



InterNational Committee for Information Technology Standards (INCITS)

Secretariat: Information Technology Industry Council (ITI)

700 K Street NW, Suite 610, Washington, DC 20001

www.INCITS.org



eb-2024-00128

Document Date: 03/25/2024

To: INCITS Members

Reply To: Lbarra@itic.org

Subject: **Public Review and Comments Register for the Approval of:**

INCITS 569-202x - Information technology - Fibre Channel – Link Services – (FC-LS-5)

Due Date: The public review is from April 5, 2024 to June 4, 2024.

Action: The InterNational Committee for Information Technology Standards ([INCITS](http://www.INCITS.org)) announces that the subject-referenced document(s) is being circulated for a 60-day public review and comment period. Comments received during this period will be considered and answered. Commenters who have objections/suggestions to this document should so indicate and include their reasons.

All comments should be forwarded not later than the date noted above to the following address:

INCITS Secretariat/ITI
700 K Street NW - Suite 600
Washington DC 20001
Email: comments@standards.incits.org (preferred)

This public review also serves as a call for patents and any other pertinent issues (copyrights, trademarks). Correspondence regarding intellectual property rights may be emailed to the INCITS Secretariat at patents@itic.org.

FIBRE CHANNEL

LINK SERVICES - 5
(FC-LS-5)

REV 5.23

INCITS working draft proposed
American National Standard
for Information Technology

December 7, 2023

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.

POINTS OF CONTACT:

Howard Johnson (T11 Chair) Broadcom Inc 1320 Ridder Park Drive San Jose, CA 95131 408-433-8000 howard.johnson@broadcom.com	Roger G. Hathorn (T11.3 Chair) IBM Corporation 9032 S. Rita Road Tucson, AZ 85744 502-799-5950 rhathorn@us.ibm.com
---	---

Patty Driever (FC-LS-5 Chair) IBM Corporation 2455 South Road Poughkeepsie, NY 12601 845-435-6442 pgd@us.ibm.com	David Peterson (FC-LS-5 Editor) Broadcom Inc 1320 Ridder Park Drive San Jose, CA 95131 408-433-8000 david.peterson@broadcom.com
---	--

Release Notes for revision 5.23:

- Incorporated T11-2023-00040-v010.

Release Notes for revision 5.22:

- Incorporated T11-2023-00040-v009.

Release Notes for revision 5.22:

- Incorporated T11-2023-00040-v009.

Release Notes for revision 5.21:

- Incorporated T11-2023-00040-v008.

Release Notes for revision 5.2:

- Incorporated T11-2022-00120-v004.
- Incorporated T11-2022-00314-v006.

Release Notes for revision 5.1:

- Incorporated T11-2021-00015-v002.
- Incorporated T11-2021-00121-v000.
- Incorporated T11-2021-00125-v004.
- Incorporated T11-2021-00135-v006.
- Incorporated T11-2021-00012-v003.
- Incorporated T11-2022-00031-v002.

Release Notes for revision 5.01:

- Incorporated T11-2020-00013-v004.
- Incorporated T11-2020-00085-v000.
- Incorporated T11-2020-00109-v003.

Release Notes for revision 5.00:

- Initial version.
- Incorporated T11-2019-00098-v002.

American National Standard
for Information Technology

**Fibre Channel —
Link Services - 5 (FC-LS-5)**

Secretariat

Information Technology Industry Council

Approved (not yet approved)

American National Standards Institute, Inc.

Abstract

This standard specifies Extended Link Services for Fibre Channel. The framing and signaling requirements are specified in the Fibre Channel - Framing and Signaling standard and the physical interface requirements are specified in the Fibre Channel - Physical Interface family of standards. This standard is recommended for new implementations, but does not obsolete the existing Fibre Channel 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.

The user's attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more such claims has been received.

By publication of this standard, no position is taken with respect to the validity of any such claim(s) or of any patent rights in connection therewith. If a patent holder has filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license, then details may be obtained from the standards developer.

Published by

American National Standards Institute
25 West 43rd Street, 4th Floor, New York, NY 10036

Copyright © 20xx 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, 700 K St NW, Suite 600 Washington, D.C. 20001.

Printed in the United States of America

Foreword (This Foreword is not part of American National Standard INCITS 569-202x.)

The Fibre Channel Link Services (FC-LS-5) standard describes in detail the Fibre Channel Link Services introduced in FC-FS-6. In addition, this document describes any ancillary functions and services required to support the Fibre Channel Link Services.

This standard was developed by the INCITS Fibre Channel T11 Technical Committee of Accredited Standards Committee INCITS. The standards approval process started in 2023. This document includes annexes which 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, 700 K St NW, Suite 600, Washington, D.C. 20001.

