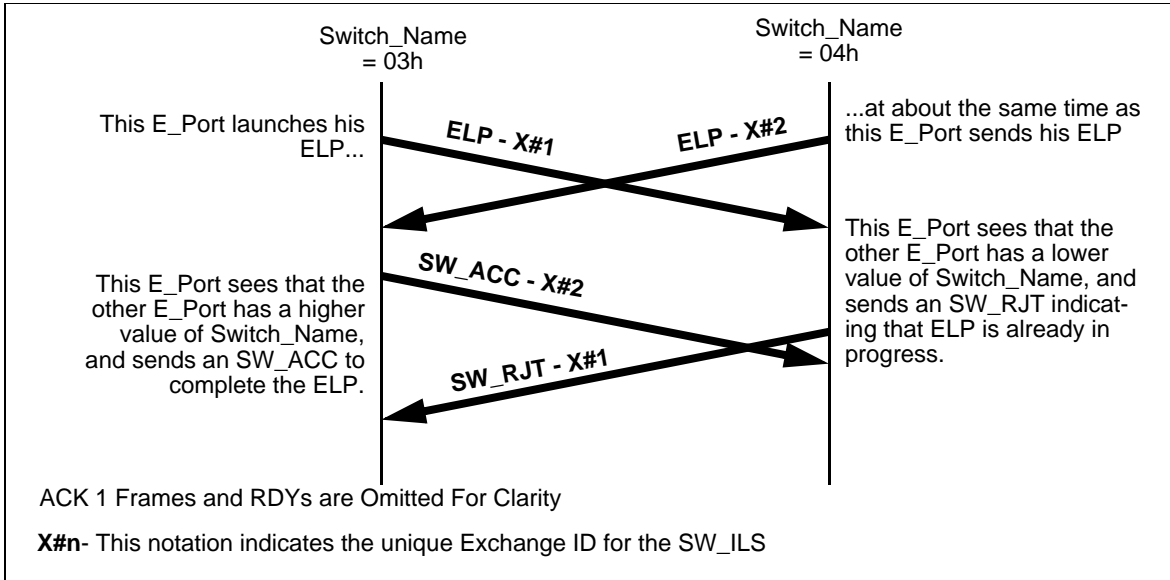
**Table 189 – Responses to ELP request for originating Interconnect\_Port (Part 2 of 2)**

Response to ELP	Indication	Originating Interconnect_Port Action
7. ELP (rcvd Switch_Name < own Switch_Name)	Both Interconnect_Ports sent ELP at the same time	Send SW_RJT <sup>b</sup> , (see figure 18 for an example of this response)
8. ELP (rcvd Switch_Name = own Switch_Name)	Interconnect_Port output is looped back to input	Remove loopback condition, Transition (P5:P9)
9. SW_RJT	Reason code/explanation: - Command already in progress <sup>d</sup>  - Logical busy  - other	(see figure 18 for an example of this response) - retry transition to P11 <sup>a</sup> , or P5 - respond accordingly, and transition to P11 if appropriate
10. FLOGI	Destination is a PN_Port	Respond accordingly <sup>c</sup> , transition to P8
11. any other frame	Indeterminate	Discard frame and retry <sup>a</sup> , transition to P11
12. E_D_TOV+4 expires	Destination is busy; or, ELP, SW_ACC, ACK_1 frame lost; or, destination is not an Interconnect_Port	Retry <sup>a</sup> , transition to P11
<p><sup>a</sup> The retry is performed following a timeout period, as defined in P11 below.</p> <p><sup>b</sup> The reason code shall be “Unable to perform command request” with a reason code explanation of “Command already in progress”.</p> <p><sup>c</sup> Response is defined in FC-FS-5.</p> <p><sup>d</sup> An SW_ACC is sent for the other ELP Exchange in progress, as described in Response #6, if the received ELP parameters are acceptable. If the received ELP parameters are not acceptable, then an SW_RJT is sent. See figure 18.</p>		

The originating Interconnect\_Port shall consider the exchange of link parameters complete, but not necessarily successful, when it has received the SW\_ACC or SW\_RJT and has transmitted the ACK\_1 for the SW\_ACC or SW\_RJT reply Sequence.

The responding Interconnect\_Port shall consider the exchange of link parameters complete when it has received the ACK\_1 for the SW\_ACC or SW\_RJT.

The exchange of link parameters shall be considered successful when the exchange of link parameters is complete, and the reply to the ELP is an SW\_ACC, and both Interconnect\_Ports agree that the parameters exchanged are acceptable.



**Figure 18 – Simultaneous ELP processing- parameters acceptable to both Switches**

**Transition P5:P5.** This transition occurs if the originating Interconnect\_Port does not agree that the parameters in the SW\_ACC are acceptable, or it receives an SW\_RJT indicating the parameters in the ELP request were not acceptable to the responding Interconnect\_Port, and it is able to originate a new ELP request Sequence with modified parameters. This transition may also occur if an SW\_RJT is received indicating a logical busy.

**Transition P5:P6.** This transition is taken by the originator of the ELP if the exchange of link parameters are complete.

**Transition P5:P8.** This transition occurs if the exchange of link parameters is unable to be completed, and FLOGI is received.

**Transition P5:P9.** This transition occurs if the originating Interconnect\_Port does not agree that the parameters in the SW\_ACC are acceptable, or it receives an SW\_RJT indicating the parameters in the ELP request were not acceptable to the responding Interconnect\_Port, and it is not able to originate a new ELP request Sequence with modified parameters (see 7.6).

**Transition P5:P11.** This transition occurs if the ELP is rejected with “unable to perform command request”, and no FLOGI is received. The Switch port performs the Link Offline protocol as defined in FC-FS-5 during the transition.

**Transition P5:P16.** This transition is taken when authorization checks that are based on data from the ELP fail.

**State P6: Link Reset.** When the exchange of link parameters has completed successfully, the value of buffer-to-buffer and end-to-end Class F Credit are initialized. In order to initialize the flow control parameters, the Switch port that originated the successful ELP SW\_ILS shall attempt the Link Reset protocol as defined in FC-FS-5.

NOTE 9 – The re-initialization of link credit is necessary since the flow control parameters in the ELP payload are intended to communicate link credit parameters for a specific credit model. The Link Reset is the common method defined by FC-FS-5 for establishing a known credit state.

**Transition P6:P11.** This transition occurs if the Link Reset fails.

**Transition P6:P12.** This transition occurs if the Link Reset is successful, the port is a B\_Port and no security checks are required.

**Transition P6:P13.** This transition occurs if the Link Reset is successful, the port determined in state P5 that it has to operate as an E\_Port, and ESC is supported.

**Transition P6:P17.** This transition occurs if the Link Reset is successful, the port determined in state P5 that it has to operate as an E\_Port, and ESC is not supported.

**Transition P6:P27.** This transition occurs if the Link Reset is successful and the port determined in state P5 that it has to operate as an A\_Port (see 17.3).

**Transition P6:P19.** Occurs for an E\_Port when the E\_Port detected, during the ELP processing in state P5, it is connected to a B\_Port, and the E\_Port requires B\_Port authentication. This transition occurs also for a B\_Port when the B\_Port supports authentication.

**Transition P6:P20.** Occurs for an E\_Port when the E\_Port detected, during the ELP processing in state P5, it is connected to a B\_Port, and the E\_Port does not require B\_Port authentication.

**State P7: Operate as an FL\_Port.** The Switch port has detected a functional Arbitrated Loop. The Switch port shall continue to operate as an FL\_Port until the next Initialization Event. If a Switch port enters the state in the non-participating mode, it shall remain in the non-participating mode until the next initialization event.

**State P8: Operate as an F\_Port.** The Switch port has detected an attached PN\_Port. The Switch port shall continue to operate as an F\_Port until the next Initialization Event.

**Transition All:P9.** This transition occurs whenever an Interconnect\_Port receives an SW\_RJT with a reason code explanation of "E\_Port is Isolated".

**State P9: Operate as an Isolated Interconnect\_Port.** The Interconnect\_Port shall become Isolated and not continue with Fabric Configuration as specified in 7.6. The Switch port shall continue to operate as an isolated Interconnect\_Port until the next Initialization Event.

**Transition P9:P5.** This transition occurs when an ELP is received by an Isolated Interconnect\_Port (see 7.6).

**State P10: Initialize as an E\_Port.** The Switch port has completed the exchange of link parameters with another E\_Port. If the link parameters exchanged were acceptable, then the E\_Port shall participate in the next phase of Fabric Configuration, described in 7.3. The Switch port shall continue to operate as an E\_Port until the next Initialization Event.

**State P11: Retry Switch Port Initialization.** The Switch port shall wait for R\_A\_TOV before retrying Switch port initialization. If the Switch port detects an Initialization Event during the timeout period, it shall not wait for the timeout period to expire.

**State P12: Operate as a B\_Port.** ELPs have been exchanged, the link reset is successful, security checks have been performed, and the port is operating as a B\_Port. Any further normal fabric configuration or routing operations are transparent to this port.

**State P13: Send ESC.** The link reset has been successful and ESC is supported. Information exchanged using ESC shall be carried through P17. The port shall perform ESC processing as described in 7.2.2.

**Transition P13:P9.** This transition occurs because of an ESC reject with reason code: “unable to perform command request”.

**Transition P11:P0.** This transition occurs if the R\_A\_TOV timeout period has expired.

**State P15: Other Operation.** The port operates in a mode other than FSPF.

**State P16: Invalid Attachment.** The port operates in Invalid Attachment mode and SW\_ILSs shall be rejected with an Invalid Attachment SW\_RJT reason code with the following exceptions. FSPF SW\_ILSs (e.g., HLO, LSU and LSA) shall be discarded and ACKs shall be sent upon receipt. Distributed Service CT\_IUs shall be rejected with an F\_RJT with a reason code of “Invalid Attachment”. Class N service frames shall be discarded and rejects shall be sent as appropriate to each Class of Service (see FC-FS-5). To leave this port state, the port shall receive OLS.

**State P17: Security Checks.** The port initiates and responds to all required security checks, if any, while in this state. If the port receives an EFP before security checks are complete, then the port shall respond with an SW\_RJT with a Logical busy SW\_RJT reason code and a SW\_RJT reason code explanation of Security Checks in Progress. The order and protocol of the security checks is defined in Fibre Channel Security Protocols (see FC-SP-2). Switch\_Name usage shall abide by the rules defined in FC-SP-2.

**Transition P17:P9.** This transition occurs when a required Policy or FC-SP Zoning check (see FC-SP-2) fails or is rejected.

**Transition P17:P10.** This transition occurs when all required security checks are successful and the port is to operate as an E\_Port. The port is to operate as an E\_Port if FSPF is the agreed upon path selection mechanism per the prior ESC exchange, or the prior ESC command was rejected with the reason code “command not supported”, or the port does not support ESC.

**Transition P17:P15.** This transition occurs if a routing protocol other than FSPF is agreed to in the prior ESC exchange.

**Transition P17:P16** This transition occurs when a required authentication or authorization check (see FC-SP-2) fails or is rejected.

**Transition P17:P17.** This transition occurs each time a required security check is successful.

**State P18: Disabled.** While in this state, the port transmits the offline sequence until either, a power-on reset condition occurs or outside intervention requests an initialization of the port.

**Transition All:P18.** The transition to this state occurs when the Switch determines that a model dependent threshold has been exceeded.

**State P19: B\_Port Security Checks.** While in this state an E\_Port shall authenticate the B\_Port by initiating a B\_AUTH\_ILS Authentication transaction (see FC-SP-2). While in this state a B\_Port shall respond to the B\_AUTH\_ILS Authentication transaction.

**State P20: CEC.** While in this state an E\_Port shall originate a CEC SW\_ILS request Sequence. The processing is as specified for the ELP SW\_ILS, except for configuration of flow control that is performed in state P5. The E\_Port that sent the CEC message with the numerically higher

Switch\_Name shall become the CEC Initiator, while the E\_Port that sent the CEC message with the numerically lower Switch\_Name shall become the CEC Responder. The CEC SW\_ILS is propagated by B\_Ports.

**Transition P19:P20.** Occurs for an E\_Port when the Authentication transaction performed in state P19 completes successfully.

**Transition P19:P12.** Occurs for a B\_Port when the Authentication transaction performed in state P19 completes successfully.

**Transition P19:P16.** Occurs when the Authentication transaction performed in state P19 fails.

**Transition P20:P13.** Occurs when the CEC Exchange performed in state P20 is successful and ESC is supported. This transition may occur also when the remote Switch port does not support CEC.

**Transition P20:P17.** Occurs when the CEC Exchange performed in state P20 is successful and ESC is not supported. This transition may occur also when the remote Switch port does not support CEC.

**Transition P20:P9.** Occurs when the CEC Exchange performed in state P20 is not successful. This transition may occur also when the remote Switch port does not support CEC.

**State P27: Initialize as an A\_Port.** The Switch port has completed the exchange of link parameters with another A\_Port. If the link parameters exchanged were acceptable, then the A\_Port shall participate in the Distributed Switch operations specified in 17.4. The Switch port shall continue to operate as an A\_Port until the next Initialization Event.

When an Inter-Switch Link is established the Switch shall request EFP and enter state F2.

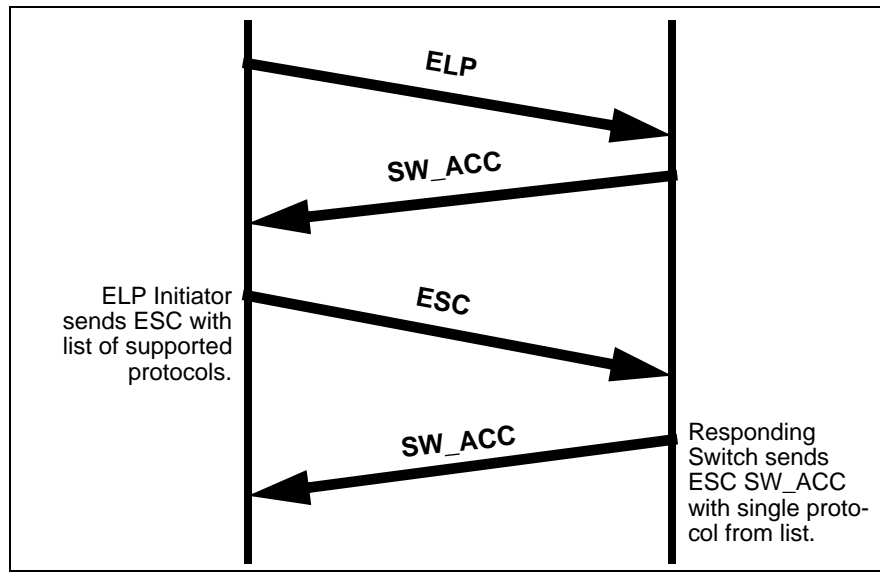
### 7.2.2 Switch\_Name usage

The Switch\_Name presented during ELP Processing in state P5 shall be identical to that used when the Switch port initializes as an E\_Port in state P10 and for any subsequent operation or protocol.

### 7.2.3 Exchange Switch Capabilities processing

Figure 19 shows a typical exchange involving the ESC SW\_ILS. In this case, the ELP Initiator (the Switch that receives the ELP SW\_ACC) initiates the ESC SW\_ILS. Contained within the payload of the ESC is a list of supported Switch-to-Switch protocols. The receiver of the ESC determines the

protocol it shall use from the list presented, and responds with that protocol in the payload of the ESC SW\_ACC.



**Figure 19 – ESC processing**

More formally, the process of exchanging Switch-to-Switch protocol capabilities shall progress as follows:

- the ELP initiator originates the ESC SW\_ILS. In the case of simultaneous ELPs from both Switches, the Switch that receives the ELP SW\_ACC shall be considered the ELP initiator. The payload of the ESC contains a list of protocols supported by the sending Switch;
- the responding Switch shall wait for a maximum of R\_A\_TOV to receive the ESC SW\_ILS request. After this time, it shall proceed in the port initialization process as if the ELP initiator does not support ESC. The responding Switch shall also proceed in the port initialization process if it receives other messages from the ELP initiator, and shall reply accordingly;
- if the receiving Switch does not support the ESC SW\_ILS, it responds with a SW\_RJT and a reason code of “Command not supported”. If the receiving Switch does support the ESC SW\_ILS, continue to the next step;
- if the receiving Switch does not support any of the protocols listed in the ESC SW\_ILS, it responds with a SW\_RJT and a reason code of “Unable to perform command request”. If the receiving Switch does support one of the protocols listed, continue to the next step; and
- the receiving Switch chooses a single protocol from the list presented in the ESC SW\_ILS and responds with this protocol in the payload of the ESC SW\_ACC.

#### **7.2.4 B\_Port impact on ESC processing**

When no B\_Ports exist between two E\_Ports, ESC processing is initiated by the ELP Initiator. This works when two E\_Ports are directly connected, but does not work when there are B\_Ports between two E\_Ports (e.g., the ELP Initiator could be a B\_Port). To accommodate the presence of B\_Ports between two E\_Ports, the ESC processing shall occur as follows:

- if two E\_Ports are directly connected, the ESC processing shall be initiated by the ELP Initiator; or
- if two E\_Ports are connected through B\_Ports, the ESC processing shall be initiated by the CEC Initiator.

### 7.2.5 Extensions to support Virtual Fabrics

The Switch port mode initialization state machine is extended to support Virtual Fabrics on the Switch. These extensions are described in 12.

### 7.3 Principal Switch Selection

If Domain\_IDs are assigned dynamically, a Principal Switch shall be selected whenever at least one Inter-Switch Link is established. The selection process chooses a Principal Switch, that is then designated as the Domain Address Manager. Figure 20 shows the state machine of the process. The recommended uses of BF and RCF are summarized in table 190.

**Table 190 – Recommended BF and RCF usage summary**

Event	BF or RCF
A Principal ISL experiences Link Failure or a transition to Offline or Isolated State	BF <sup>a</sup>
A configured Fabric is joined to another configured Fabric, and their Domain_IDs do not overlap	BF
An unconfigured Switch or Fabric is joined to a configured Fabric	neither (see figure 22)
A configured Fabric is joined to another configured Fabric, and an overlap in Domain_ID is detected	Isolate or RCF Originated by Management
Reconfiguration caused by BF fails for any reason	Isolate or RCF Originated by Management
<sup>a</sup> In lieu of BF, a Switch may attempt Principal ISL Recovery as described in 7.5.	

Non-disruptive reconfiguration of Fabrics (BF) requires that Domain\_IDs do not overlap. To ensure that the Switches being joined do not have a Domain\_ID overlap, an EFP shall be exchanged prior to either Switch issuing a Build Fabric request.






Figure 20 – Principal Switch selection state machine

**State F0: Non-disruptive.** A Switch may request a Fabric Reconfiguration by transmitting a BF on all E\_Ports that have completed Switch port initialization. Unless warranted by current conditions, a Switch shall always first attempt a non-disruptive Fabric Reconfiguration by sending a BF. If the Switch is attempting a non-disruptive Fabric Reconfiguration, the Switch shall transmit a BF to all neighbor Switches on an E\_Port that has completed Switch port initialization, and from which the Switch has not yet received a BF request. The Switch may transmit a BF on all E\_Ports that have completed Switch port initialization, and from which the Switch has not yet received a BF request.

While in this state:

- a) the Switch shall accept any BF received on any E\_Port, and shall not transmit a BF on any E\_Port from which a BF has been received;
- b) if an E\_Port from a previously unconnected neighbor completes Switch port initialization, the Switch shall transmit a BF on that E\_Port unless it has already received a BF on that E\_Port since Switch port initialization completed; and
- c) any received EFP, DIA, RDI SW\_ILSs shall result in the origination of an SW\_RJT response with a reason code of "Logical busy".

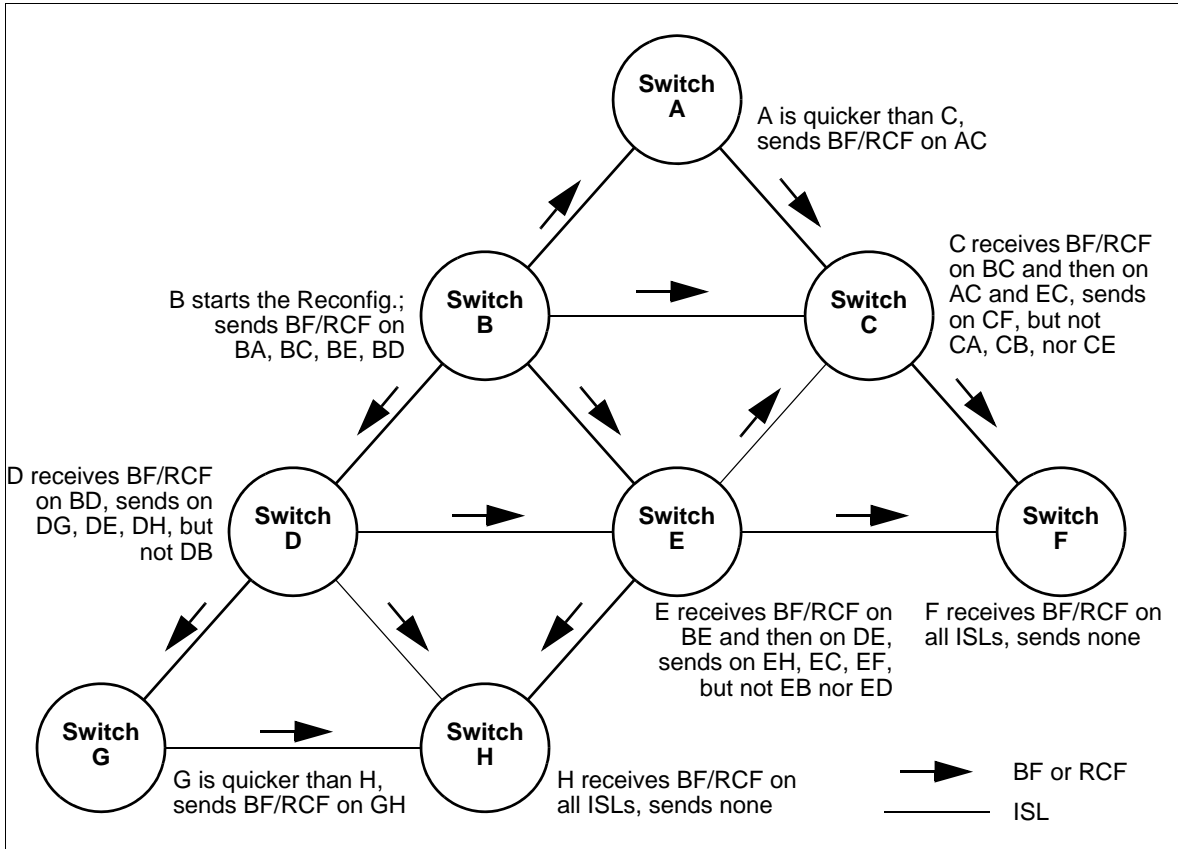
Figure 21 provides an example flow for BF requests.

**Transition All:F0.** This transition enters the state machine performing a non-disruptive Fabric Reconfiguration. This transition occurs when the Switch originates a BF, or when it receives a BF, or when the Switch receives a EFP where the received Domain\_ID\_List is non-zero, the retained Domain\_ID\_list is non-zero, the Domain\_IDs do not overlap, and the received Switch\_Priority||Switch\_Name and the retained Switch\_Priority||Switch\_Name are not the same. In addition, F\_S\_TOV shall be started when the first BF is received or when the Switch initiates non-disruptive Fabric Configuration.

**Transition F0:F0.** Occurs when a EFP, DIA, or RDI SW\_ILS is received. An SW\_RJT specifying a reason code of "logical busy" is originated.

**Transition F0:F1.** If a Switch receives and accepts an RCF request Sequence while it is in the process of attempting a non-disruptive Fabric Reconfiguration, it shall stop the non-disruptive Fabric Reconfiguration and begin processing RCF requests. Any Active or Open BF Sequences shall be abnormally terminated. In addition, F\_S\_TOV shall be started when the first RCF is accepted or when the Switch initiates disruptive Fabric Configuration.

**Transition F0:F2.** The Switch shall wait for F\_S\_TOV following the reception or origination of the first BF before originating or responding to an EFP request Sequence. At the start of a non-disruptive Fabric Reconfiguration (BF), the Domain\_ID\_List shall be empty ("zero Domain\_ID\_List"). During Fabric reconfiguration, the Switch shall retain a Switch\_Priority||Switch\_Name value that it believes is the lowest in the Fabric. This value shall be initialized at the start of Fabric Reconfiguration (caused by BF or RCF) to the Switch's value of Switch\_Priority||Switch\_Name. After the Switch is configured, it shall retain as the lowest value the Switch\_Priority||Switch\_Name of the Principal Switch.



**Figure 21 – Example Propagation of BF and RCF SW\_ILS requests**

**State F1: Disruptive.** The Switch is attempting a disruptive Fabric Reconfiguration, either originating or accepting an RCF.

Entering in this state:

- a) the Switch shall transmit an ELP on all Isolated Interconnect\_Ports and an RCF to all neighbor Switches on an E\_Port that has completed Switch port initialization, and from which the Switch has not yet received an RCF. The Switch may transmit a RCF on all E\_Ports that have completed Switch port initialization, and from which the Switch has not yet received a RCF request;
- b) any lock due to an ACA request shall be released; and
- c) the FSPF Link State Database and the associated initial message number counter shall be cleared.

NOTE 10 – The RCF processing may make the local copies of some databases related to Distributed Services out of date.

While in this state:

- a) the Switch shall respond to any RCF received on any E\_Port, and shall not transmit an RCF on any E\_Port from which an RCF has been received;
- b) if an E\_Port completes Switch port initialization, the Switch shall transmit a RCF on that E\_Port unless it has already received a RCF on that E\_Port since Switch port initialization completed;

- c) any received SW\_ILS shall result in the origination of an SW\_RJT response with a reason code of "Logical busy" except the SW\_ACC, SW\_RJT, ELP, ESC, RCF, HLO, LSU, and LSA SW\_ILSs;
- d) SW\_ILSs shall not be sent except the SW\_ACC, SW\_RJT, and RCF SW\_ILSs;
- e) the HLO, LSU and LSA SW\_ILSs shall be ignored on reception and shall not be sent;
- f) any received Class F CT frame related to Distributed Services (i.e., Type = 20h) shall result in the origination of an F\_RJT response with a reason code of "Nx\_Port not available, temporary"; and
- g) Class F CT frames related to Distributed Services (i.e., Type = 20h) shall not be sent.

Upon exiting from this state until a new domain ID is granted to the Switch (i.e., states D0 or A1):

- a) any received SW\_ILS shall result in the origination of an SW\_RJT response with a reason code of "Logical busy" except the SW\_ACC, SW\_RJT, ELP, ESC, RCF, EFP, DIA, RDI, HLO, LSU, and LSA SW\_ILSs;
- b) SW\_ILSs shall not be sent except the SW\_ACC, SW\_RJT, EFP, DIA, RDI, and RCF SW\_ILSs;
- c) the HLO, LSU and LSA SW\_ILSs shall be ignored on reception and shall not be sent;
- d) any received Class F CT frame related to Distributed Services (i.e., Type = 20h) shall result in the origination of an F\_RJT response with a reason code of "Nx\_Port not available, temporary"; and
- e) Class F CT frames related to Distributed Services (i.e., Type = 20h) shall not be sent.

Figure 21 shows an example diagram of the process to illustrate the flow of the RCF requests.

**Transition All:F1.** This transition enters the state machine performing a disruptive Fabric Reconfiguration. In this case, "All" refers to all Fx states other than F0. This transition occurs when the Switch originates an RCF, or when it receives and accepts an RCF request Sequence. In addition, F\_S\_TOV shall be started when the first RCF is received or when the Switch initiates disruptive Fabric Configuration.

**Transition F1:F1.** This transition occurs when any SW\_ILS and any Class F CT frame related to Distributed Services (i.e., Type = 20h) is received, except the SW\_ACC, SW\_RJT, ELP, ESC, RCF, HLO, LSU, and LSA SW\_ILSs. An SW\_RJT specifying a reason code of "logical busy" is originated.

**Transition F1:F2.** The Switch shall wait for F\_S\_TOV following the acceptance or origination of the first RCF before originating or responding to an EFP request Sequence. At the start of a disruptive Fabric Reconfiguration (RCF), the Domain\_ID\_List shall be empty ("zero Domain\_ID\_List"). The Switch shall retain a Switch\_Priority||Switch\_Name value that it believes is the lowest in the Fabric. This value shall be initialized at the start of Fabric Reconfiguration caused by RCF to the Switch's value of Switch\_Priority||Switch\_Name. After the Switch is configured, it shall retain as the lowest value the Switch\_Priority||Switch\_Name of the Principal Switch.

**State F2: EFP Idle.** The Switch shall remain in this state until it receives an EFP or DIA frame, or the 2xF\_S\_TOV timer expires and one of the following is true:

- a) the retained Switch\_Priority||Switch\_Name equals the Switch\_Priority||Switch\_Name of the Switch; or
- b) the retained Switch\_Priority is FFh.

In this state the Switch processes and generates EFP requests as required by the rules defined in state F3:EFP Send.

**Transition All:F2.** A Switch that is not yet configured (e.g., after initial power-on and exchange of ELPs) shall transmit an EFP SW\_ILS to all initialized E\_Ports to determine if the Switch is attached to

a configured Fabric (note that the Switch shall transition to the appropriate state and process any received BF or RCF requests as described above, as required by All:F0 and All:F1). When the first ISL to an Adjacent Switch becomes operational the Switch shall transmit an EFP on that link to determine the configuration of the Fabric that it is joining. On other ISLs the Switch may transmit an EFP. "All" in this case does not include F1:F2.

**Transition F2:F3.** When the Switch receives an EFP, or if it has not yet sent an EFP, or responded to an EFP since the reconfiguration started, it shall transition.

**Transition F2:D0.** If the retained value of Switch\_Priority||Switch\_Name does not change for twice F\_S\_TOV, and if the retained value of the Switch\_Priority||Switch\_Name is equal to the value of the Switch, then the Switch has become the Principal Switch.

**Transition F2:A0.** If the Switch receives a DIA request Sequence from the upstream Switch, then a Principal Switch has been selected. The Switch shall request a Domain\_ID as described in 7.4.

**Transition F2:P9.** If the retained value of Switch\_Priority||Switch\_Name does not change for twice F\_S\_TOV, and if the retained value of Switch\_Priority is equal to FFh, then there is no Switch capable of becoming a Principal Switch. The Switch shall cause all E\_Ports to become Isolated, as described in 7.6.

**State F3: EFP Send.** The Switch shall process all EFP payloads based on the information available at the time of processing. A Switch may receive an EFP payload either by receiving an EFP request Sequence at an E\_Port, or by receiving at an E\_Port an SW\_ACC reply Sequence in response to an EFP request Sequence. EFP Send actions shall be as follows:

- a) the Switch shall communicate its retained Switch\_Priority||Switch\_Name to neighbor Switches that it has not yet communicated that value. The Switch shall accomplish this either by originating a new EFP request Sequence, or by an SW\_ACC reply Sequence to a received EFP request;
- b) if the Switch receives in an EFP payload a non-zero Domain\_ID\_List (the list contains one or more records) and the Switch has a zero Domain\_ID\_List, then the Switch shall retain the received Switch\_Priority||Switch\_Name as the new value, and the received Domain\_ID\_List. The Switch shall also note from which neighbor Switch it received the new value, for potential use as the upstream Principal ISL during address distribution;
- c) if the Switch receives in an EFP payload a zero Domain\_ID\_List and the Switch has a non-zero Domain\_ID\_List, the Switch shall retain its current lowest Switch\_Priority||Switch\_Name value as the lowest value, without comparing with the received value. If the Switch has received a Domain\_ID, the Switch shall send a DIA to the Switch from which it received the zero Domain\_ID\_List as described in 7.4.2;
- d) if the Switch receives in an EFP payload a zero Domain\_ID\_List and the Switch has a zero Domain\_ID\_List, and the received Switch\_Priority||Switch\_Name is lower than its current retained value, it shall discard the old value and retain the new value. The Switch shall also note from which neighbor Switch it received the new value, for potential use as the upstream Principal ISL during address distribution;
- e) if the Switch receives a new lower value of Switch\_Priority||Switch\_Name before it has had a chance to communicate a prior lower value to all other E\_Ports, it shall not attempt to communicate the prior value, and shall instead attempt to communicate the new value. The Switch shall not abort or otherwise abnormally terminate an existing EFP Exchange originated by the Switch for the sole reason of the value of Switch\_Priority||Switch\_Name being adjusted lower prior to the completion of the Exchange;
- f) the Switch shall always return the lowest known value of Switch\_Priority||Switch\_Name in a SW\_ACC reply Sequence to an EFP request Sequence; and

- g) the Switch shall retain a merged Domain\_ID list after sending or receiving the SW\_ACC to the EFP.

**Transition F3:F0.** This transition is made if the received Domain\_ID List is non-zero, the retained Domain\_ID List is non-zero, and the received Switch\_Priority||Switch\_Name and the retained Switch\_Priority||Switch\_Name are not the same.

**Transition F3:F2.** This transition is made if the received Domain\_ID\_List is zero or the retained Domain\_ID\_List is zero. In this transition, the Switch\_Priority||Switch\_Name of the Switch does not change.

**Transition F3:F3.** When the Switch is in the process of sending and receiving EFP requests and responses for the most recently received EFP, and receives a new EFP that causes the retained values to change, as described in state F3, it shall re-enter state F3 and start the process over.

**Transition F3:A0.** If the Switch receives a DIA request Sequence, then a Principal Switch has been selected. The Switch shall request a Domain\_ID as described in 7.4.

**Transition F3:P9.** If the Domain\_ID\_List of the Switch is non-zero, and the Domain\_ID\_List in a received EFP payload is non-zero, and if corresponding records in the Domain\_ID\_Lists are set to the same Domain\_ID value (Domain\_ID overlap), then the E\_Port shall not continue with Fabric Configuration, and shall become Isolated, as described in 7.6.

At the completion of the Principal Switch selection process, all Switches other than the Principal Switch shall retain knowledge of the E\_Port through which was received the lowest value of Switch\_Priority||Switch\_Name. This E\_Port is the start of the first ISL in the path to the Principal Switch for the Switch; this ISL is called the upstream Principal ISL. The Switch\_Name of the Principal Switch shall be used as the Fabric\_Name.

## 7.4 Address Distribution

### 7.4.1 Address Distribution overview

If Domain\_IDs are assigned dynamically, once a Principal Switch (i.e., Domain Address Manager) has been selected, Switches that are not a principal Switch may request a Domain\_ID. The Principal Switch shall assign all Domain\_IDs. All other non-isolated Switches shall request Domain\_IDs from the Principal Switch. Figure 22 shows the state machines of each process.

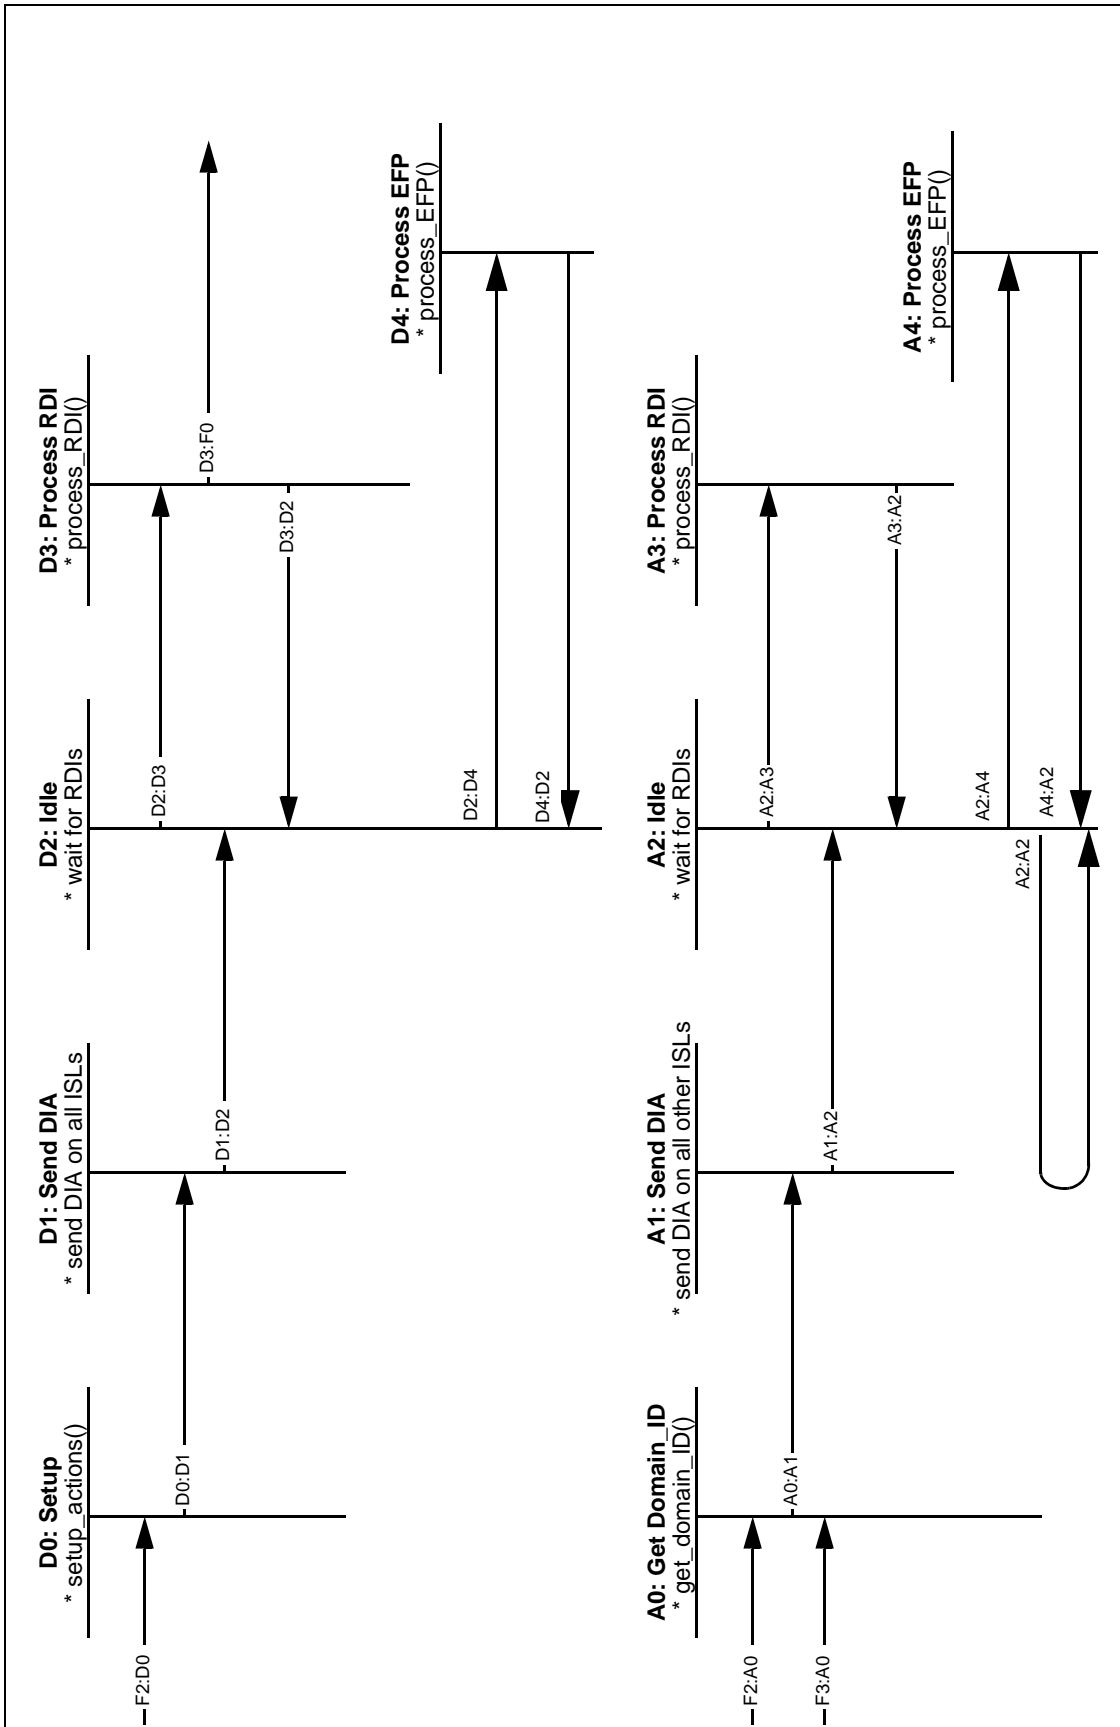


Figure 22 – Address Distribution state machines

#### 7.4.2 Domain\_ID distribution by the Principal Switch

The Principal Switch shall conduct Domain\_ID distribution as indicated in figure 22 and as described below.

**State D0: Setup.** At the completion of Principal Switch Selection, the Principal Switch shall assume the role of Domain Address Manager, and perform the following setup actions:

- a) the Principal Switch shall set its Switch\_Priority value to 02h, if the current value of its Switch\_Priority is greater than or equal to 02h. This setup action shall not cause an EFP request to be generated;
- b) the Principal Switch shall empty its Domain\_ID\_List. This setup action shall not cause an EFP request to be generated; and
- c) the Principal Switch shall then grant itself one (or more) Domain\_ID from the pool of available Domain\_IDs. This pool is maintained by the Principal Switch. If the Principal Switch had a specific Domain\_ID prior to the Reconfiguration Event, it shall grant itself that Domain\_ID. This action shall cause an EFP request to be generated as described in State D3.

**Transition F2:D0.** As defined in 7.3.

**Transition D0:D1.** This transition occurs when the setup actions described above are completed and an EFP request Sequence is sent.

**State D1: Send DIA.** The Principal Switch shall then transmit a DIA SW\_ILS request Sequence on all E\_Ports. After receiving the SW\_ACC reply, the Principal Switch may receive one or more RDI SW\_ILS request Sequences via one or more of the E\_Ports.

**Transition D1:D2.** This transition occurs when the send DIA actions described above are completed.

**State D2: Idle.** The Principal Switch shall remain in this state until it receives an RDI SW\_ILS request Sequence. Reception of RDIs and or EFPs shall be queued in this state.

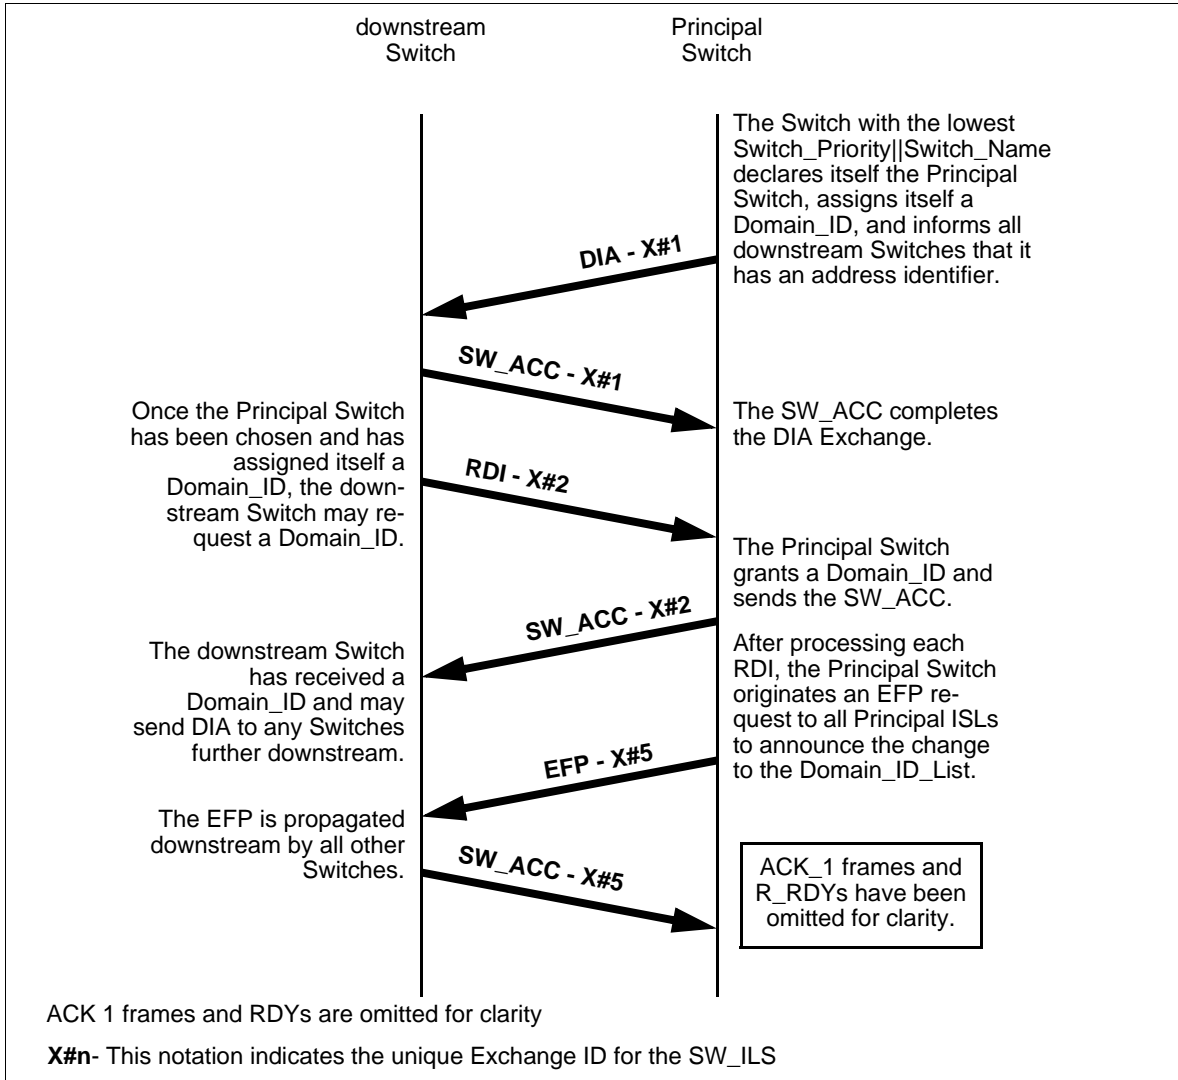
**Transition D2:D3.** This transition occurs when the Principal Switch receives an RDI SW\_ILS request Sequence via one of its E\_Ports.

**State D3: Process RDI.** The Principal Switch shall perform the following RDI processing actions:

- a) when the Principal Switch receives an RDI SW\_ILS request Sequence with a non-zero requested Domain\_ID, in the absence of any error condition preventing it, it shall allocate the requested Domain\_ID(s) to the requesting Switch, if available. If the requested Domain\_ID is zero, it shall grant an available Domain\_ID to the requesting Switch. If the requested Domain\_ID is not available, it shall either grant an available Domain\_ID to the requesting Switch or return an SW\_RJT with reason code "Domain\_ID not available". The Domain\_ID is communicated to the requesting Switch by transmitting the SW\_ACC reply Sequence via the E\_Port on which the corresponding RDI request Sequence was received;
- b) the Principal Switch shall not grant the same Domain\_ID to more than one requesting Switch;
- c) if the Principal Switch receives an RDI request for a Domain\_ID of zero, or the same requested Domain\_ID as it granted to that Switch in a previous RDI request received after Principal Switch Selection, it shall not be considered an error and the Principal Switch shall grant the Domain\_ID to the Switch using the SW\_ACC reply sequence;
- d) if a Switch that has already been granted a Domain\_ID transmits a request to the Principal Switch for a different Domain\_ID, the Principal Switch shall perform a Fabric Reconfiguration (see 7.3);



- e) if the Principal Switch receives an RDI request and no appropriate Domain\_IDs are available, the Principal Switch shall return SW\_RJT with a reason/explanation of: "Unable to perform command request", "Domain\_ID not available";
- f) all Principal ISLs via which the Principal Switch receives RDI requests shall be downstream Principal ISLs; and
- g) each time the Principal Switch grants a Domain\_ID to a Switch (including itself), it shall transmit an EFP SW\_ILS request Sequence via all Principal ISLs, with each record in the Domain\_ID\_List corresponding to a granted Domain\_ID set to the Switch\_Name granted the Domain\_ID. An example of this process is illustrated in figure 23.



**Figure 23 – RDI request processing by Principal Switch**

**Transition D3:D2.** This transition occurs when the process RDI actions described above are completed.

**Transition D3:F0.** This transition occurs when a Switch that has already been granted a Domain\_ID transmits a request to the Principal Switch for a different Domain\_ID, and the Principal Switch elects to perform a non-disruptive Fabric Reconfiguration (see 7.3).

**State D4: Process EFP.** A configured Principal Switch enters this state following the reception of an EFP request Sequence.

**Transition D2:D4.** This transition occurs when the Principal Switch receives an EFP request from an unconfigured Switch.

**Transition D4:D2.** This transition occurs when the Principal Switch transmits an EFP response and a DIA to an unconfigured Switch.

### 7.4.3 Domain\_ID requests by the Switches

The Switches shall request a Domain\_ID as indicated in figure 22, and as described below.

**Transition F2:A0.** As defined in 7.3.

**Transition F3:A0.** As defined in 7.3.

**State A0: Get Domain\_ID.** At the completion of Principal Switch Selection, the Switch receives the DIA SW\_ILS request Sequence via the upstream Principal ISL. The Switch shall reply to the request with the appropriate SW\_ACC or other response, and perform the following actions to request a Domain\_ID:

- a) the Switch shall set its Switch\_Priority value to a value greater than 02h;
- b) the Switch shall empty its Domain\_ID\_List;
- c) a DIA request Sequence received on any other ISL shall be replied to with the appropriate SW\_ACC or other response, but shall otherwise be ignored. The DIA request received via the upstream Principal ISL is the indication that the Principal Switch has assigned a Domain\_ID to all Switches between the Principal Switch and the Switch receiving the DIA request;
- d) after transmitting an SW\_ACC reply to the DIA request, the Switch shall transmit an RDI request Sequence via the upstream Principal ISL. If the Switch receives the reply SW\_ACC to the RDI request, it shall assign address identifiers to all ports within its Domain as appropriate. If the Switch receives an SW\_RJT to the RDI, it shall originate a new RDI with a different payload, or go to state P9 and become isolated; and
- e) if as a result of the RDI processing a Switch has to change its Domain\_ID, it shall perform a Link Initialization on each F\_Port and a Loop Initialization with the L bit set on the LISA Sequence on each FL\_Port. Additionally, it shall transmit an ELP on all Isolated Interconnect\_Ports, release any lock due to an ACA request, flood a LSR with the old Domain\_ID and the age field set to Max\_Age.

NOTE 11 – The change of Domain\_ID may make the local copies of some databases related to Distributed Services, or the FSPF Link State Database, out of date.

NOTE 12 – If an implementation keeps track of why a Switch port is in Isolated state, then it may avoid sending an ELP over the Interconnect\_Ports isolated for incompatible link parameters.

**Transition A0:A1.** This transition occurs when the setup actions described above are completed.

**State A1: Send DIA.** After the Switch is granted a Domain\_ID, it shall then transmit a DIA SW\_ILS request Sequence via all ISLs other than the Principal ISL. After receiving the SW\_ACC reply, the Switch may receive one or more RDI SW\_ILS request Sequences from one or more of the E\_Ports.

**Transition A1:A2.** This transition occurs when the send DIA actions described above are completed.

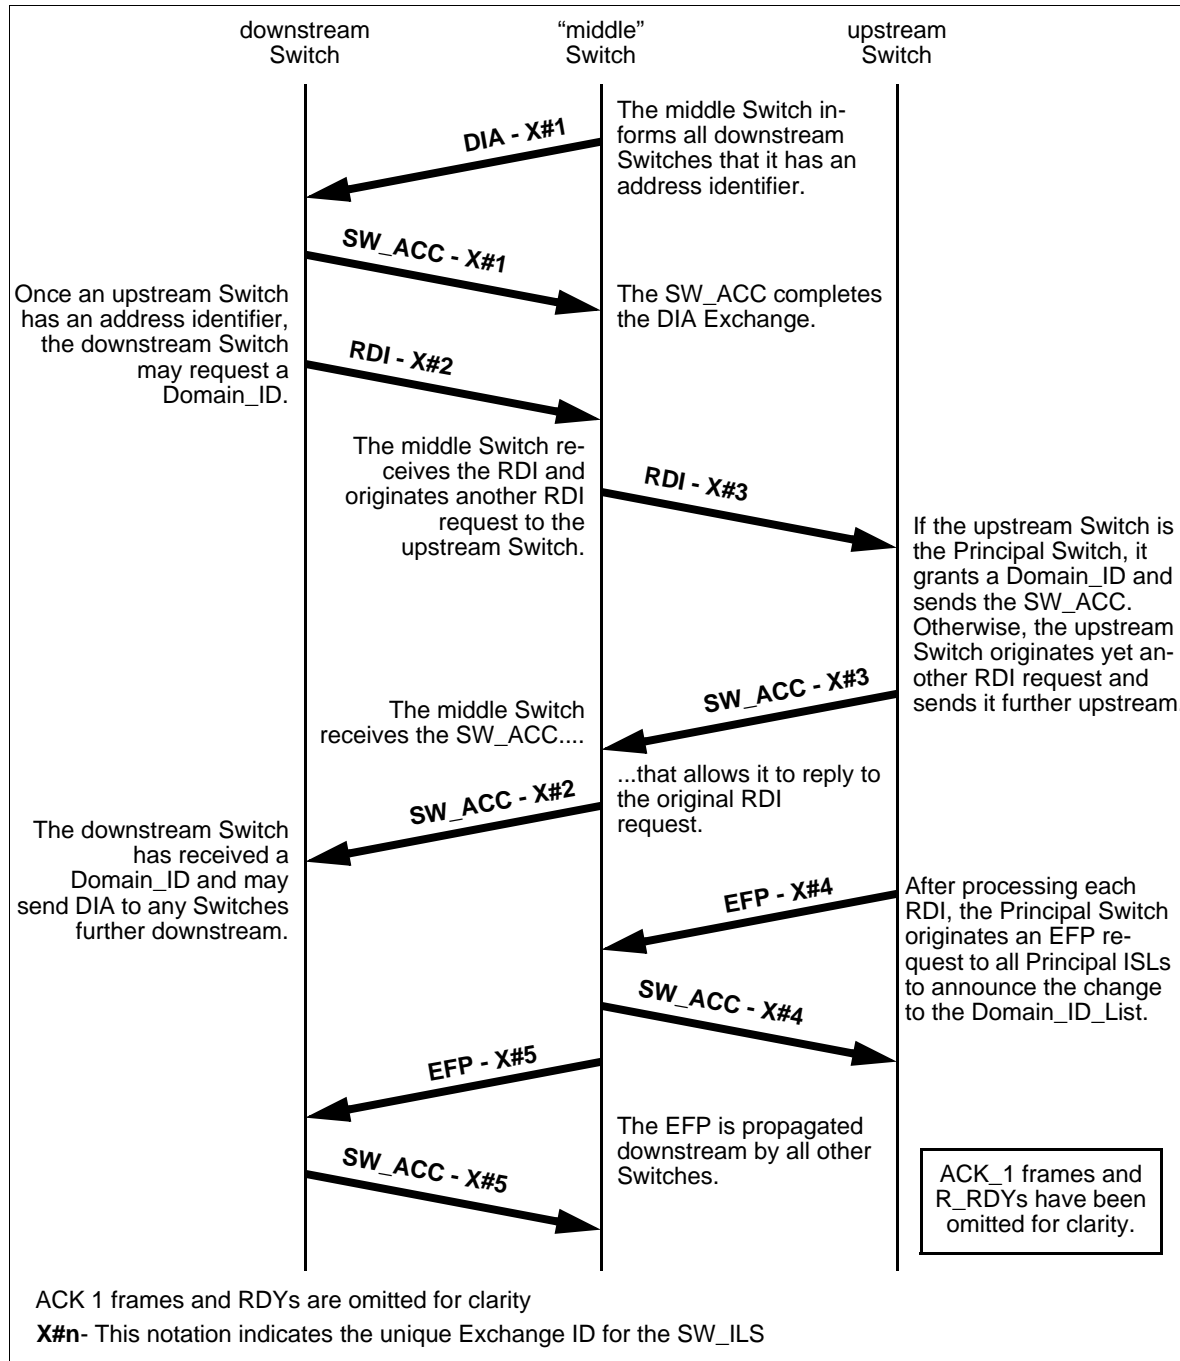
**State A2: Idle.** The Switch shall remain in this state until it receives an RDI SW\_ILS request Sequence. Reception of RDIs and or EFPs shall be queued in this state.

**Transition A2:A3.** This transition occurs when the Switch receives an RDI SW\_ILS request Sequence via one of its E\_Ports.

**Transition A2:A2.** This transition occurs when the Switch receives an EFP request from an upstream ISL and an EFP response is sent, and the EFP request is forwarded on all other ISLs.

**State A3: Process RDI.** The Switch shall perform the following RDI processing actions:

- a) all Principal ISLs via which the Switch receives valid RDI requests shall be downstream Principal ISLs. If the Switch receives an RDI request on its upstream Principal ISL, it shall return SW\_RJT with a reason/explanation of logical error, request not supported; and
- b) when the Switch receives a valid RDI request Sequence from one of its E\_Ports via a downstream Principal ISL, it shall originate an RDI request Sequence with the same payload via its upstream Principal ISL. When the reply SW\_ACC is received via the upstream Principal ISL, it shall transmit an SW\_ACC reply Sequence via the downstream Principal ISL on which the initial request was received. An example of this process is illustrated in figure 24.



**Figure 24 – RDI request Processing by non-Principal Switch**

**Transition A3:A2.** This transition occurs when the process RDI actions described above are completed.

**State A4: Process EFP.** A configured non-principal Switch enters this state following the reception of an EFP request Sequence. The Switch shall transmit EFP on all downstream principal ISLs.

**Transition A2:A4.** This transition occurs when a Switch that has already been configured receives an EFP request from a downstream unconfigured Switch.

**Transition A4:A2.** This transition occurs when a Switch that has already been configured transmits an EFP response and a DIA to a downstream unconfigured Switch.

## 7.5 Principal ISL Recovery

### 7.5.1 Overview

Failure of a Principal ISL disrupts control communication between the Principal Switch and downstream Switches. If other paths to the Principal Switch exist in the Fabric, recovery from this failure is possible with a Build Fabric operation. However, a Build Fabric operation creates a great deal of Fabric control traffic. In the case that additional ISLs exist between the two Switches that encountered the Principal ISL failure, it is possible for the two Switches to select a new Principal ISL without any impact to the remainder of the Fabric. This subclause describes this recovery process. Implementation of this process is optional.

### 7.5.2 Downstream Principal ISL discovery

If a Switch implements Principal ISL Recovery and detects a Downstream Principal ISL failure for which it has additional ISLs connecting to the same downstream Switch, the Switch shall select one of the additional ISLs to become the new Downstream Principal ISL and send an EFP on that link within F\_S\_TOV. Then the Switch shall:

- a) use the selected link as Downstream Principal ISL upon receipt of an EFP SW\_ACC; or
- b) proceed with a Build Fabric (see 7.3) upon receipt of an EFP SW\_RJT or if no response is received within F\_S\_TOV.

### 7.5.3 Upstream Principal ISL Recovery

If a Switch implements Principal ISL Recovery and has ISLs in addition to the Upstream Principal ISL to the same upstream Switch, the Switch shall:

- a) upon receiving an EFP from the upstream Switch on one of the additional ISLs:
  - 1) respond with an SW\_ACC; and
  - 2) use the ISL on which the EFP was received as the new Upstream Principal ISL; and
- b) upon detection of a failure of the Upstream Principal ISL without receiving an EFP within 2xF\_S\_TOV on any of the additional ISLs, initiate a Build Fabric (see 7.3).