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 voting members:

Howard Johnson, Chair
Patty Driever, Vice-Chair
Chris Lyon, Secretary

Company	Name
---------	------

Task Group T11.3 on Fibre Channel Protocols, which developed and reviewed this standard, had the following voting members:

Roger Hathorn, Chair
TBD, Vice-Chair
Patty Driever, Secretary

Company	Name
---------	------

Contents	Page
1 Scope	1
2 Normative References	1
2.1 Overview	1
2.2 Approved references	1
2.3 References under development	2
2.4 Other references	2
3 Definitions and conventions	4
3.1 Overview	4
3.2 Definitions	4
3.3 Editorial Conventions	15
3.4 Abbreviations and acronyms	16
3.5 Symbols	18
3.6 Keywords	18
3.7 T10 Vendor ID fields	19
4 Extended Link Services	20
4.1 Introduction	20
4.2 Link Service tag, length, value (TLV) descriptors	21
4.2.1 Overview	21
4.2.2 Link Service request with Link Service TLV descriptor payload example	21
4.2.3 Link Service accept with Link Service TLV descriptor payload example	22
4.2.4 Link Service TLV descriptors	23
4.3 Extended Link Service requests	25
4.3.1 Introduction	25
4.3.2 Sequence and Exchange management	30
4.3.3 Advise Credit (ADVC)	31
4.3.4 Echo (ECHO)	34
4.3.5 Estimate Credit (ESTC)	35
4.3.6 Establish Streaming (ESTS)	36
4.3.7 Login (FLOGI/PLOGI)	37
4.3.8 Logout (LOGO)	38
4.3.9 Read Link Error Status Block (RLS)	40
4.3.10 Read Timeout Value (RTV)	41
4.3.11 Reinstate Recovery Qualifier (RRQ)	42
4.3.12 Request Sequence Initiative (RSI)	44
4.3.13 Test (TEST)	45
4.3.14 Fabric Address Notification (FAN)	46
4.3.15 Loop Initialize (LINIT)	47
4.3.16 Loop Status (LSTS)	49
4.3.17 Registered State Change Notification (RSCN)	51
4.3.18 State Change Registration (SCR)	56
4.3.19 Process Login (PRLI)	58
4.3.20 Process logout (PRLO)	62
4.3.21 Test Process Login State (TPLS)	66
4.3.22 Request Node Identification Data (RNID)	69
4.3.23 Registered Link Incident Report (RLIR)	77
4.3.24 Link Incident Record Registration (LIRR)	84
4.3.25 Discover N_Port/Service Parameters (PDISC)	87
4.3.26 Discover F_Port Service Parameters (FDISC)	87
4.3.27 Discover Address (ADISC)	88

4.3.28	Third Party Process Logout (TPRLO)	91
4.3.29	Clock Synchronization Request (CSR)	93
4.3.30	Clock Synchronization Update (CSU)	97
4.3.31	Report Port Buffer Conditions (RPBC)	98
4.3.32	Report node FC-4 Types (RNFT)	99
4.3.33	Scan Remote Loop (SRL)	102
4.3.34	Set Bit-error Reporting Parameters (SBRP)	103
4.3.35	Query Security Attributes (QSA) Version 2	106
4.3.36	Report Port Speed Capabilities (RPSC)	108
4.3.37	Read Exchange Concise (REC)	111
4.3.38	Exchange Virtual Fabrics Parameters (EVFP)	113
4.3.39	Link Keep Alive (LKA)	119
4.3.40	Link Cable Beacon (LCB)	120
4.3.41	Registered Fabric Change Notification (RFCN)	123
4.3.42	Define FFI Domain Topology Map (FFI_DTM)	124
4.3.43	Request FFI Domain Topology Map (FFI_RTM)	126
4.3.44	FFI AE Principal Switch Selector (FFI_PSS)	127
4.3.45	FFI Map Update Registration (FFI_MUR)	128
4.3.46	FFI Registered Map Update Notification (FFI_RMUN)	130
4.3.47	FFI Suspend Map Updates (FFI_SMU)	131
4.3.48	FFI Resume Map Updates (FFI_RMU)	132
4.3.49	Read Diagnostic Parameters (RDP)	133
4.3.50	Query Fabric Priority Allocation (QFPA)	150
4.3.51	Update VE Mappings (UVEM)	152
4.3.52	Exchange Diagnostic Capabilities (EDC)	154
4.3.53	Register Diagnostic Functions (RDF)	159
4.3.54	Fabric Performance Impact Notifications (FPIN)	162
4.4	Extended Link Service Reply Sequences	170
4.4.1	Overview	170
4.4.2	LS_ACC	171
4.4.3	Reply Sequence	171
4.4.4	Link Service Reject (LS_RJT)	171
5	FC-4 Link Service	176
6	Login and Service Parameters	177
6.1	Introduction	177
6.2	Fabric Login	177
6.2.1	Introduction	177
6.2.2	Explicit Fabric Login Procedure	178
6.2.3	SOFs	180
6.2.4	Frequency	181
6.2.5	Fabric Login completion - Originator	181
6.2.6	Fabric Login completion - Responder	181
6.3	N_Port Login	181
6.3.1	Introduction	181
6.3.2	Explicit N_Port Login	182
6.3.3	SOFs	185
6.3.4	Frequency	185
6.3.5	N_Port Login completion - Originator	185
6.3.6	N_Port Login completion - Responder	186
6.4	Logout	186
6.4.1	Introduction	186
6.4.2	Explicit N_Port Logout Procedure	186