## 7.6 E\_Port and Fabric isolation

An E\_Port connected via an Inter-Switch Link to another E\_Port may determine that it is unable to communicate with the other E\_Port for one of the reasons:

- a) the two E\_Ports have incompatible link parameter requirements. For example, if one Switch has an E\_D\_TOV setting different than another, Class 2 frames sent by an N\_Port on one Switch may not receive timely F\_BSY responses from the other Switch;
- b) similarly, the two E\_Ports have incompatible Fabric Parameter requirements. For example, if an E\_Port receives an EFP that contains records it does not support, it shall Isolate;
- c) when the E\_Port receives an EFP payload and the received Domain\_ID\_List is non-zero, the retained Domain\_ID\_List is non-zero, the Domain\_IDs overlap, and the received Principal Switch\_Name is not equal to the retained Principal Switch\_Name;

- d) the two E\_Ports are a link between Switches that are not capable of performing the Domain Address Manager function, and are each also not attached via an ISL to any other Switch capable of performing the Domain Address Manager function. Since no Switch may allocate Domain\_IDs, no Class N frames may be sent between the Switches;
- e) the two E\_Ports have exchanged Zoning information via the Merge Request in an attempt to resolve a Zoning configuration. As a result of the Merge processing the Zoning configuration may not be merged (see 10.5);
- f) an SW\_RJT is received in response to an RDI request and the Switch chooses to not send a new RDI with a different payload;
- g) the E\_Port rejects an RCF SW\_ILS;
- h) an SW\_RJT with reason code explanation of "E\_Port is Isolated" is received; or
- i) a Switch configured to assign Domain\_IDs statically, on receiving an EFP, BF, RCF, DIA or RDI SW\_ILS shall reply with an SW\_RJT having reason code explanation 'E\_Port is Isolated' and shall isolate the receiving E\_Port.

When any of the above conditions occurs, the E\_Port shall Isolate itself from the other E\_Port. The following is a list of appropriate Class F frames that may be communicated between Isolated E\_Ports:

- a) an ELP SW\_ILS request may be sent by an Isolated E\_Port in an attempt to establish a working set of link parameters. This ELP SW\_ILS request may be used to support a negotiation process as outlined in annex B;
- b) an SW\_ACC response may be sent in response to any of the above SW\_ILS requests; and
- c) an SW\_RJT response may be sent in response to any of the above SW\_ILS requests, if necessary, and shall be sent as the appropriate response to any other SW\_ILS request not listed above. The SW\_RJT response shall indicate the following SW\_RJT reason/explanation code: Unable to perform command request/ E\_Port is isolated.

The buffer-to-buffer credit between the Isolated E\_Ports shall be a value of one; no alternate credit shall be in effect. No routing of Class N frames shall occur across the ISL.

A Switch may override the Isolated condition by originating an ELP, or any of the events that cause the transition ALL:P0.

## **7.7 B\_Port operation**

### **7.7.1 Differences between E\_Ports and B\_Ports**

A B\_Port supports a subset of the E\_Port Internal Link Services (ILS) and a B\_Port has the same facilities as described in this standard for an E\_Port. The underlying differences between B\_Port and E\_Port initialization are that B\_Ports perform ELP and are transparent to most other SW\_ILSs (see 5.6).

## 7.7.2 B\_Port Internal Link Services

The B\_Port shall generate a subset of the Internal Link Services defined in this standard. Table 191 details the ILS support as either being propagated or generated by the B\_Port.

**Table 191 – B\_Port ILS support (Part 1 of 2)**

<b>FC-SW-6 Internal Link Service (ILS)</b>	<b>Generated by B_Port</b>	<b>B_Port response</b>	<b>Propagated by B_Port</b>
<b>Exchange Link Parameter (ELP)</b>	Allowed	SW_ACC or SW_RJT	Prohibited
<b>Exchange Fabric Parameters (EFP)</b>	Prohibited	Propagate	Allowed
<b>Domain Identifier Assigned (DIA)</b>	Prohibited	Propagate	Allowed
<b>Request Domain_ID (RDI)</b>	Prohibited	Propagate	Allowed
<b>Hello (HLO)</b>	Prohibited	Propagate	Allowed
<b>Link State Update (LSU)</b>	Prohibited	Propagate	Allowed
<b>Link State Acknowledgment (LSA)</b>	Prohibited	Propagate	Allowed
<b>Build Fabric (BF)</b>	Prohibited	Propagate	Allowed
<b>Reconfigure Fabric (RCF)</b>	Prohibited	Propagate	Allowed
<b>Exchange Switch Capabilities (ESC)</b>	Prohibited	Propagate	Allowed
<b>Acquire Change Authorization (ACA)</b>	Prohibited	Propagate	Allowed
<b>Release Change Authorization (RCA)</b>	Prohibited	Propagate	Allowed
<b>Stage Fabric Configuration (SFC)</b>	Prohibited	Propagate	Allowed
<b>Update Fabric Configuration (UFC)</b>	Prohibited	Propagate	Allowed
<b>Registered State Change Notification (SW_RSCN)</b>	Prohibited	Propagate	Allowed
<b>Distribute Registered Link Incident Report (DRLIR)</b>	Prohibited	Propagate	Allowed
<b>Check E_Port Connectivity (CEC)</b>	Prohibited	Propagate	Allowed
<b>Exchange Switch Support (ESS)</b>	Prohibited	Propagate	Allowed
<b>Merge Request Resource Allocation (MRRA)</b>	Prohibited	Propagate	Allowed
<b>Switch Trace Route (STR)</b>	Prohibited	Propagate	Allowed
<b>Exchange Virtual Fabrics Parameters (EVFP)</b>	Prohibited	Propagate	Allowed

**Table 191 – B\_Port ILS support (Part 2 of 2)**

<b>Enhanced Acquire Change Authorization (EACA)</b>	Prohibited	Propagate	Allowed
<b>Enhanced Stage Fabric Configuration (ESFC)</b>	Prohibited	Propagate	Allowed
<b>Enhanced Update Fabric Configuration (EUFC)</b>	Prohibited	Propagate	Allowed
<b>Enhanced Release Change Authorization (ERCA)</b>	Prohibited	Propagate	Allowed
<b>Transfer Commit Ownership (TCO)</b>	Prohibited	Propagate	Allowed

**7.7.3 B\_Port initialization**


The Fabric Configuration process enables a Switch to determine its operating mode, exchange operating parameters, and provides for distribution of addresses. Changes to support Bridge devices and the B\_Port in this process are summarized in table 192

**Table 192 – B\_Port initialization summary**

Operation	Starting condition	Process	Ending condition
Establish link parameters and Switch port operating mode	Bridge port has achieved word synch.	ELPs are exchanged and the B_Port determines that it is attached to an E_Port.	Link parameters have been exchanged and Credit has been initialized, and it is known if the attached port is an E_Port.

**7.7.4 Example B\_Port configuration**

The following diagram shows an example Fabric configuration utilizing B\_Ports. In this instance, two bridge devices enable the existence of a virtual ISL between two Switches. With the exception of ELP, the B\_Port is transparent to Fabric E\_Port operation.



**Figure 25 – Example B\_Port configuration - Virtual ISL**



## 8 Fabric Shortest Path First (FSPF)

### 8.1 Overview

#### 8.1.1 Basic components

FSPF is a link state path selection protocol. FSPF keeps track of the state of the links on all Switches in the Fabric and associates a cost with each link. The protocol computes paths from a Switch to all the other Switches in the Fabric by adding the cost of all the links traversed by the path, and choosing the path that minimizes the cost. The collection of link states (including cost) of all the Switches in a Fabric constitutes the Link State Database.

FSPF has four major components:

- a) a Hello protocol, used to establish connectivity with a neighbor Switch, to establish the identity of the neighbor Switch, and to exchange FSPF parameters and capabilities;
- b) a replicated Link State Database, with the protocols and mechanisms to keep the databases synchronized across the Fabric;
- c) a path computation algorithm; and
- d) a routing table update.

The Link State Database synchronization in turn consists of two major components: an initial database synchronization, and an update mechanism. The initial database synchronization is used when a Switch is initialized, or when an Inter-Switch Link (ISL) comes up. The update mechanism is used in two circumstances:

- a) when there is a link state change, for example an ISL going down or coming up; and
- b) on a periodic basis, to prevent Switches from deleting topology information from the database.

The path computation algorithm shall be any algorithm (e.g., Dijkstra's algorithm) that yields a minimum cost path, as defined above.

The routing table update is not covered in this specification, since it is strictly implementation dependent. Different Switches may have different internal routing mechanisms, and still interoperate.

#### 8.1.2 Fabric connectivity

All the connections between Fibre Channel Switches shall be point-to-point. There are no direct connections to broadcast media, where multiple routing-capable Switches may co-exist.

#### 8.1.3 Addressing

A path selection protocol requires an addressing scheme to uniquely identify the final destination of a frame. FSPF supports the addressing scheme described in 4.9. If multiple Domain\_IDs are used by a Switch, the Switch shall use the lowest value Domain\_ID as the Originating Domain\_ID in all FSPF headers. It shall also send an LSR for each Domain\_ID that it has been assigned.

#### 8.1.4 Path selection and routing

In this standard, the term "path selection" indicates the discovery of the best path from source to destination, and the term "routing" indicates the actual forwarding of frames to a specific destination. FSPF performs hop-by-hop routing meaning that a Switch only needs to know the next hop on the best path to the destination. The replicated Link State Database insures that every Switch in the

Fabric has the same view of the Fabric itself, allowing consistent routing decisions to be made by all Switches. The replicated data base is essential to avoid routing loops.

Typically a Switch needs to know, for each destination domain in the Fabric, which path should be used to route a frame to that domain. A routing table entry minimally consists of a destination Domain\_ID, and an E\_Port to which frames are forwarded to the destination Switch.

### 8.1.5 FSPF path selection summary

Table 193 summarizes path selection via FSPF.

**Table 193 – Path selection (FSPF) operation summary**

Operation	Starting condition	Process	Ending condition
1. Perform Initial HELLO Exchange	The Switch originating the HELLO has a valid Domain_ID.	HLO SW_ISL frames are exchanged on the link until each Switch has received a HELLO with a valid neighbor Domain field.	Two way communication has been established
2. Perform Initial Database Exchange	Two way communication has been established.	LSU SW_ISL frames are exchanged containing the Initial database.	Link State Databases have been exchanged.
3. Running State	Initial Database Exchange completed.	Routes are calculated and set up within each Switch. Links are maintained by sending HELLOs every Hello_Interval. Link databases are maintained by flooding link updates as appropriate.	FSPF routes are fully functional.

## 8.2 FSPF message processing

### 8.2.1 Message transmission

FSPF information is transported using FSPF SW\_ILS messages. Details of these three FSPF SW\_ILS messages are described in 6.1.

Before sending a message, a Switch shall set the values in the header fields as follows:

- a) Command code: The value that identifies the type of message, Hello (14000000h), Link State Update (15000000h) or Link State Acknowledgement (16000000h);
- b) Version: The version number of the protocol as documented in this standard (02h);
- c) Authentication Type: No authentication is specified at this time. This field shall be set to 00h;
- d) Originating Domain\_ID: The Domain\_ID of the Switch that is transmitting this message. The Originating Domain\_ID shall be a valid value as specified in 6.2.8.2; and
- e) Authentication: No authentication is specified at this time. This field is 8 bytes long and shall be set to 0000000000000000h.

### 8.2.2 Message reception and tests

When an FSPF message is received, the following tests shall be performed on its content. These tests are:

- a) the Version number shall be 02h;
- b) the Authentication Type shall be 00h;
- c) the Originating Domain\_ID shall be checked for a valid value as specified in 6.2.8.2; and
- d) the Authentication field shall be 0000000000000000h.

If any of these tests fails, the message shall be discarded. If all the tests succeed, the message shall be passed to the relevant protocol for further processing.

### 8.3 Hello protocol

#### 8.3.1 Basic functions

The Hello protocol is used to establish two-way communication with a neighbor Switch, and determine when this communication is interrupted. An Inter-Switch Link (ISL) may be used for routing user traffic through the Fabric only if there is two-way communication between the Switches. The Hello protocol also provides some information about remote connectivity, and in particular, it allows the association of the local E\_Port with the remote E\_Port.

#### 8.3.2 Hello message transmission

A Switch is required to know that a port is connected to another Switch through an ISL, before that port may be used to route data in a multi-Switch Fabric. Prior to a Hello being sent, the following shall be true:

- a) the port shall be an E\_Port;
- b) the Switch where the E\_Port resides shall have a Domain\_ID assigned; and
- c) the Switches on the two sides of an ISL shall have agreement on a common set of link parameters and Fabric Parameters.

After a Switch determines that a port is an E\_Port and the Switch has acquired a Domain\_ID, the Switch starts sending Hello messages to the neighbor Switch. Hello messages contain the FSPF header and the parameters specific in the Hello protocol.

#### 8.3.3 Hello message parameters

The Hello\_Interval is defined to be the interval in seconds between two consecutive transmissions of a Hello message by the local Switch.

The Dead\_Interval is the interval in seconds after which the local Switch shall bring down the Adjacency, if it has not received a valid Hello message from the remote Switch. Switches may also reset this timer on reception of an FSPF LSA or LSU.

The Hello\_Interval and the Dead\_Interval are values that may be configured separately on each port. It is absolutely necessary that the two ports that are connected by an ISL on two Switches have the same value for these two variables. Default values that are appropriate for most circumstances are provided in table 292.

The Recipient Domain\_ID is the Domain\_ID of the Switch on the other side of the ISL. It is set to FFFFFFFFh in the first transmitted Hello, to indicate that the Switch has not received an Hello message from the neighbor Switch. Once an Hello message is received from the neighbor Switch, the local Switch stores the Domain\_ID of the remote Switch on that port, and from then on it sets the Recipient Domain\_ID to that value in all future Hello messages. The Recipient Domain\_ID is set back to FFFFFFFFh when the two-way communication between Adjacent Switches is disrupted. This

typically happens either because the E\_Port goes offline, or because the Dead\_Interval timer expires.

The Originating Port Index shall be set to the index of the port that transmits the Hello message.

If the Domain\_ID of a Switch changes, then the Switch shall perform a one-way Hello with FFFFFFFFh set in its Recipient Domain\_ID field.

#### **8.3.4 Hello message reception**

When a Hello message is received, the message header shall be checked according to the rules described in 8.2.2. In addition, the following checks are performed:

- a) the Hello\_Interval value shall match the value configured for the port that originated the message. If it does not, the Hello message is discarded;
- b) the Dead\_Interval value shall match the value configured for the port that originated the message. If it does not, the Hello message is discarded; and
- c) the Recipient Domain\_ID shall be either FFFFFFFFh, or the Domain\_ID of the local Switch. Any other value in this field causes the Hello message to be discarded.

When the local Domain\_ID is recognized in the incoming Hello message, a two-way communication has been established with the remote Switch, and the Neighbor FSM may proceed to the next transition. If the value FFFFFFFFh is detected in an incoming Hello message at any time after the two-way communication has been established, the neighbor shall fall back to a one-way state and the FSM transitions to that state.

The Originating Port Index does not need to be checked. Its value shall be stored in the neighbor data structure, together with the Domain\_ID of the sending Switch, the Hello\_Interval and the Dead\_Interval.

### **8.4 The Link State Database**

The Link State Database is central to the operation of FSPF. It is a replicated database where all Switches in the Fabric have the same exact copy of database at all times.

The database consists of a collection of Link State Records (LSRs). Link State Records may be of different types and have different formats and contents. This standard describes one type of LSR (i.e., Switch Link Record).

A Switch Link LSR completely describes the connectivity of a Switch to all Switches to which it is directly attached. The information contained in a LSR is a list of all the individual ISLs that the Switch may use to forward user data to a remote Switch. Each ISL is associated with a link type, the Domain\_ID of the remote Switch it is connected to, the local and remote Port ID, and the cost of the link itself.

Every Switch in the Fabric is responsible for issuing and maintaining its own LSR. An LSR is identified by a Link-State ID. For a Switch link LSR the Link-State ID is the Domain\_ID of the Switch that issues the LSR. A Switch shall not issue an Switch Link LSR with a Domain\_ID different from its own Domain\_ID. A Switch shall not generate new instances of an LSR, unless it generated the original LSR. However, a Switch shall forward LSRs that it has not generated as part of the flooding process.

Multiple instances of an LSR are issued over time. Sometimes the content of the new instance is the same as the previous instance, sometimes it is different. Every Switch is responsible for maintaining the most recent copy of its own LSR throughout the Fabric.

Multiple instances of an LSR may be temporarily present in a Fabric at the same time. Ultimately, only the most recent instance shall survive, and all Switches shall keep that instance in their Link State Database. The process of purging old instances of an LSR within the Fabric should be as fast as possible because it impacts the ability to properly route Class N frames through the Fabric.

Several fields in a LSR are used to identify the LSR and to determine which instance is the newest. The Link State Database is used by a Switch to compute the least cost path to all other Switches in the Fabric. This is why it is essential that all Switches have the same Link State Database, or different Switches may build inconsistent paths. An ISL shall be considered in the path computation only if both LSRs of the two connected Switches list this ISL (two way communication between the Switches).

The path computation is local, and the results of the computation are not distributed to other Switches, only topology information is distributed. This is a characteristic of link-state path selection protocols.

## 8.5 Usage of LSR fields

### 8.5.1 LSR Flags

The Leaf Switch bit (i.e., bit 0) in the LSR Flags field indicates a Switch is operating as a Leaf Switch (see 18). The Leaf Switch bit indicates how the Switch handles frames addressed to Domain\_IDs that do not belong to the Switch (i.e., other Domain\_IDs in the Fabric).

If the Leaf Switch bit is set to one, then the Switch does not allow frames to be routed through the Switch to other Switches in the Fabric (i.e., other Switches do not have a valid path through this Switch).

If the Leaf Switch bit is set to zero, then the Switch allows frames to be routed through the Switch to other Switches in the Fabric (i.e., other Switches have a valid path through this Switch).

### 8.5.2 LSR Age

The LSR Age field indicates how long a particular instance of an LSR has been in the database. The LSR Age field is based in seconds and is a 16-bit unsigned integer.

The LSR Age is initialized to 0000h by the advertising Switch when it is first issued. The LSR Age is incremented by one every second by every Switch in the Fabric as long as it stays in that Switch's database. It is also incremented by one every time it is transmitted during the flooding procedure.

NOTE 13 – This somewhat arbitrary increment represents the transmission time on the ISL and insures that a flooded LSR does not loop forever.

This field is also used to help determine which of two instances of an LSR is more recent, when other fields are equal.

A new instance of an LSR shall be issued when the LSR Age field of the LSR in the database reaches the value LS\_Refresh\_Time. Only the Switch that originated the LSR shall refresh it with the issue of a new instance.

The age of an LSR shall never exceed Max\_Age (i.e., 3600, 1 hour). If an LSR reaches the age of 3600, it shall be flushed from the Fabric. This operation is accomplished by flooding the LSR with the LSR Age field set to Max\_Age. Upon receiving this instance of an LSR, other Switches shall remove the LSR from the database. In order to be completely flushed from the Fabric, an aged LSR shall be removed from the database in all Switches.

Any Switch in the Fabric may flush an LSR that has reached Max\_Age from the Fabric.

### 8.5.3 LSR incarnation number

This field is a progressive number that identifies the incarnation of the LSR. It is used to determine which one of two incarnations of an LSR is more recent, in particular, the one with the larger incarnation number is the most recent.

The incarnation number is a 32-bit signed integer and is incremented in two's complement form. The lowest possible negative number is 80000000h, and it is not used. The lowest incarnation number is 80000001h. The first instance of an LSR shall have an incarnation number of 80000001h. Each new instance shall have its incarnation number incremented by one. A new instance may be issued for several reasons, but it shall always have its incarnation number increased by one, even if the content of the LSR is identical to the previous instance.

NOTE 14 – This causes the new instance to have a different checksum.

The maximum incarnation number is 7FFFFFFFh. When an LSR reaches this value as an incarnation number, the originating Switch shall flood the LSR through the Fabric with an LSR Age = Max\_Age. After the LSR is acknowledged by an LSA on all ISLs, then the originating Switch shall issue a new instance of the LSR with an incarnation number of 80000001h.

This process causes a brief interruption of service because paths to the Switch that is rolling over its incarnation number are not available until the LSR with the smallest incarnation number is installed. However, this event should be extremely rare since most of the time a new instance of an LSR is issued every 30 minutes.

### 8.5.4 LSR instance rules

Two LSR instances shall be considered identical when both of the following conditions are met:

- a) the Link State ID fields are the same; and
- b) the Link State incarnation values are the same.

For two instances of the same LSR, the LSR incarnation number, LSR Age, and LSR checksum fields shall be used to determine which instance is more recent:

- a) the LSR instance with the highest incarnation value shall be considered more recent. If both instances have the same incarnation value, then;
- b) if the LSR Age fields of only one of the two instances is equal to MaxAge, it shall be considered more recent;
- c) else, if the two instances have different LSR checksums, then the instance having the larger LSR checksum (when considered as a 16-bit unsigned integer) shall be considered more recent;
- d) else, if the LSR Age fields of the two instances differ by a value less than or equal to Max\_Age\_Diff, the instance having the smaller (i.e., younger) LSR Age shall be considered more recent; or
- e) else, the two instances shall be considered to be identical.

### 8.5.5 LSR checksum

The checksum field is used to detect data corruption in an LSR, both when it is received and when it is stored in Switch memory. When an LSR is received with a bad checksum, the LSR shall be ignored.

The integrity of the Link State Database shall be checked by calculating checksums for all the LSRs. If any of the LSRs fail this checksum, this may be an indication of a memory corruption problem, and the Switch should be reinitialized.

The LSR checksum covers the whole LSR, except the LSR Age field. The checksum algorithm is known as the Fletcher Checksum, and shall be computed byte by byte, by accumulating the sum of the payload one byte at the time.

NOTE 15 – The Fletcher algorithm is specified in RFC 905. The Nakassis algorithm, for an optimized computation of the checksum, is specified in RFC 1008.

The checksum shall be computed as specified in table 194:

**Table 194 – Checksum byte order calculation**

Word	Bits 31 to 24	Bits 23 to 16	Bits 15 to 8	Bits 7 to 0
0	LSR Type A1	Reserved A0	LSR Age	LSR Age
1	Reserved B3	Reserved B2	Reserved B1	Reserved B0
2	Link State Identifier B7	Link State Identifier B6	Link State Identifier B5	Link State Identifier B4
3	Advertising Domain_ID B11	Advertising Domain_ID B10	Advertising Domain_ID B9	Advertising Domain_ID B8
4	Link State Incarnation Number B15	Link State Incarnation Number B14	Link State Incarnation Number B13	Link State Incarnation Number B12
5	Checksum B19	Checksum B18	LSR Length B17	LSR Length B16
6	Reserved C3	Reserved C2	Number of Links C1	Number of Links C0
7	Link ID 0,D3	Link ID 0,D2	Link ID 0,D1	Link ID 0,D0
8	Reserved 0,D7	Output Port Index 0,D6	Output Port Index 0,D5	Output Port Index 0,D4
9	Reserved 0,D11	Neighbor Port Index 0,D10	Neighbor Port Index 0,D9	Neighbor Port Index 0,D8
10	Link Type 0,D15	Reserved 0,D14	Link Cost 0,D13	Link Cost 0,D12

The two checksum bytes are initialized to 00h for the checksum calculation. The checksum calculation is performed in the following order:

- 1) A0, A1, B0, B1, B2, B3, B4, B5, B6 B7, B8, B9, B10, B11, B12, B13, B14, B15, B16, B17, B18, B19, C0, C1, C2, C3;
- 2) (0, D0), (0, D1), (0, D2), (0, D3), (0, D4), (0, D5), (0, D6), (0, D7), (0, D8), (0, D9), (0, D10), (0, D11), (0, D12), (0, D13), (0, D14), (0, D15);
- 3) (1, D0), (1, D1), (1, D2), (1, D3), (1, D4), (1, D5), (1, D6), (1, D7), (1, D8), (1, D9), (1, D10), (1, D11), (1, D12), (1, D13), (1, D14), (1, D15)...;
- 4) ...(n-1, D0), (n-1, D1), (n-1, D2), (n-1, D3), (n-1, D4), (n-1, D5), (n-1, D6), (n-1, D7), (n-1, D8), (n-1, D9), (n-1, D10), (n-1, D11), (n-1, D12), (n-1, D13), (n-1, D14), (n-1, D15); and
- 5) (n, D0), (n, D1), (n, D2), (n, D3), (n, D4), (n, D5), (n, D6), (n, D7), (n, D8), (n, D9), (n, D10), (n, D11), (n, D12), (n, D13), (n, D14), (n, D15).



In this nomenclature:

- a) (x, Dy) refers to link x up to n starting at 0; and
- b) y= 0 through 15 representing the fields in each Link Descriptor that is a part of the LSR.

### 8.5.6 Link Cost

The Link Cost for each link is calculated based on the speed of the link, plus an administratively set factor. The link cost calculation is:

$$\text{Link Cost} = S * \text{DEFAULT\_COST}$$

Where *S* is an administratively defined factor. By default *S* is set to 1.0.

NOTE 16 – This value allows an administrator to adjust link cost based on a particular environment.

For each link speed, the DEFAULT\_COST value is defined in table 195.

**Table 195 – DEFAULT\_COST values**

Link speed	Values
1GFC	1000
2GFC	500
4GFC	250
8GFC	125
10GFC	83
16GFC	62
32GFC	31
64GFC	16
128GFC	8
10GFCoE	85
25GFCoE	34
40GFCoE	21
50GFCoE	17
100GFCoE	8
400GFCoE	2

This calculation shall be performed on a link by link basis. Each link in a Fabric may be advertised with a different cost. These costs shall be used by the path selection algorithm to determine the most efficient paths. If more than one least-cost path exists, use of the multiple least-cost paths is implementation specific.

## **8.6 Link State Database synchronization**

### **8.6.1 Synchronization overview**

The Link State Database shall be periodically synchronized across all Switches in the Fabric. This synchronization is required for the following reasons:

- a) LSRs in the database may change because an ISL comes up or goes down;
- b) Switches are added or removed from the Fabric;
- c) LSRs may be added or removed; and
- d) periodic issuance of new LSR instances.

Every time a new instance of an LSR or a new LSR is issued, the whole Fabric shall be informed. FSPF achieves this through reliable flooding of LSRs. A new instance of an LSR is transmitted to the directly attached Switches in a reliable fashion. The attached Switches in turn reliably forward the LSR to their attached Switches, and the process continues until all Switches in the Fabric have received the new instance of the LSR.

In addition, when an ISL between two Switches becomes operational and the Switches have successfully established two-way communication using the Hello protocol, the two Switches shall synchronize their database. Each Switch sends its full database to the other Switch, and receives the database from the other Switch. Each Switch updates its database with any received LSR that is either absent from its database, or is a newer instance of that LSR.

Generally, the initial synchronization and the ongoing database updates are implemented in the same fashion. One difference is that the initial synchronization involves the transmission of the whole database in two directions, whereas the ongoing updates typically involve only one or a few LSRs, and occur in one direction only. This characteristic of FSPF minimizes the number of different messages required by the protocol, and improves code reusability. Another difference is that the initial synchronization occurs on one ISL only, whereas the ongoing synchronization occurs on one or more ISLs.

### **8.6.2 Neighborhood and Adjacency**

Two Switches connected via an ISL shall be referred to as neighbors. Two Switches that are simply neighbors shall not use their common ISL to carry Class N frames until their Link State Databases have been synchronized. Once the Link State Databases have been synchronized, the two Switches are referred to as Adjacent on that ISL. If two Switches are connected by multiple ISLs, they may be neighbors on some ISLs and Adjacent on others at any given time.

After a Switch detects that one of its ports is connected to another Switch, it starts exchanging Hello messages with its neighbor. Initially, a Switch knows only its own Domain\_ID, and the Port Index of the port that connects the Switch to its neighbor. The Switch provides this information in the Hello message that it transmits to the neighbor. At this time the Switch does not know the Recipient Domain\_ID on the neighbor Switch. Therefore, the Switch shall set the Domain\_ID to FFFFFFFFh in the transmitted Hello message.

When a Hello message is received, the Switch stores the Domain\_ID and Port Index of the neighbor Switch and shall send the Domain\_ID in subsequent Hello messages on that ISL. When a Switch sees its own Domain\_ID as the Recipient Domain\_ID in a received Hello message, two-way communication is established on that ISL and Link State Database exchange shall be initiated. Upon detection of two-way communication, the Switch should send a Hello message immediately, rather than waiting for expiration of the Hello\_Interval time.

The next step is to synchronize the Link State Databases on the two Switches. The original databases may be already identical if the two Switches are already connected by an ISL, and the new ISL is just an additional one. The two databases may be totally different if the ISL is used to connect two previously disjointed Fabrics. In some cases it is possible for some portions of the Data Base to be identical while other portions are not.

During the process of synchronizing databases, an algorithm determines the most recent of two instances of a database entry (i.e., LSR). Both Switches shall keep in their database only the most recent version of an LSR. This algorithm is used both for the initial database synchronization process, and for any subsequent database update.

In the database synchronization phase, each Switch sends its complete Link State Database to the neighbor Switch. This topology information is transported as LSRs contained in one or more LSUs. Both Switches shall examine every LSR in the LSU, determine whether each is more recent than the associated instance in the database, and shall update the database accordingly. If the received instance of the LSR is more recent than the one contained in the database, or if the database did not contain the LSR in its database, then the received version of the LSR is stored in the database. LSRs shall be acknowledged by an LSA if they are newer or identical to the local copy, or no local copy exists. Otherwise an LSR is acknowledged by a newer LSR instance.

The receiving Switch shall acknowledge each LSR within an LSU separately. An acknowledgement for an LSR shall consist of the LSR header that uniquely identifies the instance of an LSR. Acknowledgements shall be sent in Link State Acknowledgement messages and an LSA may contain zero, one, or more acknowledgements. An LSA containing no LSR headers shall be used to acknowledge reception of the Database Complete flag from the neighbor, and shall confirm the end of the initial Link State Database synchronization process.

Unacknowledged LSRs shall be retransmitted by the sender after the Rxmt\_Interval interval expires until they are acknowledged by the neighbor.

At the end of the process, both Switches have exchanged their topology data bases and they are considered Adjacent on that ISL. Any subsequent changes shall be communicated via LSU's. The ISL itself may then be used to carry user data. Both Switches shall issue a new LSR that includes the newly Adjacent ISL. This LSR shall be flooded reliably (see 8.6.4) on the new ISL, and on all other ISLs, together with any updated LSR.

After the new LSR has been transmitted and acknowledged on all of the Switch's ISLs, the Switch shall recompute the paths to all other Switches in the Fabric, and update the routing table accordingly. All the other Switches in the Fabric shall do the same, having received the new LSR.

The process of transmitting, receiving, processing, and acknowledging LSRs is identical for the initial database synchronization process, and for ongoing or periodic updates. The same messages and the same algorithm shall be used in both cases. This characteristic of FSPF simplifies the implementation, by reducing the number of different messages, and by improving code reusability.

The Adjacency bring-up process is detailed in this standard as a finite state machine (FSM) called the Neighbor FSM (see 8.7).

### **8.6.3 Continuous Link State Database synchronization**

After initial database synchronization with its neighbors, a Switch shall maintain a synchronized database through a continuous database synchronization process. Continuous database synchronization is achieved via reliable flooding of the LSRs. This assures that the databases reflect the current topology of the Fabric.

The current topology of the Fabric may change as a result of the following:

- a) Inter-Switch Links changing state;
- b) a Switch problem that causes it not to respond to Hello messages; or
- c) an ISL becoming operational and the neighbor FSM going to the Full state.

When a Switch detects a local Fabric topology change, it shall flood the Fabric with a new LSR.

A Switch shall issue a new LSR at the LS\_Refresh\_Time to ensure that the Switch shall delete entries in its database after the Max\_Age interval expires if they are not refreshed. This allows for Switches that are permanently disconnected from the Fabric to be removed from the database. The periodic LSR update is independent of any other updates.

#### **8.6.4 Reliable flooding**

##### **8.6.4.1 Basic operation**

Reliable flooding is the mechanism by which topology changes are propagated throughout the Fabric. Reliable flooding shall be used whenever any change of a Switch link state occurs. Reliable flooding shall not be used for the initial database synchronization when an ISL between two Switches initializes. Normally, flooding occurs on all ISLs that are in the Full state at the same time, and not just between two Switches. Further, reliable flooding carries the new LSR(s) hop by hop to all Switches in the Fabric, whereas the initial database synchronization involves only two Switches.

Reliable flooding and initial database synchronization shall use LSU and LSA message structures for the updates.

##### **8.6.4.2 The flooding procedure**

The flooding procedure starts when a Switch issues a new instance of its LSR. The new Switch Link Record LSR for a Domain\_ID shall only be issued by a Switch with that same Domain\_ID.

The originating Switch shall package the LSR in an LSU and shall transmit it on all ISLs in the Full state. If there are other LSRs that are waiting to be acknowledged on an ISL, and the timer Rxmt\_Interval for that ISL has expired, all the LSRs that have been waiting longer than Rxmt\_Interval may be included in the LSU.

A receiving Switch shall acknowledge the LSR if appropriate. If it is a more recent instance than the one in its Link State Database, the Switch replaces the instance in the database with the new one. The Switch shall send the new LSR on all ISLs in the Full state, except the one from which the LSR was received. This step insures that the LSR is actually flooded throughout the Fabric.

If there are physical loops in the Fabric, a Switch may receive multiple instances of the same LSR update, even an instance that was originated by the Switch itself. A potential forwarding loop is prevented by forwarding only LSRs that are newer than the ones currently in the database. If a Switch receives an older instance of an LSR, or an LSR of the same instance as the one contained in the database, it shall acknowledge the LSR, and shall not forward the LSR to other Switches.

An LSR shall be acknowledged by sending the LSR header packaged in an LSA back to the sender. One or more LSRs may be acknowledged in the same LSA. The sender shall stop transmitting an LSR after it receives the acknowledgement.

### 8.6.4.3 Generating a new LSR

When a Switch generates an LSR, it shall set the LSR Age field to 0000h and increment the incarnation number by one. Typically different instances of an LSR have a different incarnation number that indicates a more recent instance. Under some circumstances this information is not sufficient, and other fields in the LSR shall be taken into account. The complete algorithm to determine the most recent incarnation between two LSRs is described below.

When a Switch first initializes, its Link State Database is empty because it has not recognized any neighbors yet. Link State Database information shall not be stored in non-volatile memory or be retrieved after a re-initialization. The Switch shall build a new database at every initialization or re-initialization.

A Switch generates its first LSR when the first ISL enters the Full state. The first LSR shall have an incarnation number of 80000001h. As other ISLs enter the Full state, the Switch shall generate a new incarnation of its LSR and shall increase the incarnation number by one every time. The LSR Age of a newly generated LSR shall always be 0000h.

After generating a new instance of an LSR, the Switch stores it into its Link State Database, and floods it to the rest of the Fabric.

### 8.6.4.4 Transmitting an LSR

An LSR shall be transmitted to a neighbor embedded in an LSU. Before transmission, the age of the LSR shall be incremented by 1. This value represents a nominal delay incurred by the LSR when it is transmitted.

NOTE 17 – The purpose of this increment is to prevent an LSR from being retransmitted forever (e.g., because of a software error).

The LSR shall be acknowledged by the receiving Switch. If the acknowledgement is not received within Rxmt\_Interval the LSR shall be retransmitted. The LSR shall be retransmitted until an acknowledgment is received, or until the neighbor exits the Full state.

There shall be no distinction between the first transmission and subsequent retransmissions of an LSR except for the LSR Age field. The LSRs shall be identical with the exception of the LSR Age field. If a newer instance of the LSR being retransmitted is received, the newer instance shall replace the older instance in the Link State Database. Since more than one LSR may be queued waiting for an acknowledgement, all of them may be transmitted in a single LSU for efficiency.

An LSR update shall not be sent more frequently than the Min\_LS\_Interval.

### 8.6.4.5 Receiving an LSR

An LSR is received in an LSU. After the processing of the LSU header, each LSR shall be processed separately and acknowledgements shall be provided separately for each LSR contained in the LSU. There is no specific acknowledgement to an LSU. Acknowledgements to multiple LSRs may be contained in a single LSA message.

Upon receipt, the following checks are performed in order:

- 1) the checksum is verified. If the checksum is incorrect, the LSR shall be ignored and no acknowledgement shall be returned;

- 2) the LSR Type is checked. If the type is not recognized, the LSR is ignored and no acknowledgement is returned;
- 3) if the LSR has an age equal to Max\_Age, the LSR shall be stored in the local database long enough to flood it and receive acknowledgements if already present in the database or if the neighboring Switches are in the initial database synchronization process (i.e., the states Init, Database Exchange, Database ACK Wait, Database Wait). The LSR is then deleted from the database. This ensures that when an LSR's age reaches Max\_Age in any Switch, it is removed from the databases of all Switches simultaneously;
- 4) if Min\_LS\_Arrival has not expired, then the LSR is ignored and no acknowledgement is returned;
- 5) if there is no such LSR in the Link State Database, or the received LSR is a more recent incarnation than what is stored in the database, then the new LSR is installed in the database. If an LSR existed in the database, then the older incarnation is removed from the database, and an acknowledgement is returned; and
- 6) if a Switch receives an LSR containing its own Domain\_ID in the Link State identifier field, but with an incarnation number greater than its current incarnation number, the Switch shall set the incarnation number of its current LSR to the value in the received LSR plus one, and flood the resulting LSR on all links.

## 8.7 Neighbor finite state machine (FSM)

The Neighbor FSM initializes to Down state. In this state, the FSM is waiting for the notification that the port is connected to another Switch. This notification is issued internally to the Switch when a port reaches the E\_Port status.

After the FSM receives the E\_Port input, it transitions to Init state. In this state, attempts are made to determine if the other Switch supports FSPF. If it does, these attempts may result in the establishment of two-way communication between the two Switches, which is essential for the operation of the protocol. The Link State Databases on the two Switches are unable to be reliably synchronized in the absence of two-way communication. Successful two-way communication depends on configurable parameters matching between the two Switches.

After the two-way communication is established, a Switch knows the Domain\_ID and the Port Index of the neighbor Switch on the opposite side of the ISL, and the FSM transitions to Database Exchange state.

In Database Exchange state, the two neighbor Switches share their view of the Fabric topology by exchanging their complete Link State Database. Each entry in the database is called a Link State Record (LSR). A Switch compares every LSR it receives from the neighbor to the same LSR in its database. If the received LSR is newer than the one present in the database, or if there is no LSR exists for that Domain in the database, then the received LSR is entered in the database. In the case where an LSR already exists in the database, the new LSR supersedes that LSR in the database. At the end of the process, both Switches shall have an identical Link State Database that consists of the most recent LSRs.

From Database Exchange state, the FSM may transition to two different states. If the next event is the reception of the end of database message from the neighbor, it transitions to Database Ack Wait state. If the next event is the reception of an ack to the end of database message, it transitions to Database Wait state.

From Database Ack Wait state, the FSM transitions to Full state when it receives an ack to the end of database message.

From Database Wait state, the FSM transitions to Full state when it receives the end of database message from the neighbor.

Once in Full state, the neighbor becomes an Adjacency, and the ISL that joins the two Adjacent Switches may be used to forward user data. Both Switches shall issue a new instance of their LSR to inform the rest of the Fabric about this fact.

The following aspects of the Neighbor FSM are described in detail below:

- a) state by state;
- b) the current state;
- c) an input;
- d) the new state; and
- e) actions taken in response to that input.

States are listed first, and for each state all the legal inputs are described. For ease of documentation, states are ordered. Each state in the list is considered greater than the previous one. The following states are defined, from lower to higher:

- a) Down;
- b) Init;
- c) Database Exchange;
- d) Database Ack Wait;
- e) Database Wait; and
- f) Full.

An instance of the FSM shall run on each E\_Port of the Switch.

**Table 196 – Neighbor finite state machine (Part 1 of 3)**

State	Input	Next State	Actions
Down	E_Port	Init	This input indicates that a port on the Switch is connected to another Switch. Send a Hello message to the neighbor. Start the Hello_Interval timer. The expiration of this periodic timer triggers the transmission of a Hello message to the neighbor.
Init	One-Way Received	Init	This input indicates that a Hello message that did not carry the correct Domain_ID of the local Switch has been received from the neighbor Switch. Start the Dead_Interval Timer. The expiration of the Dead_Interval Timer causes the Neighbor FSM to transition to Init State.
Init	Two-Way Received	Database Exchange	This input indicates that a Hello message carrying both the remote and the local Domain_ID has been received from the neighbor Switch. Send the Link State Database to the neighbor. The database may be sent in multiple frames, and even in multiple Fibre Channel Sequences. Restart the Dead_Interval Timer.

**Table 196 – Neighbor finite state machine (Part 2 of 3)**

State	Input	Next State	Actions
Database Exchange	Database Received	Database Ack Wait	This input indicates that an LSU with the Database Complete flag set has been received from the neighbor Switch. No action necessary, just a state transition.
Database Exchange	Database Sent	Database Exchange	This input indicates that all the Link State Records that describe the Link State Database have been sent to the neighbor. Send an LSU to the neighbor with no LSRs and the Database Complete flag set.
Database Exchange	Database Acked	Database Wait	This input indicates that an LSA with the Database Complete flag set has been received from the neighbor Switch. No action necessary, just a state transition.
Database Ack Wait	Database Sent	Database Ack Wait	This input indicates that all the Link State Records that describe the Link State Database have been sent to the neighbor. Send an LSU to the neighbor with no LSRs and the Database Complete flag set.
Database Ack Wait	Database Acked	Full	This input indicates that an LSA with the Database Complete flag set has been received from the neighbor Switch. Issue a new instance of the LSR, to reflect the new Adjacency. Compute the paths to all the Switches and program the routing tables.
Database Wait	Database Received	Full	This input indicates that an LSU with the Database Complete flag set has been received from the neighbor Switch. Issue a new instance of the LSR, to reflect the new Adjacency. Compute the paths to all the Switches and program the routing tables.
Any state except Down and Init	Two-Way Received	Same state	This input indicates that a Hello message carrying both the remote and the local Domain_ID has been received from the neighbor Switch. This is a normal periodic Hello message received from the neighbor. Restart the Dead_Interval Timer.
Any state except Down and Init	One-Way Received	Init	This input indicates that a Hello message that did not carry the correct Domain_ID of the local Switch has been received from the neighbor Switch. The retransmission lists shall be emptied, and all the timers associated with the retransmission lists shall be stopped.



**Table 196 – Neighbor finite state machine (Part 3 of 3)**

State	Input	Next State	Actions
Any state	Port Offline	Down	This input indicates that a port went offline. All the data structures related to the neighbor shall be removed. The retransmission lists shall be emptied, and all the timers associated with the neighbor shall be stopped. These include the retransmission timers, the Hello_Interval Timer and the Dead_Interval Timer. The same port may come back connected to a different Switch, or even as an F_Port, in which case the Neighbor FSM does not run. Issue a new instance of the LSR, that excludes this neighbor, to reflect the removal of an Adjacency.
Any state except Down	Hello_Interval	Same state	This input indicates that the Hello_Interval Timer has expired. Send a Hello message to the neighbor. In Down state Hello messages shall not be sent.
Any state except Down	Hello_Dead_Interval	Init	This input indicates that the Dead_Interval Timer has expired and the port is still in E_Port state. Re-initialize the data structures associated with the neighbor. Stop the Dead_Interval Timer. Issue a new instance of the LSR, that excludes this neighbor, to reflect the removal of an Adjacency.
Full	Initial Database Received	Init	This input indicates that an LSU with the Initial Database Exchange flag set has been received from the neighbor Switch. Go into Init state and send One-Way Hellos.
Full	LSU/LSA Received	Full	Indicates that a LSU or LSA was received. The Switch may reset the Dead_Interval Timer.

Figure 26 is a pictorial representation of the FSM where only the major state transitions are represented.

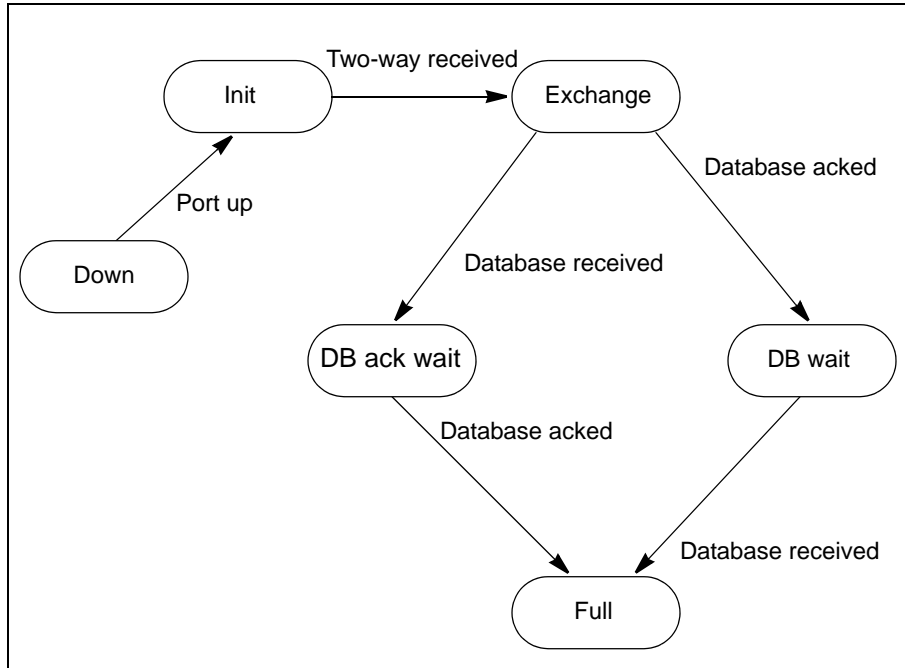


Figure 26 – Neighbor finite state machine

## 9 Distributed Services

### 9.1 Basic model

A distributed services model is used to allow a Fabric to provide consistent services to all attached N\_Ports. This standard defines a common framework by which all Distributed Services communicate. Specific mappings onto this framework are also specified for the distributed Name Server and the distributed Management Server. Please note that in the following discussion it is convenient to say that a Server is “contained” within a Switch. In this case the term “contain” does not imply that an entity is physically inside the Switch (i.e., it may be physically outside the Switch, and still operate as described below).

### 9.2 Distributed Services framework

#### 9.2.1 Goals and characteristics of the Distributed Services framework

All Distributed Services are mapped onto a common framework. The goal of this framework is three-fold:

- a) define a consistent method for distributing services across Switches in a Fabric;
- b) define a distribution method that is topology independent; and
- c) define a method that preserves processing facilities for existing frame formats.

In order to accomplish these goals this standard defines the following Distributed Server characteristics:

- a) transport;
- b) common characteristics;
- c) work categories; and
- d) frame formats.

#### 9.2.2 Distributed Service transport

##### 9.2.2.1 Required FC-2 parameters

Generic Service requests and responses are transported between Distributed Servers using the Common Transport (FC-CT) defined by FC-GS-8.

All CT frames shall be transmitted using the Class F service. The following defines the FC-2 header fields of all Distributed Services frames:

**R\_CTL:** This field shall be set to 02h for all request frames, and to 03h for all reply frames.

**CS\_CTL:** This field shall be set to 00h.

**D\_ID:** This field shall be set to the Domain Controller Identifier of the destination Switch.

**S\_ID:** This field shall be set to the Domain Controller Identifier of the source Switch.

**TYPE:** This field shall be set to 20h, indicating Fibre Channel Fabric Generic Services.

Each request shall be the first Sequence of an Exchange and the associated response shall be the last Sequence of the same Exchange. All other fields shall be set as appropriate according to the rules defined in FC-FS-5.

### 9.2.2.2 FC-CT Header usage

The following values shall be set in the FC-CT Header for all Distributed Services requests and responses:

**FC-CT revision:** Obsolete in this standard, set to 01h.

**IN\_ID:** The value of IN\_ID in a Switch related request shall be preserved in all responses to that request. This only applies to an IN\_ID value set by the Entry Switch.

**Options:** The X\_Bit shall be set to 0 to indicate a single bidirectional exchange per request/response. The Partial Response bit shall be set to zero in Switch-to-Switch requests.

NOTE 18 – Multiple requests/responses may be active using multiple bidirectional exchanges between any pair of Switches.

### 9.2.2.3 Frame distribution

It is important to note that for a Distributed Services request, a remote Switch shall never send a response directly to the requesting Nx\_Port. All responses shall be sent to the Entry Switch. It is the responsibility of the Entry Switch to send the appropriate response to the requesting Nx\_Port.

Furthermore, an Nx\_Port shall always communicate to a Distributed Service via the well-known address of the Distributed Service and Nx\_Ports shall not send Distributed Service requests to Domain Controllers. In addition, Distributed Services request and responses are transported only between Switches and not between a Switch and an Nx\_Port.

### 9.2.2.4 Domain Controller Service Parameters

The default values for Domain Controller communication should be as specified in table 197.

**Table 197 – Default Domain Controller Service Parameter values**

Item	Value
Receive Data Field Size	2112
End-to-End Credit	1
Concurrent Sequences	1
Open Sequences per Exchange	1

Optionally, the Domain Controller Service Parameters may be established using the ESS SW\_ILS Domain Controller Capability Object (see 6.2.22.4.7). When a Switch supporting Domain Controller Service Parameter establishment joins a Fabric, it shall use the ESS SW\_ILS (see 6.2.22) to determine the Domain Controller Service Parameters of the other Switches.

### 9.2.3 Common characteristics

Each Distributed Service shares a set of common characteristics. These characteristics shall be defined as follows:

- a) timeouts: for requests between Switches, the timeout value shall be D\_S\_TOV;

- b) local data copies: local data copies may be optionally allowed by a Distributed Service. If a Distributed Service allows local data copies it shall also specify the method by which the integrity of the local copied data is maintained;
- c) Exchange management: each request between Switches shall be mapped to a unique Exchange. Multiple outstanding requests are allowed between a pair of Switches up to the end-to-end credit resources specified by the receiving Switch;
- d) responses: each request sent shall receive a response. If the receiving Switch is unable to perform a requested operation, then it shall respond with a Reject CT\_IU specifying an appropriate reason code and reason code explanation. If a response is not received from all Switches to which a request was sent within the timeout period, then the request shall be considered partial and a response shall be sent back to the Nx\_Port as appropriate for the Service;
- e) partial response: for many requests even a partial response to the requesting Nx\_Port is useful. A partial response may occur for a number of reasons (e.g., one of the Switches a request is directed to is busy and unable to respond within the timeout period, or one of the Switches a request is directed to does not support the service requested). A service may allow partial responses for a subset of its requests. If the response to a request is partial, the service shall set the partial response bit in the CT Header of the response back to the Nx\_Port. This notifies the Nx\_Port that the data in the response may not be complete;

NOTE 19 – There are legacy implementations that return an LS\_RJT instead of a Reject CT\_IU when a receiving Switch is unable to perform a requested operation.

- f) data merge: describes how data from multiple responses is consolidated; and
- g) error recovery: if an error on a Distributed Services frame is detected (e.g., No ACK, P\_BSY), the frame may be retransmitted for a time interval up to D\_S\_TOV.

#### 9.2.4 Zoning considerations

If Zoning is present in a Fabric, Distributed Services may be affected. The following rules shall apply for Zoning with regard to Distributed Services:

- a) Switch-to-Switch communications shall not be Zoned. This only applies to the Class F CT Header based Distributed Services frames; and
- b) Zoning is applied by the Entry Switch. If a particular Distributed Service is affected by Zoning, it is the responsibility of the Entry Switch to make sure that a requesting Nx\_Port does not receive data for that Distributed Service that is outside of the Nx\_Port's Zone.

#### 9.2.5 Work categories

Work categories are definitions that allow consistent mapping of services to Distributed Services. These categories define how each Distributed Service maps its commands given the Distribution characteristic.

The work categories are:

##### Local

Local requests are those that may be handled entirely by the Entry Switch. A request is local for the following reasons:

- a) the data being requested is owned entirely by the Entry Switch. This situation would be dependant on the type of request; and

- b) the Entry Switch is maintaining a local copy of the data being requested. This situation may occur for any request depending on the local data copy rules of the Distributed Service to which the request belongs.

Any request that is determined to be local shall be processed as appropriate for the service as defined in FC-GS-8.

### **1-to-1**

A 1-to-1 request is a request that is unable to be handled entirely by the Entry Switch, but for which the Entry Switch has identified a single remote Switch that may handle the request. The local Switch sends the request frame directly to the Domain Controller of remote Switch.

### **1-to-Many**

A 1-to-Many request is a request that is unable to be handled entirely by the Entry Switch, but for which the Entry Switch has identified multiple remote Switches that may handle the request. The local Switch sends request frames directly to the Domain Controller of all remote Switches that it has identified to contain requested data.

### **1-to-All**

A 1-to-All request is a request that is unable to be handled entirely by the Entry Switch, but for which the Entry Switch is unable to identify the set of remote Switches to query. The Entry Switch sends request frames directly to the Domain Controller of all Switches in the Fabric.

## **9.2.6 Frame formats**

Where possible Distributed Services should use the same frame formats for Switch-to-Switch communications as are used for Nx\_Port requests.

## **9.2.7 FC-CT command restrictions**

To avoid overlap of command codes associated with FC-CT commands originated external to the Fabric with FC-CT commands originated internal to the Fabric, the following FC-CT command codes shall not be used by any well-known Server for the FC-GS-8 client/server interface.

Command codes 0400h-04FFh and E000h-EFFFh: Fabric Internal FC-CT commands.

Command codes F000h-FFFFh: Vendor specific FC-CT commands.

## **9.3 Distributed Name Server**

### **9.3.1 General behavior**

The distributed Name Service is provided as follows:

- a) each Switch contains its own resident Name Server, called a distributed Name Server (dNS);
- b) each dNS within a Switch is responsible for the name entries associated with the Domain(s) assigned to the Switch;
- c) each dNS within a Switch shall only return information associated with the Domain(s) for which the Switch is responsible;
- d) a client Nx\_Port communicates its Name Service request, as defined in FC-GS-8, to the Entry Switch via the well-known address;

- e) the dNS within the local Switch services the request by making any needed requests of other dNS's contained by the other Switches, if the required information is not available locally;
- f) a dNS may maintain local data copies. Integrity of locally copied data is maintained via SW\_RSCN notification. This implies that all Switches shall distribute SW\_RSCN notification throughout the Fabric whenever a change takes place in their local dNS database;
- g) the communication between dNS's to acquire the requested information is transparent to the original requesting client; and
- h) partial responses to dNS queries are allowed. If an Entry Switch sends a partial response back to an Nx\_Port it shall set the partial response bit in the CT Header.

### **9.3.2 FC-CT for distributed Name Servers**

#### **9.3.2.1 dNS command codes**

The command codes for FC-CT requests defined for dNS use are summarized in table 198. Codes 0100h - 0300h shall be as defined in FC-GS-8. All other requests are defined below. The format of

the Entry field used in the following commands follow the Name Server Entry format described in 9.3.3.

**Table 198 – FC-CT command codes for DNS (Part 1 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Notes
0100	Get all next	GA_NXT	1-to-All	FC-GS-8	
0101	Get Identifiers According to Scope	GID_A	1-to-1 or 1-to-Many	FC-GS-8	<sup>a</sup>
0112	Get Port_Name, based on Port Identifier	GPN_ID	1-to-1	FC-GS-8	
0113	Get Node_Name, based on Port Identifier	GNN_ID	1-to-1	FC-GS-8	
0114	Get Class of Service, based on Port Identifier	GCS_ID	1-to-1	FC-GS-8	
0117	Get FC-4 TYPEs, based on Port Identifier	GFT_ID	1-to-1	FC-GS-8	
0118	Get Symbolic Port_Name, based on Port Identifier	GSPN_ID	1-to-1	FC-GS-8	
011A	Get Port Type, based on Port Identifier	GPT_ID	1-to-1	FC-GS-8	
011B	Obsolete				
011C	Get Fabric Port_Name, based on Port Identifier	GFPN_ID	1-to-1	FC-GS-8	
011D	Get Hard Address, based on Port Identifier	GHA_ID	1-to-1	FC-GS-8	
011E	Obsolete				
011F	Get FC-4 Features - Port Identifier	GFF_ID	1-to-1	FC-GS-8	
0121	Get Port Identifier, based on Port_Name	GID_PN	1-to-All	FC-GS-8	

- <sup>a</sup> The GID\_A request may either be 1-to-1 or 1-to-Many depending on the format of the request.
- <sup>b</sup> Registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.
- <sup>c</sup> De-registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.
- <sup>d</sup> Work categories for Name Server Entry Object requests are at the discretion of the originating Switch.



**Table 198 – FC-CT command codes for dNS (Part 2 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Notes
012B	Obsolete				
0131	Get Port Identifier, based on Node_Name	GID_NN	1-to-All	FC-GS-8	
0132	Get Port_Names based on Node_Name	GPN_NN	1-to-All	FC-GS-8	
0135	Obsolete				
0136	Get Initial Process Associator, based on Node_Name	GIPA_NN	1-to-All	FC-GS-8	
0139	Get Symbolic Node_Name, based on Node_Name	GSNN_NN	1-to-All	FC-GS-8	
0153	Obsolete				
0156	Obsolete				
0171	Get Port Identifiers, based on FC-4 TYPE	GID_FT	1-to-All	FC-GS-8	
0172	Get Port_Names, based on FC-4 TYPE	GPN_FT	1-to-All	FC-GS-8	
0173	Get Node_Names, based on FC-4 TYPE	GNN_FT	1-to-All	FC-GS-8	
01A1	Get Port Identifiers, based on Port Type	GID_PT	1-to-All	FC-GS-8	
01B1	Obsolete				
01B2	Obsolete				
01C1	Get Port Identifiers, based on Fabric Port_Name	GID_FPN	1-to-All	FC-GS-8	
01D1	Get Permanent Port_Name	GPPN_ID	1-to-All	FC-GS-8	
<p><sup>a</sup> The GID_A request may either be 1-to-1 or 1-to-Many depending on the format of the request.</p> <p><sup>b</sup> Registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.</p> <p><sup>c</sup> De-registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.</p> <p><sup>d</sup> Work categories for Name Server Entry Object requests are at the discretion of the originating Switch.</p>					

**Table 198 – FC-CT command codes for dNS (Part 3 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Notes
01F1	Get Port Identifiers, based on FC-4 Features	GID_FF	1-to-All	FC-GS-8	
01F2	Get Port_Identifier	GID_DP	1-to-All	FC-GS-8	
0212	Obsolete				
0213	Register Node_Name	RNN_ID	1-to-1	FC-GS-8	b
0214	Register Class of Service	RCS_ID	1-to-1	FC-GS-8	b
0217	Register FC-4 TYPEs	RFT_ID	1-to-1	FC-GS-8	b
0218	Register Symbolic Port_Name	RSPN_ID	1-to-1	FC-GS-8	b
021A	Obsolete				
021B	Obsolete				
021D	Register Hard Address - Port Identifier	RHA_ID	1-to-1	FC-GS-8	b
021E	Obsolete				
021F	Register FC-4 Features - Port Identifier	RFF_ID	1-to-1	FC-GS-8	b
0235	Register IP Address (Node)	RIP_NN	1-to-All	FC-GS-8	b
0236	Register Initial Process Associator	RIPA_NN	1-to-All	FC-GS-8	b
0239	Register Symbolic Node_Name	RSNN_NN	1-to-All	FC-GS-8	b
0300	De-register all	DA_ID	1-to-1	FC-GS-8	c
0410	Get Entry, based on Port Identifier	GE_ID	Any	FC-SW-7	d
0420	Get Entry, based on Port_Name	GE_PN	Any	FC-SW-7	d

- <sup>a</sup> The GID\_A request may either be 1-to-1 or 1-to-Many depending on the format of the request.
- <sup>b</sup> Registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.
- <sup>c</sup> De-registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.
- <sup>d</sup> Work categories for Name Server Entry Object requests are at the discretion of the originating Switch.