6.4.3	Explicit Fabric Logout Procedure	186
6.4.4	Implicit Logout	187
6.4.5	The effects of FLOGI, FDISC, and LOGO	188
6.5	Extended Login Processing	189
6.6	Service Parameters	190
6.6.1	ELS and LS_ACC Payload	190
6.6.2	Common Service Parameters	192
6.6.3	Port_Name	202
6.6.4	Node_Switch_Fabric_Name	202
6.6.5	Class Service Parameters	203
6.6.6	Auxiliary Parameter Data	214
6.6.7	Vendor Version Level	216
6.6.8	Services Availability	216
6.6.9	Login Extension	217
6.6.10	Clock Synchronization Quality of Service	218
7	Process Login/Logout	223
7.1	Process Login	223
7.1.1	Introduction	223
7.1.2	Mode of operation	224
7.2	Process Logout	225
8	Virtual Fabrics PN_Port Support	227
8.1	Overview	227
8.2	Enabling VFT Tagging on PN_Ports	227
8.3	Configuration Information	230
8.4	Exchange Virtual Fabrics Parameters Processing	230
8.4.1	Overview	230
8.4.2	Changing the VFT Tagging Mode	233
8.4.3	Adding or Removing Virtual Fabrics	233
8.4.4	Changing the Port VF_ID	234
9	Priority Tagging	235
9.1	Overview	235
9.2	Operations	236
9.3	Scalability	237
Annex A:	Fabric notification information and examples (informative)	238
A.1	Overview	238
A.2	Causes of congestion	238
A.3	General use and device behaviors	239
A.3.1	Overview	239
A.3.2	Port name lists	239
A.3.3	FPIN Generation and Transmission	239
A.3.4	Detecting ports and attached ports	243
A.3.5	Congestion events	244
A.4	Device participation	245
A.4.1	Overview	245
A.4.2	Registration only	245
A.4.3	Logging	245
A.4.4	Device behavior modifications	245
A.5	Examples	247
A.5.1	Overview	247
A.5.2	Improving multi-path solutions	247

A.5.3 Eliminating Target Credit Stall	248
A.5.4 Speed matching to minimize oversubscription	249
A.5.5 Handling first burst congestion spikes	249

Figure		Page
Figure 1 –	Illustration of parameters	104
Figure 2 –	LS_RJT format	172
Figure 3 –	Image pairs	223
Figure 4 –	PN_Port Initialization of Virtual Fabrics	228
Figure 5 –	A Generic EVFP Transaction	231
Figure 6 –	Example of Relationship among Local VE IDs, N_Port_IDs, and PN_Ports.	235
Figure 7 –	Example of Priority Tagging of the Traffic of a VE	236
Figure 8 –	Basic port name list example	240
Figure 9 –	Typical port name list example	241
Figure 10 –	NPIV port name list example.	242
Figure 11 –	ISL port name list example	243

Table	Page
Table 1 – Comparison of numbering conventions	15
Table 2 – Extended Link Services Routing Bits and Information Categories	20
Table 3 – Link Service TLV descriptor format	21
Table 4 – Link Service request with Link Service TLV descriptor payload example	21
Table 5 – Link Service accept with Link Service TLV descriptor payload example	22
Table 6 – Link Service TLV descriptors	23
Table 7 – Link Service Request Information descriptor format	24
Table 8 – N_Port_ID descriptor format	25
Table 9 – ELS_Command codes	26
Table 10 – Responses to Received ELSs	30
Table 11 – ADVC Payload	32
Table 12 – ADVC LS_ACC Payload	33
Table 13 – ECHO Payload	34
Table 14 – ECHO LS_ACC Payload	35
Table 15 – ESTC Payload	35
Table 16 – ESTS Payload	36
Table 17 – ESTS LS_ACC Payload	37
Table 18 – LOGO Payload	39
Table 19 – LOGO LS_ACC Payload	39
Table 20 – RLS Payload	40
Table 21 – RLS LS_ACC Payload	41
Table 22 – RTV Payload	41
Table 23 – RTV LS_ACC Payload	42
Table 24 – RRQ Payload	43
Table 25 – RRQ LS_ACC Payload	44
Table 26 – RSI Payload	45
Table 27 – RSI LS_ACC Payload	45
Table 28 – TEST Payload	46
Table 29 – FAN Payload	47
Table 30 – LINIT Payload	48
Table 31 – Initialization Function	48
Table 32 – LINIT LS_ACC Payload	48
Table 33 – LINIT Status	49
Table 34 – LSTS Payload	49
Table 35 – LSTS LS_ACC Payload	50
Table 36 – Loop State	51
Table 37 – RSCN Payload	53
Table 38 – Generic affected Port_ID page	54
Table 39 – RSCN Event Qualifier values	55
Table 40 – Address Format	56
Table 41 – RSCN LS_ACC Payload	56
Table 42 – SCR Payload	57
Table 43 – Registration Function Bitmap	57
Table 44 – SCR LS_ACC Payload	58
Table 45 – PRLI Payload	58
Table 46 – PRLI service parameter page format	59
Table 47 – PRLI LS_ACC Payload	60
Table 48 – PRLI LS_ACC service parameter response page format	61
Table 49 – PRLI accept response code	62
Table 50 – PRLO Payload	63
Table 51 – PRLO logout parameter page format	63
Table 52 – PRLO LS_ACC Payload	64

Table 53 – PRLO LS_ACC logout parameter response page format	65
Table 54 – PRLO accept response code	66
Table 55 – TPLS Payload	67
Table 56 – TPLS image pair ID page format.	67
Table 57 – TPLS LS_ACC Payload	68
Table 58 – TPLS response page format	68
Table 59 – TPLS accept response code	69
Table 60 – RNID Payload	70
Table 61 – Node Identification Data Format	70
Table 62 – RNID Accept Payload	71
Table 63 – Node Identification Data Format	71
Table 64 – VN_Port Phy Type.	72
Table 65 – Common Identification Data	72
Table 66 – General Topology Specific Identification Data	73
Table 67 – Associated Type	74
Table 68 – Multi-function device bit definitions	75
Table 69 – Node Management	76
Table 70 – IP Version	76
Table 71 – RLIR Payload	78
Table 72 – Link Incident Record Format	79
Table 73 – Common Link Incident Record Data	80
Table 74 – Time Stamp Format values	81
Table 75 – Incident Qualifier	81
Table 76 – Incident Code values.	83
Table 77 – RLIR LS_ACC Payload	84
Table 78 – LIRR Payload	85
Table 79 – Registration Function	86
Table 80 – Link Incident Record-Registration Format.	86
Table 81 – LIRR LS_ACC Payload	87
Table 82 – ADISC Request payload	89
Table 83 – ADISC LS_ACC Payload	90
Table 84 – Response summary to FDISC/PDISC.	91
Table 85 – TPRLO Payload	92
Table 86 – TPRLO logout parameter page	92
Table 87 – TPRLO LS_ACC Payload	93
Table 88 – CSR Payload.	94
Table 89 – CSR Clock Sync Mode Meaning.	94
Table 90 – CSR LS_ACC Payload	95
Table 91 – CSU Clock Sync Mode Meaning.	96
Table 92 – CSU Payload.	97
Table 93 – Clock Count Field Meaning	97
Table 94 – RPBC Payload	98
Table 95 – ELS Buffer Parameters Field.	99
Table 96 – RPBC LS_ACC Payload	99
Table 97 – RNFT Payload.	100
Table 98 – RNFT LS_ACC Payload	101
Table 99 – RNFT FC-4 Entry	101
Table 100 – SRL Payload	102
Table 101 – Flag field definitions.	103
Table 102 – SRL LS_ACC Payload	103
Table 103 – SBRP Payload	104
Table 104 – SBRP LS_ACC Payload.	106
Table 105 – QSA Request Payload	107
Table 106 – Enforced Security Attribute Registration Mask	107

Table 107 – Extended Security Attribute Registration Mask	107
Table 108 – QSA LS_ACC Payload	108
Table 109 – Enforced Security Attribute Field	108
Table 110 – Extended Security Attribute Field	108
Table 111 – RPSC Payload	109
Table 112 – RPSC LS_ACC Payload	109
Table 113 – Port Speed Capabilites	110
Table 114 – Port Operating Speed	110
Table 115 – REC Payload.	112
Table 116 – REC LS_ACC Payload	112
Table 117 – EVFP Request Payload.	113
Table 118 – EVFP Message Codes	113
Table 119 – EVFP Accept Payload	114
Table 120 – LS_RJT Reason Codes for EVFP	115
Table 121 – EVFP_SYNC Message Payload	115
Table 122 – Descriptor Format	116
Table 123 – Descriptor Control Codes	116
Table 124 – Descriptor Types	116
Table 125 – Tagging Administrative Status descriptor	117
Table 126 – Administrative Tagging Modes	117
Table 127 – Tagging Mode Negotiation	117
Table 128 – Port VF_ID descriptor	118
Table 129 – Locally-Enabled VF_ID List descriptor	118
Table 130 – Vendor Specific descriptor	119
Table 131 – LKA Payload	120
Table 132 – LKA LS_ACC Payload.	120
Table 133 – LCB Payload	121
Table 134 – Subcommand field functions	121
Table 135 – Capability field functions	122
Table 136 – Status field functions	122
Table 137 – LCB LS_ACC Payload	123
Table 138 – RFCN Request Payload	124
Table 139 – Change Flags	124
Table 140 – FFI_DTM Payload	125
Table 141 – FFI_DTM LS_ACC Payload	125
Table 142 – FFI_RTM Payload	126
Table 143 – FFI_RTM LS_ACC Payload	127
Table 144 – FFI_PSS Payload	128
Table 145 – FFI_PSS LS_ACC Payload.	128
Table 146 – FFI_MUR Payload.	129
Table 147 – Registration Function	129
Table 148 – FFI_MUR LS_ACC Payload	129
Table 149 – FFI_RMUN Payload	130
Table 150 – FFI_SMU Payload	131
Table 151 – FFI_SMU LS_ACC Payload	132
Table 152 – FFI_RMU Payload.	132
Table 153 – FFI_RMU LS_ACC Payload	133
Table 154 – RDP Payload.	134
Table 155 – RDP LS_ACC Payload	135
Table 156 – Port Speed descriptor	136
Table 157 – Link Error Status Block descriptor	136
Table 158 – Px_Port Phy Type	137
Table 159 – Port Names descriptor.	137
Table 160 – SFP Diagnostic Parameters descriptor	137