**Table 198 – FC-CT command codes for dNS (Part 4 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Notes
0430	Get Entries, based on Node_Name	GE_NN	Any	FC-SW-7	<sup>d</sup>
0450	Obsolete				
0470	Get Entries, based on FC-4 TYPE	GE_FT	Any	FC-SW-7	<sup>d</sup>
04A0	Get Entries, based on Port Type	GE_PT	Any	FC-SW-7	<sup>d</sup>
04B0	Get Entries, based on Zone Member	GE_ZM	Any	FC-SW-7	
04C0	Get Entries, Based on Zone Name	GE_ZN	Any	FC-SW-7	
04D0	Obsolete				
04E0	Get Entries, Based on FC-4 Features	GE_FF	Any	FC-SW-7	
04F0	Get Entries, Based on Fabric Port_Name	GE_FPN	Any	FC-SW-7	
8001	Reject CT_IU	CT_RJT	1-to-1	FC-GS-8	
8002	Accept CT_IU	CT_ACC	1-to-1	FC-GS-8	

<sup>a</sup> The GID\_A request may either be 1-to-1 or 1-to-Many depending on the format of the request.

<sup>b</sup> Registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.

<sup>c</sup> De-registration requests may be distributed if performed by a third party. Otherwise, they are local and outside the scope of this document.

<sup>d</sup> Work categories for Name Server Entry Object requests are at the discretion of the originating Switch.

### 9.3.2.2 FC-CT Header usage for dNS

The following FC-CT Header parameters, beyond those defined in 9.2.2.2, shall be used for dNS frames:

- a) **GS\_Type:** FCh (i.e., Directory Service application);
- b) **GS\_Subtype:** 02h (i.e., Name Service); and
- c) **Command code:** see table 198.

### 9.3.3 Name Server Objects

The Name Server Objects communicated between distributed Name Servers using FC-CT are as defined in FC-GS-8 with no modification, but with several additions. The format of a Name Server Entry Object is as specified in table 199.

**Table 199 – Name Server Entry Object**

<b>Mandatory</b>	<b>Item</b>	<b>Size (bytes)</b>	<b>Comments</b>
Yes	Entry Object Format Indicator	1	
Yes	Owner Identifier	3	
Yes	Port Type	1	
Yes	N_Port_ID	3	
Yes	N_Port_Name	8	
No	Port Symbolic Name	256	<sup>a</sup>
Yes	Node_Name	8	
No	Node Symbolic Name	256	<sup>a</sup>
Yes	Initial Process Associator	8	
Yes	Obsolete	16	
Yes	Class of Service	4	
Yes	FC-4 TYPEs	32	
Yes	Obsolete	16	
Yes	F_Port_Name	8	
Yes	Reserved	1	
Yes	Hard Address	3	
No	FC-4 Features	128	<sup>b</sup>
<sup>a</sup> This field is not present in the Small Name Server Entry Object. <sup>b</sup> This field is not present in the Large or Small Name Server Entry Objects.			

All fields shall be fixed length as indicated in the table 199. The Owner Identifier shall be the Domain Controller Identifier for the Switch that owns this Entry. All other fields shall be formatted as defined in FC-GS-8.

The Entry Object Format Indicator is specified in table 200.

**Table 200 – Entry Object Format Indicator**

Bit	Description (Bit Value=1)
0	The Port Symbolic Name and Node Symbolic Name are not included in the Entry Object
1	The FC-4 Features are Included in the Entry Object
2-7	Reserved

The sizes of the Name Server Entry Object is specified in table 201.

**Table 201 – Name Server Entry Object description**

Value (Hex)	Length (bytes)	Description
00	624	Large Name Server Entry Object
01	112	Small Name Server Entry Object
02	1012	Large Name Server Entry Object + FC-4 Features
03	500	Small Name Server Entry Object + FC-4 Features

The normal response to Get Entry requests in a distributed Name Server model returns one or more Name Server Entry Objects.

When a response to a request contains either a Port Symbolic Name or Node Symbolic Name that is greater than zero in length, and does not contain FC-4 Features, the Name Server Entry Object with an Entry Object Format Indicator of 00h shall be used by the responder.

The responder may return the Name Server Entry Object with an Entry Object Format Indicator of 01h if neither a Port Symbolic Name, or Node Symbolic Name is registered for the port and would result in those Name Server Objects being of length zero, and FC-4 Features have not been registered for the port.

When a response to a request contains either a Port Symbolic Name or Node Symbolic Name that is greater than zero in length, and contains FC-4 Features, the Name Server Entry Object with an Entry Object Format Indicator of 02h shall be used by the responder.

The responder shall return the Name Server Entry Object with an Entry Object Format Indicator of 03h if it would contain FC-4 Features, and does not contain a Port Symbolic Name or Node Symbolic Name.

### 9.3.4 FC-CT requests for dNS

#### 9.3.4.1 Get Entry based on Port Identifier

The dNS shall, when it receives a GE\_ID request, return the Entry object for the specified Port Identifier. The format of the GE\_ID request is specified in table 202.

**Table 202 – GE\_ID request payload**

Item	Size (bytes)
FC-CT Header	16
Reserved	1
Port Identifier	3

The Port Identifier format shall be as defined in FC-GS-8. The dNS may reject a GE\_ID request for reasons not specified in this standard.

The format of the reply payload to a GE\_ID request is specified in table 203.

**Table 203 – GE\_ID accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry	n

Since this request returns only one Entry, the Number of Entries field shall always be set to one for this reply. The Entry field shall contain the Entry for the requested Port Identifier.

#### 9.3.4.2 Get Entry based on Port\_Name

The dNS shall, when it receives a GE\_PN request, return the Entry object for the specified Port\_Name. The format of the GE\_PN request is specified in table 204.

**Table 204 – GE\_PN request payload**

Item	Size (bytes)
FC-CT Header	16
Port_Name	8

The Port\_Name format shall be as defined in FC-GS-8. The dNS may reject a GE\_PN request for reasons not specified in this standard.

The format of the reply payload to a GE\_PN request is specified in table 205.

**Table 205 – GE\_PN accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry	n

Since this request returns only one Entry, the Number of Entries field shall always be set to one for this reply. The Entry field shall contain the Entry for the requested Port\_Name.

**9.3.4.3 Get Entries based on Node\_Name**

The dNS shall, when it receives a GE\_NN request, return the Entry object for the specified Node\_Name. The format of the GE\_NN request is specified in table 206.

**Table 206 – GE\_NN request payload**

Item	Size (bytes)
FC-CT Header	16
Node_Name	8

The Node\_Name format shall be as defined in FC-GS-8. The dNS may reject a GE\_NN request for reasons not specified in this standard.

The format of the reply payload to a GE\_NN request is specified in table 207.

**Table 207 – GE\_NN accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of Entries returned. Each Entry field shall contain an Entry for the requested Node\_Name.

#### 9.3.4.4 Get Entries based on FC-4 TYPEs

The dNS shall, when it receives a GE\_FT request, return the Entry object for the specified FC-4 TYPEs. Note that more than one FC-4 TYPE may be specified. The format of the GE\_FT request is specified in table 208.

**Table 208 – GE\_FT request payload**

Item	Size (bytes)
FC-CT Header	16
FC-4 TYPEs	32

The FC-4 TYPE format shall be as defined in FC-GS-8. The dNS may reject a GE\_FT request for reasons not specified in this standard.

The format of the reply payload to a GE\_FT request is specified in table 209.

**Table 209 – GE\_FT accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of Entries returned. Each Entry field shall contain an Entry for the requested FC-4 TYPEs.

#### 9.3.4.5 Get Entries based on Port Type

The dNS shall, when it receives a GE\_PT request, return the Entry object for the specified Port Type. The format of the GE\_PT request is specified in table 210.

**Table 210 – GE\_PT request payload**

Item	Size (bytes)
FC-CT Header	16
Reserved	3
Port Type	1



The Port Type format shall be as defined in FC-GS-8. The dNS may reject a GE\_PT request for reasons not specified in this standard.

The format of the reply payload to a GE\_PT request is specified in table 211.

**Table 211 – GE\_PT accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of Entries returned. Each Entry field shall contain an Entry for the requested Port Type.

#### 9.3.4.6 Get Entries based on Zone Member

The dNS shall, when it receives a GE\_ZM request, return the Entry objects that are in the same Zone as the Zone Member specified in the GE\_ZM request. The format of the GE\_ZM request is specified in table 212.

**Table 212 – GE\_ZM request payload**

Item	Size (bytes)
FC-CT Header	16
Zone Member	n

The Zone Member format shall be as defined in 10.4.4.6.1.

The format of the reply payload to a GE\_ZM request is specified in table 213.

**Table 213 – GE\_ZM accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of Entries returned. Each Entry field shall contain an Entry for a port that is in the same Zone as the Zone Member specified in the GE\_ZM request. Each Entry shall be only for ports local to the Switch to which the request was sent.

**9.3.4.7 Get Entries based on Zone Name**

The dns shall, when it receives a GE\_ZN request, return the Entry objects that are in the same Zone as the Zone Name indicates in the GE\_ZN request. The format of the GE\_ZN request is specified in table 214.

**Table 214 – GE\_ZN request payload**

Item	Size (bytes)
FC-CT Header	16
Zone Name	n

The Zone Name format shall be as defined in 10.4.2.3.

The format of the reply payload to a GE\_ZN request is specified in table 215.

**Table 215 – GE\_ZN accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of Entries returned. Each Entry field shall contain an Entry for a port that is in the same Zone as the Zone Name specified in the GE\_ZN request. Each Entry shall be only for ports local to the Switch to which the request was sent.

**9.3.4.8 Get Entries based on FC-4 Features**

The dNS shall, when it receives a GE\_FF request, return the Entry objects for the specified FC-4 features code specified in the GE\_FF request. The format of the GE\_FF request is specified in table 216.

**Table 216 – GE\_FF request payload**

Item	Size (bytes)
FC-CT Header	16
FC-4 Features	128

The format of the FC-4 Features value is as specified in FC-GS-8. The dNS may reject a GE\_FF for reasons not specified in this standard.

The format of the reply payload to a GE\_FF request is specified in table 217.

**Table 217 – GE\_FF accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of entries returned. Each Entry field shall contain an Entry for the requested FC-4 features code.

### 9.3.4.9 Get Entries based on Fabric Port\_Name

The dNS shall, when it receives a GE\_FPN request, return the Entry objects for the specified Fabric Port\_Name specified in the GE\_FPN request. The format of the GE\_FPN request is specified in table 218.

**Table 218 – GE\_FPN request payload**

Item	Size (bytes)
FC-CT Header	16
Fabric Port_Name	8

The Fabric Port\_Name format shall be as defined in FC-GS-8. The dNS may reject a GE\_FPN request for reasons not specified in this standard. The format of the reply payload to a GE\_FPN request is specified in table 219.

**Table 219 – GE\_FPN accept payload**

Item	Size (bytes)
FC-CT Header	16
Number of Entries	4
Entry 1	n
...	
Entry N	n

The Number of Entries field shall be set to the number of entries returned. Each Entry field shall contain an Entry for the requested Fabric Port\_Name. There is one entry for an F\_Port and more than 1 entry for an FL\_Port.

## 9.4 Distributed Management Servers

### 9.4.1 General behavior

Servers that comprise the Management Service are provided as follows:

- a) for each Management Server of the Management Service, each Switch contains its own instance of the Server. Generically, each Server instance is called a distributed Management Server (dMS);
- b) each dMS within a Switch is responsible for the entries associated with the Domain(s) assigned to the Switch;
- c) a client Nx\_Port communicates its Management Server request, as defined in FC-GS-8, to the Entry Switch via the well-known address and appropriate sub-type;
- d) a dMS within the Entry Switch services the request by making any needed requests of other dMS instances contained by the other Switches, if the required information is not available locally;

- e) a dMS may maintain local data copies, and a dMS shall notify other dMS that they should remove local data copies;
- f) the communication between each dMS to acquire the requested information is transparent to the original requesting client;
- g) partial responses for some dMS requests are allowed. Partial response support is specified in the following clauses on a per dMS basis; and
- h) the responses returned to a client for some dMS Servers are not subject to Zoning as indicated in table 220.

**Table 220 – Zoning effect on Servers of the distributed Management Service**

Server	Subject to Zoning
Fabric Configuration Server	No
Unzoned Name Server	No
Fabric Zone Server	No
Enhanced Fabric Configuration Server	No
Fabric Device Management Interface Server	No
Application Server	No
Security Information Server <sup>a</sup>	
<sup>a</sup> The impact of Zoning on the Security Information Server is separately specified for each Security Information request (see FC-GS-8).	

**9.4.2 FC-CT Header**

**9.4.2.1 FC-CT Header parameters**

FC-CT Header parameters, beyond those defined in 9.2.2.2, that shall be used for dMS frames are specified in table 221.

**Table 221 – dMS FC-CT Header parameters**

Item	Value
CT_Type	FAh (Management Service)
CT_Subtype	00h - Non-Server Specific 01h - Fabric Configuration Server 02h - Unzoned Name Server 07h - Security Information Server 08h - Enhanced Fabric Configuration Server 10h - Fabric Device Management Server 20h - Application Server others - Reserved
Command code	see tables 198, 220, 222, 227, 232, 233 and 240

### 9.4.2.2 FC-CT Header rule for Fabric internal requests

For non-Server specific requests, the GS\_Subtype value shall be set to 00h.

### 9.4.3 Fabric Configuration Service

The FC-CT command codes defined for use by Fabric Configuration Service requests of the distributed Management Server are specified in table 222.

**Table 222 – Fabric Configuration Service command codes for dMS (Part 1 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Capability See 6.2.22.4.3
0100	Get Topology Information	GTIN	1-to-1	FC-GS-8	Topology
0101	Get Interconnect Element (IE) List <sup>a</sup>	GIEL	Local	FC-GS-8	Basic
0111	Get Interconnect Element Type	GIET	Local or 1-to-1	FC-GS-8	Basic
0112	Get Domain Identifier <sup>a</sup>	GDID	Local	FC-GS-8	Basic
0113	Get Management Identifier	GMID	Local or 1-to-1	FC-GS-8	Basic
0114	Get Fabric Name <sup>a</sup>	GFN	Local	FC-GS-8	Basic
0115	Get Interconnect Element Logical Name	GIELN	Local or 1-to-1	FC-GS-8	Basic
0116	Get Interconnect Element Management Address List	GMAL	Local or 1-to-1	FC-GS-8	Basic
0117	Get Interconnect Element Information List	GIEIL	Local or 1-to-1	FC-GS-8	Basic
0118	Get Port List	GPL	Local or 1-to-1	FC-GS-8	Basic
0121	Get Port Type	GPT	Local or 1-to-All	FC-GS-8	Basic
<sup>a</sup> These requests are handled by the Entry Switch with no assistance from other Switches. <sup>b</sup> These requests function differently than the other requests (see 6.2.24 and FC-GS-8).					

**Table 222 – Fabric Configuration Service command codes for dMS (Part 2 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Capability See 6.2.22.4.3
0122	Get Physical Port Number	GPPN	Local or 1-to-All	FC-GS-8	Basic
0124	Get Attached Port_Name List	GAPNL	Local or 1-to-All	FC-GS-8	Basic
0126	Get Port State	GPS	Local or 1-to-All	FC-GS-8	Basic
0127	Get Port Speed Capabilities	GPSC	Local or 1-to-All	FC-GS-8	Basic
0128	Get Attached Topology Information	GATIN	Local or 1-to-All	FC-GS-8	Topology
0130	Get Switch Enforcement Status	GSES	Local or 1-to1	FC-GS-8	Enhanced
0140	Get Interconnect Element Attribute Group	GIEAG	Local or 1-to1	FC-GS-8	Enhanced
0141	Get Port Attribute Group	GPAG	Local or 1-to1	FC-GS-8	Enhanced
0191	Get Platform Node_Name List	GPLNL	Local or 1-to-All	FC-GS-8	Platform
0192	Get Platform Type	GPLT	Local or 1-to-All	FC-GS-8	Platform
0193	Get Platform Management Address List	GPLML	Local or 1-to-All	FC-GS-8	Platform
<p><sup>a</sup> These requests are handled by the Entry Switch with no assistance from other Switches.</p> <p><sup>b</sup> These requests function differently than the other requests (see 6.2.24 and FC-GS-8).</p>					

**Table 222 – Fabric Configuration Service command codes for dMS (Part 3 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Capability See 6.2.22.4.3
0197	Get Platform Attribute Block	GPAB	Local or 1-to-All	FC-GS-8	Platform
01A1	Get Platform Name - Node_Name	GNPL	Local or 1-to-All	FC-GS-8	Platform
01A2	Get Platform List	GPNL	Local or 1-to-All	FC-GS-8	Platform
01A4	Get Platform FCP Type	GPFCP	Local or 1-to-All	FC-GS-8	Platform
01A5	Get Platform OS LUN Mappings	GPLI	Local or 1-to-All	FC-GS-8	Platform
01B1	Get Node Identification Data - Node_Name	GNID	Local or 1-to-All	FC-GS-8	Platform
215	Register Interconnect Element Logical Name	RIELN	Local or 1-to-1	FC-GS-8	Basic
0280	Register Platform	RPL	Local or 1-to-All	FC-GS-8	Platform
0291	Register Platform Node_Name	RPLN	Local or 1-to-All	FC-GS-8	Platform
0292	Register Platform Type	RPLT	Local or 1-to-All	FC-GS-8	Platform
0293	Register Platform Management Address	RPLM	Local or 1-to-All	FC-GS-8	Platform
0298	Register Platform Attribute Block	RPAB	Local or 1-to-All	FC-GS-8	Platform
029A	Register Platform FCP Type	RPFCP	Local or 1-to-All	FC-GS-8	Platform
029B	Register Platform OS LUN Mappings	RPLI	Local or 1-to-All	FC-GS-8	Platform

<sup>a</sup> These requests are handled by the Entry Switch with no assistance from other Switches.

<sup>b</sup> These requests function differently than the other requests (see 6.2.24 and FC-GS-8).



**Table 222 – Fabric Configuration Service command codes for dMS (Part 4 of 4)**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Capability See 6.2.22.4.3
0380	Deregister Platform	DPL	Local or 1-to-All	FC-GS-8	Platform
0391	Deregister Platform Node_Name	DPLN	Local or 1-to-All	FC-GS-8	Platform
0392	Deregister Platform Management Address	DPLM	Local or 1-to-All	FC-GS-8	Platform
0393	Deregister Platform Management Address List	DPLML	Local or 1-to-All	FC-GS-8	Platform
0394	Deregister Platform OS LUN Mappings	DPLI	Local or 1-to-All	FC-GS-8	Platform
0395	Deregister Platform Attribute Block	DPAB	Local or 1-to-All	FC-GS-8	Platform
039F	De-Register All Platform Information	DPALL	Local or 1-to-All	FC-GS-8	Platform
0400	FC Trace Route	FTR	NA <sup>b</sup>	FC-GS-8	Topology
0401	FC Ping	FPNG	NA <sup>b</sup>	FC-GS-8	Topology
<p><sup>a</sup> These requests are handled by the Entry Switch with no assistance from other Switches.</p> <p><sup>b</sup> These requests function differently than the other requests (see 6.2.24 and FC-GS-8).</p>					

#### 9.4.4 Unzoned Name Service

The Distributed Service associated with the Unzoned Name Service, shall be identical to the services defined by the dMS in 9.3. These services are classified as the Basic Unzoned Name service.

#### 9.4.5 Fabric Zone Service

The Fabric Zone Service is described in clause 10.

#### 9.4.6 Fabric-Device Management Service

##### 9.4.6.1 Operational characteristics of the FDMI Server

As with the other Servers, the FDMI Server is distributed, but it has additional requirements placed upon it because an HBA may register through ports attached to different Switches. This raises the possibility that multiple Switches in the Fabric may be able to manage information for the same HBA. Since it is desirable that only one Switch manage an HBA's information, the FDMI Server requires that a mechanism be provided above and beyond the normal distribution and caching mechanisms provided for other Servers.

Switches in the Fabric may contain and manage a portion of the FDMI database. Each Switch that manages a portion of the FDMI database participates with other Switches that manage portions of the FDMI database through the exchange of a new FDMI Inter-Switch messages. Each Switch that manages a portion of the FDMI database may also maintain cached information associated with portions of the FDMI database on other Switches.

Each HBA attached to the Fabric has one Switch that functions as its Principal Manager. This ensures that only one Switch manages information on behalf of a given HBA. Through the exchange of FDMI Inter-Switch messages, Switches resolve which one becomes the Principal Manager for each HBA. This requires that each Switch with HBAs attached maintain a map that associates its attached HBA's to the Principal Manager for each HBA.

The GS client refers to the entity that issues an FDMI request to the HBA Management Server via the well-known Management Server address.

### **9.4.6.2 Registration scenarios**

#### **9.4.6.2.1 HBA attached to a single Switch**

In the simple case, an HBA is attached to only one Switch, possibly through multiple ports. The HBA attempts to register with the Switch and the Switch performs the checks as mandated by the FDMI interface specification in FC-GS-8. The Switch uses information contained in its local database and its cache entries to perform these checks. If the checks complete successfully then the HBA information is registered with the Switch and the HBA is notified of successful registration. If subsequent registrations are attempted over additional ports attached to the same Switch, the Switch would reject those requests because the HBA information is already contained in its FDMI database. Following the successful registration of HBA information with the Switch, the Switch may forward the information to other Switches in the Fabric allowing other Switches to update their caches.

#### **9.4.6.2.2 HBA attached to multiple Switches**

In the more complicated case, an HBA is attached to multiple Switches through multiple ports. Since the Switches at this point may have not had time to update each others caches between registrations, the HBA registration may pass the checks in multiple Switches. This means that multiple Switches are managing the information for the same HBA in their respective databases. This is not desirable because inconsistencies may be introduced into the FDMI database when multiple Switches manage information for the same HBA.

#### **9.4.6.2.3 Resolution of the principal HBA manager**

When multiple Switches allow the registration of information for a particular HBA, the Switches shall resolve which Switch acts as the principal manager for the HBA. The Switch with the lowest Switch\_Name shall become the principal HBA manager. Switches that have accepted registrations from an HBA exchange FDMI messages that contain the associated Switch\_Name. Switches participating in the protocol determine from the Switch\_Names which Switch serves as the principal HBA manager for the HBA. All this ensures that one and only one Switch serves as the principal HBA manager for a given HBA, even if the HBA is attached to other Switches. The protocol that determines the principal manager runs immediately following a successful HBA registration (see annex C).

### 9.4.6.3 FDMI Inter-Switch messages

#### 9.4.6.3.1 General format

FDMI Inter-Switch messages are exchanged between Switches using the Inter-Switch FC-CT. The general format of the FDMI Inter-Switch message is specified in table 223.

**Table 223 – FDMI Inter-Switch message format**

Item	Size (bytes)
FC-CT Header	See FC-GS-8
FDMI Header	28
Payload	n

#### 9.4.6.3.2 FC-CT Header

The FC-CT Header follows the format specified in FC-GS-8 and values for the GS\_Type, GS\_Subtype, and Command/Response fields are specified in table 224.

**Table 224 – FC-CT Header parameters**

Item	Value
GS_Type	FAh (Management Service)
GS_Subtype	10h - Fabric Device Management Server
Command/Response Code	see 9.4.6.4

#### 9.4.6.3.3 FDMI Header

The format of the FDMI header is specified in table 225.

**Table 225 – FDMI Header format**

Item	Size (bytes)
FDMI Version	1
Reserved	3
Switch_Name	8
Vendor Specified	16

**FDMI Version:** This field represents the version of the FDMI Header. The only value allowed is 01h. All other values are reserved.

**Switch\_Name:** This field contains the Switch\_Name for the Switch that originated the FDMI CT operation.

**Vendor Specified:**

The format of the Vendor Specified field is specified in table 226.

**Table 226 – Vendor Specified field format**

Item	Size (bytes)
Vendor ID	8
Vendor Specified Information	8

**Vendor ID:** Contains the T10 Vendor ID of the vendor that defines the content of Vendor Specified Information field.

**Vendor Specified Information:** This field contains 8 bytes of vendor specified information. The processing of the Vendor Specified information shall be subject to the following rules:

- a) if the information contained in the Vendor Specified Information field is not recognized or processed by the Server, then the command proceeds as defined; and
- b) for any FDMI command defined in the standard, the Vendor Specified information shall not cause the Server to exhibit any behavior different from that defined for the command.

**9.4.6.3.4 Payload**

The Payload field is either null or contains the GS client payload depending on the FDMI Inter-Switch request (see table 227). The GS client payload includes the entire CT request that was received by the entry Switch from the HBA, including the CT Header.

**9.4.6.4 FDMI Inter-Switch requests**

FDMI Inter-Switch requests are mapped to Request CT\_IUs. Table 227 indicates the operations performed by the HBA Management Server and indicates their associated command codes and payload contents.

**Table 227 – FDMI Fabric Internal command codes (Part 1 of 2)**

Code	Mnemonic	Description	Request attributes	Accept attributes
E100	FCRN	De-Registration Notification	FDMI Header, GS client payload	Null
E101	FRN	Registration Notification	FDMI Header, GS client payload	Null
E102	FUN	Update Notification	FDMI Header, GS client payload	Null

**Table 227 – FDMI Fabric Internal command codes (Part 2 of 2)**

Code	Mnemonic	Description	Request attributes	Accept attributes
E103	FDRF	De-Registration Forward	FDMI Header, GS client payload	Null
E104	FUF	Update Forward	FDMI Header, GS client payload	Null
E105	FETCH	Fetch	FDMI Header	HBA/Port List
E106- E10F		Reserved		

### 9.4.6.5 FDMI Inter-Switch responses

#### 9.4.6.5.1 Reject response

When the destination Switch is unable to perform a requested operation, an HBA Management Server Reject CT\_IU is sent to the originating Switch. HBA Management Server Reject CT\_IUs specify a reason code of 09h (i.e., Unable to perform command request) and a reason code explanation specified in table 228.

**Table 228 – Reason code explanation**

Encoded value (hex)	Description
00	No Additional Explanation
E0	Fetch Unsuccessful
others	Reserved

#### 9.4.6.5.2 Accept response

When the destination Switch has successfully performed the requested operation, an HBA Management Accept CT\_IU is sent to the originating Switch indicating completion of the requested operation, and containing any response information associated with the requested operation.

### 9.4.6.6 FDMI Inter-Switch operations

#### 9.4.6.6.1 Registration Notification (FRN) operation

When the HBA Management Server on the entry Switch registers HBA information in its FDMI database, the HBA Management Server shall send the FRN request to all Switches in the Fabric. The FRN request payload shall specify the FDMI header, and the original CT request from the GS client. The FRN accept payload shall be null.

**9.4.6.6.2 De-Register Notification (FDRN) operation**

When the HBA Management Server on the entry Switch de-registers HBA information in its FDMI database, the HBA Management Server shall send the FDRN request to all Switches in the Fabric. The FDRN request payload shall specify the FDMI header, and the original CT request from the GS client. The FDRN accept payload shall be null.

**9.4.6.6.3 Update Notification (FUN) operation**

When the HBA Management Server on the entry Switch updates HBA information in its FDMI database, the HBA Management Server shall send the FUN request to all Switches in the Fabric. The FUN request payload shall specify the FDMI header, and the original CT request from the GS client. The FUN accept payload shall be null.

**9.4.6.6.4 Update Forward (FUF) operation**

When the HBA Management Server on the entry Switch receives a request to update HBA information, but the Switch is not the Principal HBA Manager for the specified HBA, the HBA Management Server shall send the FUF request to the Switch that is the Principal HBA Manager for the specified HBA. The FUF request payload shall specify the FDMI header, and the original CT request from the GS client. The FUF accept payload shall be null.

**9.4.6.6.5 De-Register Forward (FDRF) operation**

When the HBA Management Server on the entry Switch receives a request to de-register HBA information, but the Switch is not the Principal HBA Manager for the specified HBA, the HBA Management Server shall send the FDRF request to the Switch that is the Principal HBA Manager for the specified HBA. The FDRF request payload shall specify the FDMI header, and the original CT request from the GS client. The FDRF accept payload shall be null.

**9.4.6.6.6 Fetch**

When a Switch becomes part of the Fabric (e.g., result of a Merge), the HBA Management Server shall send the FETCH request to all Switches in the Fabric to obtain their Registered HBA/Port lists. The FETCH request payload shall specify the FDMI header. The FETCH accept payload shall return the Registered HBA/Port list. The format of the Registered HBA/Port list is specified in table 229.

**Table 229 – Registered HBA/Port list format**

Item	Size (bytes)
Number of HBA Entries (n)	4
HBA Entry 1	x
HBA Entry 2	y
...	...
HBA Entry n	z

**Number of HBA Entries:** This field specifies the number of HBA entries contained in the Registered HBA/Port list.

**HBA Entry:** The format of the HBA Entry is specified in table 230.

**Table 230 – HBA Entry format**

Item	Size (bytes)
HBA Identifier	8
Number of Port Entries (m)	4
Port Entry 1	8
Port Entry 2	8
...	...
Port Entry n	8

**Number of Port Entries:** This field specifies the number of Port Entries for the specified HBA.

**Port Entry:** The format of the Port Entry is specified in table 231.

**Table 231 – Port Entry format**

Item	Size (bytes)
Port_Name	8

#### 9.4.6.7 GS client initiated FDMI requests

In addition to the Fabric originated FDMI operations, there are GS client initiated FC-CT commands that are forwarded to other Switches in the Fabric by the entry Switch. The FC-CT command codes

defined for use by FDMI requests of the Distributed HBA Management Server are specified in table 232.

**Table 232 – FDMI CT commands for the dms**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in
0100	Get Registered HBA List	GRHL	Local or 1 to Many	FC-GS-8
0101	Get HBA Attributes	GHAT	Local or 1-to-1	FC-GS-8
0102	Get Registered Port List	GRPL	Local or 1-to-1	FC-GS-8
0110	Get Port Attributes	GPAT	Local or 1-to-1	FC-GS-8
0120	Get Port Statistics	GPAS	Local or 1-to-1	FC-GS-8
0200	Register HBA	RHBA	Local	FC-GS-8
0201	Register HBA Attributes	RHAT	Local	FC-GS-8
0210	Register Port	RPRT	Local	FC-GS-8
0211	Register Port Attributes	RPA	Local	FC-GS-8
0300	De-Register HBA	DHBA	Local	FC-GS-8
0301	De-Register HBA Attributes	DHAT	Local	FC-GS-8
0310	De-Register Port	DPRT	Local	FC-GS-8
0311	De-Register Port Attributes	DPA	Local	FC-GS-8

### 9.4.7 Other Fabric internal services

#### 9.4.7.1 Fabric internal requests

Fabric internal FC-CT requests for the distributed Management Server are specified in table 233.

**Table 233 – Fabric internal Management Server requests**

Code	Mnemonic	Description	Request attributes	Accept attributes
E020	GCAP	Get Management Server Capabilities	None	Management Server Capabilities
E100-E10F		Reserved for Inter-Switch FDMI use (see 9.4.6.4)		



**9.4.7.2 Get Management Server Capabilities (GCAP) request**

**9.4.7.2.1 Overview**

The GCAP request allows a Management Server instance on one Switch to query the Management Server capabilities of a Management Server instance on another Switch in the Fabric.

The responding distributed Management Server shall, when it receives a GCAP request, return its capabilities. The GCAP request payload shall be null. The GCAP accept payload contains the requested Management Server Capabilities. The GCAP request payload format is specified in table 234.

**Table 234 – GCAP request payload format**

Item	Size (bytes)
Null	0

The GCAP CT\_ACC payload format is specified in table 235.

**Table 235 – GCAP CT\_ACC payload format**

Item	Size (bytes)
Number of Capability Entries (n)	4
Capability Entry 1	8
Capability Entry 2	8
...	8
Capability Entry n	8

**9.4.7.2.2 Capability Entry**

The format of the Capability Entry is specified in table 236. The Management Server Subtype indicates the service. The Capability mask designates the supported capabilities of the service.

**Table 236 – Capability Entry format**

Item	Size (bytes)
Management Server GS_Subtype	1
Vendor Specific Capability Bit Mask	3
Subtype Capability Bit Mask	4

**Management Server GS\_Subtype:** This field shall indicate the Management Server associated with the capabilities in the entry.

**Vendor Specific Capability Bit Mask:** This field shall indicate any vendor specific capabilities associated with the designated Management Server. The format of the field is not defined by this standard.

**Subtype Capability Bit Mask:** This field shall indicate capabilities associated with the designated Management Server.

### 9.4.7.2.3 Subtype Capability Bit Masks

The Subtype Capability Bit Masks for the Fabric Configuration Server are specified in table 237.

**Table 237 – Fabric Configuration Server (CT\_Subtype 01h)**

Option Bit(s)	Description (see 6.2.22.4.3)
0	Basic Configuration Services
1	Platform Configuration Services
2	Topology Discovery Configuration Services
3	Enhanced Configuration Services
4-31	Reserved

The Subtype Capability Bit Masks for the Unzoned Name Server are specified in table 238.

**Table 238 – Unzoned Name Server (CT\_Subtype 02h)**

Option Bit(s)	Description (see 9.4.4)
0	Basic Unzoned Name Services
1-31	Reserved

### 9.4.8 Security Information Server

The FC-CT command codes defined for use by Security Information Server requests are summarized in table 239.

**Table 239 – Security Information Server command codes for dMS**

Encoded Value (hex)	Description	Mnemonic	Work category	Payload defined in
0001	Get Authorization State for Port Identifier	GAS_ID	Local or 1 to 1	FC-GS-8

### 9.4.9 Application Server

The FC-CT command codes defined for use by Application Server requests are summarized in table 240.

**Table 240 – Application Server command codes for dMS**

Encoded Value (hex)	Description	Mnemonic	Work category	Payload defined in
0100	Get Application Identifier Allocations - Entity Identifier	GAPPIA_ENT	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0101	Get All Application Identifier Allocations	GALL_APPID	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0102	Get Application Identifier Allocations - N_Port_ID	GALLAPPIA_ID	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0103	Get Application Identifier Allocations - N_Port_ID and Application Identifier	GAPPIA_IDAPP	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0200	Register Application Identifier - N_Port_ID and Entity Identifier	RAPP_IDENT	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0300	Deregister Application Identifier - N_Port_ID and Entity Identifier	DAPP_IDENT	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8
0301	Deregister All Application Identifiers - N_Port_ID	DALLAPP_ID	Local, 1 to 1, 1 to many, or 1 to all	FC-GS-8

### 9.4.10 Enhanced Fabric Configuration Service

The FC-CT command codes defined for use by Enhanced Fabric Configuration Service requests of the distributed Management Server are specified in table 241.

**Table 241 – Enhanced Fabric Configuration Service command codes for dMS**

Encoded value (hex)	Description	Mnemonic	Work category	Payload defined in	Capability See 6.2.22.4.11
0100	Get Interconnect Element (IE) List <sup>a</sup>	GIEL	Local	FC-GS-8	Basic
0101	Get Interconnect Element Attribute Block	GIEAB	Local or 1-to-1	FC-GS-8	Basic
0102	Get Interconnect Element Port List	GIEPL	Local or 1-to-1	FC-GS-8	Basic
0103	Get Fabric Object	GFO	Local or 1-to-1	FC-GS-8	Basic
0110	Get Physical Object Attribute Block	GPOAB	Local or 1-to-1	FC-GS-8	Basic
0111	Get Physical Object Port List	GPOPL	Local or 1-to-1	FC-GS-8	Basic
0130	Get Logical Port Attribute Block	GLPAB	Local or 1-to-1	FC-GS-8	Basic
0131	Get Attached Port List	GAPL	Local or 1-to-1	FC-GS-8	Basic
0140	Get Physical Port Object Attribute Block	GPPOAB	Local or 1-to-1	FC-GS-8	Basic
0200	Register Interconnect Element Logical Name	RIELN	Local or 1-to-1	FC-GS-8	Basic

<sup>a</sup> These requests are handled by the Entry Switch with no assistance from other Switches.

## 9.5 Distributed Event Server

### 9.5.1 General behavior

The distributed Event Server is provided as follows:

- a) each Switch contains its own resident Event Server, called a distributed Event Server (dES);
- b) each dES within a Switch is responsible for the registrations and notifications associated with the Domain(s) assigned to the Switch;
- c) each dES within a Switch shall only originate events associated with the Domain(s) for which the Switch is responsible; and
- d) a client Nx\_Port communicates its Event Server requests, as defined in FC-GS-8, to the Entry Switch via the well-known address.

### 9.5.2 FC-CT for distributed Event Server

#### 9.5.2.1 FC-CT Header parameters

FC-CT Header parameters, beyond those defined in 9.2.2.2, that shall be used for dES frames are specified in table 242.

**Table 242 – dES FC-CT Header parameters**

Item	Value
CT_Type	F4h (Event Service)
CT_Subtype	01h

#### 9.5.2.2 dES command codes

The command codes for FC-CT requests defined for dES use are specified in table 243.

**Table 243 – FC-CT command codes for dES**

Encoded value (Hex)	Description	Work category	Payload defined in
0100	Event Registration	1-to-All	FC-GS-8
0101	Event Notification	1-to-All	FC-GS-8

## 9.6 Distributed VE Identification Server

### 9.6.1 General behavior

The distributed VE Identification Service is provided as follows:

- a) each Switch contains its own resident VE Identification Server, called a distributed VE Identification Server (dVEIS);
- b) each dVEIS within a Switch is responsible for the VE mappings associated with the Domain(s) assigned to the Switch;
- c) a client Nx\_Port communicates its VE Identification Server request, as defined in FC-GS-8, to the Entry Switch via the well-known address and appropriate sub-type;

- d) a dVEIS within the Entry Switch services the request by making any needed requests of other dVEIS instances contained by the other Switches, if the required information is not available locally;
- e) a dVEIS may maintain local VE mapping copies, and a dVEIS shall notify other dVEIS of changes in the VE mappings; and
- f) the communication between each dVEIS to acquire the requested information is transparent to the original requesting client.

### 9.6.2 FC-CT for distributed VE Identification Server

FC-CT Header parameters, beyond those defined in 9.2.2.2, that shall be used for dVEIS frames are specified in table 244.

**Table 244 – dVEIS FC-CT Header parameters**

Item	Value
CT_Type	FCh
CT_Subtype	04h

The command codes for FC-CT requests defined for dVEIS use are specified in table 245.

**Table 245 – FC-CT Command Codes for dVEIS**

Code (hex)	Mnemonic	Description	Work Category	Payload defined in
0100	G_GVID	Get Global VE ID	1-to-1	FC-GS-8
0101	G_FVID	Get Fabric VE ID	1-to-1	FC-GS-8
0102	G_VEMID	Get VEM IDs	1-to-All	FC-GS-8
0103	G_VEM	Get VE Mappings	1-to-1	FC-GS-8
0300	G_VEM_D	Get VE Mappings - Domain_ID	1-to-All	FC-SW-7
0301	U_VE_M	Update VE Mappings	1-to-All	FC-SW-7

### 9.6.3 FC-CT Requests for dVEIS

#### 9.6.3.1 Get VE Mappings - Domain\_ID (G\_VEM\_D)

The dVEIS shall, when it receives a G\_VEM\_D request, return the VE Mappings associated with the specified Domain\_ID. The format of the G\_VEM\_D CT Request is specified in table 246.

**Table 246 – G\_VEM\_D Request payload**

Item	Size (Bytes)
FC-CT Header	16
Reserved	3
Domain_ID	1

**Domain\_ID:** indicates the Domain\_ID whose VE mappings are requested.

The format of the CT Accept to a G\_VEM\_D request is specified in table 247.

**Table 247 – G\_VEM\_D Accept payload**

Item	Size (Bytes)
FC-CT Header	16
Number of VEM Mappings (n)	4
VEM Mapping entry #1	see table 248
VEM Mapping entry #2	see table 248
...	
VEM Mapping entry #n	see table 248

The format of the VEM Mapping entry is specified in table 248.

**Table 248 – VEM Mapping entry format**

Item	Size (Bytes)
VEM ID	16
Number of VE Mapping records (m)	4
VE Mapping record #1	see table 249
VE Mapping record #2	see table 249
...	
VE Mapping record #m	see table 249

**VEM ID:** contains the ID of the described VEM.

The format of the VE Mapping record is specified in table 249.

**Table 249 – VE Mapping record format**

Item	Size (Bytes)
Global VE ID	16
Number of Fabric VE IDs (t)	4
Fabric VE ID #1	see table 250
Fabric VE ID #2	see table 250
...	
Fabric VE ID #t	see table 250

**Global VE ID:** contains the global VE ID of the described VE Mapping.

**Fabric VE ID:** contains the Fabric VE ID(s) of the described VE Mapping. The Fabric VE ID format is specified in table 250..

**Table 250 – Fabric VE ID format**

Item	Size (Bytes)
N_Port_ID	3
Local VE ID	1

**N\_Port\_ID:** contains the N\_Port\_ID of the described VE.

**Local VE ID:** contains the 8-bit local VE ID of the described VE.

### 9.6.3.2 Update VE Mappings (U\_VE\_M)

If the VE mappings maintained by the VE Identification Server are updated using the UVEM ELS, the dVEIS shall communicate these updates to the other Switches in the Fabric supporting the VE Identification Server through the U\_VE\_M CT Request. The format of the U\_VE\_M CT Request is shown in table 251.

**Table 251 – U\_VE\_M Request Payload**

Item	Size (Bytes)
FC-CT Header	16
VEM ID	16
Number of instantiated VE Mapping entries (s)	4
VE Mapping entry #1	see table 252
VE Mapping entry #2	see table 252
VE Mapping entry #s	see table 252
Number of deinstantiated VE Mapping entries (t)	4
VE Mapping entry #1	see table 252
VE Mapping entry #2	see table 252
VE Mapping entry #t	see table 252

**VEM ID:** contains the ID of the VEM updating the VE mappings.



The format of the VE Mapping entry is shown in table 252.

**Table 252 – VE Mapping entry format**

Item	Size (Bytes)
Global VE ID	16
N_Port_ID	3
Local VE ID	1

**Global VE ID:** contains the global VE ID of the described VE.

**N\_Port\_ID:** contains the N\_Port\_ID of the described VE.

**Local VE ID:** contains the 8-bit local VE ID of the described VE.

Deinstantiating a VE mapping not registered in the VE Identification Server shall not be treated as an error.

A deinstantiated VE Mapping entry with a NULL Global VE ID and a NULL Local VE ID indicates the deinstantiation of all VE mappings associated with the N\_Port\_ID specified in the entry and the VEM ID specified in the U\_VE\_M request payload. An U\_VM\_E request payload containing a deinstantiated VE Mapping entry with a NULL Global VE ID and a NULL Local VE ID shall contain only that deinstantiated VE Mapping entry and no instantiated VE Mapping entries.

The format of the CT Accept to a U\_VE\_M request is shown in table 253.

**Table 253 – U\_VE\_M Accept Payload**

Item	Size (Bytes)
FC-CT Header	16

## 10 Switch Zone exchange & merge

### 10.1 Overview

This clause describes a mechanism for Switches to exchange Zoning data. FC-GS-8 contains a description of the Fabric Zoning Service architecture and management requests for administering Zoning.

When link parameters have been established for a link and the Switches have a Domain\_ID, the two Switches joined by this link exchange Zoning Configuration information to make the information consistent across the Fabric. The Fabric Management Inter-Switch messages that are addressed to a Fabric Controller are the Merge request and Merge response. These messages are used to resolve the Zoning Configuration in a Fabric when two Switches are joined. Each Switch determines if the Zoning Configuration from the Adjacent Switch may be merged with its local Zoning Configuration. The rules for merging Zoning Configurations are described in 10.5.2.

NOTE 20 – This protocol is designed to work when a single Inter-Switch Link is established. Establishing more than one Inter-Switch Link at the same time may lead to unpredictable effects over the Fabric.

### 10.2 Joining Switches

Merge request and Merge response messages are used to merge and propagate Zoning Configurations when an Inter-Switch Link becomes available. A merge operation is performed with the Adjacent Switch when an Inter-Switch Link becomes available, and with all Adjacent Switches when changes are made to the local Zoning Configuration as a result of merging the local Zoning Configuration with an Adjacent Zoning Configuration.

A Merge request message contains the local Zoning Configuration of the Switch that is generating the Merge request, together with a Protocol Version field that defines the format of the Zoning Protocol used by the Switch (e.g., Enhanced or Basic).

When a Merge request is received from an Adjacent Switch, the receiving Switch determines if the request may be accepted and executed. If the Switch is not busy, but there is a Protocol Version mismatch, or, when the Protocol version matches, the merge is unable to be executed according to the rules described in 10.5.2, a Merge response is returned indicating that the Zone Configurations are unable to be merged, and the Inter-Switch Link is isolated. If the Switch is not busy and the merge may be executed, a Merge response is returned indicating that the Zone Configurations were successfully merged. After the successful merge, the Zone Set Name is changed to "Successful Zone Set Merge: Active Zone Set Name has changed".

This information exchange may start by sending Merge Request Resource Allocation messages to allocate the resources needed to process the Merge Request Sequences. An example of the Merge data flow is provided in 10.5.1.

### 10.3 Enhanced Zoning support determination

When a Switch supporting Enhanced Zoning joins a Fabric, it shall use the ESS SW\_ILS (see 6.2.22.4.4) to determine the Enhanced Zoning capabilities of the other Switches. By doing so, the Switch also announces its Enhanced Zoning capabilities to the other Switches of the Fabric.

Each Switch supporting Enhanced Zoning shall maintain the information about the Enhanced Zoning support by all Switches in the Fabric. This information is updated whenever a Switch joins or leaves the Fabric, and it is used to reply to the GFEZ request.

If all the Switches in the Fabric support Enhanced Zoning, then the Enhanced Zoning supported bit (i.e., bit 0) of the Fabric Enhanced Zoning support flags of the GFEZ accept shall be set to one, otherwise it shall be set to 0.

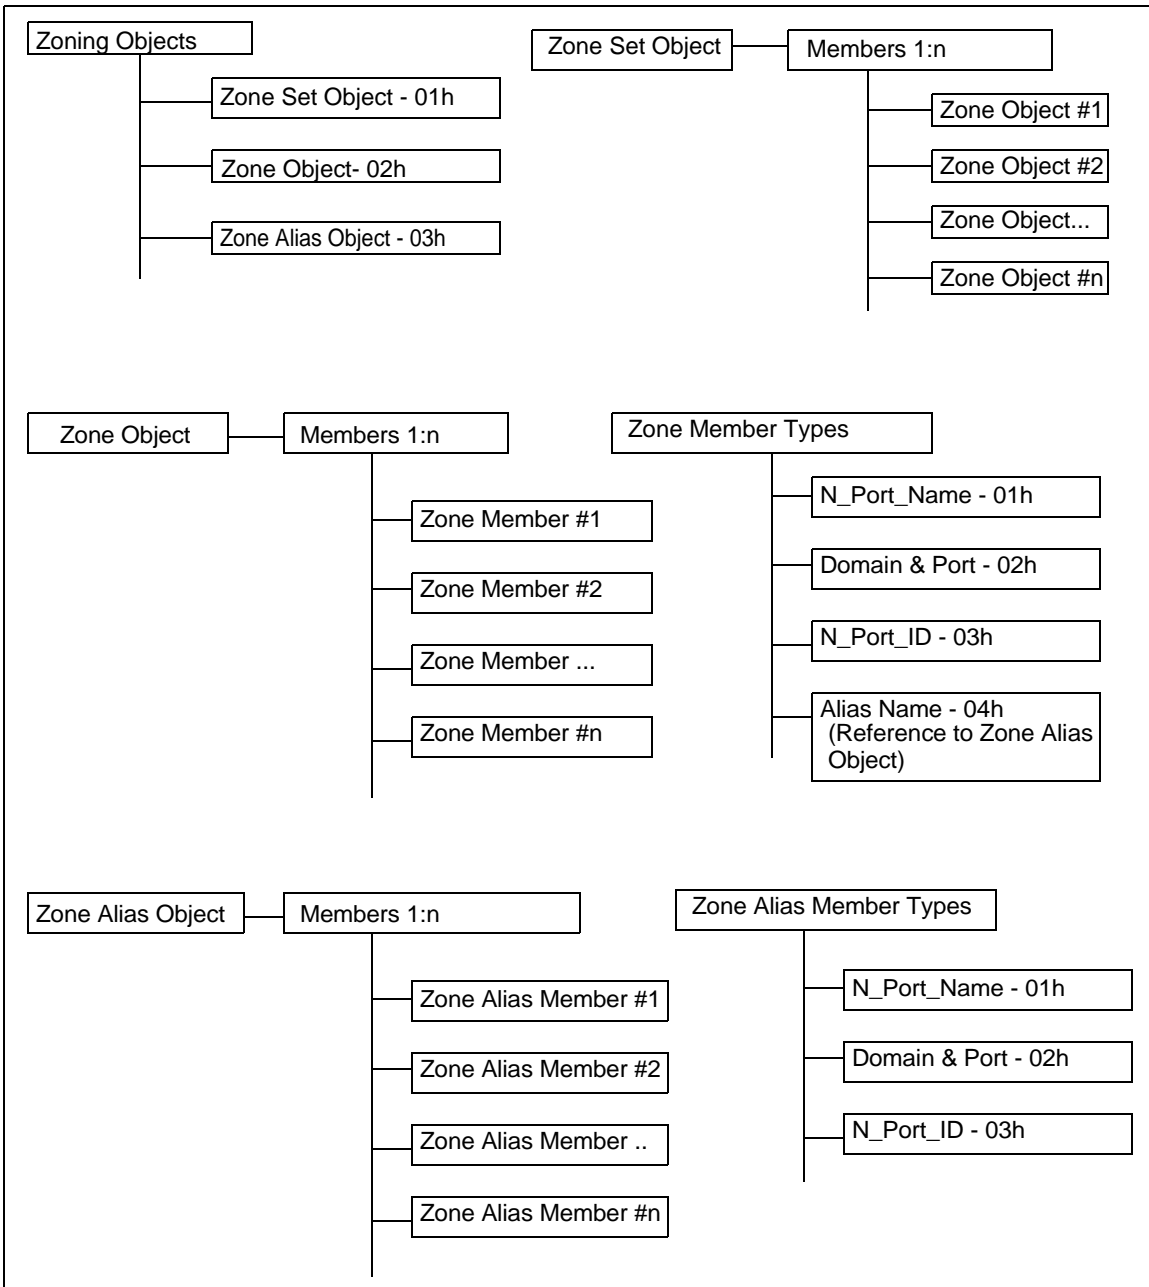
The value of the Enhanced Zoning enabled bit (i.e., bit 1) of the Fabric Enhanced Zoning support flags of the GFEZ accept is instead determined with the Zone Merge protocol because a Switch is able to join a Fabric only if it is working in the same mode of the Fabric (see 6.2.15).

If all the Switches in the Fabric support the Zone Set Database, then the Zone Set Database supported bit (i.e., bit 4) of the Fabric Enhanced Zoning support flags of the GFEZ accept shall be set to one, otherwise it shall be set to 0.

## **10.4 Zoning framework and data structures**

### **10.4.1 Basic Zoning framework**

This clause provides an overview of the Basic Zoning framework associated with the Switch Fabric. The Basic Zoning framework describes Zoning entities such as Zoning Objects, Object Members, Member types, and their relationships. The Basic Zoning framework is shown in figure 27.



**Figure 27 – Basic Zoning framework**

**Zoning Objects:** Three types of Zoning Objects are defined. They are:

- a) Zone Set;
- b) Zone; and
- c) Zone Alias.

**Zone Set Object:** The Zone Set Object defines a group of Zones. A Zone Set Object contains one or more members that are Zone Objects. In addition, the Zone Set Object has two attributes:

- a) Name; and
- b) Number of Members.

**Zone Object:** The Zone Object defines a Zone and its members. A Zone Object contains one or more Zone Members. Currently defined Zone Member Types are:

- a) N\_Port\_Name;
- b) Domain\_ID and physical port;
- c) N\_Port\_ID; and
- d) Zone Alias Name;

The Zone Object has three attributes:

- a) Name;
- b) Protocol Type; and
- c) Number of Members.

**Zone Alias Object:** A Zone Alias Object defines a Zone Alias and its members. A Zone Alias Object contains one or more Zone Alias Members. Currently defined Zone Alias Member Types are:

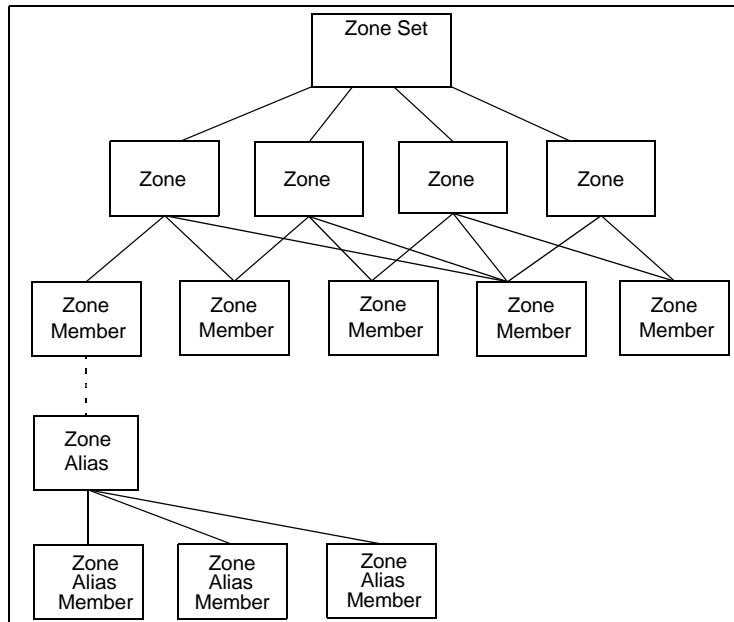
- a) N\_Port\_Name;
- b) Domain\_ID and physical port; and
- c) N\_Port\_ID;

The Zone Alias Name specified as a Zone Member serves as a reference to a Zone Alias Object.

The Zone Alias Object has two attributes:

- a) Name; and
- b) Number of Members.

The Basic Zoning hierarchy is shown in figure 28.



**Figure 28 – Basic Zoning hierarchy**

The Basic Zoning Object structure is shown in figure 29.

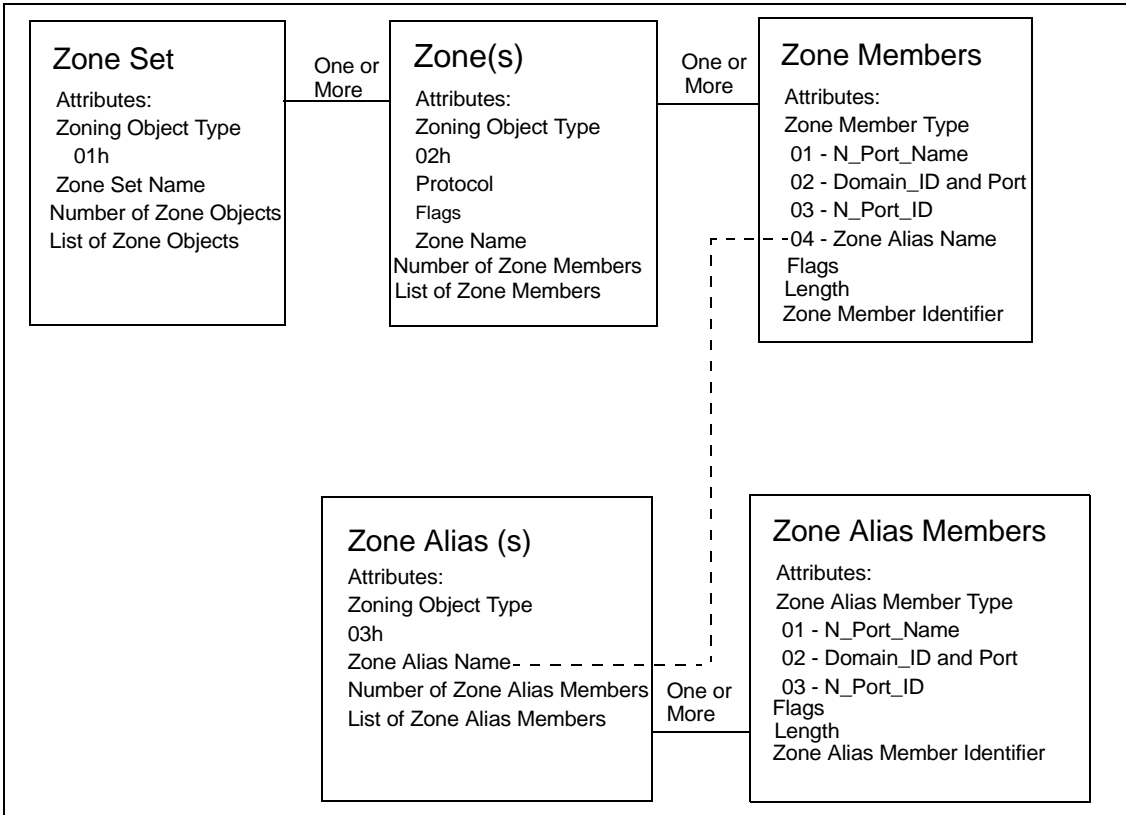


Figure 29 – Basic Zoning Object structure

## 10.4.2 Basic Zoning data structures

### 10.4.2.1 Zoning Object List

The format of the Zoning Object List is specified in table 254.

Table 254 – Zoning Object List format

Item	Size
Number of Zoning Objects	4
Zoning Object 1	x
Zoning Object 2	y
...	...
Zoning Object n	z

**Number of Zoning Objects:** This field contains a value that specifies the number of Zoning Objects contained in the Zoning Object list.

**Zoning Object:** This entry contains a Zoning Object as described in 10.4.2.2.

### 10.4.2.2 Zoning Object format

The general format of a Zoning Object is specified in table 255.

**Table 255 – Zoning Object**

Item	Size
Object Type	1
Protocol	1
Reserved	2
Object Name	a
Number of Object Members	4
Object Member 1	x
Object Member 2	y
...	...
Object Member n	z

**Object Type:** Valid values for defined Zoning Object Types are specified in table 256.

**Table 256 – Zoning Object Type values**

Description	Value (hex)
Reserved	00
Zone Set	01
Zone	02
Zone Alias	03
Reserved	04-DF
Vendor Specified	E0-FF

**Protocol:** The Protocol attribute shall be reserved for Zone Set and Zone Alias Objects. For Zone Objects, if the Protocol field is non-zero, Device\_Data and FC-4 Link\_Data frames not having the specified protocol value shall not be transmitted between members of the Zone. All other frames shall

be transmitted between members of the Zone. The values of the Protocol Format are specified in table 257.

**Table 257 – Protocol Format values**

Encoded Value (hex)	Description
00	No Protocol Zoning
01-FF	Non-zero values are taken from FC-FS-5.