Table 161 – SFP Flags	139
Table 162 – QSFP Diagnostic Parameters descriptor	140
Table 163 – QSFP Flags	141
Table 164 – FEC Status descriptor	142
Table 165 – Buffer Credit descriptor	142
Table 166 – Optical Element Data descriptor	143
Table 167 – Optical Element Function Flags	143
Table 168 – Optical Element Type Transgression Flags	144
Table 169 – Optical Product Data descriptor	144
Table 170 – Port Congestion descriptor	146
Table 171 – Counter Validity Bit Mask Values	146
Table 172 – Elapsed Time descriptor	149
Table 173 – QFPA Request Payload	150
Table 174 – QFPA LS_ACC Payload	151
Table 175 – Priority Range descriptor Format	151
Table 176 – UVEM Request Payload	152
Table 177 – VEM ID descriptor	153
Table 178 – Instantiated VE Mapping descriptor	153
Table 179 – Deinstantiated VE Mapping descriptor	154
Table 180 – UVEM LS_ACC Payload	154
Table 181 – EDC Request Payload	155
Table 182 – EDC LS_ACC Payload	155
Table 183 – Link Fault Capability descriptor	156
Table 184 – Congestion Signaling Capability descriptor	157
Table 185 – Signal Capability description	157
Table 186 – Signal Frequency description	158
Table 187 – Frame Discard TOV descriptor	159
Table 188 – Operational Flags field description	159
Table 189 – RDF Request payload	160
Table 190 – RDF LS_ACC Payload	161
Table 191 – Registered diagnostic functions	162
Table 192 – FPIN Registration descriptor	162
Table 193 – FPIN ELS format	163
Table 194 – Link Integrity notification descriptor	164
Table 195 – Link Integrity notification Event Type values	165
Table 196 – Event Modifier values for the Device Specific Event Type	165
Table 197 – Delivery Failure notification descriptor	166
Table 198 – Delivery Reason Code values	167
Table 199 – Peer Congestion notification descriptor	167
Table 200 – Congestion Event Type values	168
Table 201 – Congestion notification descriptor	169
Table 202 – Congestion notification Severity values	170
Table 203 – LS_RJT Reason Codes	172
Table 204 – LS_RJT Reason Code Explanations	173
Table 205 – Effects of FLOGI, FDISC, & LOGO on Permanent Port Name (PPN)	188
Table 206 – FLOGI, PLOGI, FDISC or LS_ACC Payload	191
Table 207 – Common Service Parameter applicability	192
Table 208 – Common Service Parameters - FLOGI	195
Table 209 – Common Service Parameters - PLOGI and PLOGI LS_ACC	196
Table 210 – Common Service Parameters - FLOGI LS_ACC	196
Table 211 – Clock Synchronization Applicability	200
Table 212 – Class Service Parameters Applicability	203
Table 213 – Class Service Parameters - FLOGI	205
Table 214 – Class Service Parameters - PLOGI and PLOGI LS_ACC	206

Table 215 – Class Service Parameters - FLOGI LS_ACC	206
Table 216 – Fabric Login Priority and Preemption Support	207
Table 217 – Class 2 and 3 Preference Bit Function	208
Table 218 – DiffServ QoS bit definition	209
Table 219 – ACK_0 Support Conditions (Initiator Control)	210
Table 220 – ACK_0 Support Conditions (Recipient Control)	211
Table 221 – Error Policy Bits Definition	212
Table 222 – Categories per Sequence Bits Definition	212
Table 223 – Auxiliary Parameter Data	214
Table 224 – NPIV_CNT field in LS_ACC	215
Table 225 – Login Extension Page format	217
Table 226 – Page Code Definitions	218
Table 227 – Vendor Specific Page format	218
Table 228 – N_Port Clock Synchronization QoS	219
Table 229 – FLOGI/PLOGI CS_QoS_Request	219
Table 230 – Fx_Port Clock Synchronization QoS	221
Table 231 – Virtual Fabrics Bit Usage	227
Table 232 – VE Identification Server Mappings	235

American National Standard for Information Technology —

Fibre Channel — Link Services - 5 (FC-LS-5)

1 Scope

FC-LS-5 describes in detail the Fibre Channel Extended Link Services.

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.

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

Approved ANSI standards; approved international and regional standards (ISO and IEC); and approved foreign standards (including JIS and DIN).

For further information, contact the ANSI Customer Service Department:

Phone +1 212-642-4900
Fax: +1 212-398-0023
Web: <http://www.ansi.org>
E-mail: ansionline@ansi.org

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

Phone 202-737-8888
Web: <http://www.incits.org>
E-mail: incits@itic.org

Additional availability contact information is provided below as needed.

2.2 Approved references

INCITS 332:1999, *Fibre Channel-Arbitrated Loop-2 (FC-AL-2)*

INCITS 332:1999/AM1-2003, *Fibre Channel-Arbitrated Loop-2 (FC-AL-2) Amendment 1*

INCITS 545:2019, *Fibre Channel - Framing and Signaling - 5 (FC-FS-5)*

INCITS 548:2020, *Fibre Channel - Generic Services - 8 (FC-GS-8)*

INCITS 547:2020, *Fibre Channel - Switch Fabric - 7 (FC-SW-7)*

INCITS 563:2023, *Fibre Channel - Protocol - 5 (FCP-5)*

INCITS 509:2014, *Fibre Channel - Backbone - 6 (FC-BB-6)*

INCITS 544:2018, *Fibre Channel - Single Byte Command Set - 6 (FC-SB-6)*

2.3 References under development

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

INCITS 562, *Fibre Channel - Framing and Signaling - 6 (FC-FS-6)*

INCITS 570, *Fibre Channel Generic Services - 9 (FC-GS-9)*

INCITS 568, *Fibre Channel - Switch Fabric - 8 (FC-SW-8)*

INCITS 576, *Fibre Channel - Non Volatile Memory Express - 3 (FC-NVMe-3)*

2.4 Other references

For more information on the current status of SFF documents, contact the SNIA SFF Technology Affiliate (SFF TA) at <http://www.snia.org/sff>.

SFF document SFF-8472 - Management Interface for SFP+ (Rev 12.4 March 31, 2021)

SFF document SFF-8636 - Management Interface for 4-lane Modules and Cables (Rev 2.11 January 3, 2023)

IETF Request for Comments (RFCs) may be obtained directly from the IETF web site at <http://www.ietf.org/rfc.html>:

RFC 768, *User Datagram Protocol*, August 1980.

RFC 791, *Internet Protocol*, September 1981.

RFC 854, *Telnet Protocol Specification*, May 1983.

RFC 1157, *A Simple Network Management Protocol (SNMP)*, May 1990.

RFC 1901, *Introduction to Community-based SNMPv2*, January 1996

RFC 4291, *IP Version 6 Addressing Architecture*, February 2006.

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*, December 2007.

RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*, July 2017.

RFC 9110, *HTTP Semantics*, June 2022.

RFC 9293, *Transmission Control Protocol*, August 2022.

3 Definitions and conventions

3.1 Overview

For FC-LS-5, the following definitions, conventions, abbreviations, acronyms, and symbols apply.

3.2 Definitions

3.2.1 acknowledged class

class of service that acknowledges a transfer

3.2.2 address identifier

address value used to identify source (i.e., S_ID) or destination (i.e., D_ID) of a frame

Note 1 to entry: See FC-FS-6.

3.2.3 AE Principal Switch

AE Switch that has no Uplinks and assumes the primary role of distributing the Domain Topology Map in an Avionics Fabric

Note 1 to entry: See FC-SW-8.

3.2.4 AE Switch

AE-Capable Switch that has activated at least one AE_Port

Note 1 to entry: AE Switches are required to implement the requirements defined for Fast Fabric Initialization in FC-SW-8.

3.2.5 Arbitrated Loop topology

Fibre Channel topology where L_Ports use arbitration to gain access to the loop

Note 1 to entry: See FC-AL-2.

3.2.6 Avionics Fabric

Fibre Channel Fabric that contains at least one AE Switch and supports all the requirements defined for Fast Fabric Initialization in FC-SW-8

3.2.7 buffer-to-buffer credit

BB_Credit

limiting value for BB_Credit_CNT in the buffer-to-buffer flow control model

Note 1 to entry: See FC-FS-6.

3.2.8 buffer-to-buffer Credit_Count

BB_Credit_CNT

counter used in the buffer-to-buffer flow control model

Note 1 to entry: See FC-FS-6.

3.2.9 B_Port

Fabric inter-element port used to connect bridge devices with E_Ports on a switch by providing a subset of E_Port functionality

Note 1 to entry: See FC-SW-8.

3.2.10 buffer

logical construct that holds a single frame

3.2.11 class of service

type of frame delivery service used by the communicating Nx_Ports

Note 1 to entry: See FC-FS-6.

3.2.12 Class 2 service

service that multiplexes frames at frame boundaries to or from one or more Nx_Ports with acknowledgement provided

Note 1 to entry: See FC-FS-6.

3.2.13 Class 3 service

service that multiplexes frames at frame boundaries to or from one or more Nx_Ports without acknowledgement

Note 1 to entry: See FC-FS-6.

3.2.14 Class F service

service that multiplexes frames at frame boundaries with acknowledgement provided to control and coordinate the internal behavior of the Fabric

Note 1 to entry: See FC-SW-8.

3.2.15 continuously increasing relative offset

condition of operation that requires frames ordered by SEQ_CNT within a Sequence to have a larger relative offset value in each frame

Note 1 to entry: See FC-FS-6.

3.2.16 Core N_Port_Name

N_Port_Name used in EVFP (see 4.3.38) that is associated with a VFT Tagging PN_Port, and not with any other PN_Port or FC_Port within the scope of its Name_Identifier format

3.2.17 Core Switch_Name

Switch_Name identifying the physical Switch (see FC-SW-8) in a Virtual Fabric capable Switch

3.2.18 Credit

number of buffers available at a recipient to receive frames from a transmitting FC_Port

Note 1 to entry: See FC-FS-6.

3.2.19 data block

ordered string of application data contained in a single Information Category

3.2.20 data frame

FC-4 Device_Data frame, an FC-4 Video_Data frame, or a Link_Data frame

Note 1 to entry: See FC-FS-6.

3.2.21 Destination_Identifier

D_ID

address identifier used to indicate the targeted destination FC_Port of the transmitted frame

Note 1 to entry: See FC-FS-6.

3.2.22 destination Nx_Port

Nx_Port designated by the Destination_Identifier of a frame

3.2.23 discard policy

error handling policy where a Sequence Recipient is able to discard Data frames received following detection of a missing frame in a Sequence

Note 1 to entry: See FC-FS-6.

3.2.24 Domain Controller

entity that controls activity within a given domain that is allocated an address

Note 1 to entry: See FC-SW-8.

3.2.25 Domain_ID

highest or most significant hierarchical level in the three-level addressing hierarchy (i.e., the most significant byte of the address identifier)

Note 1 to entry: See FC-FS-6 and FC-SW-8.

3.2.26 Domain Topology Map

entity within the Avionics Fabric (see FC-SW-8) that unambiguously describes the Domain_IDs and all of the Inter-Switch Links of the Avionics Fabric

3.2.27 end-to-end Credit

EE_Credit

limiting value for EE_Credit_CNT in the end-to-end flow control model

Note 1 to entry: See FC-FS-6.