**Object Name:** This attribute specifies the name of the object. The format of this attribute is described in 10.4.2.3.

**Number of Object Members:** This field indicates the number of Object Members.

**Object Members:** One or more Object Members are contained in the Zoning Object. Object Members may be other Zoning Objects or Zone Members.

**10.4.2.3 General name format**

All name attributes pertaining to Zoning shall use the general name format as specified in FC-GS-8. Examples include Zone Set Names, Zone Names, and Zone Alias Names.

**10.4.2.4 Zone Member format**

Zone Objects shall have Zone Members. The format of a Zone Member is specified in table 258.

**Table 258 – Zone Member format**

Item	Size
Zone Member Type	1
Reserved	1
Flags	1
Identifier Length	1
Identifier	x

**Zone Member Type:** Valid Zone Member Types are specified in table 259.

**Flags:** Implementation is vendor specific.

**Identifier Length:** The Identifier Length field value is determined by the Zone Member Type as specified in table 259.



**Identifier:** The description of the Identifier fields for valid Zone Member Types are specified in table 259.

**Table 259 – Zone Member Type and Identifier formats**

Type (hex)	Identifier	Size (bytes)
00	Reserved	
01	N_Port_Name: The format of the Zone Member information is a N_Port_Name.	8
02	Domain_ID & Physical Port Number: The format of the Zone Member Information is a combination of a Domain_ID + Physical Port Number.  (i.e., 00DDPPPPh; where DD is the Domain_ID and PPPP is the Physical Port Number).	4
03	N_Port_ID: Address identifier format (00ddaapp). Valid address identifiers are those assignable to F and FL port attached devices.	4
04	Alias Name	Variable
05-DF	Reserved	
E0-FF	Vendor specific	

Name field describes either an Alias or Zone Name in the format described in FC-GS-8.

### 10.4.3 Enhanced Zoning framework

#### 10.4.3.1 Introduction

In the Enhanced Zoning framework more Zoning Objects are defined to those defined in Basic Zoning. This subclause delineates the structures of the Enhanced Active Zone Set and of the Enhanced Zone Set Database.

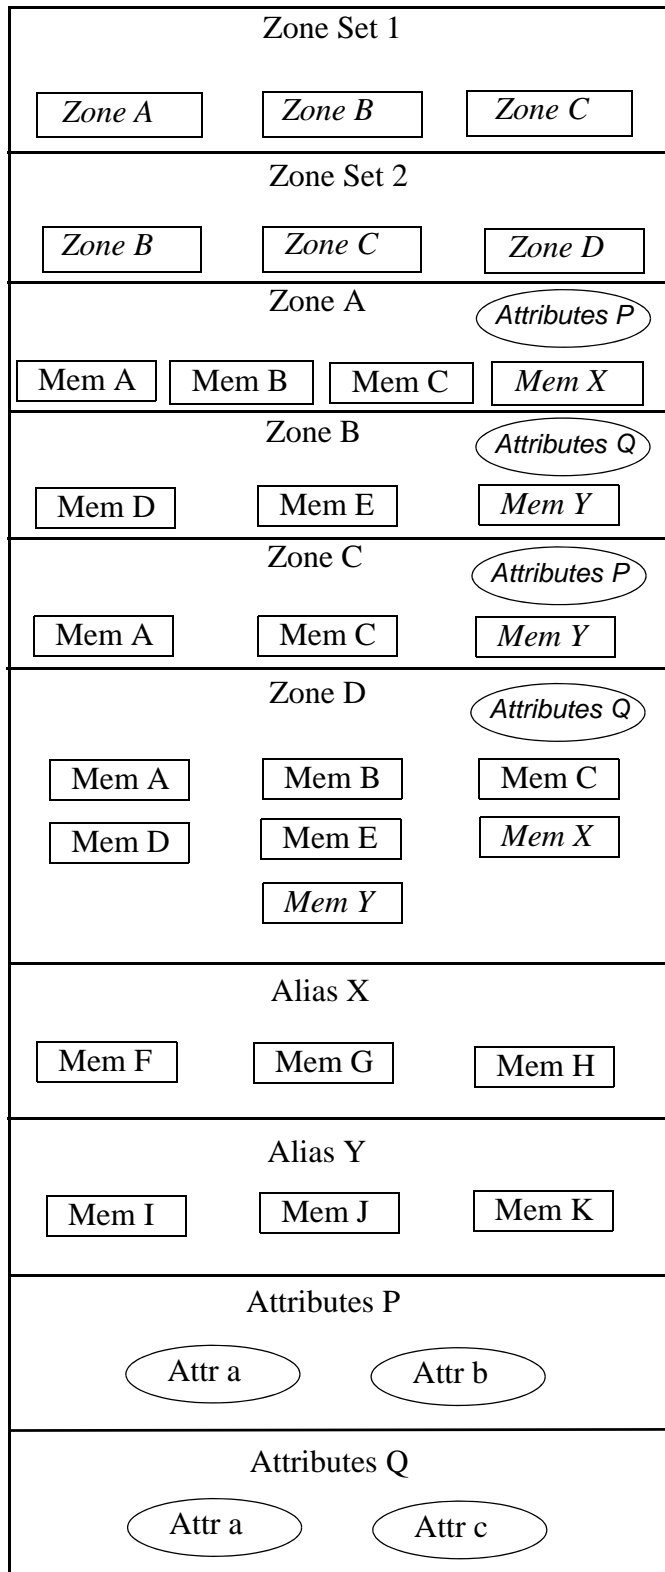
#### 10.4.3.2 Zone Set Database

The Zone Set Database may contain Zoning Object types as defined in the following subclauses. The currently defined Zoning Object types are:

- a) Zone Set;
- b) Zone;
- c) Zone Alias; and
- d) Zone Attribute.

The Zone Set Database shall not contain the Active Zone Set. The Zone Set Database may contain multiple Zoning Objects. Objects defined in the Zone Set Database need not be referenced. In the Zone Set Database the Zoning Objects may reference each other using names formatted as specified in FC-GS-8.

The logical structure of the Zone Set Database is shown in figure 30.



*Italics = Reference*

**Figure 30 – Logical structure of the Zone Set Database**

Each Zone Set definition references its Zone Objects. Each Zone may reference a Zone Attribute Object or, in the member definitions, one or more Zone Alias Objects.

### 10.4.3.3 Active Zone Set

References are not allowed in the Active Zone Set. At activation time any reference present in a Zone Set or Zone definition in the Zone Set Database shall be resolved. The resulting logical structure of the Active Zone Set is shown in figure 31.

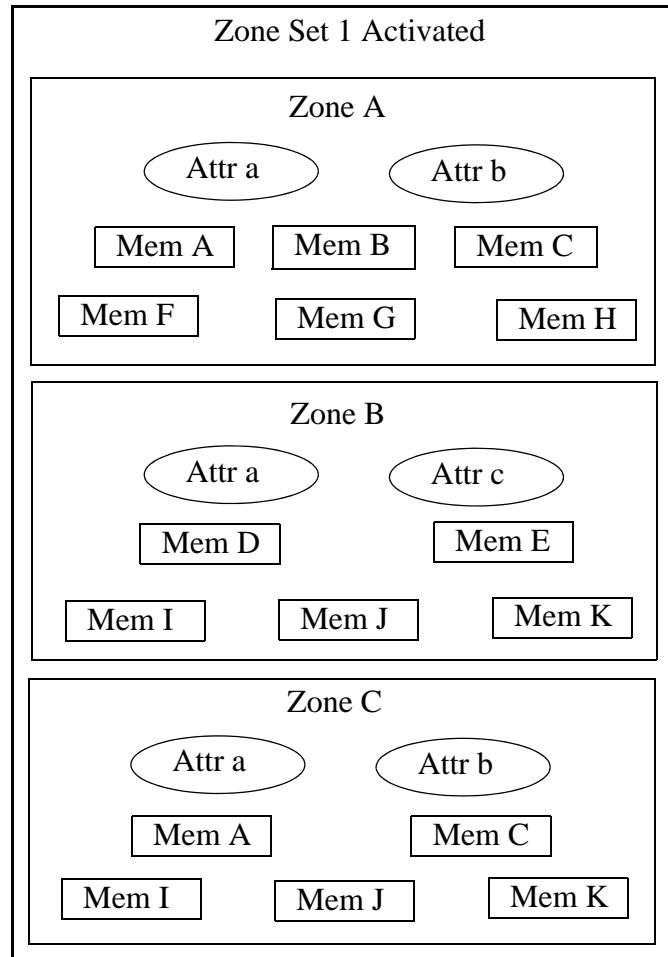


Figure 31 – Logical structure of the Active Zone Set

## 10.4.4 Enhanced Zoning data structures

### 10.4.4.1 Zoning Object List

The format of the Zoning Object List is specified in table 260.

**Table 260 – Zoning Object List format**

Item	Size
Number of Zoning Objects	4
Zoning Object 1	x
Zoning Object 2	y
...	...
Zoning Object n	z

**Number of Zoning Objects:** This field contains a value that specifies the number of Zoning Objects contained in the Zoning Object List.

**Zoning Object:** This entry contains a Zoning Object. Valid Zoning Objects are described in the following clauses.

### 10.4.4.2 Zoning Object Types

Zoning Object Type values are specified in table 261.

**Table 261 – Zoning Object Type values**

Description	Value (hex)
Reserved	00
Zone Set	01
Zone	02
Zone Alias	03
Zone Reference	04
Zone Attribute Object	05
Reserved	06-DF
Vendor Specified	E0-FF

### 10.4.4.3 Zone Set Object

#### 10.4.4.3.1 Zone Set Object in the Zone Set Database

The Zone Set Object used in the Zone Set Database has the attributes described below and the format specified in table 262.

**Table 262 – Zone Set Object format in the Zone Set Database**

Item	Size (bytes)
Zone Set Object Identifier	1
Reserved	3
Zone Set Name	a
Number of Zone References	4
Zone Reference #1	x
Zone Reference #2	y
...	...
Zone Reference #n	z

#### Zone Set Name

The Zone Set Name attribute shall follow the general name format described in FC-GS-8.

#### Number of Zone References

This attribute shall contain the integer number of Zone References in the Zone Set.

#### Zone Reference

Only Zone Reference Objects (i.e., Zone Set Member Type 04h) are allowed in the Zone Set Object when used in the Zone Set Database.

**10.4.4.3.2 Zone Set Object in the Active Zone Set**

The Zone Set Object used in the Active Zone Set has the attributes described below and the format specified in table 263.

**Table 263 – Zone Set Object format in the Active Zone Set**

Item	Size (bytes)
Zone Set Object Identifier	1
Reserved	3
Number of Zones	4
Zone #1	x
Zone #2	y
...	...
Zone #n	z

**Number of Zones**

This attribute shall contain the integer number of Zones in the Zone Set.

**Zone**

Only Zone Objects (i.e., Zone Set Member Type 02h) are allowed in the Zone Set Object when used in the Active Zone Set.

NOTE 21 – The Active Zone Set does not have a name because such a name is not significant when a Zone Merge occurs.

**10.4.4.4 Zone Reference Object**

The Zone Reference Object format is specified in table 264.

**Table 264 – Zone Reference Object format**

Item	Size (bytes)
Zone Reference Object Identifier	1
Reserved	3
Zone Name	x

**Zone Name**

The Zone Name attribute shall follow the general name format described in FC-GS-8.

#### 10.4.4.5 Zone Object in the Zone Set Database

The Zone Object used in the Zone Set Database allows references to other objects. It has the attributes described below and the format specified in table 265.

**Table 265 – Zone Object format in the Zone Set Database**

Item	Size (bytes)
Zone Object Identifier	1
Reserved	3
Zone Name	a
Zone Attribute Object Name	b
Number of Zone Members	4
Zone Member #1	x
Zone Member #2	y
...	...
Zone Member #n	z

#### **Zone Name**

The format of the Zone Name attribute shall follow the general name format described in FC-GS-8.

#### **Zone Attribute Object Name**

The format of the Zone Attribute Object Name attribute shall follow the general name format described in FC-GS-8. Its value references the Zone Attribute Object whose attributes apply for the Zone. A null value ('00 00 00 00 00 00 00 00') indicates that no attributes are associated with that Zone definition.

#### **Number of Zone Members**

This attribute shall contain the integer number of Zone Members in the Zone.

#### **Zone Member**

The Zone Member format is specified in 10.4.4.6.1.



#### 10.4.4.6 Zone Object in the Active Zone Set

The Zone Object used in the Active Zone Set does not allow references to other objects. It has the attributes described below and the format specified in table 266.

**Table 266 – Zone Object format in the Active Zone Set**

Item	Size (bytes)
Zone Object Identifier	1
Reserved	3
Zone Name	a
Number of Zone Attribute Entries	4
Zone Attribute Entry #1	b
Zone Attribute Entry #2	c
...	...
Zone Attribute Entry #m	d
Number of Zone Members	4
Zone Member #1	x
Zone Member #2	y
...	...
Zone Member #n	z

#### **Zone Name**

The format of the Zone Name attribute shall follow the general name format described in FC-GS-8.

#### **Number of Zone Attribute Entries**

This attribute shall contain the integer number of Zone Attribute Entries in the Zone.

#### **Zone Attribute Entry**

The format of the Zone Attribute Entry is specified in 10.4.4.8.1.

#### **Number of Zone Members**

This attribute shall contain the integer number of Zone Members in the Zone.

#### **Zone Member**

The Zone Member format is specified in 10.4.4.6.1.

**10.4.4.6.1 Zone Member format**

Zone Objects shall have Zone Members. The format of a Zone Member is specified in table 267.

**Table 267 – Zone Member format**

Item	Size
Zone Member Type	1
Reserved	2
Identifier Length	1
Identifier	x

**Zone Member Type**

Valid Zone Member Types are specified in table 268.

**Identifier Length**

The Identifier Length field value is determined by the Zone Member Type as specified in table 268.

**Identifier**

The description of the Identifier fields for valid Zone Member Types are specified in table 268.

**Table 268 – Zone Member Type and Identifier formats (Part 1 of 2)**

Type (hex)	Identifier	Size (bytes)
00	Reserved	
01	N_Port_Name: The format of the Zone Member information is a N_Port_Name.	8
02	Domain_ID & Physical Port Number: The format of the Zone Member Information is a combination of a Domain_ID + Physical Port Number.  (i.e., 00DDPPPPh; where DD is the Domain_ID and PPPP is the Physical Port Number).	4
03	N_Port_ID: Address identifier format (00ddaapp). Valid address identifiers are those assignable to F and FL port attached devices.	4

**Table 268 – Zone Member Type and Identifier formats (Part 2 of 2)**

Type (hex)	Identifier	Size (bytes)
04	Alias Name: The format of the Zone Member information is a general name.	a
05	Node_Name: The format of the Zone Member information is a Node_Name.	8
06	F_Port_Name: The format of the Zone Member information is a F_Port_Name.	8
07	Wildcard	8
08-DF	Reserved	
E0-FF	Vendor Specific	

**Wildcard Zone Member format**

The Wildcard Zone Member Identifier format is specified in table 269. See FC-GS-8 for how to use the Wildcard Zone Member.

**Table 269 – Wildcard Zone Member Identifier format**

Item	Size (bytes)
Subtype	2
Flags	2
Parameter	4

**Subtype:** see FC-GS-8.

**Flags:** see FC-GS-8.

**Parameter:** see FC-GS-8.

**Vendor Specified Zone Member Identifier format**

The Vendor Specified Zone Member Identifier format is specified in table 270.

**Table 270 – Vendor Specified Zone Member Identifier format**

Item	Size (bytes)
Vendor ID	8
Vendor Specified Value	n
Pad	m

**Vendor ID:** Contains the T10 Vendor ID of the vendor that defines the content of the Vendor Specified Value field.

**Vendor Specified Value:** This field contains the Vendor Specified Value.

**Pad:** Fill bytes are added as necessary to the end of the Vendor Specified Value in order to ensure that the total length of the Vendor Specified Zone Member is a multiple of four. Fill bytes shall be nulls (i.e., 00h). The number of fill bytes (m) is zero, one, two, or three depending on the length of the actual value (n).

#### 10.4.4.7 Zone Alias Object

The Zone Alias Object has the attributes described below and the format specified in table 271.

**Table 271 – Zone Alias Object format**

Item	Size (bytes)
Zone Alias Object Identifier	1
Reserved	3
Zone Alias Name	a
Number of Zone Alias Members	4
Zone Alias Member #1	x
Zone Alias Member #2	y
...	...
Zone Alias Member #n	z

#### Zone Alias Name

The format of the Zone Alias Name attribute shall follow the general name format described in FC-GS-8.

#### Number of Zone Alias Members

This attribute shall contain the integer number of Zone Alias Members in the Zone Alias.

## Zone Alias Member

The Zone Alias Member has the format specified in 10.4.4.6.1. All Zone Member Identifier Types may be used, with the exception of the Alias Name member (i.e., type '04').

### 10.4.4.8 Zone Attribute Object

The Zone Attribute Object is a variable length structure that contains extensible attributes that may be associated with a Zone. The format of the Zone Attribute Object is specified in table 272.

**Table 272 – Zone Attribute Object format**

Item	Size (bytes)
Zone Attribute Object Identifier	1
Reserved	3
Zone Attribute Object Name	x
Zone Attribute Block	w

### Zone Attribute Object Name

The format of the Zone Attribute Object Name attribute shall follow the general name format described in FC-GS-8.

### Zone Attribute Block

The format of the Zone Attribute Block is specified in table 273.

**Table 273 – Zone Attribute Block format**

Item	Size (bytes)
Number of Zone Attribute Entries	4
Zone Attribute Entry #1	x
Zone Attribute Entry #2	y
...	...
Zone Attribute Entry #n	z

### Number of Zone Attribute Entries

This field specifies the number of Zone Attribute Entries contained in the Zone Attribute Block. A value of zero in this field shall indicate that no attributes are registered.

### 10.4.4.8.1 Zone Attribute Entry format

The format of the Zone Attribute Entry is specified in table 274.

**Table 274 – Zone Attribute Entry format**

Item	Size (bytes)
Zone Attribute Type	2
Zone Attribute Length	2
Zone Attribute Value	x

**Zone Attribute Type:** This field indicates the attribute entry type. Valid Zone Attribute Types are specified in table 275 and shall be restricted to a value between 0000h and 00FFh

**Table 275 – Zone Attribute Type values**

Zone Attribute Type (hex)	Description
0001	Protocol
0002	Broadcast Zone
0003	Hard Zone
0004	IFR Zone <sup>a</sup>
0005	Peer Zone
00E0	Vendor Specific
other values	Reserved
<sup>a</sup> See FC-IFR for a definition of the IFR Zone attribute type.	

**Zone Attribute Length:** This field indicates the total length in bytes of the Zone Attribute Entry. This length shall be a multiple of four and includes the following fields:

- a) Zone Attribute Type;
- b) Zone Attribute Length; and
- c) Zone Attribute Value.

**Zone Attribute Value:** This field specifies the actual attribute value. If present, Zone Attribute Values shall be at least four bytes in length and the length shall be a multiple of four. For Zone Attribute Value fields, fill bytes are added as necessary to the end of the actual value in order to ensure that the length of the Zone Attribute Value field is a multiple of four. Fill bytes shall be nulls (i.e., 00h). The number of fill bytes (m) is zero, one, two, or three depending on the length of the actual value (n). Therefore the total length of the value field is (n+m).

#### 10.4.4.8.1.1 Protocol Attribute

When a Protocol Attribute Type is specified, the Protocol Attribute Value specifies the FC-4 type for which protocol Zoning is enforced. Valid values are 01h-FFh. If the FC-4 Type is non-zero, Device\_Data and FC-4 Link\_Data frames not having the specified FC-4 Type value shall not be transmitted between members of the Zone. All other frames shall be transmitted between members of the Zone. The format of the Protocol Attribute Value is specified in table 276.

**Table 276 – Protocol Attribute Value format**

Item	Size (bytes)
FC-4 Type	1
Reserved	3

#### 10.4.4.8.1.2 Broadcast Zone Attribute

Broadcast Zoning is enabled by setting the Broadcast Zone Attribute on Zones that contain ports that send or receive broadcast frames. Use of the Broadcast Zone Attribute allows multi-protocol devices to send broadcast frames to some devices but not to others. The Broadcast Zone Attribute only affects the processing of broadcast frames. The Broadcast Zone Attribute has no effect on Zoning enforcement for Name Server, RSCN, or Hard Zoning.

When Zoning is active, broadcast frames are delivered to all logged in Nx\_Ports that share a Broadcast Zone with the source of the frame, as indicated by the S\_ID of the frame. This implies that if Zoning is active and no Zones have the Broadcast Zone Attribute set, then no broadcast frames are delivered. Zoning is enforced at the destination Switch. Broadcast Zoning shall not have any affect on the Switch to Switch routing of frames.

If any NL\_Port attached to an FL\_Port shares a Broadcast Zone with the source of the broadcast frame, or Zoning is not active, the frame shall be sent to the all the devices on the loop using the OPNfr open replicate primitive. Use of the OPNyr selective replicate is prohibited (see FC-DA-2).

Broadcast Zoning shall be enforced on the destination Switches, using the S\_ID of the broadcast frame to determine the Zones. Some Zone Member types are unable to be directly mapped to an address identifier. For these types a Switch shall use the Name Server to get the address identifier associated with the Zone Member data. Similar discovery may be needed to implement Hard Zones. The address identifier(s) for specific Zone Member Identifier Types may be obtained as follows:

- a) N\_Port\_Name or Node\_Name Zone Member: use the GID\_PN or GID\_NN Name Server commands;
- b) Domain + Physical Port Zone Member: use the GID\_DP Name Server command; or
- c) Fabric Port\_Name Zone Member: use the GID\_FPN Name Server command.

NOTE 22 – GID\_NN and GID\_DP may return multiple address identifiers.

If address discovery finds address identifiers for non address identifier Zone Members, ports in those Broadcast Zones become accessible to the ports whose address identifiers have been discovered. If ports become accessible to new broadcast sources, an RSCN shall be issued for those ports. If a Switch is issuing an RSCN for some other reason after address discovery, and the RSCN includes newly accessible ports, a separate RSCN for address discovery for those ports is not necessary.

If Zoning is inactive in the Fabric, then broadcast frames are delivered to all logged in Nx\_Ports.

There is no value associated with the Broadcast Zone Attribute. Therefore the Zone Attribute Length shall be set to four.

#### 10.4.4.8.1.3 Hard Zone Attribute

Hard Zoning is specified by setting the Hard Zone Attribute.

When the Hard Zone Attribute is specified, the Zone configuration shall use Hard Zoning enforcement (see 3.1.53). If an implementation is unable to enforce the Zone configuration using Hard Zoning enforcement (e.g., enforcing the Zone configuration requires hardware resources that are not available), the activation of the Zone configuration shall fail.

When the Hard Zone Attribute is not specified, the Zoning configuration shall be enforced in one of the following ways:

- a) Hard Zoning enforcement; or
- b) Soft Zoning enforcement (see 3.1.89).

The Zone configuration should be enforced using Hard Zoning enforcement whenever possible.

If Zoning for an Nx\_Port is enforced using Hard Zoning enforcement for any Zone, Zoning for that Nx\_Port shall be enforced using Hard Zoning enforcement for all Zones in which it is a member.

The activation of a Hard Zone may succeed because some of the Nx\_Ports specified in the Zoning configuration are not connected to the Fabric at activation time. If a FLOGI request received from such an Nx\_Port makes it impossible to enforce the Zoning configuration using Hard Zoning enforcement, then the Fabric shall reject or discard, as appropriate for the class of service, any frames not addressed to Fabric Services to or from that Nx\_Port.

Fabrics may be composed of Switches that support Hard Zoning and Switches that support only Soft Zoning. The Fabric administrator decides which Nx\_Ports are to be managed with Hard Zoning, and which ones with Soft Zoning.

The Fabric administrator may manage the situation where some Nx\_Ports need to be restricted with Hard Zoning and others with Soft Zoning by defining two overlapping Zones, one with the Hard Zone Attribute and the other without the Hard Zone Attribute. The Zone without the Hard Zone Attribute defines which Nx\_Ports are allowed to communicate. The Zone with the Hard Zone Attribute defines the subset of these Nx\_Ports for which Zoning shall be enforced on a frame-by-frame basis.

There is no value associated with the Hard Zone Attribute. Therefore the Zone Attribute Length shall be set to four.

Hard Zoning enforcement requires a mapping of Zone Members to N\_Port\_IDs. A Switch may use the following Name Server requests to perform this mapping:

- a) N\_Port\_Name member type: use the GID\_PN Name Server command;
- b) Node\_Name member type: use the GID\_NN Name Server command, which may return multiple N\_Port\_IDs; or
- c) F\_Port\_Name member type: use the GID\_FPN Name Server command.



#### 10.4.4.8.1.4 Peer Zone Attribute

The Peer Zone Attribute is used to define a Peer Zone (see FC-GS-7). A Peer Zone is a Zone with the Peer Zone Attribute. A Peer Zone identifies a Principal member through the Peer Zone Attribute and a list of Peer members as Zone Members. The semantic of a Peer Zone is that:

- a) Peer members are allowed to communicate with the Principal member; and
- b) Peer members are not allowed to communicate among themselves, unless allowed by other Zones in the Zone Set.

The format of the Peer Zone Attribute Value is specified in table 277.

**Table 277 – Peer Zone Attribute Value format**

Item	Size (bytes)
Principal N_Port_Name	8

**Principal N\_Port\_Name:** The N\_Port\_Name of the Principal member of a Peer Zone.

#### 10.4.4.8.1.5 Vendor Specific Zone Attribute

The format of the Vendor Specific Attribute Value is specified in table 278.

**Table 278 – Vendor Specific Attribute Value format**

Item	Size (bytes)
Vendor ID	8
Vendor Specific Value	n
Pad	m

**Vendor ID:** Contains the T10 Vendor ID of the vendor that defines the content of the Vendor Specific Value field.

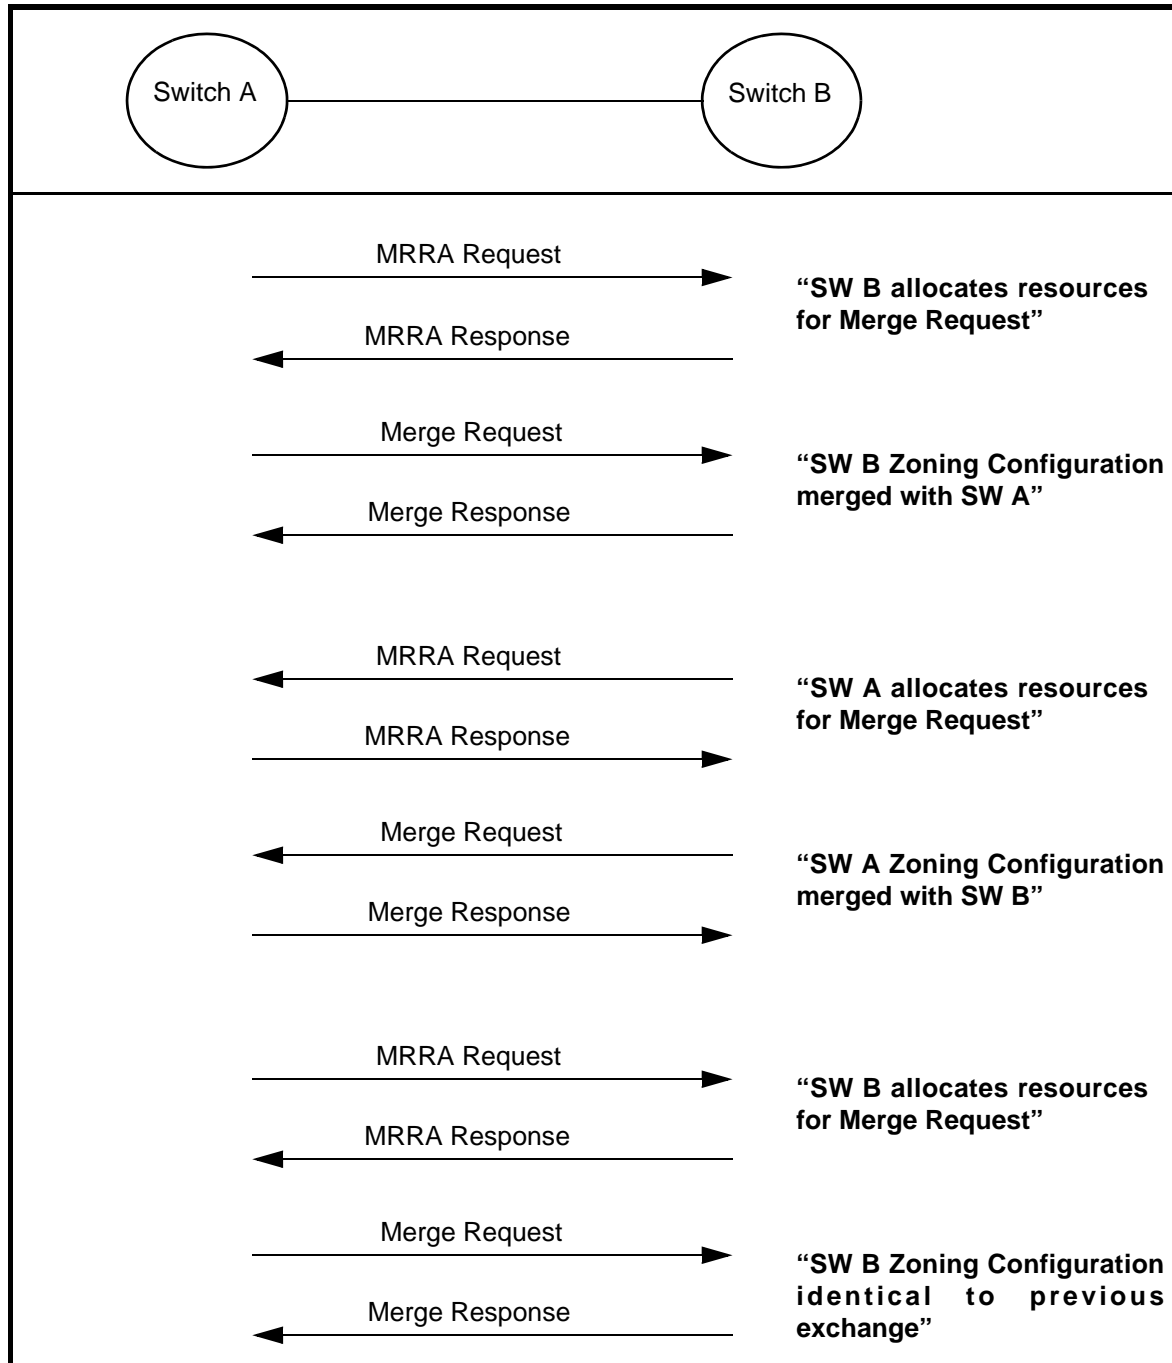
**Vendor Specific Value:** This field contains the Vendor Specific Value.

**Pad:** Fill bytes are added as necessary to the end of the Vendor Specific Value in order to ensure that the total length of the Vendor Specific Zone Member is a multiple of four. Fill bytes shall be nulls (i.e., 00h). The number of fill bytes (m) is zero, one, two, or three depending on the length of the actual value (n).

## 10.5 Merge Zone

### 10.5.1 Example merge operation

Figure 32 shows how the Zones are merged when two Switches are joined.



**Figure 32 – Merge operation between two Switches**

The use of the MRRA SW\_ILS is optional. Switch A may send a Merge Request Resource Allocation (MRRA) over the new link, to request Switch B to allocate the resource required to accept the following Merge Request. When the MRRA is accepted, a Merge Request conveying Switch A's Zoning configuration is sent to Switch B. Switch B merges the two Zoning configurations and changes the name of the Active Zone Set to "Successful Zone Set Merge: Active Zone Set Name has changed". Since Switch B's Zoning configuration has changed, it transmits its new Zoning configuration on all ISLs, including the ISL connecting with Switch A. Switch A receives Switch B's

new Zoning configuration, that it merges with its Zoning configuration and changes the Zone Set Name to "The Active Zone Set has changed due to a Zone Merge". After the merge, Switch A sends out its new Zoning configuration on all ISLs. When Switch B receives the new Zoning configuration, it confirms that it has an identical Zoning configuration and sends a Merge Response to end the Zone Merge.

Figure 33 shows the Zone merge process when more than two Switches are involved. Note that initial state is Switches connected but no merge processing has begun.




Figure 33 – Merge operation among several Switches

## 10.5.2 Merge Zone rules

### 10.5.2.1 Merge rules in Basic Zoning

In Basic Zoning, when Adjacent Switches exchange Zoning information, they shall abide by the rules in table 279.

**Table 279 – Basic Zoning merge rules**

Adjacent Zoning Configuration	Local Zoning Configuration	Result in Local Switch
Zone Set State is Deactivated.	Zone Set State is Activated or Deactivated.	No change
Zone Set State is Activated.	Zone Set State is Deactivated.	Zone Set gets the Adjacent Zone Set State (i.e., the Zone Set is activated). Zone Set gets the adjacent Zone Set.
Adjacent Zone Set is equal to the Local Zone Set		No change.
Adjacent and Local Zone Sets contain a Zone with the same name but different members, or different Protocol Types		ISL Isolated.
Adjacent Zone Set contains Zones that are not included in the Local Zone Set, and/or Local Zone Set contains Zones that are not included in the Adjacent Zone Set.		Zone Set State becomes Activated. Zone Set is the merge of the local Zones plus the Adjacent Zones. The Zone Set Name is changed to "Successful Zone Set Merge: Active Zone Set Name has changed".

To prevent potential ISL isolation, it is suggested that Zoning be inactivated and a Zone Set propagated through the Fabric by the techniques documented in clause 10.6.

If a received Zoning Configuration contains some unsupported member types, the Zone merge shall fail.

### 10.5.2.2 Merge rules in Enhanced Zoning

In Enhanced Zoning the merge behavior depends by the Merge Control setting. Merge Control is a Fabric wide setting that indicates the type of behavior two Switches exhibit as a result of a merge operation. This setting may be set as Allow or Restrict. A setting of Allow means that the two databases are merged according to the specified merge rules, and the resulting database is the union of the two databases. A setting of Restrict indicates that if the two databases on the corresponding Switches are not identical, then the links between the two Switches shall become isolated.

Before the merge rules pertaining to the Zoning Database are applied, the Zone Merge flags shall be compared.

If one of the two Switches does not support the Zone Set Database and the other one is using it, as indicated by the Zone Merge flags, then no merge shall occur and the ISL shall become isolated.

If the Merge Control or Default Zone settings are not equal for both Switches, then no merge shall occur and the ISL shall become isolated. If the Merge Control and Default Zone settings are equal, Adjacent Switches exchanging Zoning information shall abide by the rules in table 280. If the local Zoning Database is modified as a result of the merge, the new database is propagated to the neighboring Switches.

If a received Zoning Configuration contains some unsupported member types, the Zone merge shall fail.

**Table 280 – Enhanced Zoning merge rules**

<b>Condition Before Merge</b>	<b>Result in Local Switch</b>
Adjacent Zoning Database is equal to the Local Zoning database. Merge Control = Restrict or Allow	No change.
The local and Adjacent Zoning Databases contain a Zone Set, Zone, or Zone alias object with the same name but unlike members, Merge Control = Restrict or Allow	ISL Isolated.
The local and adjacent Zoning Databases contain references with the same name but with different definitions, Merge Control = Restrict or Allow	ISL Isolated
The local and adjacent Zoning Databases contain Zones with the same name but with different attributes, Merge Control = Restrict or Allow	ISL Isolated
Adjacent Zoning Database contains Zones or aliases that are not included in the local Zoning Database. Merge Control = Allow.	Zoning Database is the union of the local database plus the adjacent database
Local Zoning Database is empty, Adjacent Fabric Zone database contains data. Merge Control = Allow.	Data from Adjacent database is used to populate Local database
Adjacent Zoning Database contains Zones or aliases that are not included in the local Zoning Database. Merge Control = Restrict.	ISL Isolated
Local Zoning Database is empty, Adjacent Fabric Zone database contains data. Merge Control = Restrict.	ISL Isolated
Adjacent Zoning Database is empty, Local Fabric Zone database contains data. Merge Control = Restrict.	ISL Isolated
Adjacent Zoning Database is empty, Local Fabric Zone database contains data. Merge Control = Allow.	No Change

## 10.6 Fabric Management Session protocol

### 10.6.1 Fabric Management Session protocol overview

Requests to change a Fabric's Zoning Configuration or security policies are a result of an administration action or control requests via the Management Server. To keep consistency in the Fabric, the Fabric Management Session protocol ensures that only one management entity may change the Fabric at one time. Requests may be made to activate or deactivate a Zone Set, or

security policy. The Switch that receives a change request is referred to as the Managing Switch. The Managing Switch validates the request and exchanges Inter-Switch messages with the other Switches in the Fabric. The protocol for locking the Fabric during policy changes is the Fabric Management Session protocol.

The policy change requests serviced by the Managing Switch are used to:

- a) reserve Change Authorization in each Switch in the Fabric, blocking any other changes from taking place while this change is in progress;
- b) stage configuration changes in each Switch in the Fabric;
- c) apply the staged configuration change in each Switch in the Fabric; and
- d) release Change Authorization in each Switch in the Fabric.

Acquire Change Authorization request and response messages are used to reserve Local Change Authorization in each Switch in the Fabric (i.e., to establish Fabric Management Sessions between the Managing Switch and the Managed Switches in the Fabric). Stage Fabric Configuration request and response messages are used to distribute a Configuration update to each Switch in the Fabric in order to verify that each Switch is able to implement the change. Update Fabric Configuration request and response messages are used to modify the Configuration on each Switch in the Fabric. Release Change Authorization request and response messages are used to release Local Change Authorization in each Switch in the Fabric (i.e., end the Fabric Management Sessions).

### 10.6.2 Reserving Fabric Change Authorization

The Managing Switch shall send an Acquire Change Authorization request to each Managed Switch in the Fabric to reserve the Fabric Change Authorization. Each Acquire Change Authorization request message includes a list of the Domain\_IDs of the Managed Switches being included in the update.

Each Managed Switch that receives an Acquire Change Authorization request shall return an Acquire Change Authorization response to the Managing Switch. The Acquire Change Authorization response message is either an SW\_ACC that indicates success, or an SW\_RJT that specifies why the Acquire Change Authorization operation was unsuccessful. The Managed Switch returns an SW\_RJT indicating "Fabric Changing" if the Domain\_IDs the Managed Switch knows to be in the Fabric are not the same as the Domain\_IDs identified in the request by the Managing Switch. The Managed Switch returns an SW\_RJT indicating "Busy" if its Local Change Authorization is already reserved by another process. Otherwise, the Managed Switch reserves Local Change Authorization for the Managing Switch, and returns an SW\_ACC indicating the operation was "Successful".

The Managing Switch shall wait until an Acquire Change Authorization response has been received from each Managed Switch in the Fabric. If any of the responses indicate that a Managed Switch is Busy or that the Fabric is Changing, the Managing Switch shall initiate the process to release Fabric Change Authorization (see 10.6.5). If all the responses indicate that Local Change Authorization was successfully acquired, the Managing Switch shall initiate the process to stage the update (see 10.6.3).

### 10.6.3 Staging the Fabric Configuration

The Managing Switch shall send a Stage Fabric Configuration request to each Managed Switch in the Fabric to stage the Fabric Configuration. The Stage Fabric Configuration request includes a change command and the appropriate operation data.

Each Managed Switch that receives a Stage Fabric Configuration request shall return a Stage Fabric Configuration response to the Managing Switch. The Stage Fabric Configuration response message is either an SW\_ACC that indicates success, or an SW\_RJT that specifies why the Stage Fabric Configuration operation was unsuccessful.

If a received Zoning Configuration contains some unsupported member types, the Stage Fabric Configuration shall fail.

The Managing Switch shall wait until a Stage Fabric Configuration response has been received from each Managed Switch in the Fabric. If a staging error occurred, the process to release Fabric Change Authorization (see 10.6.5) shall be initiated. If the update was successfully staged in all the Switches in the Fabric, the process to update the Fabric Configuration (see 10.6.4) shall be initiated.

#### **10.6.4 Updating the Fabric Configuration**

The Managing Switch shall send an Update Fabric Configuration request to each Managed Switch in the Fabric to update the Fabric Configuration. There is no data included in the Update Fabric Configuration request message.

Each Managed Switch that receives an Update Fabric Configuration request shall return an Update Fabric Configuration response to the Managing Switch. The Update Fabric Configuration response message is either an SW\_ACC that indicates success, or an SW\_RJT that specifies why the Update Fabric Configuration operation was unsuccessful.

The Managing Switch shall wait until an Update Fabric Configuration response has been received from each Managed Switch in the Fabric. The Managing Switch shall then initiate the process to release Fabric Change Authorization (see 10.6.5).

#### **10.6.5 Releasing Fabric Change Authorization**

The Managing Switch shall send a Release Change Authorization request to each Managed Switch in the Fabric that has reserved Local Change Authorization for the Managing Switch. There is no data included in the Release Change Authorization request message.

Each Managed Switch that receives a Release Change Authorization request shall return a Release Change Authorization response to the Managing Switch. The Release Change Authorization response message is either an SW\_ACC that indicates success, or an SW\_RJT that specifies why the Release Change Authorization operation was unsuccessful.

The Managing Switch shall wait until a Release Change Authorization response has been received from each Managed Switch in the Fabric at which time the Fabric update is complete.

#### **10.6.6 Mapping of a Server session to a Fabric Management Session**

In the context of Enhanced Zoning Management, a management action (i.e., write access to the Zoning Database) to the Zone Server shall occur only inside a Server session. A Server session is delimited by the CT requests Server Session Begin (SSB) and Server Session End (SSE), directed to the Management Service and with GS\_Subtype specifying the Zone Server. Query requests that result in read access to the Zoning Database are not required to be issued inside a Server session, although the information which they report is only consistent inside a Server session.

The Server session on the Zone Server side is translated in a lock of the Fabric on the Fabric side, using the Fabric Management Session protocol. This ensures serialized management access to the Zoning Database by different management applications, and to guarantee a deterministic behavior.

The Switch handling the Zone Server CT requests, on receiving the SSB CT request shall try to become also the Managing Switch of the Fabric Management Session protocol. As such it shall send the ACA SW\_ILS to all the other Switches in the Fabric. The result of this action may be one of the following:

- a) each other Switch responds to the ACA with an SW\_RJT indicating failure. This means that the Fabric is already locked by another Switch, and so somebody else is managing the Zoning Database. The Zone Server shall reject the SSB CT request. The managing application may retry later;
- b) some Switches respond to the ACA with an SW\_ACC indicating success, some with an SW\_RJT indicating failure. This means that another Switch is trying to lock the Fabric at the same moment. The Switch shall send an RCA SW\_ILS to each Switch that accepted the ACA to release its attempt to lock the Fabric. The Zone Server shall reject the SSB CT request. The managing application may retry later; or
- c) each other Switch responds to the ACA with an SW\_ACC indicating success. This means that the Switch has been successful in locking the Fabric, and it owns the lock. The Switch shall prepare a copy of the Zoning Database for the subsequent management actions, and after that the Zone Server shall accept the SSB CT request.

At this point the management application may manage the Zoning Database using the Enhanced Zoning Commands. The Zone Server shall handle the requests applying them to the copy of the Zoning Database. The Zoning Database present in all the other Switches of the Fabric shall not be affected by these Enhanced Zoning Commands.

To apply the updated Zoning Database to the Fabric, the management application shall send a Commit (CMIT) CT request to the Zone Server. The Zone Server shall perform a consistency check of the updated Zoning Database. If the consistency check fails, then the Zone Server shall reject the CMIT CT request, with an appropriate reason code. Consistency checks may fail for many reasons including the following:

- a) a Zone Set Object in the Zone Set Database references Zone Objects that do not exist;
- b) a Zone Object in the Zone Set Database references Zone Alias Objects that do not exist; or
- c) a Zone Object in the Zone Set Database references Zone Attribute Objects that do not exist.

If the consistency check succeeds, then the Managing Switch shall stage the new Zoning Database to the other Switches of the Fabric, sending to each Switch the SFC SW\_ILS.

If one or more Switches reject the SFC SW\_ILS, then this may mean that these Switches are not able to support and enforce the new Zoning Database. The Fabric is unable to have a consistent Zoning Database, and so the Zone Server shall reject the CMIT CT request.

If all the Switches accept the SFC SW\_ILS, then the Fabric is able to have a consistent Zoning Database. For certain operations the Managing Switch may need to send a second SFC message to each Managed Switch (see 10.6.7). When all Switches have accepted the SFCs they are prepared to successfully update the Zoning definitions. Then the Managing Switch shall send to each other Switch the UFC SW\_ILS to make the staged Zoning Database the Fabric Zoning Database.

To terminate the Server session, the management application shall send a Server Session End (SSE) CT request to the Zone Server. Then the Managing Switch shall end the Fabric Management Session by sending an RCA SW\_ILS to every Switch in the Fabric. When every RCA has been accepted, the Managing Switch shall destroy the copy of the Zoning Database and the Zone Server shall accept the SSE CT request.



If a management application does not send a CMIT request inside a Server session, then every modification that it may have performed is not applied to the Fabric and is lost. The management application may issue the CMIT request more than one time inside the same Server session.

If after a successful SSB CT request a SSE CT request is never received (i.e, the management application locks the Fabric and then crashes), the Fabric Zone Server shall close the Server session if it does not receive any Enhanced Zoning Commands for 2 minutes with a 10% tolerance. Consequently, the Managing Switch shall release the lock over the Fabric. Management applications are expected to keep the Fabric Management Session alive even in absence of management inputs.

If instead the Managing Switch crashes while locking the Fabric, the other Switches may detect the situation and release the lock.

#### **10.6.7 Fabric behavior to handle the CT SFEZ request**

If the Fabric is functioning in Basic mode, and the SFEZ command has requested that the Zoning operational mode of the Fabric be changed to Enhanced, the Switch handling the CT SFEZ request shall initiate a Fabric Management Session. This causes a redistribution of the existing Zoning Database to the other Switches of the Fabric by using the operation request value 'Activate Zone Set Enhanced' in a Stage Fabric Configuration (SFC) SW\_ILS. In this manner the Zoning Database is distributed using the Enhanced Zoning Data structures described in 10.4.4. If this step is successful, then the Zoning Policy Flags requested by the SFEZ command are propagated to the other Switches of the Fabric by using the operation request value 'Set Zoning Policies' in a second Stage Fabric Configuration (SFC) SW\_ILS. If this step is successful, the UFC SW\_ILS is used to apply the new configuration, changing the Zoning operational mode of all Switches in the Fabric to Enhanced mode. The Fabric Management Session shall then be released.

If the Fabric is functioning in Enhanced mode, then only the SFEZ command changes the Fabric's Zoning Policies. This causes a Fabric Management Session to be initiated and an SFC SW\_ILS to be sent to all Switches. The operation request value 'Set Zoning Policies' shall be used in this SFC SW\_ILS. If a Fabric Management Session is already active, then the SFEZ command shall generate an immediate SFC SW\_ILS with operation request value 'Set Zoning Policies' and an UFC SW\_ILSs, keeping alive the existing Fabric Management Session.

NOTE 23 – Zoning structures managed within the Enhanced Zoning framework may not be subsequently managed using management operations defined in the Basic Zoning framework.

NOTE 24 – No mechanism exists to convert from Enhanced Zoning mode to Basic Zoning mode.

#### **10.6.8 Fabric behavior to handle the CT AAPZ and RAPZ requests**

The AAPZ and RAPZ CT requests (see FC-GS-8) modify the Active Zone Set, therefore they require a lock of the Fabric, however they are not processed inside a Server session. An implementation shall not wait more than one minute after receiving an AAPZ or RAPZ request before attempting to acquire a Fabric lock to change the Active Zone Set. This enables coalescing multiple AAPZ and RAPZ requests into a single change to the Active Zone Set through a single Fabric lock.

### **10.7 Switch behaviors during merge**

To facilitate interoperability between newer and older Switches, a Switch shall inspect the ELP Revision field and send the appropriate Merge Request Protocol Version. If the ELP Revision field is 2, the link should isolate or the Merge Request should have a Protocol Version = 0. If the ELP Revision field is greater than 2, the Switch may send the configured Merge Request Protocol Version or the link shall be Isolated.

## 11 Distributed broadcast

### 11.1 Overview

Distributed broadcast provides a mechanism to distribute broadcast frames without duplicating them or forming a loop. It is based on the FSPF tree with additional rules to provide one and only one broadcast path. The key is to build the same spanning tree with the same root for all Switches.

Broadcast frames shall be Class 3 frames with a D\_ID of FFFFFFFh specified.

### 11.2 Spanning tree

The following list of rules are used to determine which ISLs are broadcast ISLs, and which are not:

- 1) lowest Domain\_ID Switch becomes the root of the tree;
- 2) build the Shortest Path tree using FSPF cost as metric;
- 3) if a link is advertised with different costs in the two directions, the cost advertised by the Switch nearer to the root Switch (i.e., the upstream Switch) shall be used to determine the lowest cost path;

NOTE 25 – This is necessary when the two Switches advertise different costs for ISLs.

- 4) if a Switch has multiple equivalent paths to the root of the tree, the ISL to the upstream Switch with the lowest Domain\_ID shall be selected;
- 5) if there are multiple equivalent ISLs between a pair of Switches, the ISL connected to the upstream Switches' lowest E\_Port Index shall be selected; and
- 6) the E\_Ports selected in this process (i.e., broadcast member E\_Ports) are used to forward broadcast frames;

The following set of rules are used for forwarding broadcast frames through the Fabric once the broadcast member E\_Ports are identified:

- a) a broadcast frame received on any Fx\_Port is forwarded on all broadcast member E\_Ports. In addition if Zoning is enabled the broadcast frame is forwarded to all other Fx\_Ports in the Broadcast Zone on that Switch, or if Zoning is not enabled on all other Fx\_Ports on that Switch;
- b) a broadcast frame received on any broadcast member E\_Port is forwarded on all other broadcast member E\_Ports on that Switch. In addition if Zoning is enabled the broadcast frame is forwarded to all other Fx\_Ports in the Broadcast Zone on that Switch, or if Zoning is not enabled on all other Fx\_Ports on that Switch; and
- c) a broadcast frame received on any other port is discarded.

#### 11.2.1 Spanning tree example

In the example in figure 34, the Switch with Domain\_ID 5 is the root of the broadcast tree. From Domain\_ID 7 to Domain\_ID 5, there are 2 equal cost ISLs, so the path from Domain\_ID 7 port 5 to Domain\_ID 5 port 1 becomes the broadcast ISL. From Domain\_ID 30 to the root, there are multiple equal cost paths. Domain\_ID 7 is the lowest upstream Switch, therefore, this path is chosen. Secondly, there are multiple equal cost ISLs to Domain\_ID 7, the ISL from Domain\_ID 30 port 1 to Domain\_ID 7 port 6 is chosen because this is the lowest E\_Port index on the upstream Switch.

A broadcast frame received on an Fx\_Port on Domain\_ID 30 is forwarded to Domain\_ID 7 E\_Port index 6. From there it is forwarded to Domain\_ID 5 E\_Port index 1, and from there to Domain\_ID 12 E\_Port index 2.




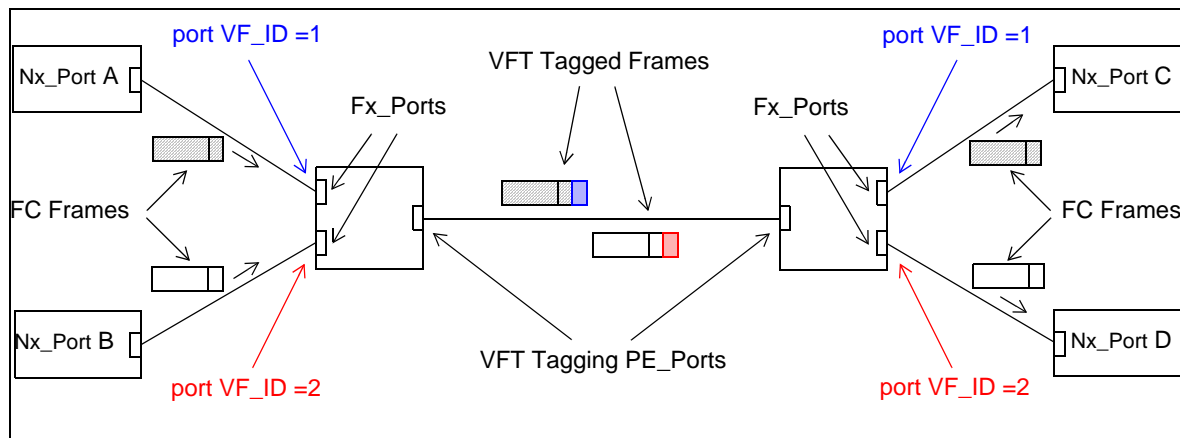
Figure 34 – Broadcast path selection example

## 12 Virtual Fabrics Switch support

### 12.1 Overview

The Virtual Fabric Tagging Header (VFT\_Header, see FC-FS-5) allows Fibre Channel frames to be tagged with the Virtual Fabric Identifier (VF\_ID) of the Virtual Fabric (VF) to which they belong. Tagged frames (i.e., frames with a VFT\_Header) belonging to different Virtual Fabrics may be transmitted over the same physical link. By combining VFT-Headers and other features, Virtual Fabrics provide compartmentalization of access and management. The VFT\_Header may be supported by PN\_Ports, PF\_Ports and PE\_Ports.

The use of VFT\_Header between PE\_Ports allows implementation of Virtual Fabrics without requiring any change in the PN\_Port/Fx\_Port interface, as shown in figure 35.



**Figure 35 – Virtual Fabrics**

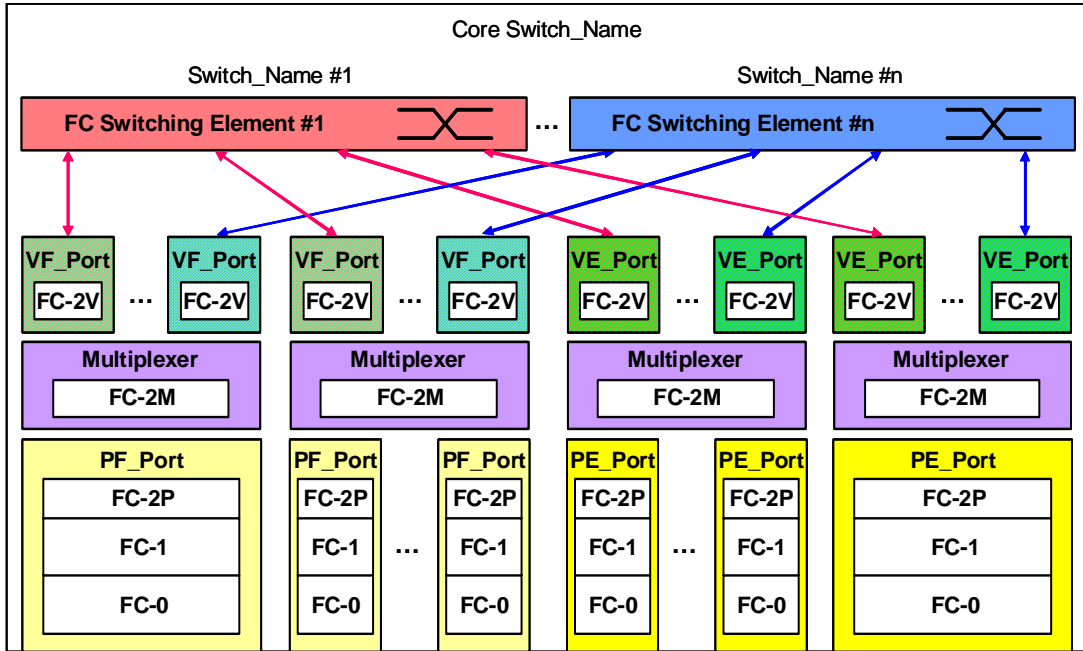
Each Fx\_Port of a VF capable Switch (i.e., a Switch supporting Virtual Fabrics) shall have a configurable port VF\_ID, so that Nx\_Ports may access Virtual Fabric features without modifications. The port VF\_ID shall be associated to any untagged Fibre Channel frame received by the Fx\_Port. A VF capable Switch shall perform frame forwarding by considering the Virtual Fabric the frame belongs to, as identified by the VF\_ID. When transmitted between a pair of tagging PE\_Ports (i.e., PE\_Ports belonging to FC\_Ports processing the VFT\_Header), each Fibre Channel frame shall be tagged with a VFT\_Header. When a VFT\_Header tagged frame is received by a tagging PE\_Port of a VF capable Switch, the VF\_ID carried in the VFT\_Header shall be used to perform frame forwarding, together with the D\_ID carried in the Frame\_Header.

As shown in figure 35, the Fibre Channel frames sent by Nx\_Port A are associated with the Virtual Fabric having VF\_ID 1 when received by the Fx\_Port. The VF\_ID is used by the Switch to perform frame forwarding. Frames transmitted over the tagging PE\_Port are tagged with a VFT\_Header carrying the VF\_ID, and their CRC is recomputed (see FC-FS-5). The receiving tagging PE\_Port retrieves from the VFT tagged frame the VF\_ID and uses it together with the D\_ID carried in the Frame\_Header to route the frames to Nx\_Port C. The Fx\_Port connected to the destination Nx\_Port C removes the VFT\_Header, recomputes the CRC (see FC-FS-5) and delivers the original Fibre Channel frames.

NOTE 26 – Under rare circumstances VF\_IDs may be aliased, resulting in multiple VF\_IDs referring to the same Virtual Fabric.

## 12.2 VF Capable Switch functional model

The functional model of a VF capable Switch is shown in figure 36.



**Figure 36 – Functional model of a VF capable Switch**

A VF capable Switch is functionally a collection of multiple Switching Elements hosted in the same Core Switch. There is one Switching Element per each Virtual Fabric hosted on the Core Switch.

Each Switching Element is identified by a unique Switch\_Name. In addition, the Core Switch is identified by a unique Core Switch\_Name. Each Virtual Fabric is identified by a 12-bit Virtual Fabric Identifier (VF\_ID). VF\_ID Identifiers are administratively set using a management interface.

Each Switching Element is connected to one or more VF\_Ports or VE\_Ports. Each VF\_Port or VE\_Port is connected to a single Switching Element. Multiple VF\_Ports belonging to different Virtual Fabrics may share one or more PF\_Ports through the multiplexing and tagging functions of the Multiplexer. Multiple VE\_Ports belonging to different Virtual Fabrics may share one or more PE\_Ports through the multiplexing and tagging functions of the Multiplexer.

Physical links are shared across multiple Virtual Fabrics using the VFT\_Header. The Multiplexer functions of multiplexing and tagging logic are driven by the VF\_ID in the VFT\_Header. Upon receiving a VFT tagged frame from a PF\_Port or PE\_Port, the Multiplexer logic delivers the frame to the appropriate VF\_Port or VE\_Port connected with the appropriate Switching Element. (i.e., the Switching Element associated with the Virtual Fabric whose VF\_ID is carried in the VFT\_Header).

When transmitted between a pair of tagging E\_Ports, each Fibre Channel frame shall be tagged with a VFT\_Header. A VF capable Switch shall perform frame forwarding by considering the Virtual Fabric the frame belongs to, as identified by the VF\_ID.

Each Switch port of a VF capable Switch shall have a configurable port VF\_ID. The port VF\_ID shall be associated to any untagged Fibre Channel frame received by the Switch port. This allows the interconnection of VF capable Switches with non VF capable Switches. Any untagged Fibre Channel

frame received by an E\_Port or Fx\_Port on a VF capable Switch shall be implicitly associated with the port VF\_ID for processing. The port VF\_ID is then used by the tagging logic to deliver the frame to the appropriate Switching Element. In absence of any explicit configuration, the value 001h should be used as default port VF\_ID.

Switches supporting Virtual Fabrics may not receive VFT tagged frames on all Switch ports. This may occur for the following reasons:

- a) a Switch port is administratively configured to not use VFT\_Headers;
- b) the port at the far end of a link is administratively configured to not use VFT Headers; or
- c) the port at the far end of a link is not capable of processing VFT\_Headers.

### 12.3 Switch\_Names usage

The Switch\_Names of the Switching Elements and the Core Switch\_Name shall be used as follows:

- a) in state P5 (see figure 16), the Switch\_Name of the Switching Element associated with the port VF\_ID shall be used when transmitting the ELP SW\_ILS;
- b) in state P17 (see figure 17), the Switch\_Name shall abide by the rules in FC-SP-2; and
- c) when a Switch port initializes as an E\_Port in state P10 (see figure 17), the Switch\_Name of the Switching Element associated with the port VF\_ID shall be used for any subsequent operation or protocol.

### 12.4 Configuration information

A VF capable Switch shall maintain the following configuration parameters per each Switch port:

- a) Tagging Administrative Status, used to negotiate the VFT tagging operational mode of the Switch port (see 6.2.25.2.2);
- b) port VF\_ID (see 12.2 and 6.2.25.2.3); and
- c) Locally-Enabled VF\_ID List, used to negotiate the list of Virtual Fabrics operational over the Switch port (see 6.2.25.2.4).

### 12.5 Enabling VFT tagging on Switch ports

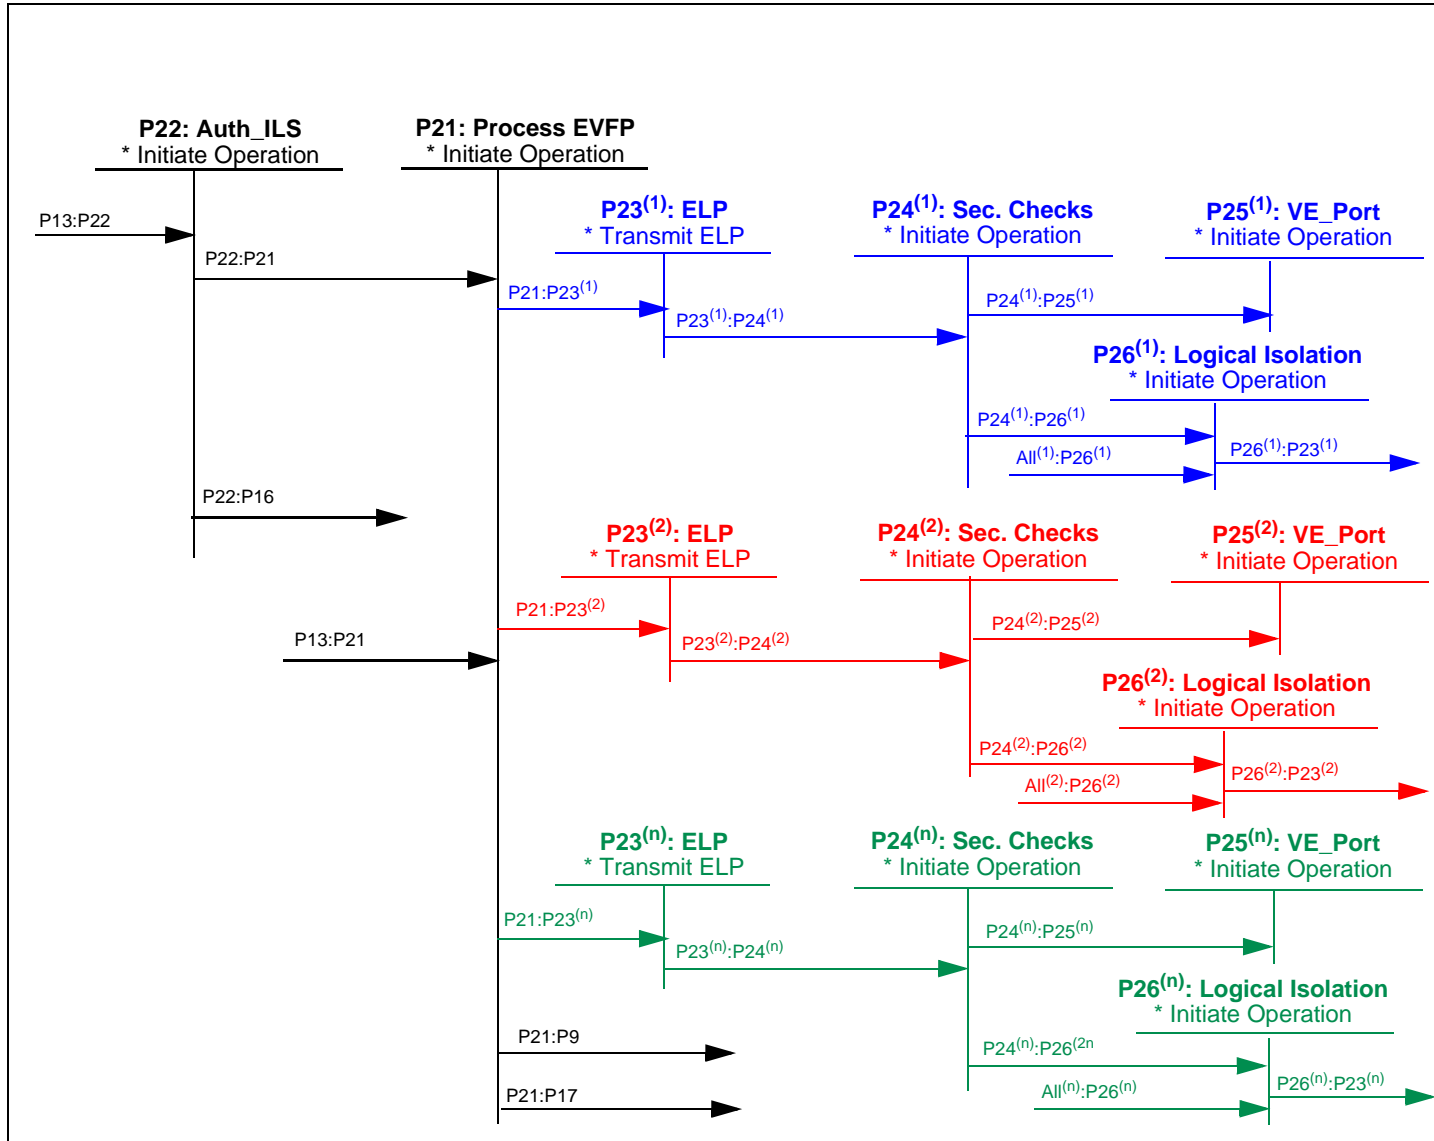
Figure 37 shows the extended Switch port mode initialization state machine enhanced to enable Virtual Fabrics. In state P13 (see figure 17), two Switch ports may negotiate to perform the EVFP processing (see 12.6) if both of them support Virtual Fabrics. The support for Virtual Fabrics is indicated by the 'Virtual Fabrics Supported' Protocol ID value specified in table 72.

The Switch port sending the ESC request indicates support for Virtual Fabrics by including the 'Virtual Fabrics Supported' Protocol ID specified in table 72 in the ESC payload. The replying Switch port selects to negotiate Virtual Fabrics parameters by choosing the 'Virtual Fabrics Supported' Protocol ID specified in table 72 in the ESC SW\_ACC payload. Then the Switch port initialization proceeds to state P21 or to state P22.

A Switch port connected to a B\_Port shall not indicate nor select the support for Virtual Fabrics in the ESC protocol if the directly connected B\_Port did not announce support for Virtual Fabric Tagging by setting to one the 'Bridge Virtual Fabrics' flag in the ELP SW\_ILS.

If one of the two Switch ports does not support Virtual Fabrics, the Switch port initialization proceeds to state P17 (see figure 17).

When VFT tagging is enabled on a link, a Link Reset (see FC-FS-5) shall not change the tagging process, while a Link Initialization (see FC-FS-5) shall stop the tagging process and bring the involved Switch ports to state P0 (see figure 16).



**Figure 37 – Switch port mode initialization state machine - Virtual Fabric support**

**Transition P13:P21.** Occurs when the two Switch ports negotiated to perform the EVFP processing, and neither Switch port requires Authentication in state P22.

**Transition P13:P22.** Occurs when the two Switch ports negotiated to perform the EVFP processing, and authentication is required by at least a Switch port in state P22.

**State P22: AUTH\_ILS.** While in this state an Authentication transaction (see FC-SP-2) shall be performed. If the port receives an EVFP before authentication checks are complete, the port shall respond with an SW\_RJT, indicating a reason code of Logical Busy and a reason code explanation of Security Checks in Progress. Switch\_Name usage shall abide by the rules defined in FC-SP-2.

**Transition P22:P16.** Occurs when the Authentication transaction performed in state P22 fails.



**Transition P22:P21.** Occurs when the Authentication transaction performed in state P22 completes successfully.

**State P21: Process EVFP.** The Switch port shall perform EVFP processing as described in 12.6.

**Transition P21:P17.** Occurs when the EVFP processing determined that VFT tagging is not performed and the two Switch ports have the same port VF\_ID.

**Transition P21:P9.** Occurs when the EVFP processing determined that VFT tagging is not performed and the two Switch ports have a different port VF\_ID.

**Transition P21:P23<sup>(k)</sup>.** Occurs when the EVFP processing determined that VFT tagging is performed. There is a different state for each Virtual Fabric negotiated to be used on the link. The state for Virtual Fabric K is denoted P23<sup>(k)</sup>.

**State P23<sup>(k)</sup>: ELP.** In this state the Fibre Channel frames transmitted by the Switch port are tagged with the VFT\_Header carrying VF\_ID K. An ELP, tagged with VF\_ID K, is transmitted. This ELP shall carry the Switch Name of the Switching Element associated with VF\_ID K and the operational parameters (e.g., timeout values, Classes of service) of Virtual Fabric K. No flow control configuration is required in this state, because it is performed in state P5.

**Transition P23<sup>(k)</sup>:P24<sup>(k)</sup>.** Occurs when the ELP processing in state P23 is completed.

**State P24<sup>(k)</sup>: Security Checks.** In this state the Fibre Channel frames transmitted by the Switch port are tagged with the VFT\_Header carrying VF\_ID K. The Switch port initiates and responds to all required security checks (see FC-SP-2), if any. If the port receives an EFP before authentication checks are complete, the port shall respond with an SW\_RJT, indicating a reason code of Logical Busy and a reason code explanation of Security Checks in Progress.

**Transition P24<sup>(k)</sup>:P25<sup>(k)</sup>.** Occurs when the Security Checks performed in state P24<sup>(k)</sup> complete successfully.

**State P25<sup>(k)</sup>: VE\_Port.** In this state the Switch port operates as VFT tagging PE\_Port. Fibre Channel frames transmitted by the Switch port are tagged with the VFT\_Header carrying VF\_ID K. The VE\_Port shall participate in the next phase of Fabric Configuration in Virtual Fabric K. The Switch\_Name of the Switching Element associated with VF\_ID K shall be used for any subsequent operation or protocol in Virtual Fabric K.

**State P26<sup>(k)</sup>: Logical Isolation.** In this state the VE\_Port corresponding to Virtual Fabric K becomes logically Isolated (i.e., in Virtual Fabric K no Class N traffic flows and only the SW\_ILSs specified in 7.6 may be communicated).

**Transition P24<sup>(k)</sup>:P26<sup>(k)</sup>.** Occurs when the Security Checks performed in state P24<sup>(k)</sup> complete unsuccessfully.

**Transition All<sup>(k)</sup>:P26<sup>(k)</sup>.** Occurs when a protocol in Virtual Fabric K causes the corresponding VE\_Port to go in Logical Isolation state for any of the reasons listed in 7.6.

**Transition P26<sup>(k)</sup>:P23<sup>(k)</sup>.** Occurs when the Logically Isolated VE\_Port corresponding to Virtual Fabric K receives or transmits an ELP.

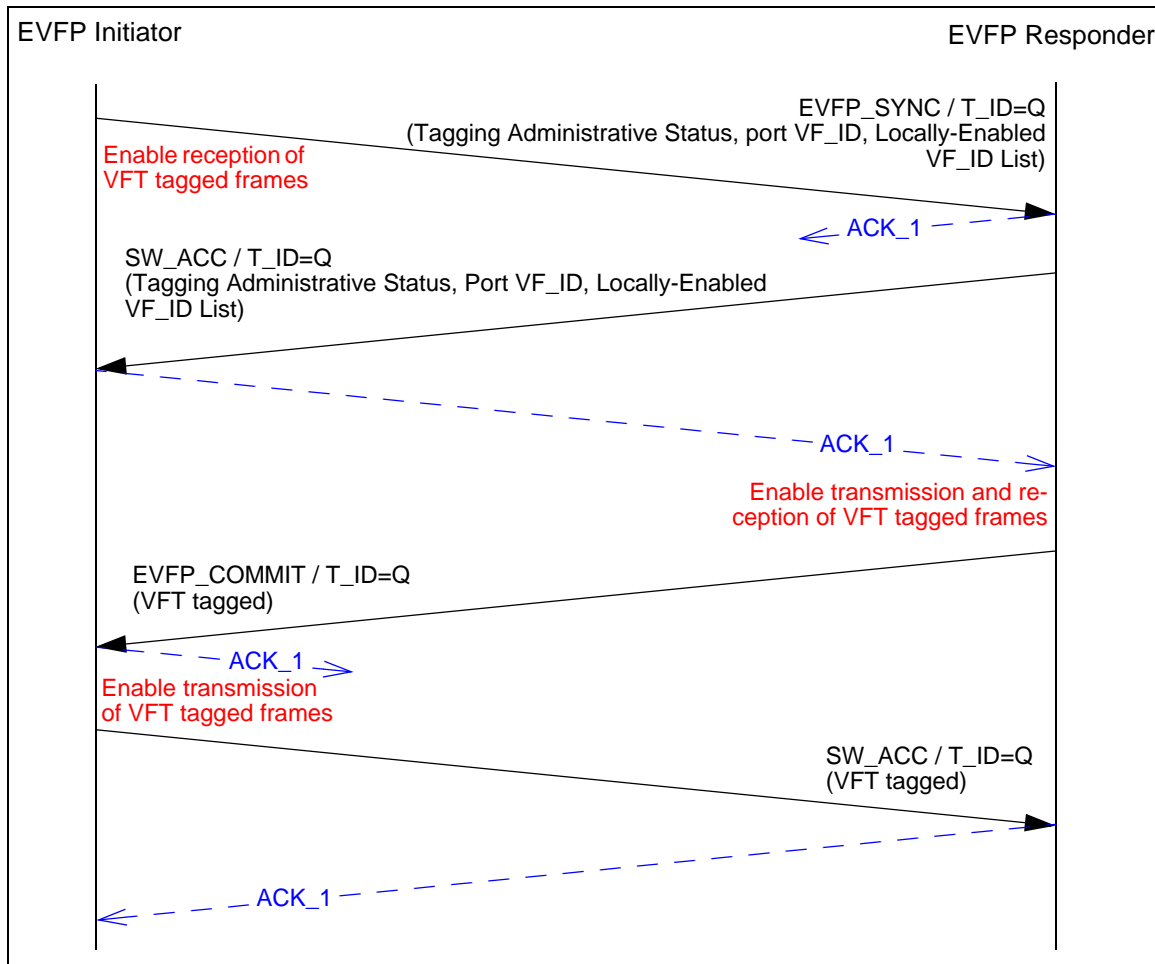
## 12.6 Exchange Virtual Fabrics parameters processing

### 12.6.1 Overview

The Exchange Virtual Fabrics Parameters (EVFP) protocol allows peers of Interconnect\_Ports belonging to VF capable Switches to:

- a) negotiate the VFT Tagging operational mode;
- b) verify the consistency of the two port VF\_IDs; and
- c) establish the list of operational Virtual Fabrics across the Inter-Switch Link.

An EVFP transaction occurs between an EVFP Initiator and an EVFP Responder. An EVFP transaction (see figure 38) is identified by a unique Transaction Identifier (T\_ID), and consists of a synchronizing phase (EVFP\_SYNC) followed by a commit phase (EVFP\_COMMIT).



**Figure 38 – A generic EVFP transaction**

The VF\_ID value FEFh is used by the EVFP protocol for certain operations and is referred to as Control VF\_ID. The EVFP protocol, during the Switch port initialization, proceeds as follows:

- 1) the EVFP Initiator shall start the EVFP transaction by sending the EVFP\_SYNC message (see 6.2.25.2) to the EVFP Responder. In the EVFP\_SYNC message, the EVFP Initiator shall specify the Transaction Identifier, and shall send its Core Switch\_Name, together with its Tagging Administrative Status (see 6.2.25.2.2), port VF\_ID (see 6.2.25.2.3) and Locally-Enabled VF\_ID List (see 6.2.25.2.4). On sending the EVFP\_SYNC message the EVFP Initiator enables the reception of VFT tagged frames;
- 2) the EVFP Responder shall reply with an EVFP\_SYNC SW\_ACC carrying its Tagging Administrative Status, port VF\_ID and Locally-Enabled VF\_ID List. Then the EVFP Responder shall determine if VFT Tagging is to be enabled on the link, according to table 108. If VFT Tagging is to be enabled, the EVFP Responder shall go to step 3. If VFT Tagging is not to be enabled, the EVFP Responder shall check the received peer's port VF\_ID:
  - a) if the peer's port VF\_ID is not equal to the local port VF\_ID, on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) the EVFP protocol terminates and the EVFP Responder goes in Isolated state (transition P21:P9, see 12.5); or
  - b) if the peer's port VF\_ID is equal to the local port VF\_ID, on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) the EVFP protocol terminates and the EVFP Responder goes in state P17 (transition P21:P17, see 12.5).

On receiving the EVFP\_SYNC SW\_ACC, the EVFP Initiator shall determine if VFT Tagging is to be enabled on the link, according to table 108. If VFT Tagging is to be enabled, on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) the EVFP Initiator shall enable the reception of VFT tagged frames in its port VF\_ID and shall go to step 4. If VFT Tagging is not to be enabled, the EVFP Initiator disables the reception of VFT tagged frames and shall check the received peer's port VF\_ID:

- a) if the peer's port VF\_ID is not equal to the local port VF\_ID, on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) the EVFP protocol terminates and the EVFP Initiator goes in Isolated state (transition P21:P9, see 12.5); or
  - b) if the peer's port VF\_ID is equal to the local port VF\_ID, on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) the EVFP protocol terminates and the EVFP Initiator goes in state P17 (transition P21:P17, see 12.5);
- 3) on completion of the EVFP\_SYNC Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC), the EVFP Responder shall enable both transmission and reception of VFT tagged frames for the Virtual Fabrics operational on the link, computed as explained in 6.2.25.2.4. Transmission and reception of VFT tagged frames for the Control VF\_ID shall be implicitly enabled. Transmission and reception of VFT tagged frames for the EVFP Initiator's port VF\_ID shall be also enabled on the link, to allow a successful completion of the EVFP protocol. Then the EVFP Responder shall send an EVFP\_COMMIT message (see 6.2.25.3), tagged with the EVFP Initiator's port VF\_ID; and
  - 4) on receiving the VFT tagged EVFP\_COMMIT, the EVFP Initiator shall enable both transmission and reception of VFT tagged frames for the Virtual Fabrics operational on the link, computed as explained in 6.2.25.2.4. Transmission and reception of VFT tagged frames for the Control VF\_ID shall be implicitly enabled. Transmission and reception of VFT tagged frames for the EVFP Initiator's port VF\_ID shall be also enabled on the link, to allow a successful completion of the EVFP protocol. Then the EVFP Initiator shall send an EVFP\_COMMIT SW\_ACC message tagged with its port VF\_ID.

When tagging is enabled the EVFP transaction completes successfully on completion of the EVFP\_COMMIT Exchange, for both the EVFP Initiator and EVFP Responder. If the computed set of VF\_IDs operational on the link does not include the EVFP Initiator's port VF\_ID, transmission and

reception of VFT tagged frames for such VF\_ID shall be disabled on the link upon completion of the EVFP transaction. When the EVFP transaction is completed the processing continues independently for each Virtual Fabric operational on the link, as specified by transitions P21:P23<sup>(k)</sup> (see 12.5). If the computed set of VF\_IDs operational on the link is NULL, the involved Switch ports remain in state P21 (see 12.5) until a new EVFP transaction is performed in the Control VF\_ID.

If two Interconnect\_Ports start an EVFP transaction at the same time, or if an Interconnect\_Port is acting as an EVFP Initiator and receives an EVFP\_SYNC message from the designated EVFP Responder, one of the two EVFP transactions shall be aborted. The Interconnect\_Port that sent the EVFP\_SYNC message with the numerically higher Core Switch\_Name shall remain the EVFP Initiator, while the Interconnect\_Port that sent the EVFP\_SYNC message with the numerically lower Core Switch\_Name shall become the EVFP Responder. The Interconnect\_Port that remains the EVFP Initiator shall reply to the received EVFP\_SYNC message with a 'EVFP collision' SW\_RJT (see 6.2.25.1). The Interconnect\_Port that becomes the EVFP Responder shall reply to the received EVFP\_SYNC message and abort its own transaction upon receipt of the SW\_RJT.

The EVFP protocol is used also when some Switch port configuration information (see 12.4) are changed by a management action. The EVFP messages may be carried in Fibre Channel frames tagged with the port VF\_ID if the EVFP protocol begins while the link is not performing VFT tagging (see 12.6.1). The EVFP messages are carried in Fibre Channel frames tagged with the Control VF\_ID if the EVFP protocol begins while the link is performing VFT tagging (see 12.6.2 and 12.6.3).

### 12.6.2 Changing the VFT tagging mode

When a management action changes the Administrative Tagging Mode of an E\_Port belonging to a VF capable Switch that determined during initialization the peer supports the EVFP protocol, the E\_Port shall determine if the link has to change its VFT Tagging mode (i.e., if it has to transition from tagging to untagging mode or from untagging to tagging mode) by acting as EVFP Initiator as follows. If the E\_Port is currently performing tagging, all EVFP protocol messages shall be tagged with the Control VF\_ID. If the E\_Port is currently not performing tagging, all EVFP protocol messages shall be untagged.

- 1) the EVFP Initiator shall start the EVFP transaction by sending the EVFP\_SYNC message to the EVFP Responder. The EVFP\_SYNC message shall carry the updated Tagging Administrative Status (see 6.2.25.2), port VF\_ID, and the Locally-Enabled VF\_ID List; and
- 2) the EVFP Responder shall reply with an EVFP\_SYNC SW\_ACC carrying its Tagging Administrative Status, port VF\_ID and Locally-Enabled VF\_ID List. The EVFP Responder shall determine if VFT Tagging has to be changed on the link, according to table 108. The EVFP Responder:
  - a) if VFT Tagging has not to be changed, on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) terminates the EVFP protocol; or
  - b) if VFT Tagging has to be changed, on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) shall perform a link initialization.

On receiving the EVFP\_SYNC SW\_ACC, the EVFP Initiator shall determine if VFT Tagging has to be changed on the link, according to table 108. The EVFP Initiator:

- a) if VFT Tagging has not to be changed, on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) terminates the EVFP protocol; or
- b) if VFT Tagging has to be changed, shall participate in the link initialization initiated by the EVFP Responder.

### 12.6.3 Adding or removing Virtual Fabrics

When a management action changes the Locally-Enabled VF\_ID List over a tagging E\_Port, the E\_Port shall initiate the EVFP protocol by acting as EVFP Initiator as follows. All EVFP protocol messages shall be tagged with the Control VF\_ID.

- 1) the EVFP Initiator shall start the EVFP transaction by sending the EVFP\_SYNC message to the EVFP Responder. The EVFP\_SYNC message shall carry the Tagging Administrative Status, port VF\_ID, and the updated Locally-Enabled VF\_ID List (see 6.2.25.2.4);
- 2) the EVFP Responder shall reply with an EVFP\_SYNC SW\_ACC carrying its Tagging Administrative Status, port VF\_ID and Locally-Enabled VF\_ID List. The EVFP Responder, depending on the resulting operational VF\_ID List (see 6.2.25.2.4):
  - a) if the operational VF\_ID List did not change, terminates the EVFP protocol on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) in the Control VF\_ID; or
  - b) if the operational VF\_ID List did change, performs step 3 on completion of the Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) in the Control VF\_ID.

On receiving the EVFP\_SYNC SW\_ACC in the Control VF\_ID, the EVFP Initiator, depending on the resulting operational VF\_ID List (see 6.2.25.2.4):

- a) if the operational VF\_ID List did not change, terminates the EVFP protocol on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) in the Control VF\_ID; or
  - b) if the operational VF\_ID List did change, performs step 4 on completion of the Exchange (i.e., on sending the ACK\_1 for the EVFP\_SYNC SW\_ACC) in the Control VF\_ID.
- 3) on completion of the EVFP\_SYNC Exchange (i.e., on receiving the ACK\_1 for the EVFP\_SYNC SW\_ACC) in the Control VF\_ID, the EVFP Responder shall apply the updated operational VF\_ID List, enabling the added Virtual Fabrics and disabling the removed Virtual Fabrics. Then the EVFP Responder shall send an EVFP\_COMMIT message; and
  - 4) on receiving the EVFP\_COMMIT message, the EVFP Initiator shall apply the updated operational VF\_ID List, enabling the added Virtual Fabrics and disabling the removed Virtual Fabrics. Then the EVFP Initiator shall send an EVFP\_COMMIT SW\_ACC message.

When the operational VF\_ID List changes, the EVFP transaction completes successfully on completion of the EVFP\_COMMIT Exchange for both the EVFP Initiator and EVFP Responder. When the EVFP transaction is completed, the updated operational VF\_ID List is operative.

### 12.6.4 Changing the port VF\_ID

When a management action changes the port VF\_ID of a tagging PE\_Port, no changes are applied to the link.

When a management action changes the port VF\_ID of a non-tagging PE\_Port, the PE\_Port shall perform a link initialization.

When a management action changes the port VF\_ID of a Switch port in Isolated state, the Switch port shall go in state P0 (see figure 16).

## 13 Enhanced Commit Service

### 13.1 Overview

The Enhanced Commit Service (ECS) builds on the Fabric Management Session Protocol defined in clause 10 to provide a general mechanism to manage the serialization and updating of Fabric resources. ECS provides:

- a) serialization and locking of resources on a per fabric application basis;
- b) error recovery; and
- c) transaction semantics.

ECS may operate in assisted mode or in autonomous mode. When operating in assisted mode the protocol processing is controlled by a fabric application (see 13.2). In this case assisted Mode does not perform or enable the error recovery defined by the enhanced commit service. When operating in autonomous mode the protocol processing proceeds as described in 13.3 and provides error recovery during the commit process.

The enhanced commit service internal link services (EACA, ESFC, EUFC, ERCA, and TCO) are defined in this standard and are based on the Zoning update commands (ACA, SFC, UFC, RCA) also defined in this standard (see 6.1).

### 13.2 Assisted mode protocol operations

When operating in assisted mode (see 6.2.26.1.3) the ECS protocol processing begins, proceeds and terminates under the control of a fabric application.

A Switch begins the ECS protocol processing when requested by the fabric application (e.g., when an SSB CT request is received by the Security Server, see FC-SP-2). The Switch begins the ECS protocol attempting to lock the Fabric for the specified fabric application, by sending an EACA request to the Switches specified in the ECS Switch List, following the order of the ECS Switch List. If all these Switches accept the EACA request, then the sending Switch becomes the ECS Managing Switch for the transaction and shall return the control to the fabric application with a success indication (e.g., by replying with an SSB Accept CT\_IU). If one or more Switches rejects the EACA request, then the sending Switch shall send an ERCA request to the Switches that accepted the EACA request and shall return the control to the fabric application with an error indication (e.g., by replying with an SSB Reject CT\_IU).

Once the Fabric is locked for the specified fabric application, the ECS protocol processing proceeds under the control of the fabric application. According to the requests received from the fabric application, the ECS Managing Switch may distribute and commit information to the Switches participating in the ECS transaction by sending one or more ESFC and EUFC requests to the Switches specified in the ECS Switch List, following the order of the ECS Switch List. A success indication shall be returned to the fabric application if all Switches participating in the ECS transaction accept the ESFC or EUFC request. An error indication shall be returned to the fabric application if one or more of the Switches participating in the ECS transaction reject an ESFC or EUFC request.

The ECS protocol processing terminates when the fabric application requests to unlock the Fabric (e.g., when an SSE CT request is received by the Security Server, see FC-SP-2). The ECS Managing Switch unlocks the Fabric for the specified fabric application by sending an ERCA request to the Switches specified in the ECS Switch List, following the order of the ECS Switch List.

The TCO SW\_ILS is not used by the ECS protocol when operating in assisted mode.

### **13.3 Autonomous mode protocol operations**

#### **13.3.1 Protocol phases**

##### **13.3.1.1 Overview**

When operating in autonomous mode (see 6.2.26.1.3) the ECS protocol allows Fabric resources associated with a specific fabric application to be locked in each managed Switch. An ECS operation consists of one transaction bounded by an EACA request that begins the transaction, and an ERCA request that ends the transaction. Between the acceptance of the EACA by all managed Switches and the generation of the ERCA by the Managing Switch, the ESFC and EUFC requests are used to update and commit application resources in each managed Switch.

An ECS operation affects a subset of Switches in the Fabric. The Switch List indicates the list of managed Switches for the ECS operation and indicates the order that the Managing Switch sends ECS commands to the managed Switches. The Switch List also specifies which Switches are authorized to participate in ECS recovery processing.

##### **13.3.1.2 Phase one**

The Managing Switch initiates the commit process by sending an EACA request to all Switches specified in the ECS Switch List. Once the EACA is accepted, the Managing Switch and all the managed Switches of EACA, reserve the resources associated with the specified application. Once all managed Switches accept the EACA, the Managing Switch transitions to phase two of the commit process.

If any Switch rejects the EACA, then the Managing Switch aborts the commit process by sending an ERCA to all the Switches that accepted the EACA.

##### **13.3.1.3 Phase two**

The Managing Switch initiates the second phase of the commit process by sending an ESFC request to all managed Switches. The application specific data is validated and staged by the managed Switch receiving the ESFC. If all the managed Switches successfully complete the ESFC processing, then the Managing Switch transitions to phase three of the commit process.

If the application data is invalid or consistency checks fail on the managed Switch, then the Switch rejects the ESFC. If a reject is received for an ESFC then the Managing Switch performs an ERCA on all managed Switches and the commit process is aborted.

##### **13.3.1.4 Phase three**

The Managing Switch initiates phase three of the commit process by sending a EUFC request to all managed Switches. The data received from the prior ESFC request is committed by the managed Switch according to the requirements of the application. When all the managed Switches successfully complete the EUFC, then the Managing Switch enters phase four of the commit process.

##### **13.3.1.5 Phase four**

The Managing Switch initiates phase four of the commit process by sending an ERCA request to all managed Switches. This causes the Switches to release the resources reserved by the relinquishes EACA and the commit process to conclude.



### 13.3.2 Handling Fabric changes

If a Switch that was specified in the ECS Switch List leaves the Fabric after the EACA has been accepted, it is removed from the ECS Switch List and only those that are currently in the ECS Switch List are involved in the commit process. After the first ESFC is sent, the Managing Switch will continue the commit process regardless of changes to the Fabric. If a new Switch joins the Fabric during the commit process it would not participate in the commit process for this transaction. All operations relating to the application shall be held off until the commit process successfully completes or is aborted.

NOTE 27 – In certain cases, there is a slight chance that isolation of Switches in the Fabric during ECS may result in inconsistencies between Switches in the Fabric. These inconsistencies could prevent the Fabric from merging and user intervention may be required before the Fabric is able to merge.

### 13.3.3 Error recovery

#### 13.3.3.1 Overview

When the ECS protocol operates in autonomous mode, a mechanism is provided to select a new Managing Switch without interaction with the fabric application. Only Switches in the ECS Switch List are authorized to participate in the ECS recovery processing.

#### 13.3.3.2 Managing Switch not functional

When the EACA is accepted by each managed Switch, each authorized Switch in the ECS Switch List begins monitoring for a domain unreachable condition associated with the Managing Switch. When an authorized managed Switch determines that the Managing Switch is unreachable, a new Managing Switch shall be selected.

##### 13.3.3.2.1 Dead Man timer

When each authorized managed Switch determines that the Managing Switch is unreachable and a transaction has been initiated (i.e., EACA received), a Dead Man timer is started that is based on the following equation:

Timer Value = 30sec \* 'Switch position in list'.

The first authorized Switch in the ECS Switch List following the Managing Switch shall assume a Switch position of zero. Other authorized Switches after that assume positions in increasing order from zero. The Dead Man timer in each authorized Switch is set according to its position in the ECS Switch List. The Dead Man timer in each authorized Switch is kept alive until an ECS request is received from a Switch different than the previous Managing Switch with the same Transaction\_ID.

##### 13.3.3.2.2 Basic procedure

Typically, the first authorized Switch in the ECS Switch List following the Managing Switch assumes the role of the new Managing Switch. However, if that Switch is also non-operational then the next authorized Switch will assume the role of Managing Switch and so forth. This will occur when the timer expires on a given Switch and no ECS requests are received with the same Transaction\_ID.

When the new Switch assumes the role of Managing Switch, it shall either issue an ERCA, or replay the current ECS phase with all remaining Switches specified in the ECS Switch List. When all Switches complete the current phase under the control of the new Managing Switch, then the Managing Switch proceeds to the next phase.

In the case where the new Managing Switch determines that the process is in the EACA phase or the ERCA phase, then the new Managing Switch shall send an ERCA to all remaining Switches in the ECS Switch List.

### **13.3.3.3 Resolution of multiple Managing Switches**

#### **13.3.3.3.1 Two Managing Switches - same commit phase**

When two Switches have assumed the role of Managing Switch and they are in the same commit phase, the Managing Switch with the lowest Domain\_ID assumes the role of Managing Switch and the other Switch relinquishes its role as Managing Switch. This condition is detected when multiple ECS requests are received with the same Transaction\_ID.

#### **13.3.3.3.2 Two Managing Switches - different commit phases**

When two Switches have assumed the role of Managing Switch, and one Switch is in a higher phase than the other, the Managing Switch in the higher commit phase retains the role of Managing Switch. To accomplish this the Switch in the higher commit phase returns an SW\_RJT reason code of "Unable to perform command request and an SW\_RJT reason code explanation of "In Advanced Phase". When the Switch receives such a reject for an ECS request, it sends a TCO request to the Switch that rejected the ECS request for this phase. The Switch that receives the TCO then retains the role of the Managing Switch and continues with the commit process. The Managing Switch that originated the TCO relinquishes its role as a Managing Switch and becomes a managed Switch.

### 13.3.4 Ladder diagrams

#### 13.3.4.1 Normal case

Figure 39 shows the interactions for the successful case.

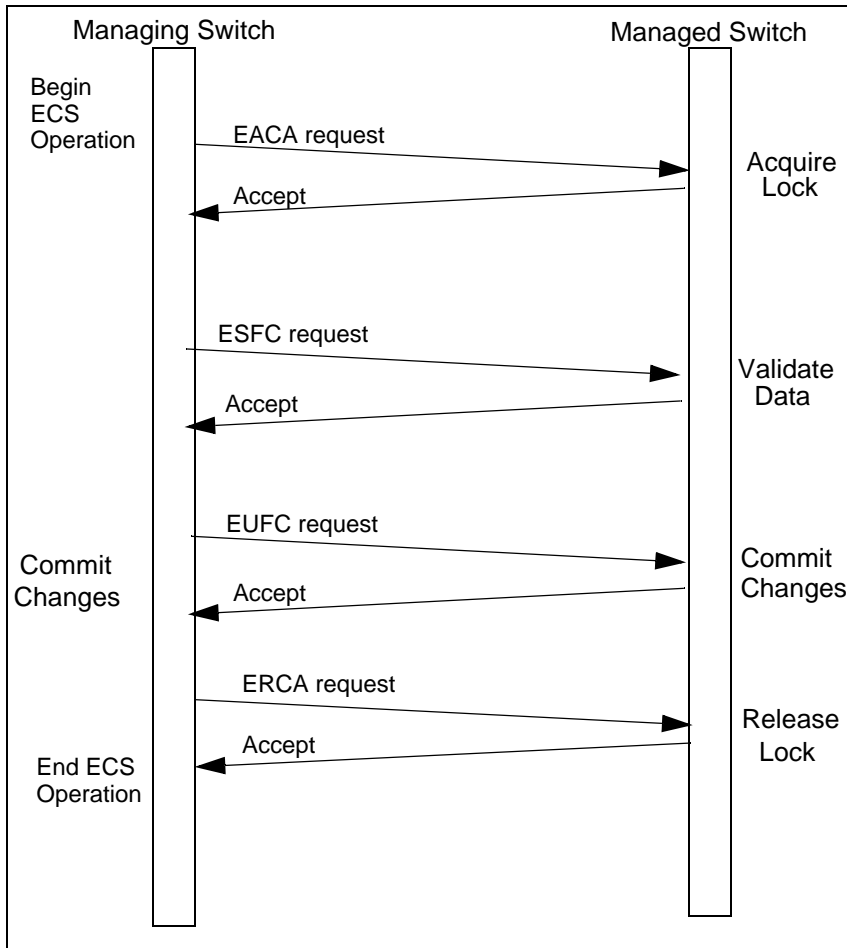



Figure 39 – Normal commit ladder diagram

### 13.3.4.2 Unsuccessful case

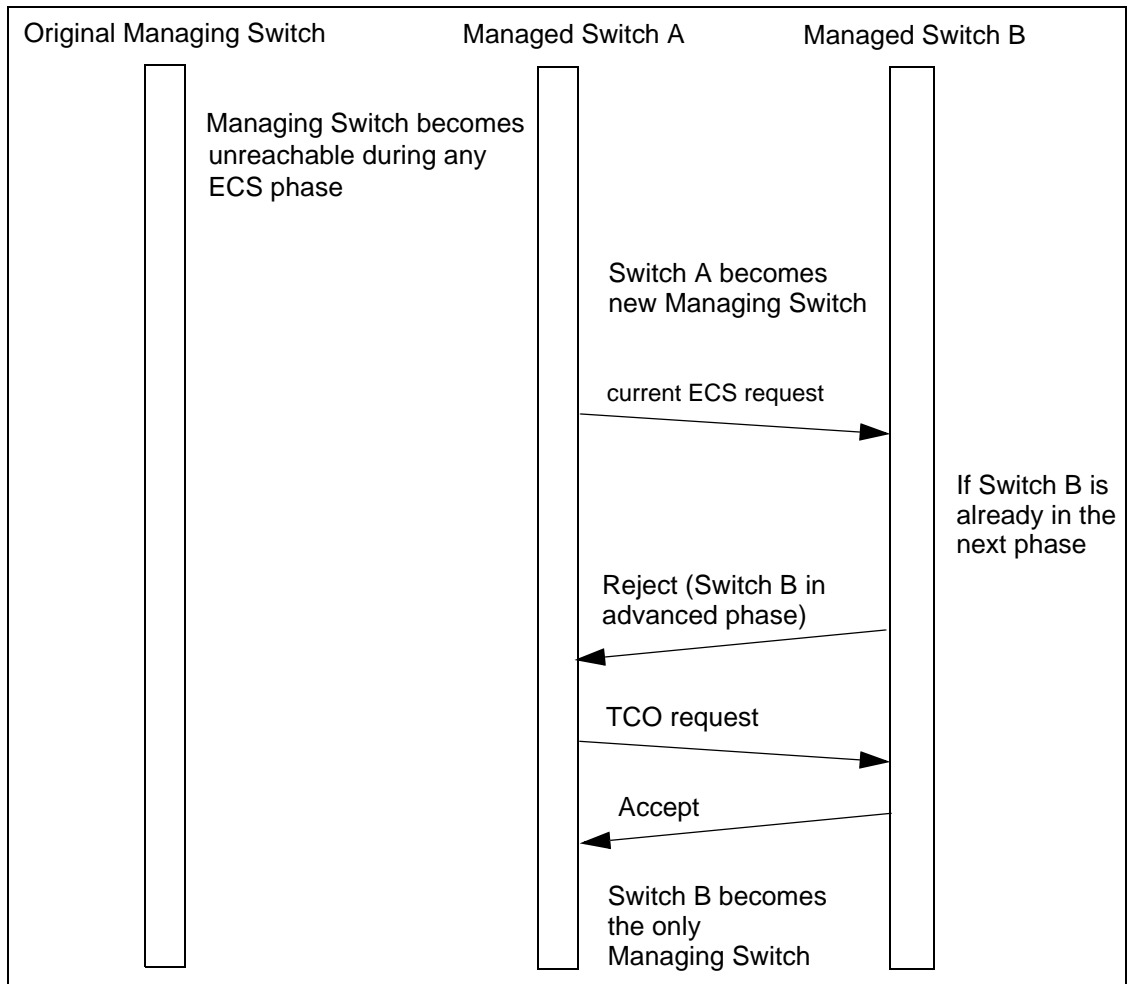
Figure 40 shows the interactions for the unsuccessful case.



**Figure 40 – Unsuccessful commit ladder diagram**

**13.3.4.3 Transfer ownership case - recovery processing enabled**

Figure 41 shows the interactions for the transfer ownership case.



**Figure 41 – Transfer ownership ladder diagram**

**13.3.5 State machines**


**13.3.5.1 Overview**

State machines are specified for both a Managing Switch and a managed Switch. State Machines are in the context of a specific service or application.

Figure 42 shows the state machine for a Managing Switch. Figure 43 shows the state machine for a managed Switch. Figure 44 shows the state machine for the transfer commit ownership case.

**13.3.5.2 States and transitions for the Managing Switch**

The states and transitions for the Managing Switch are shown in figure 42.



**Figure 42 – ECS Managing Switch state machine**

**State P0: ECS Begin.** This state marks the beginning of an ECS operation due to ECS being invoked. ECS may be invoked due to internal distribution requirements or an externally initiated Fabric Management Session. The Managing Switch determines whether the operation can be accepted or not.

**Transition P0:P1.** The Managing Switch is not busy with another ECS operation for the same service or application.

**Transition P0:P5.** The Managing Switch is busy with another ECS operation for the same service or application.

**State P1: EACA Phase.** The Managing Switch initiates EACAs to all Switches participating in the ECS operation and waits for responses.

**Transition P1:P1.** This transition occurs whenever a response to an EACA is received.

**Transition P1:P2.** This transition occurs when responses to the EACAs have been received from all Switches and all the responses indicated that the EACAs were accepted.

**Transition P1:P4.** This transition occurs if any of the responses to the EACAs indicated that the EACA was rejected.

**State P2: ESFC Phase.** The Managing Switch initiates ESFCs to all Switches participating in the ECS operation and waits for responses.

**Transition P2:P2.** This transition occurs whenever a response to an ESFC is received.

**Transition P2:P3.** This transition occurs when responses to the ESFCs have been received from all Switches and all the responses indicated that the ESFCs were accepted.

**Transition P2:P4.** This transition occurs if any of the responses to the ESFCs indicated that the ESFC was rejected.

**State P3: EUFC Phase.** The Managing Switch initiates EUFCs to all Switches participating in the ECS operation and waits for responses.

**Transition P3:P3.** This transition occurs whenever a response to an EUFC is received.

**Transition P3:P4.** This transition occurs when responses to the EUFCs have been received from all Switches and there are no other ESFC requests to be performed as part of the ECS operation.

**State P4: ERCA Phase.** The Managing Switch initiates ERCAs to all Switches participating in the ECS operation and waits for responses.

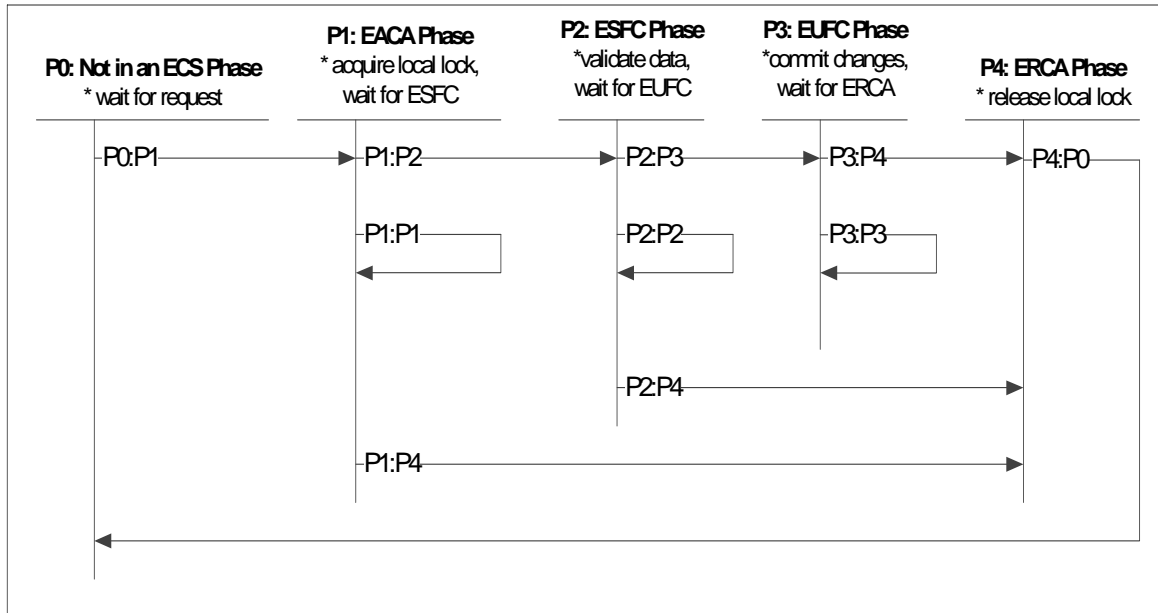
**Transition P4:P4.** This transition occurs whenever a response to an ERCA is received.

**Transition P4:P5.** This transition occurs when responses to the ERCAs have been received from all Switches.

**State P5: ECS End.** This state marks the end of an ECS operation. The Managing Switch reports any errors that occurred during the ECS operation to the requesting service or application.

### 13.3.5.3 States and transitions for the managed Switch

The states and transitions for the managed Switch are specified in figure 43.



**Figure 43 – ECS managed Switch state machine**

**State P0: Not in an ECS Phase.** The Switch is not a managed Switch for a given service or application. The Switch is waiting for an ECS request for a given service or application.

**Transition P0:P1.** This transition occurs whenever an EACA for a given service or application is received. It becomes a managed Switch with respect to the specified transaction.

**State P1: EACA Phase.** The managed Switch acquires the appropriate lock and waits for the ESFC request.

**Transition P1:P1.** This transition occurs when an ECS request other than an ESFC or ERCA is received (see table 281).

**Table 281 – EACA phase - events and actions**

Event	Action
1. EACA Received for same transaction	EACA accepted. If the request is from a new Managing Switch, the identity of the new Managing Switch is noted.
2. EUFC Received	Reject EUFC - Reason "Invalid phase transition within transaction"
3. EACA Received from a Switch that is not Authorized	Reject EACA - "Switch Not Authorized"



**Transition P1:P2.** This transition occurs when an ESFC for the given transaction is received and the data is valid.

**Transition P1:P4.** This transition occurs when an ERCA for the current transaction is received.

**State P2: ESFC Phase.** The managed Switch validates the received data in the ESFC and waits for the EUFC request.

**Transition P2:P2.** This transition occurs when an ECS request other than an EUFC or ERCA is received (see table 282).

**Table 282 – ESFC phase - events and actions**

Event	Action
1. ESFC Received for same transaction	ESFC accepted. If the request is from a new Managing Switch, the identity of the new Managing Switch is noted.
2. EACA Received for same transaction	EACA Rejected - Reason "In Advanced Phase"
3. ESFC Received from a Switch that is not Authorized	Reject ESFC - "Switch Not Authorized"

**Transition P2:P3.** This transition occurs when an EUFC for the given transaction is received.

**Transition P2:P4.** This transition occurs when an ERCA for the current transaction is received.

**State P3: EUFC Phase.** The managed Switch commits the application data and waits for the ERCA request.

**Transition P3:P3.** This transition occurs when an ECS request other than an ERCA is received (see table 283).

**Table 283 – EUFC phase - events and actions**

Event	Action
1. EUFC Received for same transaction	EUFC accepted. If the request is from a new Managing Switch, the identity of the new Managing Switch is noted.
2. ESFC Received for same transaction	ESFC Rejected - Reason "In Advanced Phase"
3. EACA Received for same transaction	EACA Rejected - Reason "In Advanced Phase"
4. EUFC Received from a Switch that is not authorized.	Reject EUFC - "Switch Not Authorized"

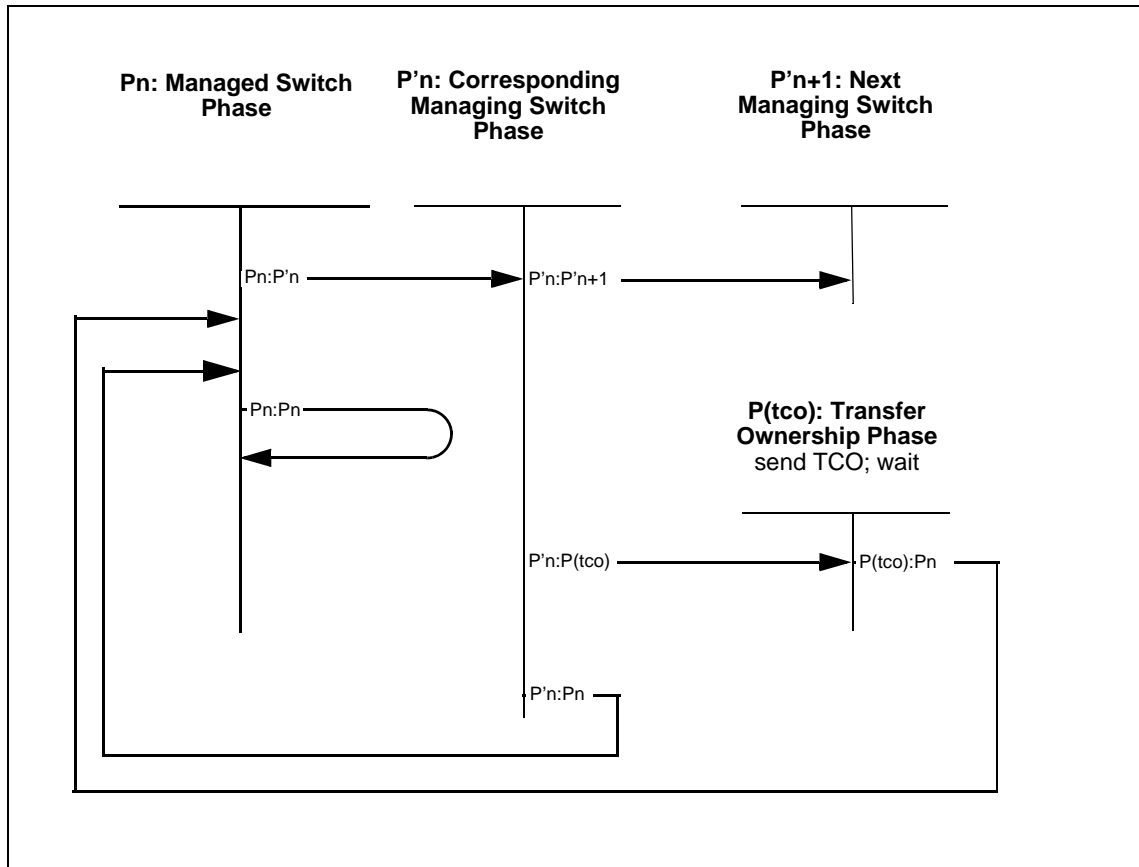
**Transition P3:P4.** This transition occurs when an ERCA for the current transaction is received.

**State P4: ERCA Phase.** The managed Switch releases the appropriate lock and concludes the ECS operation. Current transaction ends.

**Transition P4: P0.** This transition concludes ECS operation processing.

### 13.3.5.4 States and transitions for transfer commit ownership

The state machine for the transfer commit ownership case is shown in Figure 44.



**Figure 44 – ECS transfer commit ownership (TCO) state machine**

All managed Switches currently involved in the ECS operation monitor for domain unreachable for the Managing Switch. If domain unreachable is detected for the Managing Switch then the detecting managed Switch starts a deadman timer with the appropriate value based on its own position in the ECS operation’s Switch List. Managed Switches keep track of the current Managing Switch by examining the source of ECS requests.

**State Pn: Managed Switch Phase.** This is the current state of the managed Switch. This can be P1: EACA Phase, P2: ESFC Phase, or P3: EUFC Phase.

**Transition Pn:P’n.** This transition occurs whenever the deadman timer expires, or a TCO is received due to an “In Advanced Phase” response to an ECS request. The managed Switch assumes the role of Managing Switch and transitions to the corresponding Managing Switch phase.

**Transition Pn:Pn.** This transition occurs whenever an ECS request that is older than the current request is received. SW\_RJT reason code explanation “In Advanced Phase” is returned.

**State P'n: Corresponding Managing Switch Phase.** This is the Managing Switch state that corresponds to the managed Switch state that the new Managing Switch was in. The new Managing Switch initiates ECS requests corresponding to that phase and waits for responses.

**Transition P'n:P'n+1.** This is the normal transition that occurs when responses to the ECS requests have been received from all Switches and all the responses indicated that the ECS requests were accepted.

**Transition P'n:Pn.** This transition occurs if an ECS request for this phase is received from a Switch with a lower Domain\_ID. The Managing Switch goes back to its role as managed Switch.

**Transition P'n:P(tc).** This transition occurs if any of the responses to the ECS requests indicated "In Advanced Phase."

**State P'n+1: Next Managing Switch Phase.** This is the next state that the Managing Switch transitions to if all ECS requests are accepted.

**State P(tc): Transfer Commit Ownership Phase.** The Managing Switch initiates TCO to Switch that returned "In Advanced Phase" to the ECS operation and waits for response.

**Transition P(tc):Pn.** This transition occurs if the TCO was accepted by the managed Switch that indicated it was "In Advanced Phase." The Managing Switch goes back to its role as managed Switch.

## 14 Virtual Channels for Switched Fabric

### 14.1 Overview

The Virtual Channel (VC) architecture enables different data flows to be identified that in turn allows them to be differentiated and subjected to different service policies. Virtual Channels are allocated to group traffic based on destination, control traffic, and intra-fabric traffic. Virtual Channels operate on individual ISLs and are not required to operate across the entire fabric.

### 14.2 Assignment of Virtual Channels

#### 14.2.1 Overview of assignment

Frames are assigned to Virtual Channels based on the assignment scheme and the number of virtual channels for that assignment scheme as determined during link initialization. The assignment schemes are described in 14.2.2, 14.2.3, and 14.2.4.

#### 14.2.2 Simple

In the Simple assignment scheme, frames are assigned to Virtual Channels as specified in table 284. Class F frames shall always be assigned to VC 0 and all other classes of frames shall be assigned to VC 1.

**Table 284 – Simple assignment scheme**

Virtual Channel Number	Description
0	Class F frames
1	All other frames

#### 14.2.3 Fixed

In the Fixed assignment scheme, Class F frames shall be assigned to VC 0 and broadcast frames shall be assigned to the highest VC as determined during link initialization. Frames other than Class F and broadcast shall be assigned to Virtual Channels based on D\_ID as specified in table 285.

**Table 285 – Fixed assignment scheme**

Virtual Channel Number	Description
0	Class F frames
1	Reserved
2 - (n-3)	Based on D_ID (see table 286)
n-2	Reserved
n-1	Broadcast frames

VC assignments for the Fixed assignment scheme is specified in table 286.

**Table 286 – VC assignments - Fixed**

Number of VCs (n)	VC mapped to D_ID bits
8	D_ID(9:8)
12	D_ID(10:8)
20	D_ID(11:8)
36	D_ID(12:8)
68	D_ID(13:8)
132	D_ID(14:8)

If Virtual Channels are supported, implementation of the Fixed assignment scheme shall be mandatory when eight or greater Virtual Channels are available.

#### 14.2.4 Variable

In the Variable assignment scheme, frames are allocated to VCs according to the D\_ID. In this scheme Class F and broadcast frames are not allocated to dedicated Virtual Channels. Class F and broadcast frames are transported over Virtual Channels according to D\_ID and are given priority over other frames within each Virtual Channel. The Variable assignment scheme is specified in table 287.

**Table 287 – Variable assignment scheme**

Virtual Channel Number	Description (see table 288)
0	Based on D_ID
1	Based on D_ID
...	
n-2	Based on D_ID
n-1	Based on D_ID

VC assignments for the Variable assignment scheme is specified in table 288.

**Table 288 – VC assignments - Variable**

Number of VCs (n)	VC mapped to D_ID bits
4	D_ID(9:8)
8	D_ID(10:8)
16	D_ID(11:8)
32	D_ID(12:8)
64	D_ID(13:8)
128	D_ID(14:8)
256	D_ID(15:8)

If Virtual Channels are supported and less than eight Virtual Channels are available, implementation of the Variable assignment scheme shall be mandatory.

### 14.3 VC parameter negotiation

#### 14.3.1 Agreement of assignment schemes

If VC\_RDY flow control mode is specified, then the originator of the ELP and the recipient of the ELP shall agree on the assignment scheme. If the recipient of the ELP does not support the assignment scheme specified in the ELP by the originator, then the recipient rejects the ELP request with a reason code explanation of "Invalid Flow Control Parameters". The originator of the ELP may then request another assignment scheme defined for VC\_RDY flow control mode. In the event that an assignment scheme cannot be agreed upon, the originator shall request that R\_RDY flow control mode be used.

#### 14.3.2 Negotiation of number of VCs

Once the assignment scheme has been agreed upon, the recipient checks the number of VCs specified by the VC value. If the recipient of the ELP supports a VC value less than that specified by the originator, it shall transmit the lower value that it supports in the SW\_ACC to the ELP. The originator will then use this value representing the number of VCs and their assignment for operation. The originator may use the VC\_Credit assigned to the valid VCs in the original ELP, or it may originate a new ELP with the reduced number of VCs and a different VC\_Credit distribution.

The recipient of the ELP shall not send an SW\_ACC indicating a higher VC Value than what was advertised in the originator's ELP.

In the event that the originator and responder cannot agree upon the number of VCs, the originator shall request that R\_RDY flow control be used.

**14.4 Credit management**


**14.4.1 Overview**

Each VC maintains a set of VC Buffer credits. The credit for each VC is set to the value exchanged in the ELP SW\_ILS at the completion of Link Initialization for the ISL. The VC Credit shall be decremented for every frame transmitted that is associated with that VC. VC Credit shall be incremented for each VC\_RDY received for that VC. Frames shall not be transmitted for a VC for which there is no VC Credit even if there is Buffer-to-buffer Credit. Frames shall not be transmitted on any VC if there is not Buffer-to-buffer Credit for that ISL. If a frame is received on a Virtual Channel with no buffer-to-buffer Credit, the frame shall be discarded and no VC\_RDY shall be returned.

After negotiating VC\_RDYs in ELP and following the link reset, the VC\_RDYs shall be sent and R\_RDYs shall not be sent.

**14.4.2 VC\_RDY Primitive Signals**

VC\_RDY primitive signals are used for buffer-to-buffer flow control on ISLs that support Virtual Channels. The general format is shown in figure 45.



**Figure 45 – VC\_RDY Primitive Signal format**

Table 289 provides the VC\_ID values used in the VC\_RDY Primitive Signals for each virtual channel number.