3.2.28 end-to-end Credit_Count

EE_Credit_CNT

counter used in the end-to-end flow control model

Note 1 to entry: See FC-FS-6.

3.2.29 E_Port

Fabric expansion port that connects to another E_Port or B_Port to create an Inter-Switch Link

Note 1 to entry: See FC-SW-8.

3.2.30 Exchange

unit of protocol activity, identified by an OX_ID and RX_ID (see FC-FS-6), that transfers information between a specific Originator FC_Port and specific Responder FC_Port using one or more related non-concurrent Sequences that may flow in the same or opposite directions

3.2.31 Exchange_Identifier**X_ID**

collective reference to OX_ID and RX_ID

Note 1 to entry: See FC-FS-6.

3.2.32 Exchange Status Block

logical construct that contains the status of an Exchange

Note 1 to entry: See FC-FS-6.

3.2.33 F_Port

FC_Port within the Fabric that attaches to a PN_Port through a link, and is addressable by the Nx_Ports communicating through the PN_Port attached to the F_Port by the F_Port Controller well-known address (i.e., FFFFFFFEh)

Note 1 to entry: See FC-SW-8.

3.2.34 Fabric

entity that interconnects Nx_Ports attached to it and is capable of routing frames by using the D_ID information in a Frame_Header

Note 1 to entry: See FC-FS-6.

3.2.35 Fabric Controller

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

Note 1 to entry: See FC-SW-8.

3.2.36 F_Port Controller

entity at the well-known address FFFFFFFEh

Note 1 to entry: See FC-SW-8.

3.2.37 Fabric_Name

Name_Identifier associated with a Fabric

Note 1 to entry: See FC-FS-6.

3.2.38 Fabric VE ID

{N_Port_ID, Local VE ID} pair of values used to identify a VE within a Fabric

3.2.39 Fast Fabric Initialization**FFI**

technique that provides accelerated initialization of an Avionics Fabric through the distribution of the Domain Topology Map using the FFI request Sequence

Note 1 to entry: See FC-SW-8.

3.2.40 FC-4 TYPE

value in the Type field in the header of a data frame associated with an FC-4 protocol

Note 1 to entry: See FC-FS-6.

3.2.41 FC_Port

port that is capable of transmitting and receiving Fibre Channel frames according to the FC-0, FC-1, and FC-2 levels of the Fibre Channel architecture

Note 1 to entry: See FC-FS-6.

3.2.42 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

Note 1 to entry: See FC-SW-8.

3.2.43 FFI SW_ILS

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

Note 1 to entry: See FC-SW-8.

3.2.44 Fibre Channel interaction space

set of Fibre Channel ports, devices, and Fabrics that are connected by Fibre Channel links or are accessible by a common instance of an administrative tool or tools

3.2.45 FL_Port

F_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology

Note 1 to entry: See FC-AL-2.

3.2.46 frame

indivisible unit of information used by FC-2

Note 1 to entry: See FC-FS-6.

3.2.47 F_Port_Name

Name_Identifier associated with an F_Port

Note 1 to entry: See FC-FS-6.

3.2.48 Fx_Port

switch port capable of operating as an F_Port or FL_Port

Note 1 to entry: See FC-FS-6 and FC-AL-2.

3.2.49 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).

3.2.50 Hypertext Transfer Protocol

protocol for communicating various formats of text with embedded links and display controls

Note 1 to entry: See RFC 2616.

3.2.51 Infinite buffer

terminology to indicate that at FC-2 level, the amount of buffer available at the Sequence Recipient is unlimited

3.2.52 Information Category

category to which the frame Payload belongs (e.g., Solicited Data, Unsolicited Data, Solicited Control and Unsolicited Control)

Note 1 to entry: See FC-FS-6.

3.2.53 Information Unit

organized collection of data specified by an upper level to be transferred as a single Sequence by FC-2

3.2.54 initial relative offset

relative offset value specified at the sending end by an upper level for a given data block and used by the sending FC-2 in the first frame of that data block

Note 1 to entry: Initial relative offset value may be zero or non-zero (see FC-FS-6).

3.2.55 Internet Protocol

protocol for communicating data packets between identified endpoints on a multipoint network

Note 1 to entry: See RFC 791, RFC 2460, RFC 4291, and RFC 5095.

3.2.56 IP Address

identifier of an endpoint in Internet Protocol

3.2.57 link

one or more pairs of unidirectional fibres transmitting in opposite directions and their associated transmitters and receivers

3.2.58 Link Control Facility

LCF

hardware facility that attaches to an end of a link and manages transmission and reception of data

Note 1 to entry: See FC-FS-6.

3.2.59 local Fx_Port

Fx_Port to which an Nx_Port is directly attached by a link or an Arbitrated Loop

3.2.60 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 FC header.

3.2.61 Loop Fabric Address

LFA

address identifier used to address an FL_Port for the purpose of loop management

Note 1 to entry: See FC-SW-8.

3.2.62 Loss-of-signal

indication by an FC_Port that it is unable to detect an input signal on its physical interface

Note 1 to entry: See FC-FS-6.

3.2.63 L_Port

FC_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology

Note 1 to entry: See FC-AL-2.

3.2.64 Multi-function device

device that provides more than one function

3.2.65 Multiplexer

entity that provides the functions of the FC-2M sublevel

Note 1 to entry: See FC-FS-6.

3.2.66 Name_Identifier

64-bit identifier used to identify entities in Fibre Channel (e.g., Nx_Port, node, F_Port, or Fabric)

Note 1 to entry: See FC-FS-6.