**Table 289 – VC\_ID values for VC\_RDY Primitive Signals**

Virtual Channel Number	VC_ID Value
00h	D0.0
01h	D1.0
02h	D2.0
...	see FC-FS-5
FDh	D29.7
FEh	D30.7
FFh	D31.7

## 15 Inter-Fabric Routing support

### 15.1 F\_RJT and F\_BSY processing for Class 2/F

#### 15.1.1 Overview

Switches may support either or both of:

- a) Inter-Fabric Routers operating according to the FC-IFR simple mode of operation (see FC-IFR); or
- b) Inter-Fabric Routers operating according to the FC-IFR NAT mode of operation (see FC-IFR).

When a Switch supporting Inter-Fabric Routers operating according to the FC-IFR simple mode of operation (see FC-IFR) determines a Class 2/F F\_RJT or Class 2/F F\_BSY needs to be generated in response to a received frame, the Switch shall:

- a) generate an encapsulated Class 2/F F\_RJT or Class 2/F F\_BSY as specified in 15.1.2 if the received frame contains an Enc\_Header; or
- b) generate a Class 2/F F\_RJT or Class 2/F F\_BSY as specified in FC-FS-5 if the received frame does not contain an Enc\_Header.

When a Switch not supporting Inter-Fabric Routers operating according to the FC-IFR simple mode of operation (see FC-IFR) determines a Class 2/F F\_RJT or Class 2/F F\_BSY needs to be generated in response to a received frame, the Switch shall generate a Class 2/F F\_RJT or Class 2/F F\_BSY as specified in FC-FS-5.

#### 15.1.2 Encapsulated Class 2 F\_RJT or Class 2 F\_BSY frame format

##### 15.1.2.1 Overview

An encapsulated Class 2/F F\_RJT or Class 2/F F\_BSY frame shall be formatted as specified in figure 46.

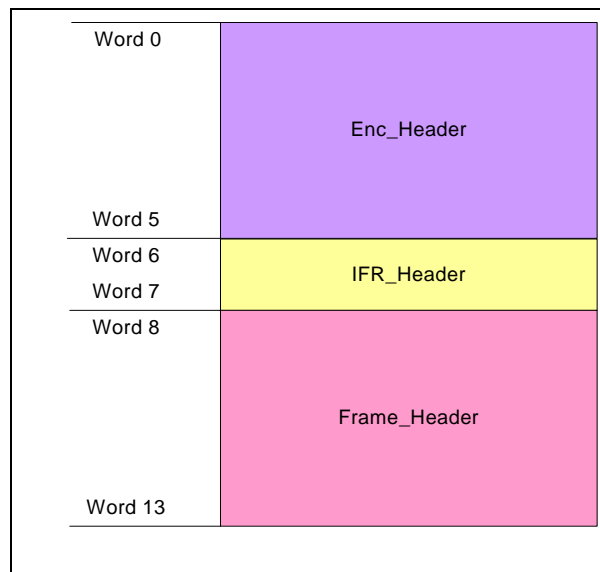


Figure 46 – Encapsulated Class 2/F F\_RJT and Class 2/F F\_BSY frame format



**15.1.2.2 Encapsulated Enc\_Header field values**

The Enc\_Header field values for an encapsulated Class 2/F F\_RJT or Class 2/F F\_BSY frame shall be set as specified in table 290.

**Table 290 – Encapsulated Class 2/F F\_RJT/F\_BSY Enc\_Header field values**

<b>Field</b>	<b>Value</b>
R_CTL	52h
D_ID	S_ID field in received Enc_Header
CS_CTL	CS_CTL field in received Enc_Header
S_ID	D_ID field in received Enc_Header
TYPE	TYPE field from Frame_Header (see 15.1.2.4)
F_CTL	F_CTL field from Frame_Header (see 15.1.2.4)
SEQ_ID	SEQ_ID field in received Enc_Header
DF_CTL	DF_CTL field from Frame_Header (see 15.1.2.4)
SEQ_CNT	SEQ_CNT field in received Enc_Header
OX_ID	OX_ID field in received Enc_Header
RX_ID	RX_ID field in received Enc_Header
Parameter	Parameter field from Frame_Header (see 15.1.2.4)

**15.1.2.3 Encapsulated IFR\_Header field values**

The IFR\_Header field values for an encapsulated Class 2/F F\_RJT or Class 2/F F\_BSY frame shall be set as specified in table 291.

**Table 291 – Encapsulated Class 2/F F\_RJT/F\_BSY IFR\_Header field values**

Field	Value
R_CTL	51h
DF_ID	SF_ID field in received IFR_Header
Exp_Time	a) zero; or b) valid expiration time if port is IFR capable
SF_ID	DF_ID field in received IFR_Header
Ver	Ver field in received IFR_Header
Pri	Pri field in received IFR_Header
ETV	a) zero; or b) one if valid expiration time
HCV	a) zero; or b) one if valid Hop_Cnt
Hop_Cnt	a) zero; or b) valid Hop_Cnt if port is IFR capable

**15.1.2.4 Encapsulated Frame\_Header field values**

The Frame\_Header field values for an encapsulated Class 2/F F\_RJT or Class 2/F F\_BSY frame shall be set as specified in FC-FS-5 and this standard.

## 16 Timers and constants

### 16.1 General timers and constants

General timers and constants referenced in FC-SW-7 are specified in table 292.

**Table 292 – FC-SW-7 timers and constants**

Timer/Constant	Value	Description
F_S_TOV	5 Seconds	Fabric Stability timeout value. Ensures that Fabric stability has been achieved during various aspects of Fabric Configuration
D_S_TOV	5 Seconds	Distributed Service timeout value. A value that indicates the maximum time that a Distributed Service requestor shall wait for a response.
Min_LS_Arrival	1 seconds	The minimum amount of time that shall pass before a Switch shall accept updates of any given LSR via flooding.
Min_LS_Interval	5 seconds	The minimum amount of time that shall pass before a Switch is allowed to send an LSR update via flooding.
Check_Age	5 minutes	The minimum amount of time between verification checks of LSRs in a Switch's database.
Max_Age_Diff	15 minutes	If the age of two instances of the same LSR differ by more than this amount, they are considered to come from different incarnations. The LSR with the smaller age field is considered to be more current.
LS_Refresh_Time	30 minutes	The maximum interval between transmission of refresh LSRs.
Max_Age	1 hour	The maximum age that an LSR may reach. When an LSR reaches this age, it is removed from the database.
Initial_Message_Number	80000001h	The initial value for the LSR Incarnation field.
Max_Message_Number	7FFFFFFFh	The maximum value for the LSR Incarnation field.
Rxmt_Interval	5 seconds	The maximum time period for which an LSR may go unacknowledged. If an LSR is not acknowledged within this time period, it shall be retransmitted.
Hello_Interval	20 seconds	The minimum interval between HELLOs sent by a Switch on a link to verify link health.
Dead_Interval	80 seconds	The maximum interval for which no HELLO may be received on a link. If no HELLO is received on a link after this time period the link is considered broken and thus removed from the database.

**16.2 SW\_ILS timeout values**

SW\_ILS timeout values and the recommended actions for them are specified in table 293.

**Table 293 – SW\_ILS timeout values**

<b>SW_ILS</b>	<b>Timeout (SW_ACC or SW_RJT)</b>	<b>Recommended action when timeout expires</b>
ELP	E_D_TOV + 4 secs	Go to state P11
ESC	E_D_TOV + 4 secs	Go to state P11
BF	F_S_TOV + E_D_TOV	Restart BF
RCF	F_S_TOV + E_D_TOV	Restart RCF
EFP	2*F_S_TOV	Restart BF or Retransmit
DIA	F_S_TOV	Restart BF or Retransmit
RDI	R_A_TOV	Retransmit
MR	R_A_TOV + 70 secs	Retransmit
MRRA	R_A_TOV	Retransmit
ACA	R_A_TOV	Retransmit
RCA	R_A_TOV	Retransmit
SFC	R_A_TOV + 20 secs	Retransmit
UFC	R_A_TOV + 20 secs	Retransmit
EACA	R_A_TOV	Retransmit
ERCA	R_A_TOV	Retransmit
ESFC	R_A_TOV + 20 secs	Retransmit
EUFC	R_A_TOV + 20 secs	Retransmit
TCO	R_A_TOV	Retransmit
SW_RSCN	R_A_TOV	Retransmit
DRLIR	R_A_TOV	Retransmit
ESS	R_A_TOV	Retransmit
CEC	R_A_TOV	Retransmit
EVFP	2*R_A_TOV	Retransmit
STR	R_A_TOV	Retransmit

## 17 Distributed Switch

### 17.1 Overview

A Distributed Switch is a set of FCDFs associated with at least one Controlling Switch that controls the operations of the set of FCDFs. Figure 47 shows an example of a Distributed Switch composed of a Controlling Switch and two FCDFs.

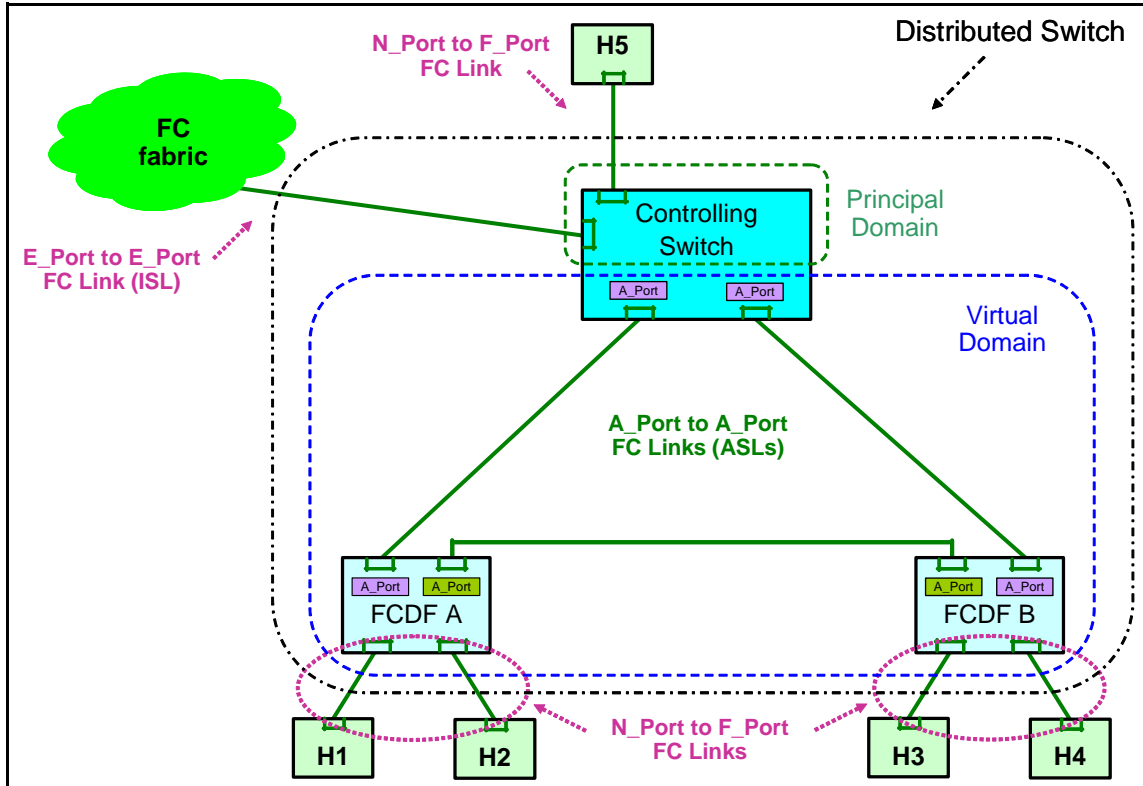


Figure 47 – Distributed Switch example - one Controlling Switch and two FCDFs


From an external point of view (i.e., outside the dotted and dashed black line in figure 47), a Distributed Switch behaves as a Fibre Channel Switch. In particular, a Distributed Switch supports the instantiation of N\_Port to F\_Port links and of E\_Port to E\_Port links (ISLs). N\_Port to F\_Port links are supported by both FCDFs and Controlling Switches, while ISLs are supported only by Controlling Switches. This means that it is possible to connect a Distributed Switch to another Switch only through a Controlling Switch, not through an FCDF.

From an internal point of view (i.e., inside the dotted and dashed black line in figure 47), A\_Port to A\_Port links (ASLs) enable FC frame forwarding between the Controlling Switch and FCDFs, as well as between FCDFs. ASLs are also used to exchange control information between a Controlling Switch and FCDFs.

The Controlling Switch uses one Virtual Domain\_ID to perform N\_Port\_ID allocations for N\_Ports connected to the FCDFs of the Distributed Switch (i.e., a Virtual Domain\_ID is used as the most significant byte in the N\_Port\_IDs allocated to N\_Ports that are attached to an FCDF of the Distributed Switch). The Controlling Switch also uses another Domain\_ID, called a Principal Domain\_ID, for its normal functions as a Fibre Channel Switch, including allocation of N\_Port\_IDs for directly attached N\_Ports. As a result, a Distributed Switch as the one shown in figure 47 uses two

Domain\_IDs: one for the Principal Domain and one for the Virtual Domain. To properly support the operations of a Virtual Domain, a Controlling Switch shall have at least one unique Switch\_Name to associate with the Virtual Domain, in addition to its own unique Switch\_Name.

For a Distributed Switch as the one shown in figure 47 the Controlling Switch generates the FSPF LSR describing the Virtual Domain in the Distributed Switch. In addition, the Controlling Switch lists the Virtual Domain as a directly attached Domain in its FSPF LSR. The resulting FSPF topology is shown in figure 48, where Z is the Domain\_ID belonging to the Virtual Domain and P is the Domain\_ID of the Principal Domain of the Controlling Switch.



**Figure 48 – FSPF Topology of figure 47**

FCDFs are not able to operate without a Controlling Switch, therefore the Controlling Switch is a single point of failure in a Distributed Switch configuration with only one Controlling Switch, as the one shown in figure 47. To avoid this issue, Distributed Switches may support a redundant configuration consisting of two Controlling Switches, with one Controlling Switch elected as a Primary Controlling Switch, and one Controlling Switch elected as a Secondary Controlling Switch. The Primary Controlling Switch controls the operation of the entire Distributed Switch, including the allocation of N\_Port\_IDs to N\_Ports connected to the FCDFs of the Distributed Switch. The Primary Controlling Switch and Secondary Switch keep their state synchronized to allow the Secondary Controlling Switch to become the Primary Controlling Switch if the Primary Controlling Switch fails as specified by the Controlling Switch redundancy protocol (see 17.5).

An example of a redundant Distributed Switch with two Controlling Switches and two FCDFs is shown in figure 49.

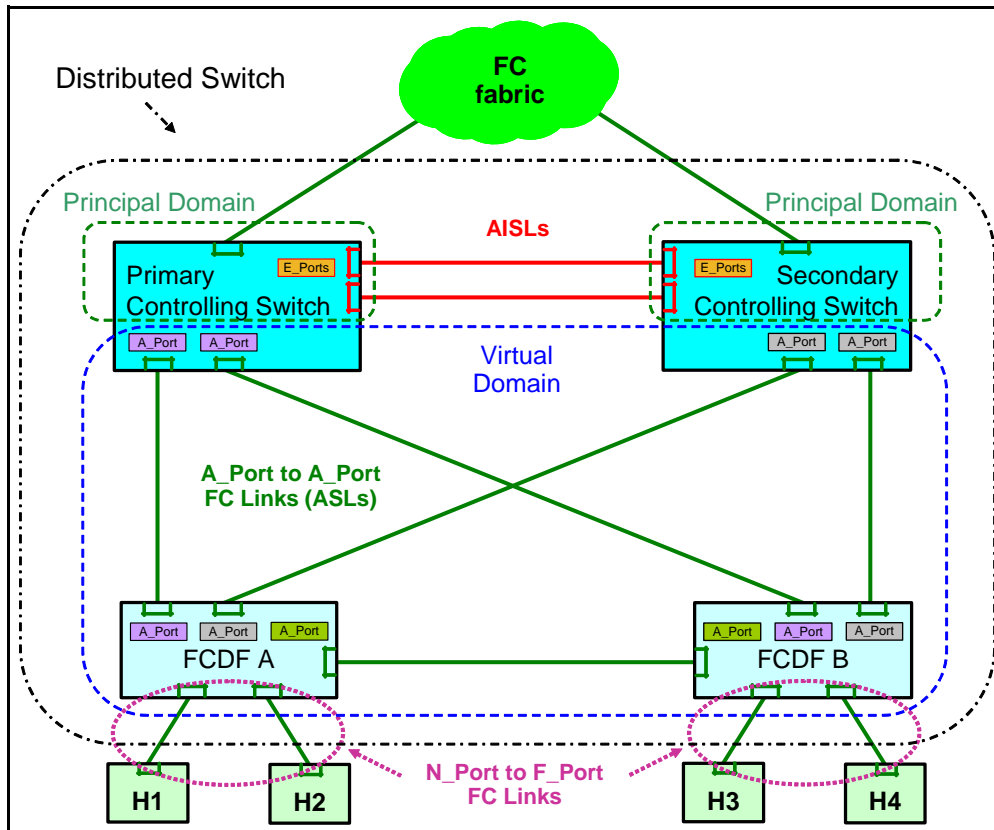


Figure 49 – Distributed Switch example - two Controlling Switches - two FCDFs

The Primary Controlling Switch and the Secondary Controlling Switch in a redundant Distributed Switch should be connected by at least two links to be used as Augmented ISLs (AISLs), where the term 'augmented' indicates that link is used for the Controlling Switch redundancy protocol, in addition to normal E\_Port operation.

The Primary Controlling Switch uses one Virtual Domain\_ID to perform N\_Port\_ID allocations for N\_Ports connected to the FCDFs of the Distributed Switch (i.e., a Virtual Domain\_ID is used as the most significant byte in the N\_Port\_IDs allocated to N\_Ports that are attached to an FCDF of the Distributed Switch). Using a Virtual Domain\_ID to assign N\_Port\_IDs enables continued operation in case of failures of the Primary Controlling Switch. Each Controlling Switch also uses another Domain\_ID, called a Principal Domain\_ID, for its normal functions as a Fibre Channel Switch, including allocation of N\_Port\_IDs to directly attached N\_Ports. As a result, a redundant Distributed Switch as the one shown in figure 49 uses three Domain\_IDs: one for each Principal Domain and one for the Virtual Domain. To properly support the operations of a Virtual Domain, a Controlling Switch shall have at least one unique Switch\_Name to associate with the Virtual Domain, in addition to its own unique Switch\_Name.

The Controlling Switches instantiate ASLs to enable the forwarding of FC frames and the communication of control information between Controlling Switches and FCDFs. In a redundant configuration, FCDFs instantiate ASLs to each of the Controlling Switches and between themselves.

A Distributed Switch may have a cascaded FCDF configuration. An example of a redundant Distributed Switch with two Controlling and four cascaded FCDFs is shown in figure 50.

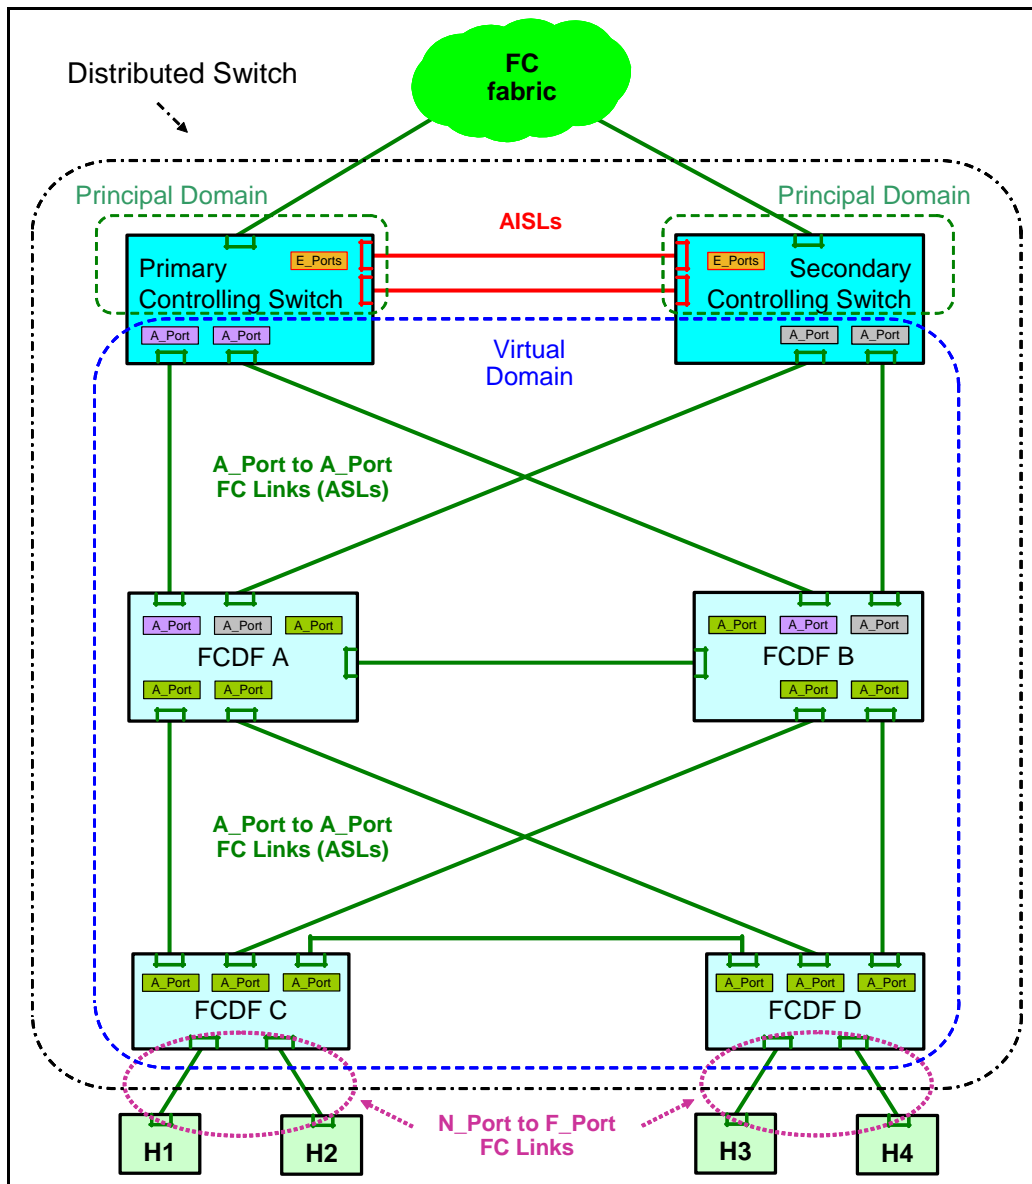



Figure 50 – Distributed Switch example - two Controlling Switches - four cascaded FCDFs

For Distributed Switches as the ones shown in figure 49 and in figure 50 the Primary Controlling Switch generates the FSPF LSR describing the Virtual Domain in the Distributed Switch. In addition, both Primary and Secondary Controlling Switch list the Virtual Domain as a directly attached Domain in their FSPF LSR. The resulting FSPF topology is shown in figure 51, where Z is the Domain\_ID



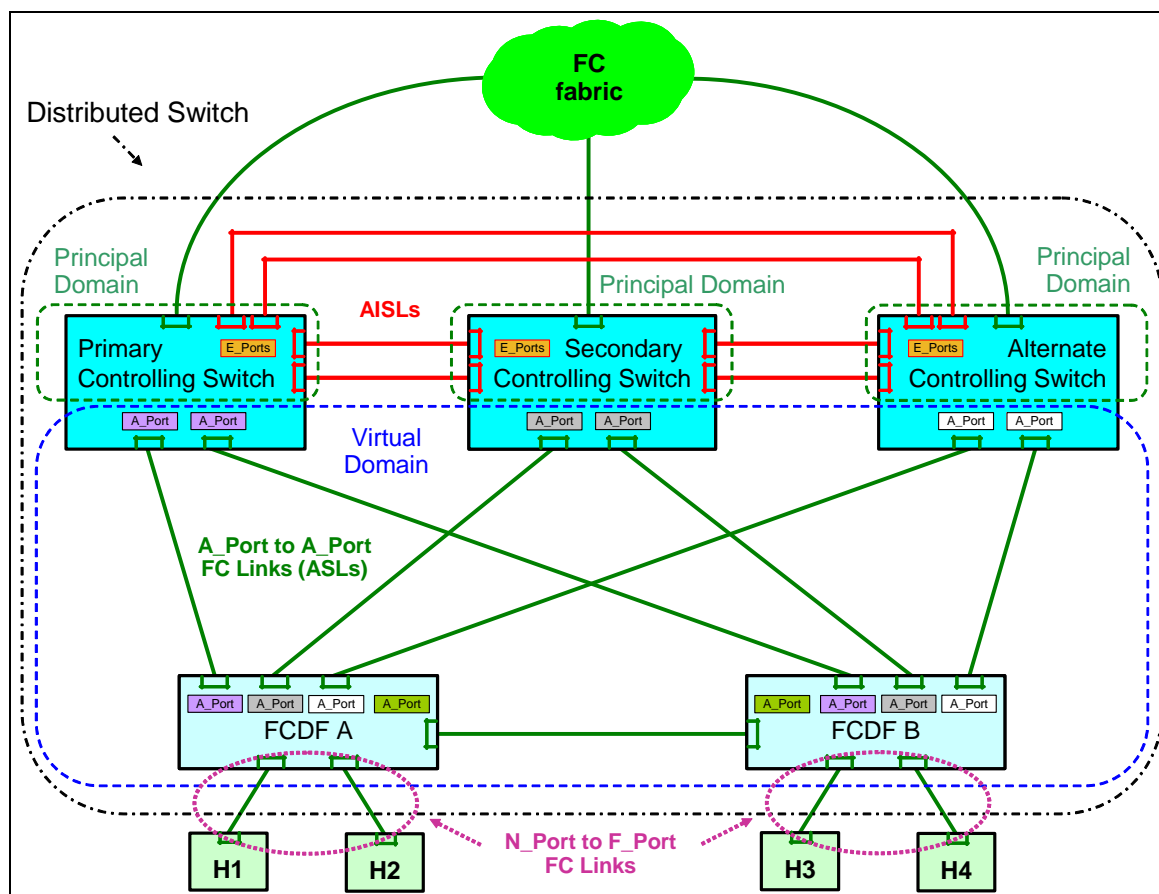
belonging to the Virtual Domain and P and S are the Domain\_IDs of the Principal Domains of the two Controlling Switches. P and S are also connected between themselves by virtue of the AISLs.



**Figure 51 – FSPF topology of figure 49 and figure 50**

To further increase the availability of a Distributed Switch, Distributed Switches may optionally support a redundant configuration of more than two Controlling Switches, a Primary Controlling Switch, a Secondary Controlling Switch, and one or more Alternate Controlling Switches. The Primary Controlling Switch controls the operation of the entire Distributed Switch, including the allocation of N\_Port\_IDs to N\_Ports connected to the FCDFs of the Distributed Switch. The Primary Controlling Switch and Secondary Switch keep their state synchronized to allow the Secondary Controlling Switch to become the Primary Controlling Switch if the Primary Controlling Switch fails as specified by the Controlling Switch redundancy protocol (see 17.5). Alternate Controlling Switches are able to take the place of the Secondary Controlling Switch in case of its failure as specified by the Controlling Switch redundancy protocol (see 17.5).

An example of a redundant Distributed Switch with three Controlling Switches and two FCDFs is shown in figure 52.



**Figure 52 – Distributed Switch example - three Controlling Switches - two FCDFs**

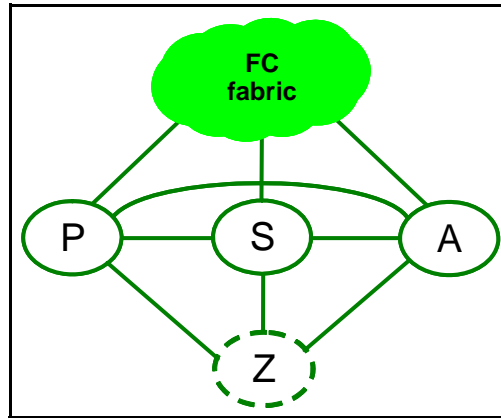
For proper operation of a redundant Distributed Switch with more than two Controlling Switches, the set of Controlling Switches shall be a connected set (i.e., between any pair of Controlling Switches there shall be a path composed of only Controlling Switches) and it is recommended for these Controlling Switches to be connected in a fully meshed topology.

The Controlling Switches in a redundant Distributed Switch with more than two Controlling Switches should be connected by at least two links to be used as Augmented ISLs (AISLs), where the term 'augmented' indicates that link is used for the Controlling Switch redundancy protocol, in addition to normal E\_Port operation.

The Primary Controlling Switch uses one Virtual Domain\_ID to perform N\_Port\_ID allocations for N\_Ports connected to the FCDFs of the Distributed Switch (i.e., a Virtual Domain\_ID is used as the most significant byte in the N\_Port\_IDs allocated to N\_Ports that are attached to an FCDF of the Distributed Switch). Using a Virtual Domain\_ID to assign N\_Port\_IDs enables continued operation in case of failures of the Primary Controlling Switch. Each Controlling Switch also uses another Domain\_ID, called a Principal Domain\_ID, for its normal functions as a Fibre Channel Switch, including allocation of N\_Port\_IDs to directly attached N\_Ports. As a result, a redundant Distributed Switch such as the one shown in figure 52 uses four Domain\_IDs, one for each Principal Domain and one for the Virtual Domain. To properly support the operations of a Virtual Domain, a Controlling

Switch shall have at least one unique Switch\_Name to associate with the Virtual Domain, in addition to its own unique Switch\_Name.

For a Distributed Switch as the one shown in figure 52 the Primary Controlling Switch generates the FSPF LSR describing the Virtual Domain in the Distributed Switch. In addition, all Controlling Switches list the Virtual Domain as a directly attached Domain in their FSPF LSR. The resulting FSPF topology is shown in figure 53, where Z is the Domain\_ID belonging to the Virtual Domain and P, S, and A are respectively the Domain\_IDs of the Principal Domains of the Primary Controlling Switch, Secondary Controlling Switch, and Alternate Controlling Switch. P, S, and A are also connected between themselves by virtue of the AISLs.



**Figure 53 – FSPF topology of figure 52**

For proper operation of a redundant Distributed Switch it is recommended that the FCDFs directly connected to at least one Controlling Switch be directly connected to all Controlling Switches.

A Controlling Switch is uniquely identified by its Switch\_Name Name\_Identifier. An FCDF is uniquely identified by its Switch\_Name Name\_Identifier. A Distributed Switch is defined by an administrative configuration on the Controlling Switches, listing:

- a) the Switch\_Names of the Controlling Switches that are part of that Distributed Switch (i.e., the Controlling Switch Set); and
- b) the Switch\_Names of the FCDFs that are part of that Distributed Switch (i.e., the FCDF\_Set).

## 17.2 FCDF handling of well-known addresses

N\_Ports use well-known addresses (WKAs) and Domain Controller address identifiers to exchange information with the Fabric, either through ELSs or through the Common Transport protocol.

An FCDF supports VF\_Ports, therefore it shall terminate FC frames destined to the F\_Port Controller WKA. This implies local processing by the FCDF of the FLOGI, FDISC, LOGO, and RLS ELSs (see FC-LS-3).

The handling of other WKAs and Domain Controllers address identifiers is performed by the Primary Controlling Switch, therefore an FCDF shall forward all FC frames having as D\_ID the address identifiers listed in table 294 to the Primary Controlling Switch through a VA\_Port. The NPRD

SW\_ILS (see 6.3.3.5) provides to FCDFs the routing information needed to reach the Primary Controlling Switch.

**Table 294 – Forwarded Domain Controller and well-known address identifiers**

Address Value	Description
FFFC01h .. FFFCFEh	Domain Controller address identifiers
FFFFFF4h	Event Service WKA
FFFFFF6h	Clock Synchronization Service WKA
FFFFFF7h	Security Key Distribution Service WKA
FFFFFFAh	Management Service WKA
FFFFFFBh	Time Service WKA
FFFFFFCh	Directory Service WKA
FFFFFFDh	Fabric Controller WKA

The AISLs between Controlling Switches are used as paths to reach the Primary Controlling Switch if an FCDF is connected to the Secondary Controlling Switch or to an Alternate Controlling Switch but not to the Primary Controlling Switch. In order to do so, the Secondary Controlling Switch and the Alternate Controlling Switches shall forward to the Primary Controlling Switch over the appropriate AISLs any FC frame destined to the address identifiers shown in table 294 if they are received from

an ASL or from an AISL. An example configuration and associated WKA frame processing for an FCDF with no link to the Primary Controlling Switch is shown in figure 54.

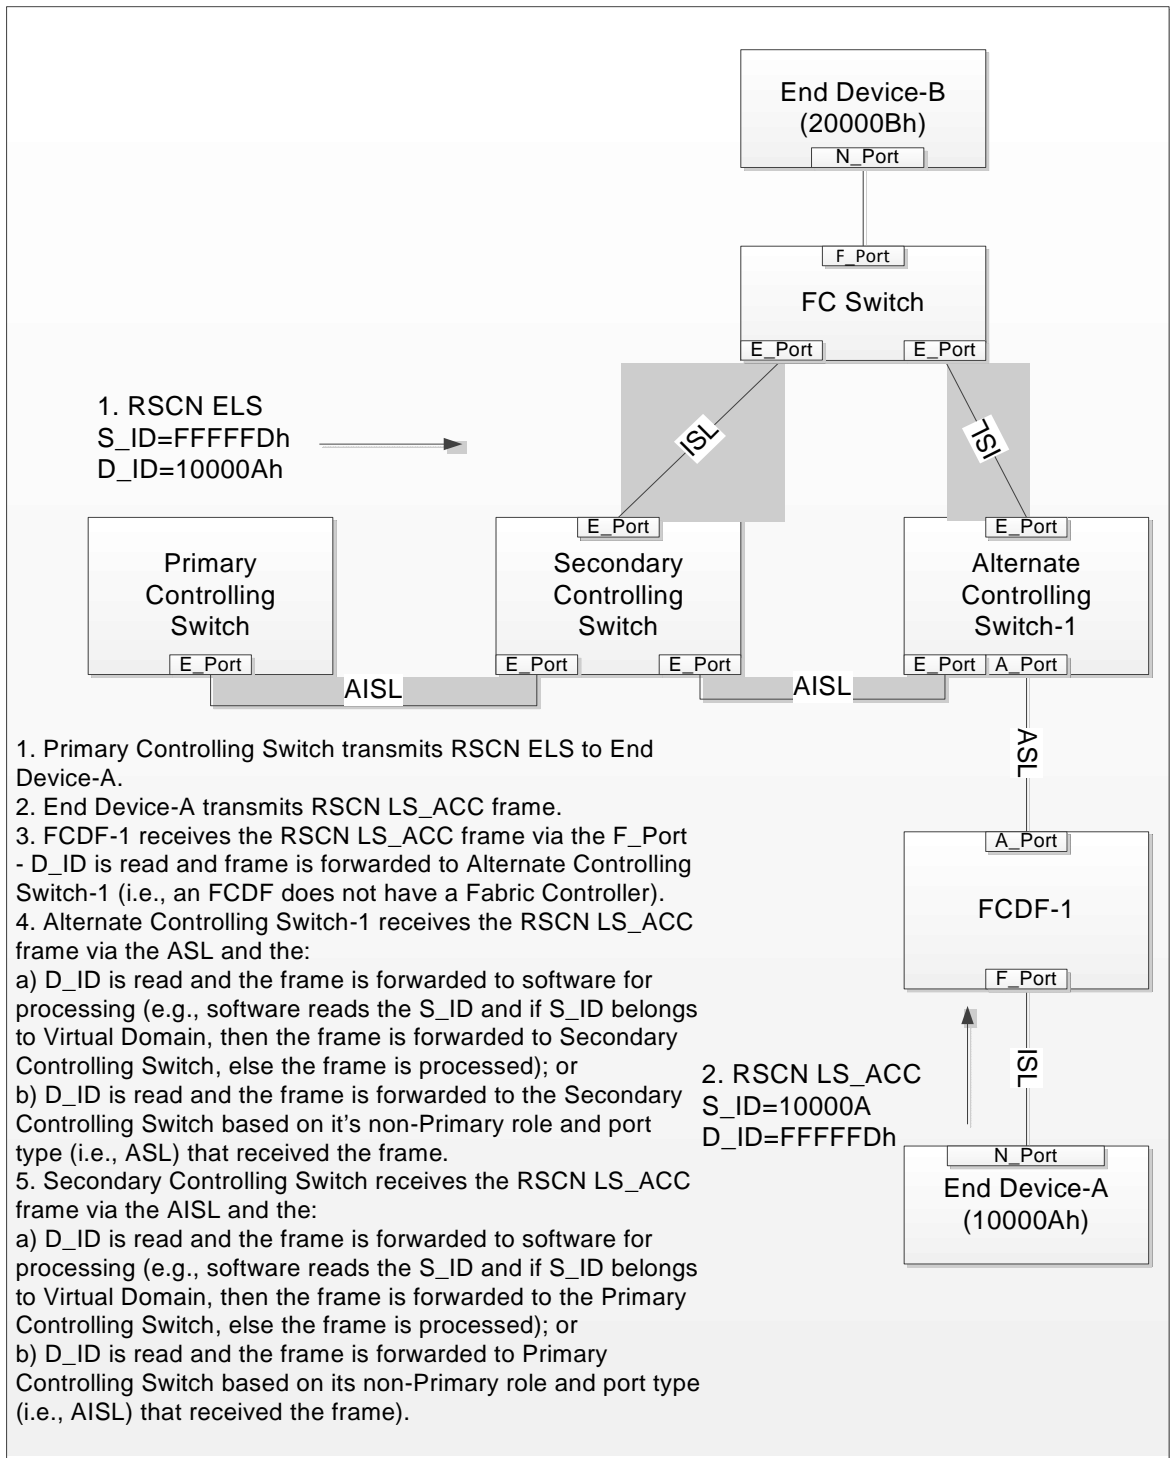


Figure 54 – FCDF with no link to Primary Controlling Switch WKA frame processing example

### 17.3 A\_Port to A\_Port links (ASLs)

An ASL becomes operational on successful completion of an ELP Exchange between a Controlling Switch and an FCDF or between two FCDFs. Bits 13 (i.e., the Controlling FCF/Switch bit) and 12 (i.e., the FDF/FCDF bit) in the Flags field of the ELP payload indicate if the originator of the ELP request or SW\_ACC is a Controlling Switch or an FCDF (see 6.2.4).

A received ELP request having both these bits set to one is invalid, shall be rejected, and the link shall be Isolated. A received SW\_ACC having both these bits set to one is invalid and the link shall be Isolated. Table 295 specifies the meaning of the values of these bits.

**Table 295 – VA\_Port ELP Flag bits**

Bit 13 value	Bit 12 value	Description
0b	0b	The originator of the ELP request or SW_ACC is a FC Switch or FCF that is not a Controlling Switch or Controlling FCF.
0b	1b	The originator of the ELP request or SW_ACC is an FCDF or an FDF.
1b	0b	The originator of the ELP request or SW_ACC is a Controlling Switch or a Controlling FCF.
1b	1b	Invalid combination

A port of a Controlling Switch shall transmit an ELP request after completing Link Initialization. This ELP request shall have the Controlling FCF/Switch bit set to one and the FDF/FCDF bit set to zero. The Controlling Switch port behaves as an E\_Port (i.e., as ISL is instantiated) if the ELP is accepted by the neighbor and:

- a) the received ELP SW\_ACC has both the Controlling FCF/Switch bit and the FDF/FCDF bit set to zero; or
- b) the received ELP SW\_ACC has the Controlling FCF/Switch bit set to one and the neighbor Switch is not part of the Controlling Switch Set of this Distributed Switch.

If the ELP is accepted and the received ELP SW\_ACC has the Controlling FCF/Switch bit set to one and the neighbor Switch is part of the Controlling Switch Set of this Distributed Switch, then the Controlling Switch port behaves as an Augmented E\_Port for this Distributed Switch (i.e., an AISL is instantiated and used by the Controlling Switch redundancy protocol (see 17.5)).

If the ELP is accepted and the received ELP SW\_ACC has the FDF/FCDF flag set to one and the neighbor FCDF is part of the FCDF\_Set of this Distributed Switch, and the Controlling Switch has established its role (i.e., when in state P2, S2, or A2 of the Controlling Switch redundancy protocol (see 17.5)), then the Controlling Switch port behaves as an A\_Port (i.e., an ASL is instantiated).

If the ELP is accepted and the received ELP SW\_ACC has the FDF/FCDF bit set to one and the neighbor FCDF is part of the FCDF\_Set of this Distributed Switch, and the Controlling Switch has not yet established its role, then the Controlling Switch port shall transition to the Isolated state.

If the ELP is accepted and the received ELP SW\_ACC has the FDF/FCDF bit set to one and the neighbor FCDF is not part of the FCDF\_Set of this Distributed Switch, then the Controlling Switch port shall transition to the Isolated state.

A Controlling Switch port shall reject a received ELP request with the FDF/FCDF bit set to one with reason code 'Protocol Error' and reason code explanation 'Invalid Request'.

A Controlling Switch port shall reply to a received ELP request with the FDF/FCDF bit set to zero according to the normal ELP rules (i.e., acceptance or rejection includes considering the involved Switch\_Names). The Controlling Switch port behaves as an E\_Port (i.e., an ISL is instantiated) if the ELP request is accepted and:

- a) the received ELP request has both the Controlling FCF/Switch flag and the FDF/FCDF bit set to zero; or
- b) the received ELP request has the Controlling FCF/Switch bit set to one and the neighbor Switch is not part of the Controlling Switch Set of this Distributed Switch.

If the ELP is accepted and the received ELP request has the Controlling FCF/Switch bit set to one and the neighbor Switch is part of the Controlling Switch Set of this Distributed Switch, then the Controlling Switch port behaves as an Augmented E\_Port for this Distributed Switch (i.e., an AISL is instantiated and used for the Controlling Switch redundancy protocol (see 17.5)).

If an FCDF has not received from the Primary Controlling Switch the Distributed Switch's FCDF\_Set and Controlling Switch Set via a DSMD SW\_ILS (see 6.3.3.8), then its ports shall not transmit an ELP request after completing Link Initialization.

After having received the Distributed Switch's FCDF\_Set and Controlling Switch Set via a DSMD SW\_ILS (see 6.3.3.8), and routing information via an NPRD SW\_ILS (see 6.3.3.5), FCDF ports that have completed Link Initialization, except the port that received the DSMD request, shall transmit an ELP request with the FDF/FCDF bit set to one.

On receiving a valid ELP request with the Controlling FCF/Switch bit set to one or the FDF/FCDF bit set to one, the FCDF port shall process it irrespective of the value of the Switch\_Name field in the ELP request payload (i.e., acceptance or rejection shall be based on the other ELP parameters, not on the involved Switch\_Names). If the ELP is accepted then the FCDF port behaves as an A\_Port (i.e., an ASL is instantiated).

NOTE 28 – These rules enable an ordered establishments of ASLs from the Controlling Switch(es) to the peripheral FCDFs in a Distributed Switch with cascaded FCDFs.

An FCDF does not support E\_Ports, therefore an FCDF port shall reject a received ELP request with both the Controlling FCF/Switch bit and FDF/FCDF bit set to zero (i.e., an ELP request coming from a Switch that is not a Controlling Switch or an FCDF) with reason code 'Protocol Error' and reason code explanation 'Invalid Request'. After an FCDF has received the Distributed Switch's FCDF\_Set and Controlling Switch Set from the Primary Controlling Switch, that FCDF shall reject received ELP requests coming from a Controlling Switch other than the Controlling Switches in the Controlling Switch Set, with reason code 'Logical Error' and reason code explanation 'Not Authorized'.

Upon instantiating an ASL with another FCDF, an FCDF may receive a DSMD request over that ASL. If the received DSMD request contains a Virtual Domain Switch\_Name different than its Virtual Domain Switch\_Name, then the DSMD request shall be rejected and the ASL Isolated.

## 17.4 Distributed Switch operations

### 17.4.1 Overview

In a Distributed Switch, the Primary Controlling Switch defines the routes for the FCDF topology and performs N\_Port\_ID allocations and deallocations for all its controlled FCDFs. In a redundant Distributed Switch with two Controlling Switches or with more than two Controlling Switches, the Primary Controlling Switch and the Secondary Controlling Switch keep their state synchronized.

## 17.4.2 FCDF routing

A Controlling Switch establishes its role when in state P2, S2, or A2 of the Controlling Switch redundancy protocol (see 17.5). When a Controlling Switch established its role, it instantiates ASLs with the FCDFs that are directly connected and are part of its FCDF\_Set.

Upon instantiating an ASL with an FCDF, the Primary Controlling Switch shall initiate an FCRN Exchange (see 6.3.3.3) describing that link with the Secondary Controlling Switch, if available, to keep the state synchronized. Upon completion of this FCRN Exchange, the Primary Controlling Switch shall provide to that FCDF the Distributed Switch Membership information through a DSMD Exchange (see 6.3.3.8).

The Primary Controlling Switch shall recompute the N\_Port\_ID routes and distribute them to each FCDF and Alternate Controlling Switch, if any, belonging to the Distributed Switch through NPRD Exchanges. At this point the instantiated ASL becomes part of the Distributed Switch internal topology (i.e., the set of ASLs internal to the Distributed Switch).

Upon de-instantiating an ASL with an FCDF, the Primary Controlling Switch shall initiate an FCUN Exchange (see 6.3.3.4) describing that deinstantiated link with the Secondary Controlling Switch, if available, to keep the state synchronized. Upon completion of this FCUN Exchange, the Primary Controlling Switch shall recompute the N\_Port\_ID routes and distribute them to each FCDF and Alternate Controlling Switch, if any, belonging to the Distributed Switch through NPRD Exchanges.

When becoming operational, an FCDF waits for a Controlling Switch or another FCDF to initiate an ELP Exchange with it, in order to set up a ASL. Upon completing the DSMD and NPRD Exchanges with the Primary Controlling Switch, the FCDF becomes able to initiate ELP requests to instantiate other ASLs with other FCDFs. At this point the FCDF enables its ports for logins from Nodes. Any FLOGI received on a FCDF port before this point is responded to by the FCDF with a LS\_RJT having reason code 'Logical busy' and reason code explanation 'No additional explanation'.

Upon instantiating an ASL with an FCDF, an FCDF, the Secondary Controlling Switch, or an Alternate Controlling Switch, if any, shall perform an FCRN Exchange with the Primary Controlling Switch to inform it of the new link. If the FCDF reported in the FCRN request is not known to the Primary Controlling Switch, then upon completing this FCRN Exchange the Primary Controlling Switch shall provide to the newly reported FCDF the Distributed Switch Membership information through a DSMD Exchange. At this point the instantiated ASL becomes part of the Distributed Switch internal topology (i.e., the set of ASLs internal to the Distributed Switch). Upon completion of this DSMD Exchange, the Primary Controlling Switch shall recompute the N\_Port\_ID routes and distribute them to each FCDF and Alternate Controlling Switch, if any, belonging to the Distributed Switch through NPRD Exchanges.

NOTE 29 – This behavior enables an FCDF connected only to the Secondary Controlling Switch or to an Alternate Controlling Switch to become part of the Distributed Switch.

Upon de-instantiating an ASL with an FCDF, an FCDF, the Secondary Controlling Switch, or an Alternate Controlling Switch, if any, shall perform an FCUN Exchange with the Primary Controlling Switch to inform it of the deinstantiated link. Upon completion of this FCUN Exchange, the Primary Controlling Switch shall recompute the N\_Port\_ID routes and distribute them to each FCDF and Alternate Controlling Switch, if any, belonging to the Distributed Switch through NPRD Exchanges.

The distribution of NPRD requests shall take precedence over the distribution of AZAD (see 6.3.3.7) and NPZD (see 6.3.3.6) requests.



### 17.4.3 N\_Port\_ID handling

Upon receiving on a port a FLOGI request or a NPIV FDISC request from a Node, an FCDF shall send a VNRN request (see 6.3.3.1) to the Primary Controlling Switch to inform it of the newly reachable VN\_Port. If the Primary Controlling Switch rejects the VNRN request, the FCDF shall also reject the FLOGI request or NPIV FDISC request with reason code 'Unable to Perform Command Request' and reason code explanation 'No Additional Explanation'. Upon receiving the VNRN request, the Primary Controlling Switch performs the following processing:

- a) if the VNRN request carried a FLOGI request and that VN\_Port was not already logged in or if the VNRN request carried a NPIV FDISC request, then the Primary Controlling Switch shall allocate to the newly reachable VN\_Port an N\_Port\_ID from the Virtual Domain\_ID; or
- b) if the VNRN request carried a FLOGI request and that VN\_Port was already logged in, then the Primary Controlling Switch shall implicitly log out that VN\_Port and all the VN\_Ports associated to the VF\_Port that VN\_Port was associated with and then allocate to that VN\_Port an N\_Port\_ID from the Virtual Domain\_ID.

The Primary Controlling Switch shall also recompute the Zoning ACLs for the affected N\_Port\_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling Switch shall send NPZD requests indicating N\_Port\_ID allocation to the Secondary Controlling Switch, if present, to the Alternate Controlling Switches, if any, and to each FCDF belonging to the Distributed Switch (see 6.3.3.6). The Primary Controlling Switch shall wait to receive the NPZD SW\_ACC from the Secondary Controlling Switch, if present, and the NPZD SW\_ACC from the FCDF that sent the VNRN request before sending the VNRN SW\_ACC containing the FLOGI/NPIV FDISC LS\_ACC parameters to the FCDF that sent the VNRN request. If administratively configured to do so, the Primary Controlling Switch shall wait to receive the NPZD SW\_ACC from the Secondary Controlling Switch, if present, the NPZD SW\_ACC from all Alternate Controlling Switches, if any, and the NPZD SW\_ACC from all FCDFs before sending the VNRN SW\_ACC containing the FLOGI/NPIV FDISC LS\_ACC parameters to the FCDF that sent the VNRN request.

Upon receiving the VNRN SW\_ACC containing the FLOGI / NPIV FDISC LS\_ACC Parameters, the FCDF that sent the VNRN request shall accept the FLOGI request or NPIV FDISC request and complete the N\_Port Fabric login. If the FLOGI or NPIV FDISC Exchange that triggered the VNRN request has been terminated, then an FLOGI LS\_ACC or an NPIV FDISC LS\_ACC shall not be generated upon receiving the VNRN SW\_ACC. In this case, if upon receiving the VNRN SW\_ACC or upon termination of the VNRN Exchange there is no Fabric Login in progress or established for the VN\_Port that began the terminated FLOGI or NPIV FDISC Exchange, then the FCDF shall perform a VNUN Exchange with the Primary Controlling Switch to inform the Primary Controlling Switch that the VN\_Port is now unreachable.

If while performing the processing of a VNRN request the Primary Controlling Switch receives a second VNRN request for the same VN\_Port (i.e., the second VNRN request is received while the first VNRN Exchange is still open) and the processing of the second VNRN request results in NPZD requests having the same payloads as the ones generated for the first VNRN request, then the Controlling Switch may skip sending the second set of NPZD requests and reply to the second VNRN request once NPZD processing for the first VNRN request is considered completed.

When a VN\_Port is logged out or when a VF\_Port is de-instantiated, an FCDF shall perform a VNUN Exchange (see 6.3.3.2) with the Primary Controlling Switch to inform it that the VN\_Port is now unreachable or that all the VN\_Ports associated with that VF\_Port are unreachable. Upon completing a VNUN Exchange, the Primary Controlling Switch shall deallocate the N\_Port\_ID(s) assigned to the affected VN\_Port(s), recompute the Zoning ACLs for the affected N\_Port\_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling Switch shall send NPZD requests indicating N\_Port\_ID(s) deallocation to the Secondary Controlling Switch, if present,

to the Alternate Controlling Switches, if any, and to each FCDF belonging to the Distributed Switch (see 6.3.3.6).

The Primary Controlling Switch maintains a sequence number for each FCDF in the FCDF\_Set. The sequence number is incremented by one and included in the NPZD, NPRD, or AZAD sequence number descriptor each time an NPZD, NPRD, or AZAD request is sent.

Upon receipt of an NPZD, NPRD, or AZAD request, an FCDF compares the sequence number in the received sequence number descriptor to that of the last processed NPZD, NPRD, or AZAD request, or to 00000000 00000001h if none of these commands (i.e., NPZD, NPRD, AZAD) has previously been processed. If the received sequence number is lower, except in the case where a sequence number wrap condition has been detected, the received NPZD, NPRD, or AZAD request shall be discarded and a SW\_RJT shall be sent with reason code of 'Logical Error' and reason code explanation of 'Out of Order'. If the received sequence number is higher, or a wrap condition has been detected, then the NPZD, NPRD, or AZAD is processed.

An FCDF considers an N\_Port\_ID to be allocated when it has successfully received the N\_Port\_ID in an Allocation Entry of the current or previous NPZD request. If an NPZD request contains a peering entry with a Principal N\_Port\_ID that has not been allocated, that entire peering entry shall be ignored. If an NPZD request contains a peering entry with a Principal N\_Port\_ID that is currently allocated, but that peering entry contains Peer N\_Port\_ID(s) that have not been allocated, then those Peer N\_Port\_ID(s) shall be ignored.

Whenever an NPZD request is retransmitted for any reason (e.g., timeout) the Zoning ACLs for the affected N\_Port\_IDs shall be recomputed and a new NPZD request including a new sequence number and the newly computed peering entries shall be sent.

If a Primary Controlling Switch receives a SW\_RJT with a reason code of 'Logical Error' and reason code explanation of 'Out of Order' in response to an NPZD request, the Primary Controlling Switch shall retransmit the NPZD request.

When a new Zone Set is activated in the Fabric, the Primary Controlling Switch shall recompute the Zoning ACLs for all N\_Port\_IDs allocated in the Virtual Domain and distribute them to the FCDFs of the Distributed Switch through AZAD Exchanges.

If the Primary Controlling Switch has to send an AZAD request to an FCDF, any NPZD or NPRD requests outstanding to that FCDF shall first be completed. Any AZAD requests outstanding shall also be completed prior to initiating any subsequent NPZD or NPRD requests with that FCDF.

If the Primary Controlling Switch has to send an NPRD request to an FCDF, any NPZD or AZAD requests outstanding to that FCDF shall first be completed. Any NPRD requests outstanding shall also be completed prior to initiating any subsequent NPZD or AZAD requests with that FCDF.

Upon receiving on a port a FLOGI request or a NPIV FDISC request from a Node, a Controlling Switch shall allocate to the newly reachable VN\_Port an N\_Port\_ID from the Principal Domain\_ID if it accepts the received FLOGI or NPIV FDISC request.

If Zoning is enforced in the Fabric and a fabric login outside the Virtual Domain or a fabric logout outside the Virtual Domain results in the Zoning ACLs for some FCDFs to change, the Primary Controlling Switch shall send an NPZD request carrying the updated Zoning ACLs to each affected FCDF. If Zoning is not enforced in the Fabric, NPZD requests shall not be generated as a result of a fabric login outside the Virtual Domain or a fabric logout outside the Virtual Domain (see 6.3.3.6).

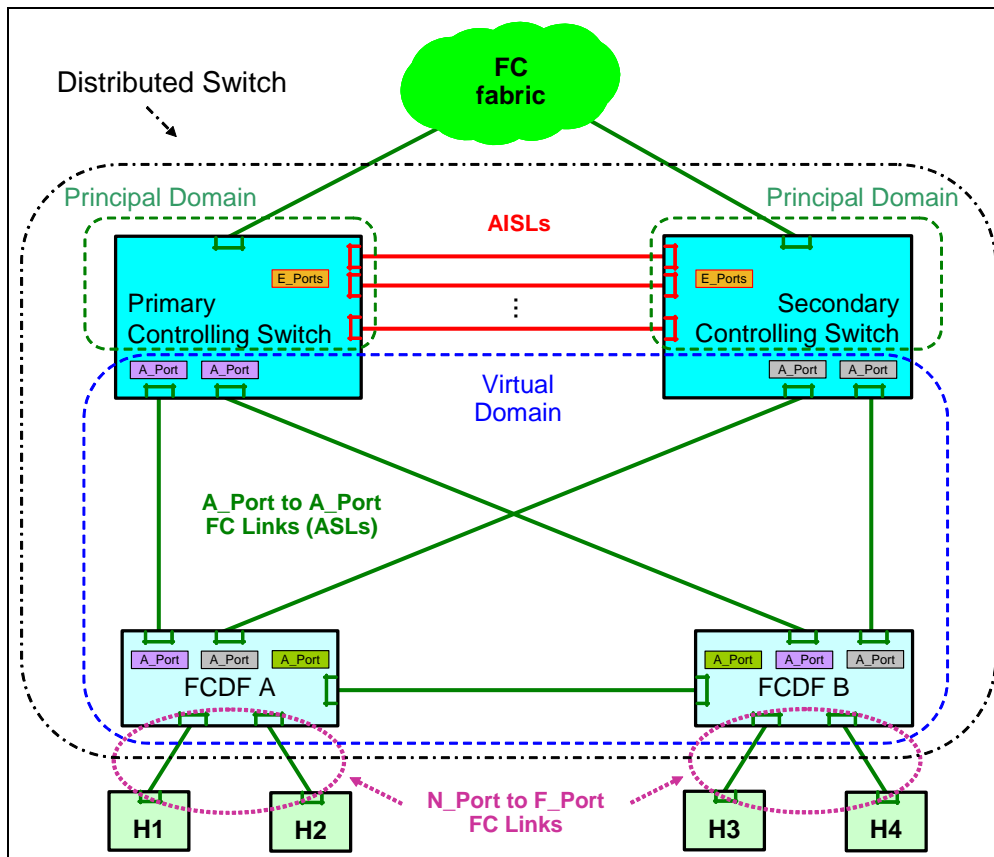
**17.4.4 Distribution tree**

The Primary Controlling Switch shall compute a distribution tree to distribute VA\_Port SW\_ILSs to the FCDFs. The distribution tree information is encoded in the NPRD request (i.e., only the first path listed in a N\_Port\_ID Reachability Entry is used to relay VA\_Port SW\_ILSs).

**17.5 Controlling Switch redundancy protocol**

**17.5.1 Controlling Switch redundancy protocol overview**

The purpose of the Controlling Switch redundancy protocol is to avoid any single point of failure in a Distributed Switch. Figure 55 shows an example of a redundant Distributed Switch with a Primary Controlling Switch and a Secondary Controlling Switch.



**Figure 55 – Redundant Distributed Switch example**

The Controlling Switch redundancy protocol uses Augmented E\_Port to E\_Port links (AISLs) between Controlling Switches for its operations. It is strongly recommended to deploy at least two AISLs between a pair of Controlling Switches, in order to distinguish the case of an AISL failure from the case of a Controlling Switch failure. Additional AISLs provide additional resiliency. The set of AISLs of a Controlling Switch is referred to as the AISL\_Set of that Controlling Switch.

**17.5.2 Controlling Switch redundancy protocol state machine**

The Controlling Switch redundancy protocol state machine reacts to AISLs failures in a timed fashion. To this end, the Controlling Switch redundancy protocol state machine uses indications from the

physical layer to determine if a link failed together with periodic Redundancy Hello messages (RHello) to verify the health of the Controlling Switches. The Controlling Switch redundancy protocol state machine uses the following time intervals and timers:

- a) **RHello\_Interval:** time interval between RHellos, expressed in milliseconds. The default value is 200 ms;
- b) **Down\_Interval:** time interval for the Primary Controlling Switch to declare the Secondary Controlling Switch down or for the Secondary Controlling Switch to declare the Primary Controlling Switch down. Calculated as  $2.5 * RHello\_Interval$ ; and
- c) **SPCS\_Timer:** timer that expires upon SPCS\_TOV and is used for selecting a Primary Controlling Switch if the Controlling Switch Set is composed of more than two Switch\_Names. The default value for SPCS\_TOV is 2000 ms.

To determine which Controlling Switch behaves as Primary and which one as Secondary, the Controlling Switch redundancy protocol uses a priority value associated to each Controlling Switch. Controlling Switch priority values are shown in table 296.

**Table 296 – Controlling Switch priority values**

Value	Description
00h	Reserved
01h	Highest priority value. This value is administratively configured to force the election of a Controlling Switch to Primary.
02h <sup>a</sup>	Primary Controlling Switch priority. This value is used by the Controlling Switch redundancy protocol to identify a Controlling Switch as Primary.
03 .. FEh	Higher to lower priority values. The default value is 128.
FFh <sup>a</sup>	This value indicates that a Controlling Switch is not willing to operate as Primary. This is used by the Primary Controlling Switch to trigger a transition of the Secondary Controlling Switch to Primary without having to wait for the current Primary to timeout, if appropriate.
<sup>a</sup> These values are used by the Controlling Switch redundancy protocol and not available to an administrator.	

Figure 56 shows the Controlling Switch redundancy protocol state machine.

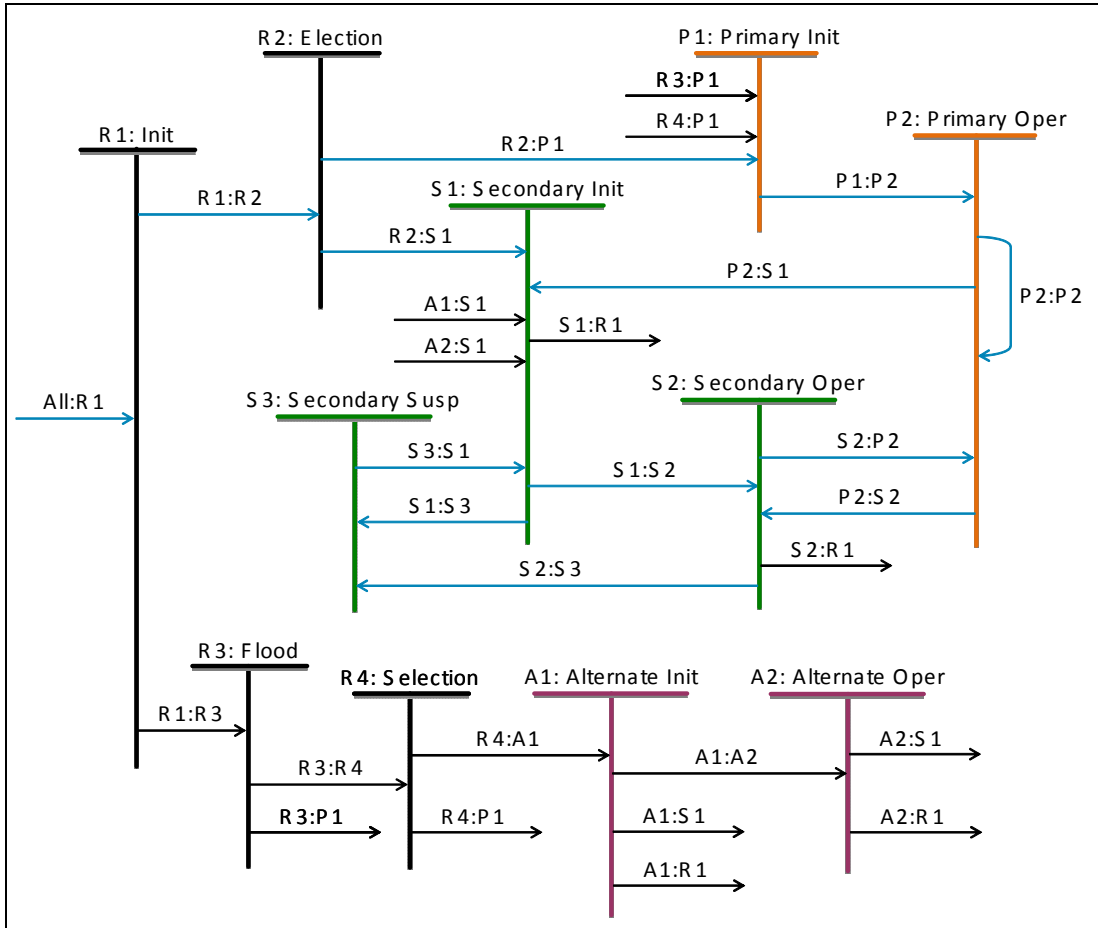


Figure 56 – Controlling Switch redundancy protocol state machine

**Transition All:R1.** Occurs when the Controlling Switch redundancy protocol is re-initialized.

**State R1:Init.** In this state a Controlling Switch clears its state and waits to begin the processing for the Controlling Switch redundancy protocol.

**Transition R1:R2.** Occurs when processing for the Controlling Switch redundancy protocol begins and the Controlling Switch Set is composed of two Switch\_Names. The Controlling Switch redundancy protocol processing begins when:

- a) the Controlling Switch redundancy protocol is enabled;
- b) the Controlling Switch Set and the FCDF\_Set are configured; and
- c) Fabric configuration is completed.

**Transition R1:R3.** Occurs when processing for the Controlling Switch redundancy protocol begins and the Controlling Switch Set is composed of more than two Switch\_Names. The Controlling Switch redundancy protocol processing begins when:

- a) the Controlling Switch redundancy protocol is enabled;
- b) the Controlling Switch Set and the FCDF\_Set are configured; and
- c) Fabric configuration is completed.

**State R2:Election.** A Controlling Switch enters this state if the Controlling Switch Set is composed of two Switch\_Names. In this state a Controlling Switch determines if it operates as Primary or Secondary. If its AISL\_Set is NULL, then the Controlling Switch exits this state. If its AISL\_Set is not NULL, then an ERP Exchange (see 6.4.3.1) is performed.

NOTE 30 – In this state the ERP payload does not contain N\_Port\_ID Ranges, nor a non-zero Virtual Domain Switch\_Name given that the Controlling Switch cleared its state in state R1.

If the ERP Exchange shows that the local Controlling Switch priority is 01h and the remote Controlling Switch priority is 01h (i.e., both Controlling Switches are manually configured to be Primary) then the Controlling Switch redundancy protocol is disabled and an error is logged.

**Transition R2:P1.** Occurs when:

- a) the Controlling Switch AISL\_Set is NULL;
- b) the Controlling Switch AISL\_Set is not NULL and the ERP Exchange showed that the local Controlling Switch priority is lower than the remote Controlling Switch priority; or
- c) the Controlling Switch AISL\_Set is not NULL, the ERP Exchange showed that the local Controlling Switch priority is equal to the remote Controlling Switch priority, and the local Switch\_Name is lower than the remote Switch\_Name.

**Transition R2:S1.** Occurs when:

- a) the Controlling Switch AISL\_Set is not NULL and the ERP Exchange showed that the local Controlling Switch priority is higher than the remote Controlling Switch priority; or
- b) the Controlling Switch AISL\_Set is not NULL, the ERP Exchange showed that the local Controlling Switch priority is equal to the remote Controlling Switch priority, and the local Switch\_Name is greater than the remote Switch\_Name.

**State P1:Primary Initialization.** In this state a Controlling Switch performs the operations to become the Primary Controlling Switch of the Distributed Switch. The Controlling Switch sets its priority to 02h, selects the Virtual Domain Switch\_Name (i.e., a Switch\_Name for the Virtual Domain), and obtains an additional Domain\_ID value (i.e., a Virtual Domain\_ID) from the Principal Switch of the fabric by generating an RDI request on behalf of the Virtual Domain Switch\_Name.

**Transition P1:P2.** Occurs when the Virtual Domain\_ID is available.

**State P2:Primary Operational.** In this state the Controlling Switch is operational as Primary. On entering this state the Controlling Switch:

- a) sets its priority to 02h;
- b) initiates an ERP Exchange with the Secondary Controlling Switch, if available;
- c) performs a DSMD Exchange with all reachable FCDF of the FCDF\_Set and with all reachable Alternate Controlling Switches of the Controlling Switch Set, if any, declaring itself as Primary Controlling Switch;
- d) performs an NPRD Exchange with all reachable FCDF of the FCDF\_Set and with all reachable Alternate Controlling Switches of the Controlling Switch Set;
- e) performs an AZAD Exchange with all reachable FCDFs of the FCDF\_Set;
- f) on native Fibre Channel links that were Isolated because connected to FCDFs, if any, it performs an ELP; and
- g) on FCoE interfaces, it establishes VA\_Port to VA\_Port Virtual Links with neighbor FDFs belonging to the FDF Set to which no VA\_Port to VA\_Port Virtual Links has been established, if any.

While in this state, the Controlling Switch:

- a) performs the Distributed Switch operations (see 17.4);
- b) generates the FSPF LSR describing the Virtual Domain in the Distributed Switch and lists the Virtual Domain as a directly attached Domain in its FSPF LSR;
- c) on receiving an SSA SW\_ILS (see 6.4.3.4) (i.e., when the Secondary Controlling Switch completed its state synchronization) performs a DSMD Exchange with all reachable FCDF of the FCDF\_Set and with all reachable Alternate Controlling Switches of the Controlling Switch Set, if any, declaring itself as Primary and the Secondary as Secondary;
- d) if the Secondary Controlling Switch is available sends RHello requests (see 6.4.3.5) every RHello\_Interval over each of its AISLs connected to the Secondary Controlling Switch and over each ASL through which the Secondary is reachable;
- e) resets the Down\_Timer to Down\_Interval everytime an RHello request from the Secondary Controlling Switch is received over at least one AISL or ASL;
- f) when the Secondary Controlling Switch is not anymore available (i.e., when Down\_Timer expires) and there are no neighbor Alternate Controlling Switches, it performs a DSMD Exchange with all reachable FCDF of the FCDF\_Set and with all reachable Alternate Controlling Switches of the Controlling Switch Set, if any, declaring itself as Primary;
- g) when the Secondary Controlling Switch is not anymore available (i.e., when Down\_Timer expires) and there are neighbor Alternate Controlling Switches, it performs an ERP Exchange with the neighbor Alternate Controlling Switch selected to become the new Secondary Controlling Switch (i.e, the neighbor Alternate Controlling Switch having the lowest value of Controlling Switch priority || Switch\_Name) followed by a DSMD Exchange with all reachable FCDF of the FCDF\_Set and with all reachable Alternate Controlling Switches of the Controlling Switch Set, if any, declaring itself as Primary and the new Secondary as Secondary;
- h) if its AISL\_Set goes from NULL to not-NULL (i.e., a Controlling Switch becomes available), it performs an ERP Exchange with the other Controlling Switch; and
- i) if an Alternate Controlling Switch becomes available, it performs a DSMD Exchange with that Alternate Controlling Switch.

**Transition P2:P2:** Occurs following the ERP Exchange performed when the Controlling Switch AISL\_Set went from NULL to not-NULL if:

- a) the received Controlling Switch State Descriptor contains a zero Virtual Domain Switch\_Name and no N\_Port\_ID Ranges;
- b) the received Controlling Switch State Descriptor contains the Virtual Domain Switch\_Name of the Distributed Switch, there is no allocated N\_Port\_ID conflict between the two Controlling Switches, and this Switch is the one selected to remain Primary;
- c) the received Controlling Switch State Descriptor contains the Virtual Domain Switch\_Name of the Distributed Switch and there is an allocated N\_Port\_ID conflict between the two Controlling Switches. In this case the AISL shall be Isolated; or
- d) the received Controlling Switch State Descriptor contains a Virtual Domain Switch\_Name different than the one of the Distributed Switch. In this case the AISL shall be isolated.

**Transition P2:S1:** Occurs following the ERP Exchange performed when the Controlling Switch AISL\_Set went from NULL to not-NULL if the received Controlling Switch State Descriptor contains the Virtual Domain Switch\_Name of the Distributed Switch, there is no allocated N\_Port\_IDs conflict between the two Controlling Switches, and this Switch is the one selected to become Secondary.

**State S1:Secondary Initialization.** In this state a Controlling Switch performs the operations to become the Secondary Controlling Switch of the Distributed Switch. The Controlling Switch has to synchronize its state with the one of the Primary Controlling Switch. To this end the Controlling Switch:

- 1) requests to the Primary the FCDF topology through the GTFS SW\_ILS (see 6.4.3.2);
- 2) requests to the Primary the Virtual Domain\_IDs and N\_Port\_IDs Allocation state in the Distributed Switch through the GFNS SW\_ILS (see 6.4.3.3);
- 3) obtains the information associated with each N\_Port\_ID in the Name Server through the GE\_ID CT request; and
- 4) communicates the achieved state synchronization to the Primary through the SSA (see 6.4.3.4) SW\_ILS.

While in this state the Controlling Switch:

- a) processes FCUN, FCRN, and NPZD requests from the Primary Controlling Switch;
- b) sends RHello requests every RHello\_Interval over each of its AISLs and over each ASL through which the Primary is reachable; and
- c) resets the Down\_Timer to Down\_Interval everytime an RHello request is received over at least one AISL or ASL.

**Transition S1:S2.** Occurs when the Secondary Controlling Switch has synchronized its state with the Primary.

**State S2:Secondary Operational.** In this state the Controlling Switch is operational as Secondary. On entering this state the Controlling Switch:

- a) sets its priority to its configured value;
- b) initiates an ERP Exchange with the Primary Controlling Switch;
- c) on native Fibre Channel links that were Isolated because connected to FCDFs, if any, it performs an ELP; and
- d) on FCoE interfaces, it establishes VA\_Port to VA\_Port Virtual Links with neighbor FDFs belonging to the FDF Set to which no VA\_Port to VA\_Port Virtual Links has been established, if any.

While in this state, the Secondary Controlling Switch shall:

- a) participate in the Distributed Switch operations (see 17.4);
- b) list the Virtual Domain as a directly attached Domain in its FSPF LSR if at least one AISL path with the Primary Controlling Switch is available;
- c) not list the Virtual Domain as a directly attached Domain in its FSPF LSR if no AISL path with the Primary Controlling Switch is available;
- d) send RHello requests every RHello\_Interval over each of its AISLs and over each ASL through which the Primary is reachable; and
- e) reset the Down\_Timer to Down\_Interval everytime an RHello request is received over at least one AISL or ASL.

**Transition S2:P2.** Occurs when the Secondary Controlling Switch becomes Primary. This occurs when:

- a) the Secondary Controlling Switch Down\_Timer expires and the Primary Controlling Switch is no longer part of the fabric FSPF topology (i.e., the Primary Controlling Switch is no longer available); or
- b) the Priority field in a received ERP request has a value of FFh. This is an indication that the Primary Controlling Switch determined to become Secondary.

**Transition S2:S3.** Occurs when the Secondary Controlling Switch Down\_Timer expires, the Primary Controlling Switch is still part of the fabric FSPF topology (i.e., the Secondary Controlling Switch is



unable to maintain synchronization with the Primary but the Primary is still part of the fabric), and no Alternate Controlling Switches are available.

**State S3. Secondary Suspended.** In this state the Controlling Switch suspends its operations as Controlling Switch to avoid split-brain scenarios. On entering this state the Controlling Switch shall Isolate all of its ASLs.

**Transition S3:S1.** Occurs when an AISL with the Primary Controlling Switch is instantiated.

**Transition S1:S3.** Occurs when the Secondary Controlling Switch Down\_Timer expires, the Primary Controlling Switch is still part of the fabric FSPF topology (i.e., the Secondary Controlling Switch is unable to maintain synchronization with the Primary but the Primary is still part of the fabric), and at least one Alternate Controlling Switch is available.

**Transition P2:S2.** Occurs when the Primary Controlling Switch determines to become Secondary by setting its priority to FFh. This may happen as result of an administrative action.

**Transition S1:R1.** Occurs when the Secondary Controlling Switch Down\_Timer expires, the Primary Controlling Switch is still part of the Fabric FSPF topology (i.e., the Secondary Controlling Switch is unable to maintain synchronization with the Primary but the Primary is still part of the Fabric), and at least one Alternate Controlling Switch is available.

**Transition S2:R1.** Occurs when the Secondary Controlling Switch Down\_Timer expires, the Primary Controlling Switch is still part of the Fabric FSPF topology (i.e., the Secondary Controlling Switch is unable to maintain synchronization with the Primary but the Primary is still part of the Fabric), and at least one Alternate Controlling Switch is available.

**State R3: Flood.** In this state the Controlling Switch starts the SPCS\_Timer and transmits a SPCS SW\_ILS (see 6.4.3.6) to all neighbor Controlling Switches on all AISLs from which the Controlling Switch has not yet received a SPCS SW\_ILS, if its AISL\_Set is not NULL. The SPCS SW\_ILS is then flooded to all reachable Controlling Switches. If its AISL\_Set is NULL, then the Controlling Switch exits this state.

**Transition R3:P1.** Occurs if the Controlling Switch AISL\_Set is NULL.

**Transition R3:R4.** Occurs after SPCS\_TOV following the reception or origination of the first SPCS SW\_ILS.

**State R4: Selection.** In this state the Controlling Switch having the lowest value of Controlling Switch priority || Switch\_Name is selected as Primary Controlling Switch. To this end:

- a) upon entering this state the Controlling Switch shall transmit a ECSP SW\_ILS (see 6.4.3.7) to all neighbor Controlling Switches on all AISLs from which the Controlling Switch has not yet received an ECSP SW\_ILS request. If the Controlling Switch has not yet received any ECSP SW\_ILS request, it shall list its Controlling Switch priority || Switch\_Name value in the Primary Controlling Switch priority || Switch\_Name record (i.e., list itself as a candidate Primary Controlling Switch, see 6.4.2.6) and in the list of Controlling Switch priority || Switch\_Name records (i.e., list itself as one of the Controlling Switches of the Distributed Switch, see 6.4.2.6) of the ECSP SW\_ILS request it generates;
- b) the Controlling Switch shall retain a Controlling Switch priority || Switch\_Name value that it believes is the lowest among the Controlling Switches belonging to the Controlling Switch Set a candidate Primary Controlling Switch Priority || Switch\_Name value. The Controlling Switch shall also retain a merged list of Controlling Switch Priority || Switch\_Name records including the Controlling Switches it learns about during the ECSP exchanges; and

- c) the Controlling Switch shall communicate its retained Primary Controlling Switch priority || Switch\_Name value to the neighbor Controlling Switches that it has not yet communicated that value. The Controlling Switch shall also communicate its retained list of Controlling Switch priority || Switch\_Name records to the neighbor Controlling Switches that it has not yet communicated that list. The Controlling Switch accomplishes this either by originating a new ECSP SW\_ILS request or by an SW\_ACC to a received ECSP SW\_ILS request. If the Controlling Switch receives a new lower value of Primary Controlling Switch priority || Switch\_Name before it has had a chance to communicate a prior lower value, it shall not attempt to communicate the prior value, and shall instead attempt to communicate the new value.

**Transition R4:P1.** Occurs after SPCS\_TOV following the reception or origination of the first ECSP SW\_ILS and the lowest Controlling Switch priority || Switch\_Name is equal to the its Controlling Switch priority || Switch\_Name (i.e., the Controlling Switch is selected to be the Primary Controlling Switch).

**Transition R4:A1.** Occurs after SPCS\_TOV following the reception or origination of the first ECSP SW\_ILS and the lowest Controlling Switch priority || Switch\_Name is not equal to the its Controlling Switch priority || Switch\_Name.

**State A1: Alternate Init.** In this state a Controlling Switch performs the operations to become an Alternate Controlling Switch of the Distributed Switch. The Controlling Switch waits until it receives an ERP request or a DSMD request from the Primary Controlling Switch.

**Transition A1:R1.** Occurs when no AISL path with the Primary Controlling Switch is available.

**Transition A1:A2.** Occurs when the Controlling Switch receives a DSMD request from the Primary Controlling Switch.

**Transition A1:S1.** Occurs when the Controlling Switch receives an ERP request from the Primary Controlling Switch.

**State A2: Alternate Oper.** In this state the Controlling Switch is operational as Alternate. On entering this state the Controlling Switch:

- a) on native Fibre Channel links that were Isolated because connected to FCDFs, if any, it performs an ELP; and
- b) on FCoE interfaces, it establishes VA\_Port to VA\_Port Virtual Links with neighbor FDFs belonging to the FDF Set to which no VA\_Port to VA\_Port Virtual Links has been established, if any.

While in this state, the Alternate Controlling Switch shall:

- a) participate in the Distributed Switch operations (see 17.4); and
- b) list the Virtual Domain as a directly attached Domain in its FSPF LSR if at least one AISL path with the Primary Controlling Switch is available.

**Transition A2:S1.** Occurs when the Controlling Switch receives an ERP request from the Primary Controlling Switch.

**Transition A2:R1.** Occurs when no AISL path with the Primary Controlling Switch is available.



## 18 Leaf Switch


### 18.1 Overview

A Leaf Switch is a Switch operating at the edge of a Fabric that cannot be used to access other Switches in the Fabric. It functions as a terminating point in a Fabric and other Switches in the Fabric shall not include paths through a Leaf Switch to reach any other Switch in the Fabric.

A Leaf Switch is identified by the Leaf Switch bit set to one in the Link State Record LSR Flags field. All Switches in a Fabric receive the Leaf Switch LSR during the FSPF LSR database synchronization, which informs the other Switches in a Fabric that the Switch operating as a Leaf Switch.

Paths through a Leaf Switch are excluded from the path selection results, such that any path between Switches through a Leaf Switch is ignored.

Figure 57 shows an example of a Leaf Switch and two Switches.



**Figure 57 – Leaf Switch example**

In figure 57, if the “dotted green” ISL between Switch-1 and Switch-2 is not active and the “red” ISL paths connecting the Leaf Switch to Switch-1 and Switch-2 are active, the path that passes through the Leaf Switch from Switch-1 to Switch-2 is not a valid path and is excluded from the FSPF calculation. In this example, Switch-1 and Switch-2 do not have a valid path to each other, but they each have a valid path to the Leaf Switch. The Leaf Switch also has a valid path to both Switch-1 and Switch-2. If the “green” ISL is active, Switch-1 and Switch-2 have a valid path to each other, distribute their device information, and send RSCNs to the attached devices.

## Annex A (informative)


### Examples of Switch port initialization

#### A.1 Introduction

This annex presents some example scenarios that may occur during Switch port initialization (see 7.2). It is expected that the reader is familiar with Loop Initialization as defined in FC-AL-2, and with Link Initialization as defined in FC-FS-5. Loop Initialization states as defined in FC-AL-2 are referenced here to facilitate the understanding of the process.

#### A.2 Example 1: two E/F/FL\_Port-capable Switch ports

In this example, two Switch ports that are E/F/FL\_Port-capable are attached to each other. Figure A.1 illustrates this example.



**Figure A.1 – Initialization example 1**

According to the initialization algorithm, since each Switch port is E/F/FL\_Port-capable, they start the process with Loop Initialization, as defined in FC-AL-2. LIP Primitive Sequences are sent and received, and each Switch port starts sending LISM frames. When Switch port X receives LISM from Switch port Y, it sees that its Port\_Name is lower than the Port\_Name in the payload, and continues sending the same LISM.

On the other hand, when Switch port Y receives LISM from Switch port X, it sees that its Port\_Name is higher than the Port\_Name in the payload. This causes Switch port Y to start sending the LISM it received, with the Port\_Name belonging to Switch port X. Switch port Y also transitions to the MONITORING state with PARTICIPATE = FALSE (0), because only one FL\_Port may be participating on a loop.

Switch port X receives its LISM and assumes the role of Loop Master. Switch port X then proceeds to send all of the other Loop Initialization Sequences, and by the end of Loop Initialization, discovers that it is the only L\_Port on the Loop.

Because there may be a Non-Participating Switch port on the Loop, Switch port X knows it is required to attempt Link Initialization. Switch port X begins Link Initialization by REQ(old-port). Switch port X transitions to the OLD-PORT-REQ state and begins transmitting LIP; this causes Switch port Y to begin Loop Initialization. Switch port Y transmits a minimum of 12 of the received LIPs in the OPEN-INIT-START state and transitions to the OPEN-INIT-SELECT-MASTER state. When Switch port X recognizes LIP, it transitions to the OLD-PORT state and transmits OLS for minimum(2xAL\_TIME).

After a maximum (1xAL\_TIME), Switch port Y recognizes Primitive Sequences (OLS, NOS) defined in FC-FS-5 and transitions from the FL\_Port operating mode to E/F\_Port mode. The Link protocol continues to completion and a point-to-point Link is now active.

Switch port X and Switch port Y may now attempt to exchange link parameters and establish an Inter-Switch Link.

### A.3 Example 2: two E/F/FL\_Port-capable Switch ports and one PN\_Port

In this example, two Switch ports that are E/F/FL\_Port-capable are attached to each other as in the first example, but there is also a PN\_Port on the loop. Figure A.2 illustrates this example.

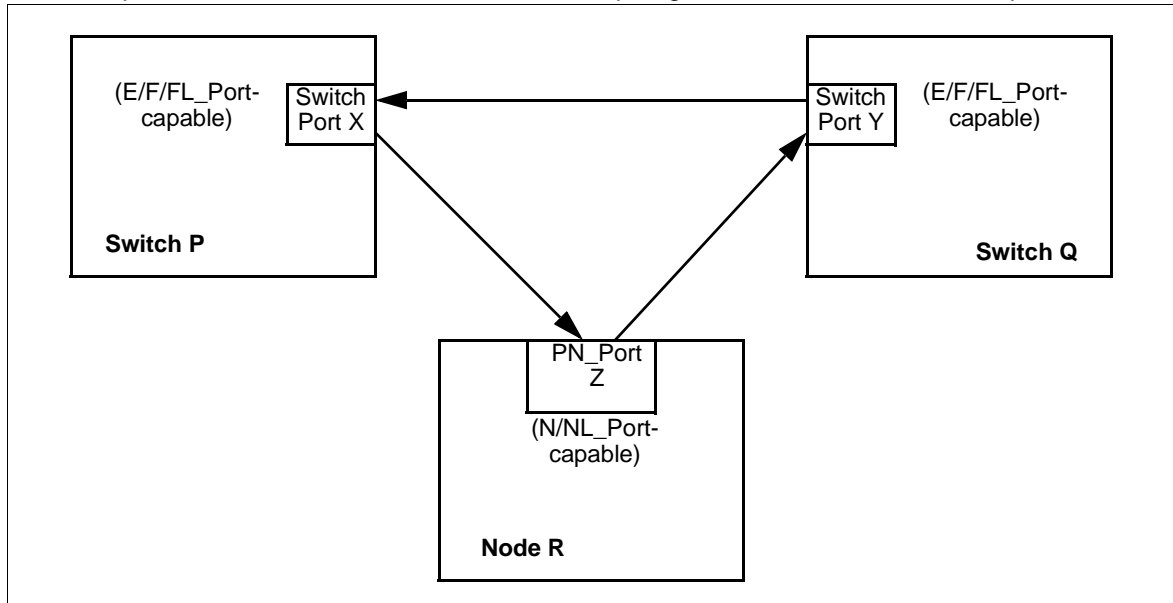


Figure A.2 – Initialization example 2

According to the initialization algorithm, since each Switch port is E/F/FL\_Port-capable and PN\_Port Z is N/NL\_Port-capable, they start the process with Loop Initialization, as defined in FC-AL-2. LIP Primitive Sequences are sent and recognized, and each Switch port and the PN\_Port start sending LISM frames. As in the first example, Switch port X receives LISM from Switch port Y, it sees that its Port\_Name is lower than the Port\_Name in the payload, and continues sending the same LISM.

When PN\_Port Z receives the LISM from Switch port X, it finds a D\_ID of zero, meaning that the originator is an FL\_Port. Since an FL\_Port always wins Loop Master, the PN\_Port begins sending the received LISM from Switch port X.

When Switch port Y receives Switch port X's LISM from Port Z, it sees that its Port\_Name is higher than the Port\_Name in the payload. This causes Switch port Y to start sending the LISM it received, with the Port\_Name belonging to Switch port X. Switch port Y also transitions to the MONITORING state with PARTICIPATE = FALSE (0), because only one FL\_Port may be participating on a loop.

Switch port X receives its LISM and assumes the role of Loop Master. Switch port X then proceeds to send all of the other Loop Initialization Sequences, and by the end of Loop Initialization, discovers that there is only one other L\_Port on the loop. Because that one other port may be capable of point-to-point operation, Switch port X knows it may attempt Link Initialization.

Switch port X begins Link Initialization by asserting REQ (old-port) causing the transmission of LIP in the OLD-PORT-REQ state, and causes PN\_Port Z to begin Loop Initialization. PN\_Port Z transmits a minimum of 12 received LIPs in the OPEN-INITSTART state causing Switch port Y to begin Loop Initialization and transitions to either the OPEN-INIT-SELECTMASTER or the SLAVE-WAIT-FOR-MASTER state. Switch port Y transmits a minimum of 12 received LIPs in the OPENINIT-START state and transitions to the OPEN-INIT-SELECT-MASTER state. Switch port X recognizes LIP, transitions to the OLD-PORT state and transmits OLS for minimum (2xAL\_TIME). If after minimum (1xAL\_TIME) PN\_Port Z recognizes OLS and reacts to it, it transitions to the OLD-PORT state and transmits LR in response. Switch port Y being in the OPEN-INIT-SELECT-MASTER state does not recognize LR and continues with the INITIALIZATION process thereby blocking LR to Switch port X. When Switch port X fails Link Initialization it should remove REQ (old-port) to allow Loop Initialization to complete. When Loop Initialization completes successfully, Switch port X operates as an FL\_Port, and PN\_Port Z operates as an L\_Port. Switch port Y remains Non-Participating.

NOTE 31 – If PN\_Port Z had been bypassed, the process would have completed as in Example 1, because the Primitive Sequences would have been ignored by PN\_Port Z. At a later time, if PN\_Port Z is then enabled, Loop Initialization begins (i.e., PN\_Port Z starts sending LIP to get an AL\_PA), and things sort themselves out as described for Example 2. If Switch port Y had been bypassed, then Switch port X would have become an F\_Port in a point-to-point Link with PN\_Port Z.

NOTE 32 – If PN\_Port Z was L\_Port capable only when it went to the OPEN-INIT-START state, it would stall in either the OPEN-INITSELECT-MASTER or SLAVE-WAIT-FOR-MASTER state transmitting LISM or waiting for an ARB(F0). This would cause Switch port X to fail at Link Initialization, and then go back to Loop Initialization. Again, Switch port Y stays Non-Participating.

#### A.4 Example 3: one E/F/FL\_Port-capable Port and one E/F\_Port-capable Port

In this example, a Switch port that is E/F/FL\_Port-capable is attached to a Switch port that is E/F\_Port-capable. Figure A.3 illustrates this example.




Figure A.3 – Initialization example 3

According to the initialization algorithm, the Switch port that is E/F/FL\_Port-capable starts the process with Loop Initialization as defined in FC-AL-2. However, the Switch port that is E/F\_Port-capable starts the process with Link Initialization as defined in FC-FS-5. Switch port X sends LIP Primitive Sequences, and Switch port Y sends OLS Primitive Sequences. If Switch port X in the NORMAL-INITIALIZE state does not receive LIP before expire(3xAL\_TIME), it transitions to the OLD-PORT state and completes Link Initialization.

Switch port X and Switch port Y may now attempt to exchange link parameters and establish an Inter-Switch Link.

**Annex B**  
(informative)

**ELP Negotiation Example**


**B.1 Introduction**

This annex presents an example of how ELP negotiation may be performed.

**B.2 ELP exchange protocol**

The following description is an extension of the ELP exchange described this standard. It allows for a negotiation of link parameters.

NOTE 33 – In the following discussion related to the ELP Protocol, the Reference Configuration given in figure B.1 is used.



**Figure B.1 – Reference ELP configuration**

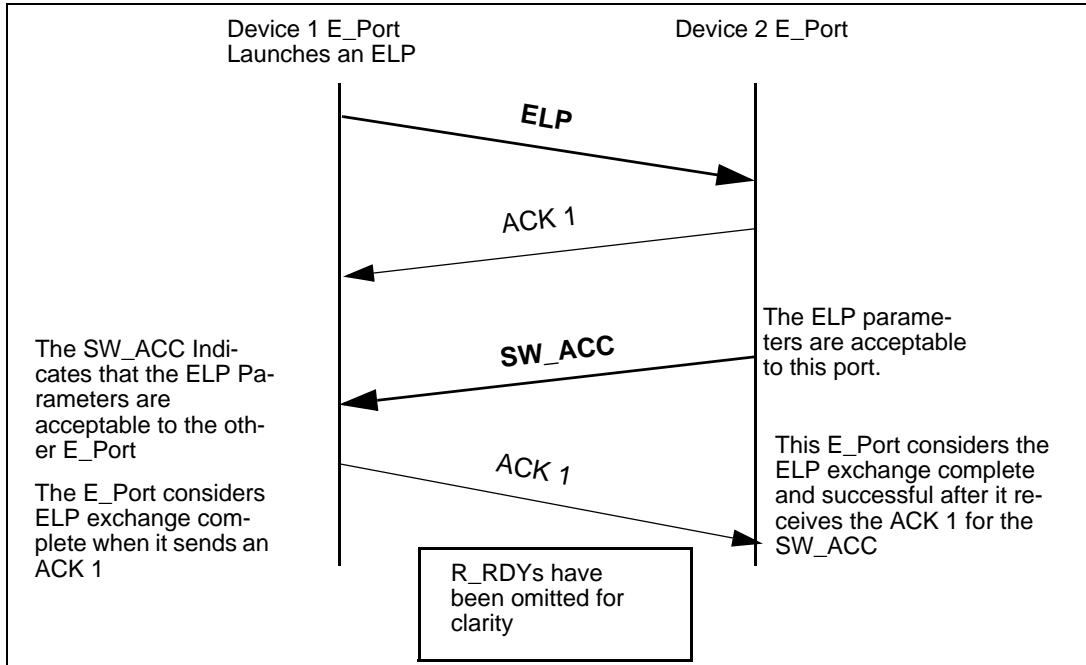
Following is a summary of the resulting state at each E\_Port after the ELP exchange with and without Parameter Negotiation.

**B.2.1 ELP exchange without parameter negotiation**

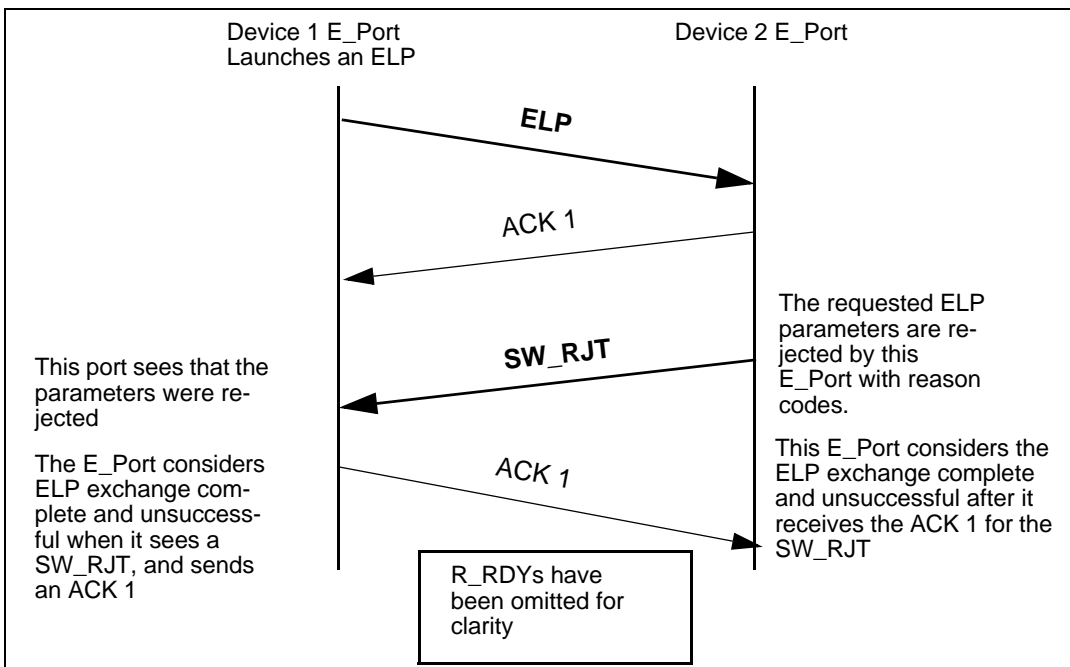
- The ELP originating port (Device 1) shall consider the exchange of link parameters complete and successful, when it has transmitted an ACK\_1 for the SW\_ACC it has received.
- The ELP originating port (Device 1) shall consider the exchange of link parameters complete but not successful, when it has transmitted an ACK\_1 for the SW\_RJT it has received. The originating port now goes into isolation.
- The responding port (Device 2) shall consider the exchange of link parameters complete and successful, when it has received the ACK\_1 for the SW\_ACC it has transmitted.
- The responding port (Device 2) shall consider the exchange of link parameters complete but not successful, when it has received the ACK\_1 for the SW\_RJT it has transmitted. The responding port now goes into isolation.

Figures B.2 and B.3 illustrate two complete ELP exchanges between two E\_Ports, one successful and the other unsuccessful without negotiation. This is the current ELP operation.





**Figure B.2 – A successful and complete ELP exchange**

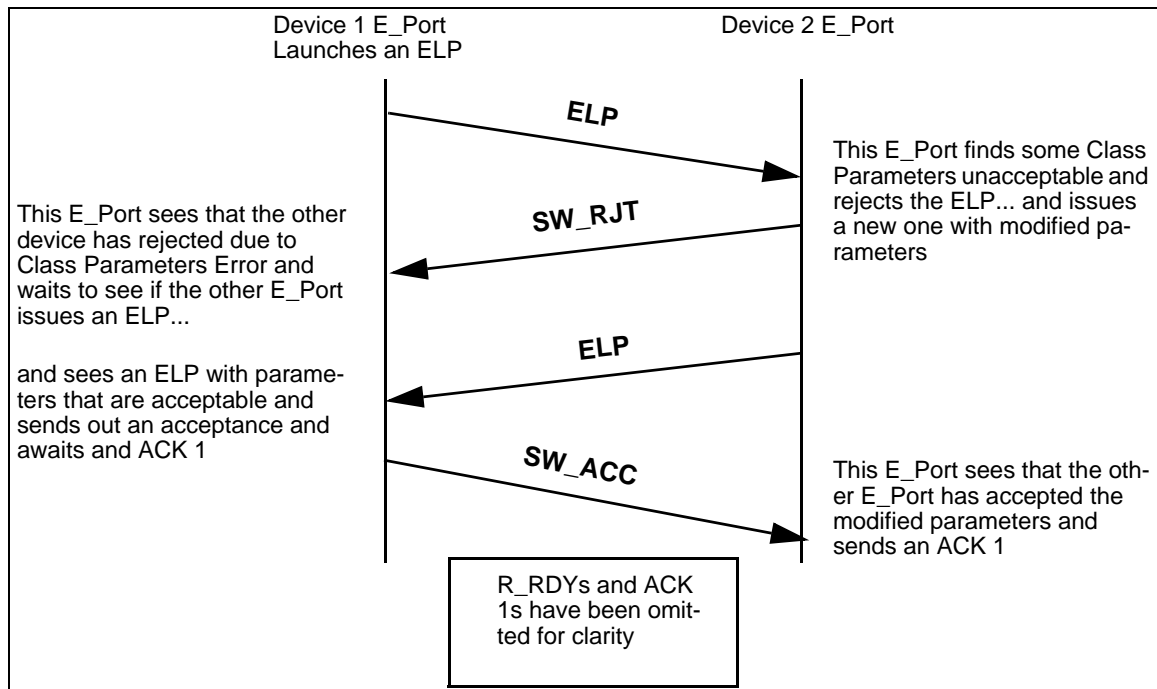


**Figure B.3 – An unsuccessful but complete ELP exchange**

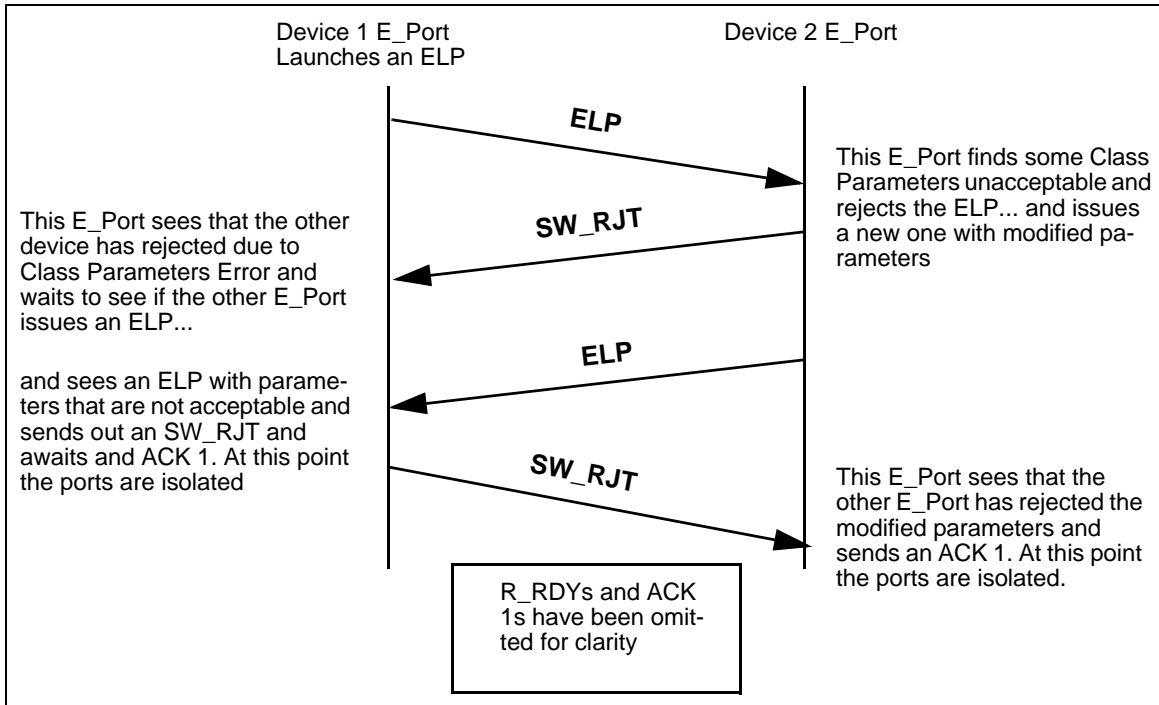
### B.2.2 ELP exchange with parameter negotiation

- The ELP parameter negotiation process always converges with the final reply being a SW\_ACC for a successful exchange or a SW\_RJT for an unsuccessful exchange.
- If a responding port (Device 2) is unable to agree to the received parameters, it sends out a SW\_RJT after which it may issue a new ELP with modified parameters. This responding port now becomes the new ELP originator (Device 2).
- The old originating port (Device 1) after receiving a SW\_RJT shall send an ACK\_1 and waits for the possible arrival of a new ELP with modified parameters; this device (Device 1) now becomes the new responder. Until a new ELP is received this device is isolated on this link.
- If the new responder (Device 1) finds the parameters acceptable, then it sends out a SW\_ACC and waits for an ACK\_1.
- If the new responder (Device 1) finds the parameters unacceptable, then it sends out a SW\_RJT and waits for an ACK\_1 after which it goes into isolation.
- There is a maximum of 2 exchanges of link parameters, after which the ports are either operational or isolated.

Figures B.4 and B.5 illustrates a complete ELP exchanges between two E\_Ports with negotiation, one successful and the other unsuccessful with negotiation:



**Figure B.4 – A successful ELP exchange protocol parameter negotiation**



**Figure B.5 – An unsuccessful ELP exchange protocol parameter negotiation**

**Annex C**  
(informative)

**Fabric Device Management Interface-Sample Flows**

**C.1 Introduction**


This annex presents sample flows for the Fabric Device Management Interface (FDMI).

**C.2 Sample flows**

**C.2.1 HBA registration - single Switch**

In Figure C.1 below, the Switch interactions are shown for HBA registration to a single Switch.

- 1) The HBA attempts registration by sending an RHBA to Switch 1. Since the HBA has not already registered with Switch 1, the registration completes successfully and Switch 1 becomes the HBA's primary manager.
- 2) Switch 1 sends a Registration Notification to Switches 2 and 3. Switches 2 and 3 now update their caches such that subsequent queries to Switches 2 and 3 regarding the HBA may be handled locally.
- 3) Some time later the HBA requests registration by sending an RHBA through another port to Switch 1. This time the registration fails because the HBA is already registered with Switch 1. Since the HBA has not registered with any other Switches, Switch 1 becomes the HBA's primary manager.

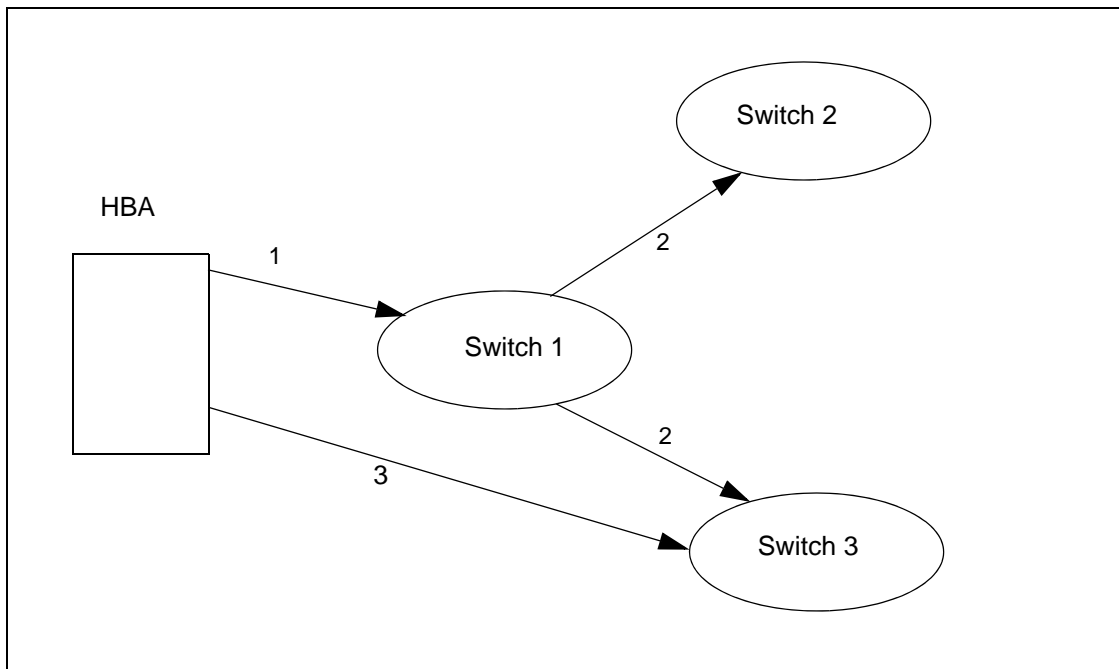


**Figure C.1 – Registration of HBA information - single Switch**

### C.2.2 HBA registration - multiple Switches - caches updated

In Figure C.2 below, the Switch interactions are shown for HBA registration to multiple Switches when the caches are updated.

- 1) The HBA attempts registration by sending an RHBA to Switch 1. Since the HBA has not already registered with Switch 1, the registration completes successfully.
- 2) Switch 1 sends a Registration Notification to Switches 2 and 3. Switches 2 and 3 now update their caches such that subsequent queries to Switches 2 and 3 regarding the HBA may be handled locally.
- 3) Some time later the HBA requests registration by sending an RHBA through another port to Switch 3. Switch 3 performs the checks against its FDMI database and its cached information. The cached information indicates that the HBA has already been registered in Switch 1 and the registration is rejected. Switch 1 becomes the HBA's primary manager.



**Figure C.2 – Registration of HBA information - multiple Switches caches updated**

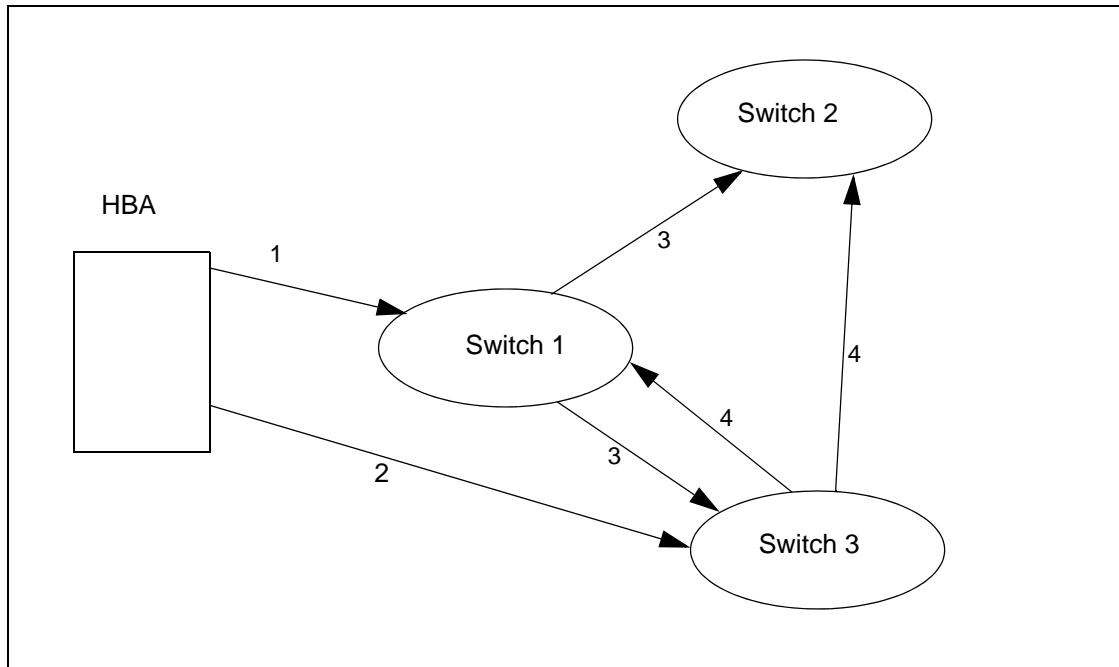
### C.2.3 HBA registration - multiple Switches - caches not updated

In Figure C.3 below, the Switch interactions are shown for HBA registration to multiple Switches when the caches are not updated:

- 1) the HBA attempts registration by sending an RHBA to Switch 1. Since the HBA has not already registered with Switch 1, the registration completes successfully;
- 2) before Switch 1 sends a Registration Notification to Switches 2 and 3, the HBA requests registration by sending an RHBA through another port to Switch 3. Since Switch 3 does not have an updated cache, Switch 3 assumes that the HBA is not registered and accepts the registration;
- 3) Switch 1 sends a Registration Notification to Switches 2 and 3; and

4) Switch 3 send a Registration Notification to Switches 1 and 2.

In this case Switch 1 has the lowest Switch\_Name. Switch 2 receives the Registration Notification from Switches 1 and 3 and notes that Switch 1 has the lower Switch\_Name. Switch 2 updates his cache with the information received from Switch 1 and designates Switch 1 has the primary manager for the HBA. Switch 3 receives the Registration Notification from Switch 1 and notes that Switch 1 has the lower Switch\_Name. Switch 3 deletes the HBA information from its FDMI database and updates his cache with the information received from Switch 1. Switch 3 designates Switch 1 as the primary manager for the HBA.




**Figure C.3 – Registration of HBA Information - multiple Switches caches not updated**

#### C.2.4 HBA de-registration - primary HBA manager

In Figure C.4 below, the Switch interactions are shown for HBA de-registration to the primary HBA manager:

1) the HBA attempts de-registration by sending a DHBA to Switch 1. Since Switch 1 is the primary HBA manager, the de-registration completes successfully and the HBA information is removed from Switch 1's FDMI database; and

2) Switch 1 sends a De-registration Notification to Switches 2 and 3. Switches 2 and 3 now delete the HBA's information from their caches.

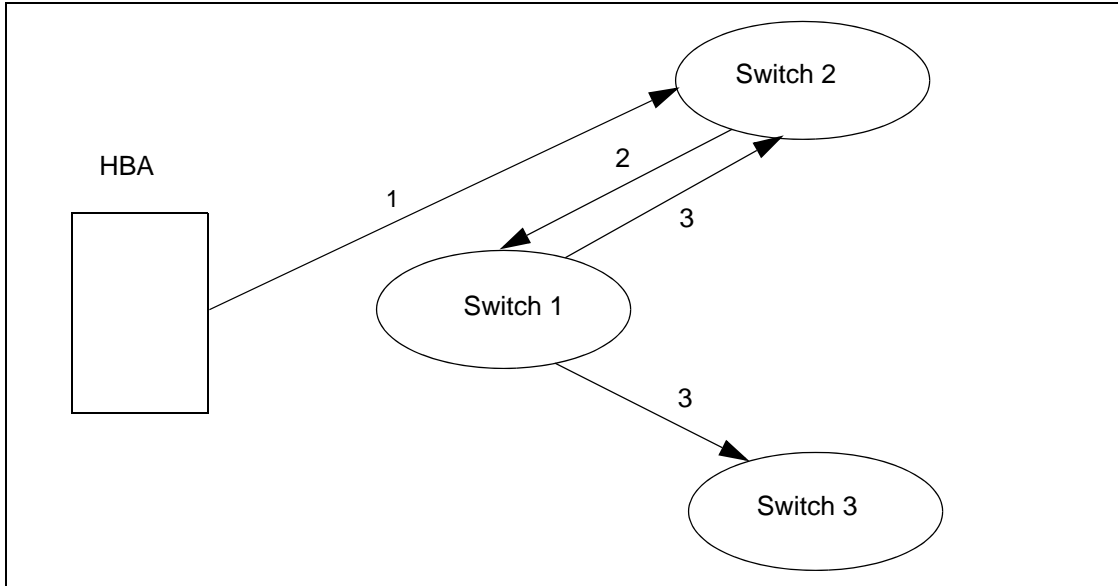


**Figure C.4 – HBA de-registration - primary HBA manager**

### **C.2.5 HBA de-registration - non-primary HBA manager**

In Figure C.5 below, the Switch interactions are shown for HBA de-registration to a Switch that is not the primary HBA manager. In this case, Switch 1 is the primary manager for the HBA and the HBA is connected to Switch 1 and Switch 2:

- 1) the HBA attempts de-registration by sending a DHBA to Switch 2. Since Switch 2 is not the primary manager for the HBA Switch 2, it is required to inform Switch 1 of the de-registration request;
- 2) Switch 2 sends a De-Registration Forward to Switch 1 which is the HBA's primary manager. Switch 1 de-registers the HBA information and deletes the information from its FDMI database; and
- 3) Switch 1 sends a De-registration Notification to Switches 2 and 3. Switches 2 and 3 now delete the HBA's information from their caches.



**Figure C.5 – HBA de-registration - non-primary HBA manager**



## Annex D (Normative)

### Fast Fabric Initialization for AE-Capable Equipment

#### D.1 Overview

This annex defines the minimal requirements for Fast Fabric Initialization (FFI). The FFI requirements apply to Switches that are specified to be “AE-Capable”. An AE-Capable Switch is a Switch that supports at least one Avionics Expansion Port (AE\_Port).

#### D.2 Background

Fibre Channel is growing in industry acceptance in the Avionics Environment. Initial systems were relatively simple, employing architectures containing no more than two Switches. Systems being designed today for tomorrow’s applications will be complex by comparison, employing many Switches. The real-time nature of modern avionics demands stringent requirements, especially during fabric initialization, for determinism, low latency, predictability, reliability, and fault tolerance.

Avionics Fibre Channel equipment differs from typical commercial equipment in several important ways. Foremost, the Avionics Fabric domain topology by definition will be predefined. The location of all domains and their associated Inter-Switch Links within the fabric topology, whether currently active or inactive, do not change for the duration of a mission.

Another important distinction in requirements between avionics systems and commercial systems is that commercial systems are plug-and-play with self-discovery, while avionics systems are inherently well-known systems with fixed configurations. The plug-and-play requirement in the commercial industry results in significant protocol mechanisms not necessary in the Avionics Environment. In Avionics all the Switches, their Domain\_IDs, Inter-Switch Links, and the topology and routing maps are well known per mission and not subject to change. By eliminating the protocol tools utilized in the commercial industry for discovering entities, Avionics systems can significantly simplify initialization protocols and lower system latency from power interruption to an active state.

This Annex formalizes the methods for accelerating the initialization of an Avionics Fabric by allowing certain AE-Capable Switches to have implicit knowledge of the entire domain topology before fabric initialization, which is then distributed throughout the fabric. Techniques used in commercial fabric initialization, such as Principal Switch Selection, and the Fabric Shortest Path First protocol are undesirable and omitted since they add overhead and uncertainty to the initialization time of Avionics systems. Avionics Fabrics must be initialized quickly, and in a deterministic and repeatable fashion.

#### D.3 Definitions

**AE (Avionics Environment):** Avionics Environment refers to hi-reliability applications in harsh environmental conditions.

**AE-Capable Switch:** A Fibre Channel Fabric Switch that is capable of supporting at least one AE\_Port.

**AE Principal Switch:** An AE Switch has no Uplinks and assumes the primary role of distributing the Domain Topology Map in an Avionics Fabric.

**AE Secondary Principal Switch:** An AE Switch that is capable of becoming the AE Principal Switch.

**AE Switch:** An AE-Capable Switch that has activated at least one AE\_Port. AE Switches are required to implement the requirements set forth in this Annex.

**AE\_Port:** A Fabric Avionics Expansion Port that connects to another Avionics Expansion Port to create an Inter-Switch Link. AE\_Ports are required to implement the requirements set forth in this Annex.

**Active AE\_Port:** An AE\_Port that has reached the AE0 state or subsequent states.

**Avionics Fabric:** A Fibre Channel Fabric that contains at least one AE Switch and supports all the requirements of this Annex. An Avionics Fabric may or may not support the requirements of the rest of the FC-SW-5 standard.

**Domain Topology Map:** An entity within the Avionics Fabric that unambiguously describes the Domain\_IDs and all of the Inter-Switch Links of the Avionics Fabric. The Domain\_IDs and all of the Inter-Switch Links shall remain unchanged for the duration of a mission.

**Downlink:** An ISL connected to an Active AE\_Port that is not part of at least one path that leads to the AE Principal Switch.

**FFI (Fast Fabric Initialization):** A technique that provides accelerated initialization of an Avionics Fabric through the distribution of the Domain Topology Map. The Domain Topology Map is distributed to all AE Switches via the AE Principal Switch using the FFI request Sequence.

**FFI Incarnation Number:** A number within FFI request Sequence that uniquely identifies each version of the Domain Topology Map. The FFI Incarnation Number is required to be managed solely by the AE Principal Switch.

**FFI Link Descriptor:** A description of an individual AE\_Port-to-AE\_Port connection within the Avionics Fabric.

**FFI Link State Record:** For an individual AE Switch, a description of the Domain and all the AE\_Port Inter-Switch Link connections of that Switch.

**FFI SW\_ILS:** An AE specific SW\_ILS command that distributes the Domain Topology Map throughout the Avionics Fabric or for reporting changes in link status and error conditions.

**Inactive AE\_Port:** A Switch port that is designated in the Domain Topology Map to be an AE\_Port, but the Switch port has not yet reached the AE0 state in its Port Mode Initialization.

**Uplink:** An ISL connected to an Active AE\_Port that is part of at least one path that leads to the AE Principal Switch.

## **D.4 Characteristics of Avionics Fabrics**

### **D.4.1 Overview**

An Avionics Fabric differs from a normal Fibre Channel Fabric in that it shall support a method for initialization called Fast Fabric Initialization (FFI). Specified in this Annex, FFI amends the requirements defined in clause 7 for Avionics Fabrics.

Fast Fabric Initialization specifies a set of coherent behaviors for the establishment of communication between multiple Switches within an Avionics Fabric. The complete definition of the fabric topology is defined in the Domain Topology Map of the system. The AE Principal Switch shall be responsible for initiating the distribution the Domain Topology Map to all AE Switches in the Avionics Fabric.

The Domain Topology Map shall be distributed to all AE Switches in the Avionics Fabric using the FFI request Sequence. The FFI request Sequence shall use AE\_Ports for distribution. In addition to the Domain Topology Map, the FFI request Sequence payload shall include the link status of all AE\_Port links and the Switch status of all AE Switches that are defined in the Domain Topology Map.

The method for the AE Principal Switch to obtain the Domain Topology Map can be implicit or through the reception of the FFI\_DTM ELS Command from an Nx\_Port.

## **D.4.2 AE Switch port mode initialization**

### **D.4.2.1 Overview**

Switch ports that are AE-Capable shall negotiate their Port mode via the Switch port mode initialization state machine defined in 7 and as modified by this clause. Switch ports that are AE-Capable shall negotiate to become an AE\_Port by following the rules of this clause.

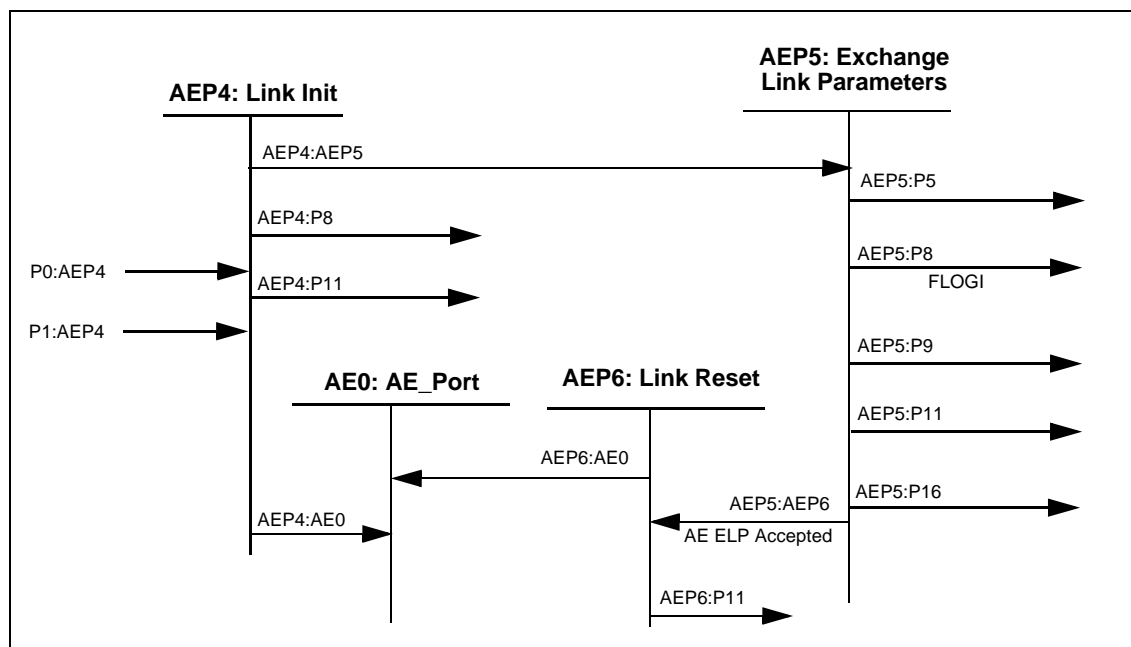
Switch ports that are AE-Capable may be connected to other Switch port types, but AE\_Port functionality shall only be used when an AE\_Port is directly connected to another AE\_Port.

### **D.4.2.2 Switch port mode initialization state machine modifications**

Figure D.1 shows the modifications needed to the Switch port mode initialization state machine specified in 7 for AE-Capable ports.

The AEP4 state replaces the P4 state that is described in clause 7. Therefore the transitions for P0:P4 and P1:P4 are identical except that they now point to state AEP4 as P0:AEP4 and P1:AEP4. The AEP4 state is described fully herein.

There are new Switch port mode initialization states called AEP5, AEP6 and AE0. They are described fully herein. All other Switch port mode initialization state machine states, including P5, remain unmodified.



**Figure D.1 – Modifications to Switch port mode initialization**

The following text specifies modifications to the Switch port mode initialization state machine for AE-Capable ports.

**Transition P0:AEP4.** The Switch port is not capable of becoming an FL\_Port. Attempt Link Initialization. (AEP4 replaces P4.)

**Transition P1:AEP4.** This transition occurs if the Loop Initialization does not complete successfully. This may occur if the Switch port is attached to a non-L\_Port capable port, so the next thing to try is a Link Initialization. (AEP4 replaces P4.)

**State AEP4: Link Initialization.** The Switch port shall attempt Link Initialization as defined in FC-FS-5.

**Transition AEP4:AE0.** This transition occurs if the Link Initialization procedure succeeds and the AE-Capable Port is programmed to immediately become an AE\_Port.

**Transition AEP4:AEP5.** This transition occurs if the Link Initialization procedure succeeds and the AE-Capable Port is not programmed to immediately become an AE\_Port.

**Transition AEP4:P8.** This transition occurs if the Link Initialization procedure succeeds and the AE-Capable Port is programmed to immediately become an F\_Port.

**Transition AEP4:P11.** This transition occurs when the Link Initialization procedure fails.

**State AEP5: Exchange Link Parameters.** An AE-Capable Switch port shall originate an AE-specific ELP SW\_ILS request Sequence by setting the ISL Flow Control Mode field in the ELP to a Vendor Specific value of AE02h. The minimum requirements of the ELP payload for AE-Capable ports are provided in D.4.3.

Table D.1 describes the responses an Originator of an ELP may receive, and further actions and state transitions the originator shall make.

**Table D.1 – Responses to ELP request for originating Interconnect\_Port (Part 1 of 2)**

Response to ELP	Indication	Originating Interconnect_Port action
1. R_RDY	Request received at destination	Wait E_D_TOV+1 second for response frame. Do not transition
2. ACK_1	Request received at destination	Wait E_D_TOV+1 second for response frame. Do not transition.
3. SW_ACC (Flow Control = "AE02")	Destination is an AE_Port and accepts all ELP parameters	Send ACK_1, Transition (AEP5:AEP6)
4. F_BSY or P_BSY	Destination is busy	Retry <sup>a</sup> , Transition (AEP5:P11)
5. F_RJT or P_RJT	The frame is not acceptable	Respond accordingly <sup>c</sup> , Transition (AEP5:P11)
6. ELP (rcvd Switch_Name > own Switch_Name)	Both Interconnect_Ports sent ELP at the same time. (Destination has control.)	Send SW_ACC or SW_RJT based on the values of the received ELP parameters. Transition (AEP5:AEP6)
7. ELP (rcvd Switch_Name < own Switch_Name)	Both Interconnect_Ports sent ELP at the same time. (Own Switch has control.)	Send SW_RJT <sup>b</sup> , Do not transition.
8. ELP (rcvd Switch_Name = own Switch_Name)	Interconnect_Port output is looped back to input	Remove loopback condition, Transition (AEP5:P9)
9. SW_RJT with reason code of "Command already in progress <sup>d</sup> "	Both Interconnect_Ports sent ELP at the same time. Destination has control (has greater Switch_Name) and has rejected own Switch's ELP.	Send SW_ACC or SW_RJT based on the values of the received ELP parameters. Do not transition.
<p><sup>a</sup> The retry is performed following a timeout period, as defined in P11. The R_A_TOV used for retry may be specifically set for Avionics Equipment.</p> <p><sup>b</sup> The reason code shall be "Unable to perform command request" with a reason code explanation of "Command already in progress".</p> <p><sup>c</sup> Response is defined in FC-FS-5.</p> <p><sup>d</sup> An SW_ACC is sent for the other ELP Exchange in progress if the received ELP parameters are acceptable. If the received ELP parameters are not acceptable, then an SW_RJT is sent.</p>		

**Table D.1 – Responses to ELP request for originating Interconnect\_Port (Part 2 of 2)**

Response to ELP	Indication	Originating Interconnect_Port action
10. SW_RJT with any other reason code	Both Interconnect_Ports sent ELP at the same time. Own Switch has control, but the destination Switch has rejected the ELP parameters	Try different ELP for E_Port if supported, Transition (AEP5:P5) Otherwise Isolate, Transition (AEP5:P9)
11. FLOGI	Destination is an N_Port	Respond accordingly <sup>c</sup> , Transition (AEP5:P8)
12. any other frame	Indeterminate	Discard frame and Retry <sup>a</sup> , Transition (AEP5:P11)
13. E_D_TOV+ 1 second expires	No response within timeout period	Retry <sup>a</sup> , Transition (AEP5:P11)
<p><sup>a</sup> The retry is performed following a timeout period, as defined in P11. The R_A_TOV used for retry may be specifically set for Avionics Equipment.</p> <p><sup>b</sup> The reason code shall be “Unable to perform command request” with a reason code explanation of “Command already in progress”.</p> <p><sup>c</sup> Response is defined in FC-FS-5.</p> <p><sup>d</sup> An SW_ACC is sent for the other ELP Exchange in progress if the received ELP parameters are acceptable. If the received ELP parameters are not acceptable, then an SW_RJT is sent.</p>		

The originating AE\_Port shall consider the exchange of link parameters complete (but not necessarily successful) when it has received the SW\_ACC or SW\_RJT and has transmitted the ACK\_1 for the SW\_ACC or SW\_RJT reply Sequence.

The responding AE\_Port shall consider the exchange of link parameters complete when it has received the ACK\_1 for the SW\_ACC or SW\_RJT.

The exchange of link parameters shall be considered successful when the exchange of link parameters is complete, and the reply to the ELP is an SW\_ACC, and both AE\_Ports agree that the parameters exchanged are acceptable.

**Transition AEP5:P5.** This transition occurs if the responding Interconnect\_Port does not agree that the AE\_Ports parameters are acceptable, and the Interconnect\_Port is capable of becoming an E\_Port. The responding Interconnect\_Port shall return an SW\_RJT reply Sequence with the reason code of “Unable to perform Command Request” and the reason code explanation of “Class F Service Parameter Error”, and shall perform the entry conditions for State P5.

This transition may also occur if the originating AE-Capable Port does not agree that the parameters in the SW\_ACC are acceptable, or it receives an SW\_RJT indicating the parameters in the ELP request were not acceptable to the responding Interconnect\_Port, and it is capable of originating a new ELP request Sequence for state P5.

**Transition AEP5:AEP6.** This transition is taken by the originator of the ELP exchange when the exchange of link parameters is complete. In order for this transition to occur, the received ISL Flow Control Mode field in the ELP SW\_ACC shall be set to the Vendor Specific value of AE02h and all other ELP parameters shall be set to acceptable values.

**Transition AEP5:P8.** This transition occurs if the exchange of link parameters is unable to be completed and an explicit FLOGI ELS request is received and the AE-Capable Switch port is capable of F\_Port operation.

**Transition AEP5:P9.** This transition occurs if the originating AE-Capable Port does not agree that the parameters in the SW\_ACC are acceptable, or it receives an SW\_RJT indicating the parameters in the ELP request were not acceptable to the responding Interconnect\_Port, and it is not capable of originating a new ELP request Sequence with modified parameters. (Note: The Switch port remains isolated until the next initialization event.)

**Transition AEP5:P11.** This transition occurs if the ELP is rejected with “unable to perform command request” reason code, and no FLOGI is received. The Switch port performs the Link Offline protocol as defined in FC-FS-5 during the transition.

**Transition AEP5:P16.** This transition is taken when authorization checks that are based on data from the ELP fail.

**State AEP6: Link Reset.** The AE\_Port shall perform the Link Reset Protocol.

**Transition AEP6:AE0.** This transition occurs if the Link Reset Protocol is successful.

**Transition AEP6:P11** This transition occurs if the Link Reset Protocol fails.

**State AE0: AE\_Port.** The Switch port has completed becoming an AE\_Port and shall continue to operate as an AE\_Port until the next initialization event. The AE\_Port shall then participate in the next phase of Fast Fabric Initialization.

**D.4.3 ELP payload requirements**

Table D.2 provides the minimum requirements of the ELP request and Accept payload for AE-Capable ports. Fields not specified below shall follow the definitions of the ELP SW\_ILS in clause 6.

The flow control model requires the use of the R\_RDY Primitive Signal to manage BB\_Credit. The format of the flow control parameters for an AE-Capable AE\_Port shall be the same as the flow control parameters for the R\_RDY mode of flow control (0002h). The Compatibility Parameters of the Flow Control Parameters shall be set to 0000h.

**Table D.2 – ELP required payload values for AE\_Ports (Part 1 of 2)**

ELP request/Accept payload	Value	Notes
Revision	3h	
Flags	0000h	B_Ports Prohibited
Class F Service Parameters		
VAL (Class Valid)	1b	
XII (X_ID Interlock)	0b	
Max BB Receive Data Field Size	≥ 2048	
Class 1 Interconnect_Port Parameters		

**Table D.2 – ELP required payload values for AE\_Ports (Part 2 of 2)**

ELP request/Accept payload	Value	Notes
VAL (Class 1 Valid)	X	
Max BB Receive Data Field Size	≥ 2048	Valid only when Class 1 validity is set to 1b
Class 2 Interconnect_Port Parameters		
VAL (Class 2 Valid)	X	
SEQ (Sequential Delivery)	1b	Valid only when Class 2 validity is set to 1b
Max BB Receive Data Field Size	≥ 2048	Valid only when Class 2 validity is set to 1b
Class 3 Interconnect_Port Parameters		
VAL (Class 3 Valid)	1b	
SEQ (Sequential Delivery)	1b	
Max BB Receive Data Field Size	≥ 2048	
ISL Flow Control Mode	AE02h	Uniquely defines AE-Capable ports
Flow Control Parameter Length	20	Follows the R_RDY flow control model
Compatibility Parameters	0000h	Not used
Legend: X = Don't Care		

#### D.4.4 AE Principal Switch

##### D.4.4.1 AE Principal Switch initialization process

An AE Principal Switch shall be required in order to support FFI.

The AE Principal Switch shall initiate fabric initialization through distribution of the Domain Topology Map. The Domain Topology Map is communicated to the AE Principal Switch either implicitly or through the reception of the FFI\_DTM ELS Command from an Nx\_Port. When the AE Principal Switch receives a valid FFI\_DTM ELS Command from an Nx\_Port, the AE Principal Switch shall update its own Domain Topology Map to the Domain Topology Map that is specified in the payload of the FFI\_DTM ELS.

The AE Principal Switch shall be capable of retaining or reacquiring its Domain Topology Map through a power cycling event.

The AE Principal Switch initiates the distribution of the Domain Topology Map throughout the Avionics Fabric. The AE Principal Switch shall send the FFI SW\_ILS request Sequence to each of its AE\_Port ISLs after the ELP exchange has been successfully completed on that link.



**D.4.4.2 Map update process**

The AE Principal Switch shall also collect link status information from all AE Switches and distribute this information to all AE Switches using the Map Update process.

The AE Principal Switch receives Link Change Notification Flags and Problem Detected Notification Flags from downstream AE Switches via the FFI request Sequences. The AE Principal Switch collates this information, and then initiates a Map Update FFI request Sequence to all downstream Switches in the Avionics Fabric. When the AE Principal Switch initiates a new Map Update FFI request Sequence, it shall also increment the FFI Incarnation Number in order to uniquely identify the Map Update. When the FFI Incarnation Number reaches the value of "FFFF FFFF", it shall be allowed to roll over with the understanding that the value of "0000 0000" is greater than "FFFF FFFF".

Only the AE Principal Switch is allowed to change the FFI Incarnation Number. The lower AE Switches shall use their current FFI Incarnation Number to identify any change in status. The AE Principal Switch may combine several Link Change Notifications before a new Map Update FFI request Sequence is initiated.

The procedure for distributing the Domain Topology Map throughout the Avionics Fabric is fully described in clause D.4.5.

**D.4.4.3 AE Principal Switch update process**

It is permitted that more than one Switch have the capability to become the AE Principal Switch. Only one AE Principal Switch shall be active at any time.

The selection of the AE Principal Switch is beyond the scope of this Annex. Similarly, the selection of a replacement AE Principal Switch is beyond the scope of this Annex. In the case where an alternate AE Principal Switch is selected, this clause describes the mechanism for replacing the AE Principal Switch and for updating the AE Principal Uplink on each Switch.

Any AE Switch that is capable of becoming the AE Principal Switch shall be marked as such using the AE Secondary Principal Switch Flag in the FFI Link State Record in the Domain Topology Map. Any AE Switch that is capable of becoming the AE Principal Switch shall be capable of retaining or reacquiring its Domain Topology Map through a power cycling event.

The selection of a replacement AE Principal Switch is accomplished implicitly or through the usage of the FFI\_PSS ELS command. The FFI\_PSS ELS command is addressed to an AE Secondary Principal Switch.

If the AE Secondary Principal Switch receives a valid FFI\_PSS ELS command, it shall reply to the FFI\_PSS ELS with the LS\_ACC reply Sequence and then become the new AE Principal Switch. The new AE Principal Switch shall send an FFI request Sequence with the iPrincipal Updateî flag set, indicating that this FFI request Sequence is informing the Avionics Fabric of the change in the AE Principal Switch. The new FFI Incarnation number shall be considered valid. After each lower Switch resets its own links to Downlinks as appropriate, the lower Switch then sends a new FFI request Sequence with the Principal Update flag to its AE\_Port Downlinks. In this manner, all lower Switches will relearn their Uplinks and Downlinks in accordance with the new AE Principal Switch.

## D.4.5 FFI Domain Topology Map distribution

### D.4.5.1 Overview

Domain\_ID Assignment shall follow the requirements of this clause. This clause supersedes the Address Distribution procedure described in clause 7 of this document.

The Domain\_IDs and all of the AE\_Port Inter-Switch Links are defined in the Domain Topology Map. The Domain Topology Map is required to be sent in full in the payload of every FFI request Sequence. From the payload, any AE Switch that does not yet have a Domain\_ID can immediately obtain one. In addition, when any AE Switch initiates a new New FFI request Sequence, it is also required to identify in the payload the intended Domain\_ID recipient before beginning transfer of the Domain Topology Map. In this manner the recipient shall adopt or confirm its own Domain\_ID before processing the Domain Topology Map.

### D.4.5.2 FFI Domain Topology Map distribution state machine diagram

Figure D.2 and figure D.3 represent the FFI Domain Topology Map Distribution state machines for non-principal AE Switches and for the AE Principal Switch, respectively.

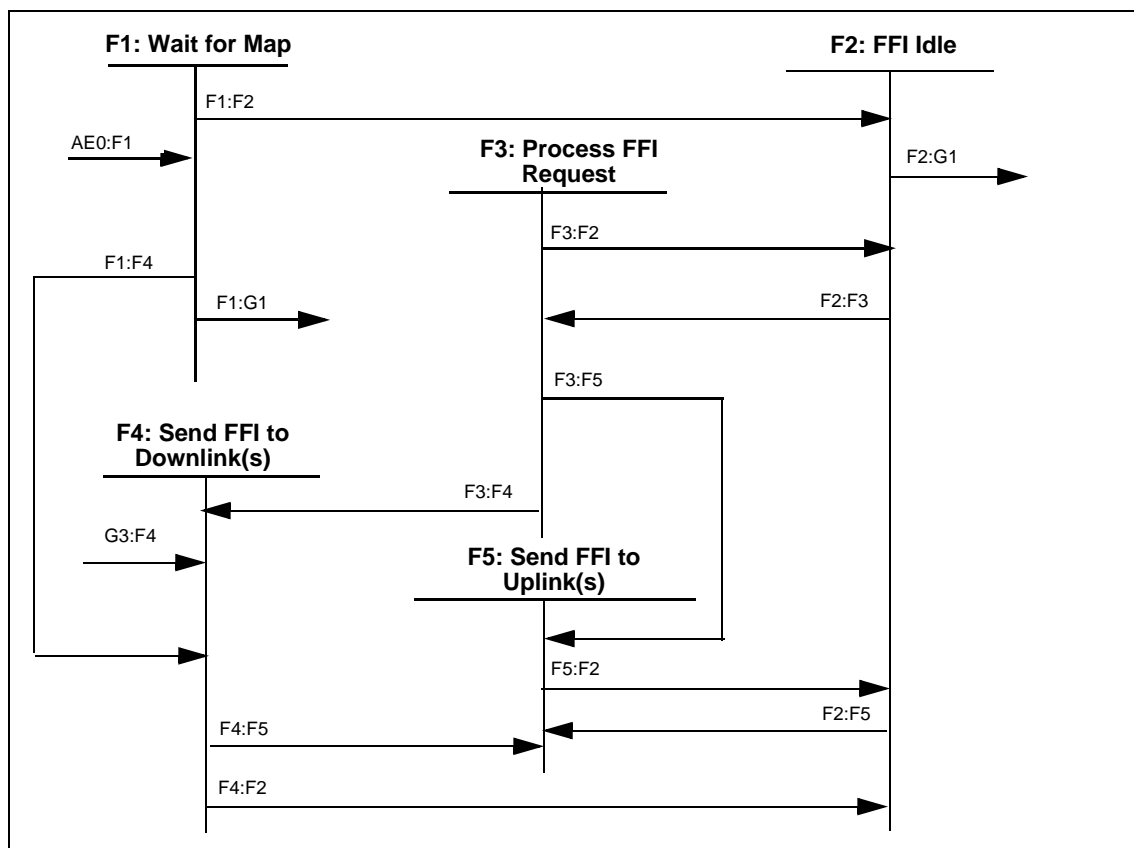


Figure D.2 – FFI Domain Topology Map distribution state machine, non-principal AE Switches




Figure D.3 – FFI Domain Topology Map distribution state machine, AE Principal Switch

**D.4.5.3 FFI Domain Topology Map distribution state machine text**

The following text represents the FFI Address Distribution State Machine for AE Switches. A prerequisite to any state in the FFI Address Distribution State Machine is that the AE Capable Switch must have at least one Active AE\_Port.

The FFI Address Distribution State Machine consists of two distinct sections. The "F" and "G" states are substates of State AE0. The "F" states shall be used when the AE Switch is not the AE Principal Switch. The "G" states shall be used when the AE Switch is the AE Principal Switch. If an AE Switch is not capable of becoming the AE Principal Switch, it is not required to implement the "G" states of the state machine.

**Transition AE0:F1.** This transition occurs after the AE\_Port Mode initialization has been completed on the first active AE\_Port and it is not pre-determined that the AE Switch is the AE Principal Switch.

**Transition AE0:G1.** This transition occurs after the AE\_Port Mode initialization has been completed on the first active AE\_Port and it is pre-determined that the AE Switch is the AE Principal Switch. In order to be the AE Principal Switch, the AE Switch must have an implicit Domain\_ID and an implicit Domain Topology Map, and the Domain Topology Map must indicate that its own Domain\_ID is designated as the AE Principal Switch.

**State F1: Wait for Map.** The AE Switch waits for the first FFI request Sequence to be received on one of its AE\_Ports, or for an implicit Domain Topology Map to be received that indicates that this AE Switch is the AE Principal Switch.

Upon entry to this state, the AE Switch shall declare all of its AE\_Ports to be Downlinks.

The received FFI request Sequence shall be checked for validity. This includes checking that either the Map Update Flag or the Override Incarnation Number Map Update Flag or the Override Domain Topology Map Update Flag or the Principal Update Flag is set. If the received FFI request Sequence is valid, the AE Switch declares this AE\_Port Inter-Switch Link (ISL) to be an Uplink for FFI. The AE Switch shall reply with the SW\_ACC.

If the AE Switch already has a Domain Topology Map, and an FFI request Sequence is received with either the Map Update Flag or Principal Update Flag or Override Incarnation Number Map Update Flag set, it shall compare the received Domain Topology Map to its own Domain Topology Map. If the Domain Topology Maps do not match, then the AE Switch shall reply with SW\_RJT and shall remain in the F1 state. The reason code shall be "Logical Error" and the reason code explanation shall be "Invalid Data".

If the AE Switch receives an FFI request Sequence with the Override Domain Topology Map Update Flag set, it shall accept the Domain Topology Map from the FFI request Sequence. The AE Switch shall reply with SW\_ACC.

Each FFI request Sequence that is received shall be checked for Map consistency. This includes checking the Originator Domain, Recipient Domain, Originator Port Index, and Recipient Port Index fields against the Domain Topology Map. If an inconsistency is detected, then the AE Switch shall send an SW\_RJT and shall remain in the F1 state. The reason code shall be "Logical Error" and the reason code explanation shall be "Invalid Data".

If the FFI request Sequence has the Link Change Notification Flag or Problem Detected Notification Flag set, then the AE Switch shall send an SW\_RJT and shall remain in the F1 state. The reason code shall be "Unable to Perform Command Request" and the reason code explanation shall be "Unable to Verify Connection".

In the case where the AE Switch receives a Domain Topology Map that indicates its FFI LSR has the AE Secondary Principal Switch Flag set, the AE Switch shall first verify that it can perform the role required by the AE Principal Switch. If not, then the AE Switch shall send an SW\_RJT and shall remain in the F1 state. The reason code shall be "Logical Error" and the reason code explanation shall be "Invalid Data".

In all cases where the FFI request Sequence is acceptable, the AE Switch shall set its Domain\_ID to the value in the FFI request Sequence. The AE Switch shall set its Domain Topology Map to be the map defined in the FFI request Sequence. The AE Switch shall set its FFI Incarnation Number to the value in the FFI request Sequence.

**Transition F1:F2.** This transition occurs after a SW\_ACC reply to a valid FFI request Sequence has been sent, and there are no other active AE\_Port Downlinks.

**Transition F1:F4.** This transition occurs after a SW\_ACC reply to a valid FFI request Sequence has been sent, and there is at least one active AE\_Port Downlink.

**Transition F1:G1.** This transition occurs when the AE Switch receives a Domain Topology Map that indicates that this AE Switch is the new AE Principal Switch. The AE Switch may learn this Domain Topology Map implicitly or by receiving a valid FFI\_DTM ELS Command from an external Nx\_Port.

**State F2: FFI Idle.** The AE Switch is responsible for continually monitoring the status of all of its AE\_Ports, and for maintaining the consistency of the Domain Topology Map database. In this state, the AE Switch monitors all of its ports and takes the appropriate action as they become Active AE\_Ports or Inactive AE\_Ports. Also in this state, the AE Switch monitors all of its AE\_Ports for any FFI request Sequences. Additionally, the AE Switch maintains the latest FFI Incarnation Number.

In the event that an AE Secondary Principal Switch receives a valid FFI\_PSS ELS Command, then the Switch shall set the AE Principal Update Flag before transitioning to state G1.

**Transition F2:F3.** This transition occurs when the AE Switch receives an FFI request Sequence on any of its AE\_Ports.

**Transition F2:F5.** This transition occurs when the AE Switch has determined that a link status change has occurred on at least one of its AE\_Ports.

**Transition F2:G1.** This transition occurs when the AE Switch that is designated as an AE Secondary Principal Switch receives a valid FFI\_PSS ELS Command.

**State F3: Process FFI Request** The AE Switch processes the FFI request Sequence that has been received on an AE\_Port. The following table describes the actions to be taken by the AE Switch.

Each FFI request Sequence that is received shall first and foremost be checked for Map consistency. This includes checking the Originator Domain, Recipient Domain, Originator Port Index, and Recipient Port Index fields against the Domain Topology Map.

Each FFI request Sequence that is received shall then be checked to insure that one and only one FFI Type Flag is set.

Other conditions detected for each FFI request Sequence and actions to be taken are described in table D.3

**Table D.3 – Actions taken by non-Principal AE Switch for an FFI request Sequence (Part 1 of 3)**

Condition detected	Actions taken
1. Map Update Flag set and Incarnation Number is incremented by one	Reply SW_ACC Mark link as Uplink Update AE Switch's local Incarnation Number Update AE Switch's local Domain Topology Map Create FFI request Sequence(s) using received payload Set Map Update Flag Transition F3:F4
2. Map Update Flag set and Incarnation Number matches	Reply SW_ACC Mark link as Uplink Transition F3:F2
3. Map Update Flag set and Incarnation Number is incremented by more than one	Reply SW_RJT <sup>a</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to Incarnation Number Error Set Problem Detected Notification Flag Transition F3:F5
<sup>a</sup> The reason code shall be "Protocol Error" with a reason code explanation of "Unable to Merge".	
<sup>b</sup> The reason code shall be "Logical Error" with a reason code explanation of "Invalid Data".	

**Table D.3 – Actions taken by non-Principal AE Switch for an FFI request Sequence (Part 2 of 3)**

Condition detected	Actions taken
4. Map Update Flag set and Incarnation Number is less than current	Reply SW_RJT <sup>a</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to Incarnation Number Error Set Problem Detected Notification Flag Transition F3:F5
5. Override Incarnation Number Map Update Flag set and Incarnation Number is different	Reply SW_ACC Mark link as Uplink Update AE Switch's local Incarnation Number Update AE Switch's local Domain Topology Map Create FFI request Sequence(s) using received payload Set Override Incarnation Number Map Update Flag Transition F3:F4
6. Override Incarnation Number Map Update Flag set and Incarnation Number matches	Reply SW_ACC Mark link as Uplink Transition F3:F2
7. Link Change Notification Flag set	Reply SW_ACC Mark link as Downlink Create FFI request Sequence(s) using received payload Set Link Change Notification Flag Transition F3:F5
8. Problem Detected Notification Flag set	Reply SW_ACC Mark link as Downlink Create FFI request Sequence(s) using received payload Copy FFI Problem Detected Reason Code from received payload Set Problem Detected Notification Flag Transition F3:F5
9. Principal Update Flag Set and Incarnation Number is different	Reply SW_ACC Mark link as Uplink Mark all other links as Downlinks Update AE Switch's local Incarnation Number Update AE Switch's local Domain Topology Map Create FFI request Sequence(s) using received payload Set Principal Update Flag Transition F3:F4
10. Principal Update Flag Set and Incarnation Number matches	Reply SW_ACC Mark link as Uplink Transition F3:F2
11. Override Domain Topology Map Update Flag Set and Incarnation Number is different	Reply SW_ACC Mark link as Uplink Update AE Switch's local Incarnation Number Update AE Switch's local Domain Topology Map Create FFI request Sequence(s) using received payload Set Override Domain Topology Map Update Flag Transition F3:F4
<p><sup>a</sup> The reason code shall be "Protocol Error" with a reason code explanation of "Unable to Merge".</p> <p><sup>b</sup> The reason code shall be "Logical Error" with a reason code explanation of "Invalid Data".</p>	

**Table D.3 – Actions taken by non-Principal AE Switch for an FFI request Sequence (Part 3 of 3)**

Condition detected	Actions taken
12. Override Domain Topology Map Update Set and Incarnation Number matches	Reply SW_ACC Mark link as Uplink Transition F3:F2
13. Map inconsistency	Reply SW_RJT <sup>b</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to Map Inconsistent Set Problem Detected Notification Flag Transition F3:F5
14. More than one flag set	Reply SW_RJT <sup>b</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to Multiple Flags Set Set Problem Detected Notification Flag Transition F3:F5
15. No flags set	Reply SW_RJT <sup>b</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to No Flags Set Set Problem Detected Notification Flag Transition F3:F5
16. Other	Reply SW_RJT <sup>b</sup> Create FFI request Sequence(s) using AE Switch's local Domain Topology Map as payload Set FFI Problem Detected Reason Code to Unexpected Condition Set Problem Detected Notification Flag Transition F3:F5
<p><sup>a</sup> The reason code shall be "Protocol Error" with a reason code explanation of "Unable to Merge".</p> <p><sup>b</sup> The reason code shall be "Logical Error" with a reason code explanation of "Invalid Data".</p>	

**Transition F3:F4.** This transition occurs when the processing of the required actions of the FFI request Sequence has been completed, and there is a need to forward new FFI request Sequences on all active AE\_Port Downlinks.

**Transition F3:F2.** This transition occurs when the processing of the FFI request Sequence has been completed, and there is no need to propagate a new FFI request sequence.

**Transition F3:F5.** This transition occurs when the processing of the required actions of the FFI request Sequence has been completed, and there is a need to forward new FFI request Sequences on all active AE\_Port Uplinks.

**State F4: Send FFI to Downlink(s):** The Switch initiates a separate FFI request Sequence in a new Exchange on each active AE\_Port Downlink and waits for a response or timeout from each active AE\_Port Downlink.

If the AE Switch determines that there are no Active AE\_Port Downlinks, then this state shall perform no action.

If one or more responses are not received within 2xE\_D\_TOV from when the last FFI request Sequence was sent, the AE Switch shall construct an FFI request Sequence with the Problem

Detected Notification Flag set. The AE Switch shall set the Problem Detected Reason Code to Downlink Timeout Occurred. The FFI Timeout Detected Flag shall be set in the appropriate Link Descriptor(s) in the Domain Topology Map.

**Transition F4:F2.** This transition occurs when the Switch has received an SW\_ACC or SW\_RJT response to every FFI request Sequence. This transition also occurs if there are no Active AE\_Port Downlinks.

**Transition F4:F5.** This transition occurs if one or more responses are not received within 2xE\_D\_TOV from the last FFI request Sequence for which a reply was not received.

**State F5: Send FFI Uplink(s):** The Switch initiates a separate FFI request Sequence in a new Exchange on each active AE\_Port Uplink and waits for a response or timeout from each active AE\_Port Uplink.

**Transition F5:F2.** This transition occurs when the Switch has received an SW\_ACC or SW\_RJT response to every FFI request Sequence. This transition also occurs if one or more responses are not received within 2xE\_D\_TOV from the last FFI request Sequence for which a reply was not received.

**State G1: Send Initial Map** The AE Principal Switch shall send an FFI request Sequence with the Override Incarnation Number Map Update Flag set or the Override Domain Topology Map Update Flag set on all active AE\_Ports. The AE Principal Switch shall wait for a reply or timeout from each active AE\_Port. All AE\_Ports links shall be set to Downlinks.

If the AE Switch determines that there are no Active AE\_Port Downlinks, then this state shall perform no action.

If a SW\_RJT reply is received on any AE\_Port ISL, then the FFI Reject Detected Flag shall be set in the FFI LSR Link Descriptor of the Domain Topology Map before transitioning.

If a timeout occurs (response is not received within 2xE\_D\_TOV) on any AE\_Port ISL, then the FFI Time-out Detected Flag shall be set in the appropriate Link Descriptor(s) of the Domain Topology Map before transitioning.

**Transition G1:G2.** This transition occurs when the Switch has received an SW\_ACC response to every FFI request Sequence. This transition also occurs if there are no Active AE\_Port Downlinks.

**Transition G1:G5.** This transition occurs when the Switch has received an SW\_RJT response to at least one of its FFI request Sequence(s), or if one or more responses are not received within 2xE\_D\_TOV from the last FFI request Sequence for which a reply was not received.

**State G2: FFI Principal Idle.** The AE Principal Switch monitors all of its ports and takes the appropriate action as they become Active AE\_Ports or Inactive AE\_Ports. Also in this state, the AE Switch monitors its AE\_Ports for any FFI request Sequences received.

**Transition G2:G3.** This transition occurs when the AE Principal Switch receives an FFI request Sequence.

**Transition G2:G5.** This transition occurs when the AE Principal Switch has detected a link status change on at least one of its own AE\_Ports.



**State G3: Process FFI Request** The AE Principal Switch processes an FFI request Sequence that has been received on an AE\_Port. Table D.4 describes the action to be taken by the AE Principal Switch on the AE\_Port that it received the FFI request Sequence.

**Table D.4 – Action taken by AE Principal Switch for an FFI request Sequence**

Condition Detected	Action taken
1. Link Change Notification Flag set	Reply SW_ACC, Update the master Domain Topology Map Transition G3:G5
2. Problem Detected Notification Flag set	Reply SW_ACC, Mark problems in the master Domain Topology Map <sup>a</sup> Transition G3:G5
3. Map Update Flag set	Reply SW_RJT <sup>b</sup> Transition G3:G2
4. Override Incarnation Number Map Update Flag set	Reply SW_RJT <sup>b</sup> Transition G3:G2
5. Override Domain Topology Map Update Flag set	Reply SW_RJT <sup>b</sup> Transition G3:G2
6. Principal Update Flag Set	Reply SW_ACC Mark link as Uplink Mark all other links as Downlinks Update AE Switch's local Incarnation Number Update AE Switch's local Domain Topology Map Create FFI request Sequence(s) using received payload Set Principal Update Flag Transition G3:F4
<sup>a</sup> Subsequent actions to be taken are implementation specific and are beyond the scope of this Annex.	
<sup>b</sup> The reason code shall be "Logical Error" with a reason code explanation of "Invalid Data".	

**Transition G3:G5.** This transition occurs when the processing of the FFI request Sequence has been completed, and a change has occurred to the Domain Topology Map.

**Transition G3:G2.** This transition occurs when the processing of the FFI request Sequence has been completed, and no change has occurred to the Domain Topology Map.

**Transition G3:F4.** This transition occurs when the AE Switch has received a valid FFI request Sequence with the Principal Update Flag set. This AE Switch is therefore no longer acting as the AE Principal Switch in the Avionics Fabric.

**State G4: Send Map Update.** The AE Principal Switch sends a new FFI request Sequence on all active AE\_Port Downlinks. The AE Principal Switch shall wait for a reply or timeout from each active AE\_Port.

If the AE Switch determines that there are no Active AE\_Port Downlinks, then this state shall perform no action.

The FFI request Sequence payload, the FFI Incarnation Number, and the FFI Type Flags that were previously set before this state are used to create the FFI request Sequence.

If a SW\_RJT reply is received on any AE\_Port ISL, then the FFI Reject Detected Flag shall be set in the FFI LSR Link Descriptor of the Domain Topology Map before transitioning.

If a timeout occurs on any AE\_Port ISL, then the FFI Timeout Detected Flag shall be set in the FFI LSR Link Descriptor of the Domain Topology Map before transitioning.

**Transition G4:G2.** This transition occurs when the Switch has received an SW\_ACC response to every FFI request Sequence. This transition also occurs if there are no Active AE\_Port Downlinks.

**Transition G4:G5.** This transition occurs when the Switch has received an SW\_RJT response to at least one of its FFI request Sequence(s), or if one or more responses are not received within 2xE\_D\_TOV from the last FFI request Sequence for which a reply was not received.

**State G5: Update Map Status.** The AE Principal Switch shall update the Domain Topology Map according to all of the link changes that have been detected or reported.

After all the changes are processed, the AE Principal Switch determines if it is necessary to distribute the Domain Topology Map to any registered Nx\_Ports. If any Nx\_Ports have registered for Map updates via the FFI\_MUR ELS Command, then the AE Principal Switch shall initiate an FFI\_RMUN Command to each of those Nx\_Ports.

After all the changes are processed, the AE Principal Switch determines if it is necessary to redistribute the Domain Topology Map to other AE Switches. If the AE Principal Switch has received a valid FFI\_SMU ELS Command, and has not subsequently received a valid FFI\_RMU ELS Command, then AE Principal Switch shall not redistribute the Domain Topology Map before making the transition to State G2.

If the AE Principal Switch has not received a valid FFI\_SMU ELS Command, or has last received a valid FFI\_RMU ELS Command, then the AE Principal Switch shall construct the FFI request Sequence payload using the latest FFI Domain Topology Map, increment the FFI Incarnation Number by one, and set the Map Update Flag before making the transition to state G4.

**Transition G5:G2.** This transition occurs when the AE Principal Switch determines that it will not update the Domain Topology Map to all of its active AE\_Ports.

**Transition G5:G4.** This transition occurs when the AE Principal Switch determines that it will update the Domain Topology Map to all of its active AE\_Ports.

## **D.4.6 Fast Fabric Initialization (FFI) SW\_ILS definition**

### **D.4.6.1 Overview**

The Fast Fabric Initialization Switch Internal Link Service (FFI SW\_ILS) provides a common mechanism for distributing the Domain Topology Map and AE\_Port link status throughout the Avionics Fabric. During initialization, the Domain Topology Map of the Avionics Fabric is distributed by the AE Principal Switch to all AE Switches in the Avionics Fabric using the FFI SW\_ILS request Sequence. After initialization, the FFI SW\_ILS request Sequence is used to communicate the link status of the Domain Topology Map within the Avionics Fabric.

#### **Protocol:**

The Fast Fabric Initialization (FFI) request Sequence  
Accept (SW\_ACC) reply Sequence

**Addressing:** For use in Fabric Configuration, the S\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Switch. The D\_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Switch.

**Payload:** The format of the FFI request Sequence payload is shown in table D.5.

**Table D.5 – FFI request payload**

Item	Size (bytes)
50000000h	4
Originator Switch Domain	1
Originator Port Index	1
Responder Switch Domain	1
Responder Port Index	1
FFI Incarnation Number	4
Reserved	2
FFI Type Flags	1
FFI Problem Detected Reason Code	1
Number of FFI Link State Records	4
FFI Link State Records	n

**Originator Switch Domain:** This field shall contain the Domain\_ID of the Switch that originated the FFI request Sequence.

**Originator Port Index:** This field shall contain the port index of the Switch that originated the FFI request Sequence.

**Responder Switch Domain:** This field shall contain the Domain\_ID of the neighboring Switch that responds to the FFI request Sequence.

**Responder Port Index:** This field shall contain the port index of the neighboring Switch that responds to the FFI request Sequence.

**FFI Incarnation Number:** This field contains the current incarnation of the FFI Domain Topology Map.

**FFI Type Flags:** This field shall define the type of FFI request Sequence. Each FFI request Sequence is required to set one and only one of these flags. The types are listed in table D.6.

**Table D.6 – FFI Type Flags definition**

Bit	Description
0	Map Update Flag
1	Override Incarnation Number Map Update Flag
2	Override Domain Topology Map Update Flag
3	Principal Update Flag
4	Link Change Notification Flag
5	Problem Detected Notification Flag
6-7	Reserved

The Map Update Flag is used to indicate that this FFI request Sequence is being sent downstream to update the link status of the system.

The Override Incarnation Number Map Update is used to indicate that the normal Incarnation Number checks on this Map Update are to be ignored on this downstream FFI request Sequence.

The Override Domain Topology Map Update Flag is used to indicate that this FFI request Sequence is being sent downstream to replace the previous Domain Topology Map of the Avionics Fabric.

The Principal Update Flag is used to indicate that the originator of this specific FFI request Sequence is a new AE Principal Switch, and that all Uplinks and Downlinks need to be recalculated.

The Link Change Notification Flag is used to indicate that this FFI request Sequence is being sent upstream in order to indicate that a change in link status of one or more AE\_Port ISLs has been detected. The change in link status is defined as an AE\_Port ISL that has become active or has become inactive.

The Problem Detected Notification Flag is used to indicate that this FFI request Sequence is being sent upstream in response to some error that has been detected in the Domain Topology Map.

**FFI Problem Detected Reason Code:** This field shall contain a code for the error that was detected in the FFI request Sequence. This field is not meaningful unless the Problem Detected Notification Flag is set. The following table defines the codes: The types are listed in table D.6.

**Table D.7 – FFI Problem Detected Reason Codes**

Value	Description
0	No information
1	Multiple Flag Bits detected
2	No Flag Bit detected
3	Invalid Incarnation Number detected
4	Downlink Timeout Occurred
5	Uplink Timeout Occurred
6	Map Inconsistency
7	Multiple AE Principal Switches detected
8	Unexpected Condition
All others	Reserved

**Number of FFI Link State Records:** This field shall specify the number of FFI Link State Records that follow this field.

**FFI Link State Records:** This field contains all of the individual FFI Link State Records that describe the Domain Topology Map of the Avionics Fabric. The format of the FFI Link State Record is described in clause D.4.6.2.

**Reply Switch Fabric Internal Link Service Sequence:**

- Service Reject (SW\_RJT)  
Signifies the rejection of the FFI request Sequence
- Accept (SW\_ACC)  
Signifies acceptance of the FFI request Sequence
- Accept payload

**Payload:** The format of the FFI accept payload is shown in table D.8

**Table D.8 – FFI accept payload**

Item	Size (bytes)
02000000h	4
Originator Switch Domain	1
Originator Port Index	1
Recipient Switch Domain	1
Recipient Port Index	1

**D.4.6.2 Fast Fabric Initialization Link State Record (FFI LSR) format**

**FFI LSR:** There is one format for the FFI LSR. The format is shown in table D.9. One or more FFI Link Descriptors may be contained in a single FFI LSR.

**Table D.9 – FFI Link State Record - Link Descriptor format**

Item	Size (bytes)
FFI Link State Record Identifier	1
FFI Link State Record Flags	1
Number of FFI Link Descriptors	2
Link Descriptor #1	4
...	4
...	4
Link Descriptor #n	4

**FFI Link State Record Identifier:** This field contains the Domain\_ID of the Switch that owns the FFI LSR.

**FFI Link State Record Flags:** This field shall contain the FFI Link State Record Flags to be used within the FFI request Sequence. Table D.10 defines the FFI LSR Flags:

**Table D.10 – FFI LSR Flags definition**

Bit	Description
0	FFI LSR Active Flag
1	FFI LSR Inconsistency Flag
2	AE Principal Switch Flag
3	AE Secondary Principal Switch Flag
4-7	Reserved

The FFI LSR Active Flag is used to indicate the current status of the particular AE Switch. When set to 1b, this indicates that the particular AE Switch is currently active and has correctly initialized. When set to 0b, this indicates that the particular AE Switch has not correctly initialized.

The FFI LSR Inconsistency Flag is used to indicate a conflict between two FFI request Sequences that have been received on different AE\_Ports. If the Domain\_ID of the topology does not properly match, or if the number and position of AE\_Port ISLs defined for a particular Domain are not consistent, then

this bit is set to indicate the location of the inconsistency. In addition, the AE Switch is required to forward this inconsistency by initiating an FFI request Sequence to its AE Principal Uplink.

The AE Principal Switch Flag is used to identify which Switch in the fabric is currently performing the role of the AE Principal Switch.

The AE Secondary Principal Switch Flag is used to identify which Switches in the system have the capability to perform the role of the AE Principal Switch.

**Number of FFI Link Descriptors:** This field specifies the number of FFI Link Descriptors contained in the FFI Link State Record.

**FFI Link Descriptor:** The format of the FFI Link Descriptor is described in clause D.4.6.3.

#### D.4.6.3 Fast Fabric Initialization Link Descriptor format

**FFI Link Descriptor:** The FFI Link Descriptor is a description of an individual AE\_Port to AE\_Port ISL within the Avionics Fabric. The format of the FFI Link Descriptor is shown in table D.11.

**Table D.11 – FFI Link Descriptor format**

Item	Size (bytes)
FFI Link Descriptor Flags	1
FFI Neighbor Link ID	1
Output Port Index	1
Neighbor Port Index	1

**FFI Link Descriptor Flags:** This field shall contain the link descriptor flags to be used within the FFI request Sequence. Table D.12 defines the FFI Link Descriptor flags:

**Table D.12 – FFI Link Descriptor Flags definition**

Bit	Description
0	FFI Link Active Flag
1	FFI Timeout Detected Flag
2	FFI Reject Detected Flag
3-7	Reserved



The FFI Link Active Flag is used to indicate the current status of the particular ISL. When set to 1b, this indicates that the particular ISL is currently active and has correctly initialized. When set to 0b, this indicates that the particular ISL is either inactive or has not yet correctly initialized.

The FFI Timeout Detected Flag is used to indicate that an FFI request Sequence has not received the required SW\_ACC or SW\_RJT within the specified timeout period.

The FFI Reject Detected Flag is used to indicate that an SW\_RJT response was received during the initialization process of that specific ISL.

**FFI Neighbor Link Identifier:** This field identifies the link and contains the Domain\_ID of the neighbor Switch at the other end of the ISL, relative to the owning Switch.

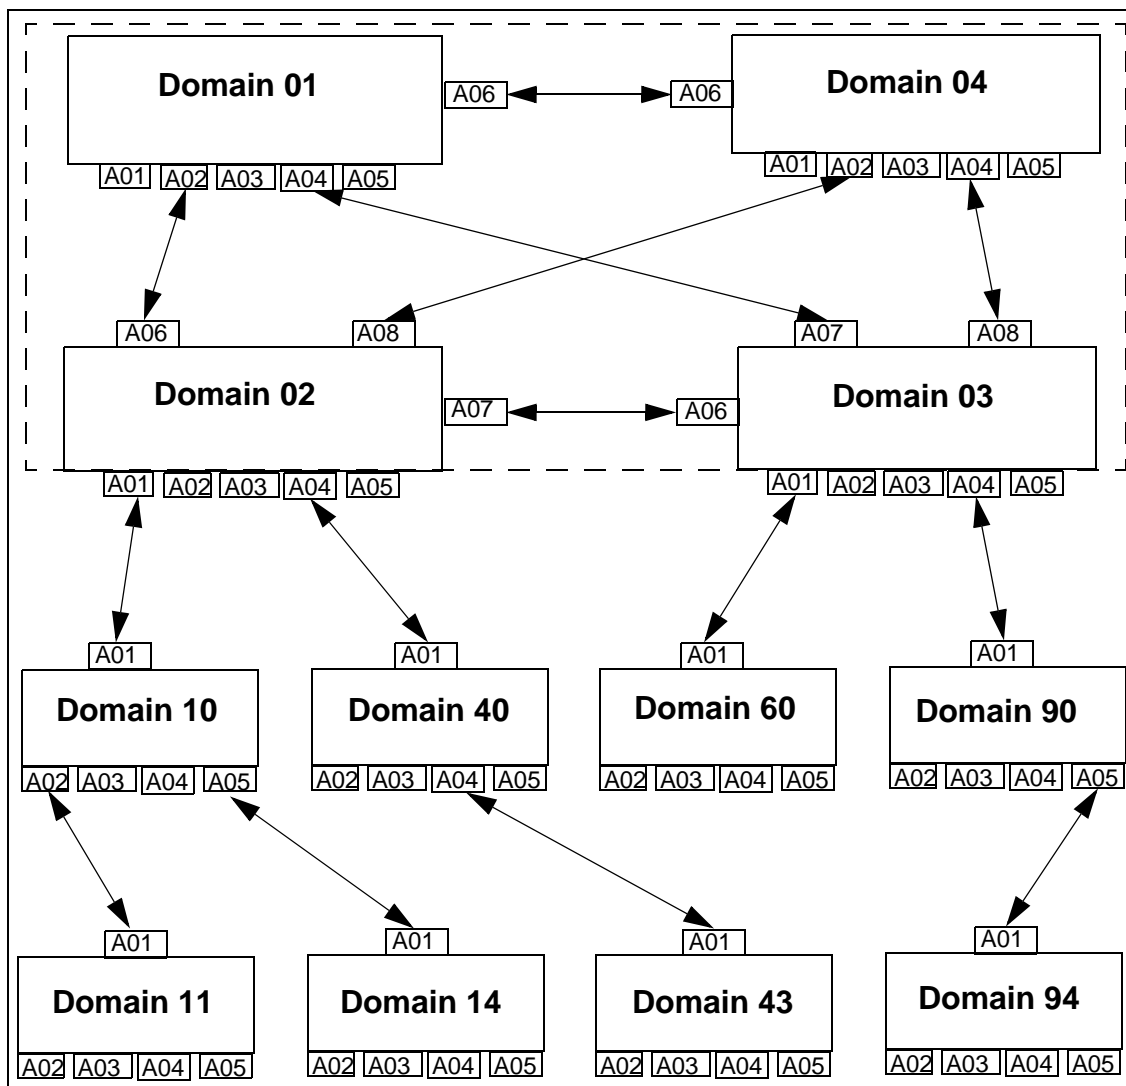
**Output Port Index:** This field shall specify the source AE\_Port Index.

**Neighbor Port Index:** The field shall specify the destination AE\_Port Index.

## **D.5 FFI Domain Topology Map Distribution (Informative)**

### **D.5.1 Sample configuration**

Figure D.4 shows an example Avionics Fabric. All the Switches shown are AE-Capable Switches. In order to improve readability, some of the SW\_ILS reply Sequences (i.e., SW\_ACC or SW\_RJT with Reason Codes) and all of the Acknowledgment frames have been omitted.



**Figure D.4 – Example Avionics Fabric**

In this example, Domains 01, 02, 03 and 04, know their Domain\_IDs implicitly upon power up. All other Switches do not know their domains upon power up. Domain 01 is the designated AE Principal Switch.

### D.5.2 Initialization procedure example

The example begins with Domains 01, 02, 03 and 04 being powered up simultaneously. The other Switches will be powered up later. These four Switches will complete Switch port mode initialization implicitly or explicitly (i.e., it does not matter for this example) and the inter-connected Switch ports arrive at State AE0 (AE\_Port). The next phase of FFI, Domain Topology Map distribution, can then begin.

Domains 02, 03 and 04 transition to State F1 (Wait for Map). All of their connected AE\_Ports are initially defined to be Downlinks.

Domain 01 (the AE Principle Switch) transitions to State G1 (Send Initial Map). Domain 01 then sends out the “master” Domain Topology Map with the Incarnation Number set to 1 to Domains 02, 03 and 04 via three separate FFI request Sequences. Domains 02, 03, and 04 receive an FFI request Sequence, mark those Inter-Switch Links (ISLs) as Uplinks, store the Incarnation Number, and compare the received Domain\_ID and Topology Map with their implicitly defined Domain\_ID and Topology Map. If there are no discrepancies, Domains 02, 03, and 04 send back an SW\_ACC to Domain 01, otherwise they send back an SW\_RJT.

If there are no discrepancies in the received Domain Topology Map (i.e., “the Map”), Domains 02, 03, and 04 then send the Map out on all of their Active AE\_Port Downlinks via FFI request Sequences. Because of this, Domains 02, 03, and 04 may receive multiple copies of the Map. If this occurs, the recipient marks these links as Uplinks and compares the received Map with its own Map. If there are no errors and the received Incarnation Number is the same as the recipient’s current Incarnation Number, the recipient replies with an SW\_ACC. If there are any discrepancies, the recipient responds with an SW\_RJT.

Now assume that power is applied to the AE-Capable Switches labeled Domain 10, Domain 11, and Domain 14 simultaneously. These Switches have no prior knowledge of their Domain\_IDs or the Topology Map so adjacent AE\_Ports exchange parameters and go through Switch port mode initialization explicitly and arrive at State AE0. The AE\_Ports are all initialized as Downlinks and these three Switches transition to State F1 (Wait for Map).

Domain 02, sensing a change of status on its ISL to the Switch at Domain 10, sends out an FFI request Sequence on all of its Active AE\_Port Uplinks with:

1. the Link Change Notification Flag set;
2. the FFI LSR Active Flag set for this Link State Record (i.e., Domain); and
3. the FFI Link Active Flag set for this specific Link Descriptor.

One copy of this FFI request Sequence will reach Domain 01 (the AE Principle Switch) directly. Other copies may arrive indirectly via Domains 03 and 04.

Domain 01 receives this FFI request Sequence, updates the master Domain Topology Map to show that this Domain is now active and that the ISL between Domain 02, Port Index 01 and Domain 10, Port Index 01 is now active. Domain 01 increments the Incarnation Number by 1 and sends out the new Map on all of its Active AE\_Port Downlinks.

Domain 02 receives the FFI request Sequence with the new Map and Incarnation Number, updates its internal Map and Incarnation Number and proceeds to send the Map to all of its Active AE\_Port Downlinks (in this case just Switch at Domain 10) via another FFI request Sequence.

When the Switch at Domain 10 receives the Map, it marks the ISL as an Uplink, adopts the Domain\_ID of 10, stores the Map and Incarnation Number (without question since there is nothing to compare it with yet) and responds with SW\_ACC.

Now that Domain 10 is active and has an Uplink, it is obligated to report the change in status on its ISLs to the Switches at Domain 11 and Domain 14. An FFI request Sequence will be sent from Domain 10 up to Domain 02 followed by another FFI request Sequence from Domain 02 up to Domain 01 with the appropriate Flags set to indicate that these Switches and links are now active. As before, Domain 01 will update the master Domain Topology Map, increment the Incarnation Number by 1, and send the new map out on all of its Active AE\_Port Downlinks. The new map is received by Domain 02, then by Domain 10, and finally by the Switches at Domain 11 and Domain 14.

This procedure is repeated whenever a new AE-Capable Switch joins the Avionics Fabric.

A similar procedure is followed whenever an AE\_Switch or one of its ISLs become inactive. In that case, the Link Change Notification Flag would be set on the FFI request Sequence, and the corresponding FFI LSR Active Flag and/or the FFI Link Active Flag would not be set.

### **D.5.3 AE Principal Switch update example**

This example describes an orderly transfer of the role of AE Principal Switch from one Switch to another.

In this example, assume the Avionics Fabric described above has been fully initialized and is stable. Assume that a problem has been identified with Domain 01 and Domain 04 must now become the new AE Principal Switch. Note that the Domain Topology Map does not change when a new AE Principal Switch is selected. Also note that Domain 04 must have previously indicated that it is capable of being the AE Principal Switch by having the AE Secondary Principal Switch Flag set in its FFI Link State Record.

The process begins with Domain 04 receiving an updated Domain Topology Map that has its own FFI Link State Record marked with the AE Principal Switch Flag, or by receiving a valid FFI\_PSS ELS Command from an external Nx\_Port. Domain 04 immediately updates its own map, adopts the role of AE Principal Switch, and proceeds to State G1 (Send Initial Map). At this time Domain 04 redefines all of its ISLs to be Downlinks and then sends out an FFI request Sequence with the new Map, the Incarnation Number set to 1, and the AE Principal Update Flag set.

Similar to what was done in the previous example at Initialization, Domains 02 and 03 will receive the FFI request Sequence from Domain 04, see that the AE Principal Update Flag is set, and perform the following actions:

1. redefine these ISLs to be Uplinks and their other ISLs to be Downlinks;
2. store the Incarnation Number; and
3. compare the received Domain\_ID and Topology Map with their implicitly defined Domain\_ID and Topology Map.

If there are no discrepancies in the received Map, Domains 02 and 03 send back an SW\_ACC to Domain 04; otherwise, they send back an SW\_RJT. If there are no discrepancies, Domains 02 and 03 will then send the Map out, with the AE Principal Update Flag set, on all of their Active AE\_Port Downlinks via FFI request Sequences.

The Switch at Domain 01, if active, will see the FFI request Sequence from Domain 04 with the AE Principle Update Flag set indicating that Domain 04 has become the new AE Principal Switch. Note that this is the only Update Flag that can be accepted by the AE Principle Switch. Upon receiving this FFI request Sequence, Domain 01 will revert to a non-principal AE Switch. It will declare this ISL to be an Uplink and propagate the FFI request Sequence to all of its Active AE\_Port Downlinks.

When the other AE Switches receive the new Map with the AE Principal Update Flag set, they redefine the receiving ISL to be an Uplink and propagate the Map to all of their Downlinks.

This process continues until all Switches have been informed that the AE Principal Switch has been changed.