3.2.67 NAS server

device that connects to a network and provides file access services

3.2.68 Network_Address_Authority

NAA

organization such as IEEE that administers network addresses

3.2.69 Network_Address_Authority identifier

four-bit identifier defined to indicate a Network_Address_Authority

Note 1 to entry: See FC-FS-6.

3.2.70 NL_Port

Nx_Port communicating through a PN_Port that is operating a Loop Port State Machine

Note 1 to entry: See FC-AL-2.

3.2.71 node

collection of one or more Nx_Ports controlled by a level above FC-2

Note 1 to entry: See FC-FS-6.

3.2.72 Node_Name

Name_Identifier associated with a node

Note 1 to entry: See FC-FS-6.

3.2.73 N_Port

Nx_Port communicating through a PN_Port

Note 1 to entry: See FC-FS-6.

3.2.74 N_Port_ID

address identifier of an Nx_Port used in the S_ID and D_ID fields of a frame

Note 1 to entry: See FC-FS-6.

3.2.75 N_Port_ID Virtualization

NPIV

ability of the multiplexer of a PF_Port or a PN_Port to support more than one VN_Port

3.2.76 N_Port_Name

Name_Identifier associated with an Nx_Port

3.2.77 Nx_Port

end point for Fibre Channel frame communication that is used in this standard to specify behavior of either N_Ports or Public NL_Ports (see FC-AL-2)

Note 1 to entry: See FC-FS-6.

3.2.78 open Exchange

period of time starting when an Exchange is initiated until that Exchange is normally or abnormally terminated (see FC-FS-6)

3.2.79 open Sequence

period of time starting when a Sequence is initiated until that Sequence is normally or abnormally terminated (see FC-FS-6)

3.2.80 Originator

logical function associated with an Nx_Port that originates an Exchange

3.2.81 Originator Exchange_ID

OX_ID

identifier assigned by an Originator to identify an Exchange (see FC-FS-6)

3.2.82 Payload

contents of the Data_Field of a frame, excluding Optional Headers and fill bytes, if present

Note 1 to entry: See FC-FS-6.

3.2.83 Permanent Port Name

Name_Identifier associated with a PN_Port

Note 1 to entry: See FC-GS-9.

3.2.84 PF_Port

LCF within a Fabric that attaches to a PN_Port through a link

Note 1 to entry: See FC-SW-8.

3.2.85 Platform

container for one or more nodes and one or more LCFs

Note 1 to entry: Any additional characteristics of a platform are outside the scope of this standard (e.g., see FC-GS-9).

3.2.86 PN_Port

LCF that supports only Nx_Ports (see FC-FS-6)

3.2.87 Policy

rule or rules used to determine how frames not received are handled during error recovery

Note 1 to entry: See FC-FS-6.

3.2.88 Port VF_ID

configurable VF_ID that is associated with any untagged frame received by a VF capable PN_Port or F_Port

Note 1 to entry: See FC-FS-6.

3.2.89 Private NL_Port

NL_Port that does not attempt a Fabric Login and does not transmit OPN(00,x)

Note 1 to entry: See FC-AL-2.

3.2.90 Public NL_Port

NL_Port that attempts a Fabric Login

Note 1 to entry: See FC-AL-2.

3.2.91 random relative offset

relationship specified between relative offset values contained in frame (n) and frame (n+1) of an Information Category within a single Sequence (see FC-FS-6), in which the relative offset of frame (n+1) is not related to the relative offset of frame (n)

3.2.92 relative offset

displacement, expressed in bytes, of the first byte of a Payload related to an upper level defined origin for a given Information Category (see FC-FS-6)

3.2.93 relative offset space

address space defined by the sending upper level for a set of information carried in one or more information units

3.2.94 Responder

logical function in an Nx_Port responsible for supporting the Exchange initiated by the Originator in another Nx_Port

3.2.95 Responder Exchange_ID

RX_ID

identifier assigned by a Responder to identify an Exchange

3.2.96 Secured Hypertext Transfer Protocol

protocol for communicating various formats of text with embedded links and display controls used in combination with a subordinate protocol that provides security features (see RFC 2818)

3.2.97 Sequence

set of one or more Data frames within an Exchange with a common Sequence_ID (SEQ_ID), transmitted unidirectionally from one Nx_Port to another Nx_Port

Note 1 to entry: See FC-FS-6.

**3.2.98 Sequence_ID
SEQ_ID**

identifier used to identify a Sequence (see FC-FS-6)

3.2.99 Sequence Initiator

FC_Port that initiates a Sequence and transmits Data frames to the destination FC_Port

Note 1 to entry: See FC-FS-6.

3.2.100 Sequence Recipient

FC_Port that receives Data frames from the Sequence Initiator

Note 1 to entry: See FC-FS-6.

3.2.101 Sequence Status Block

logical construct that tracks the status of a Sequence

Note 1 to entry: See FC-FS-6.

3.2.102 Simple Network Management Protocol

SNMP

protocol for communicating simply structured management information (see IETF STD 62)

3.2.103 Source_Identifier

S_ID

address identifier used to indicate the source Nx_Port of the transmitted frame

Note 1 to entry: See FC-FS-6.

3.2.104 source Nx_Port

Nx_Port where a frame originated

3.2.105 Storage access device

device that provides storage management and access for heterogeneous hosts and heterogeneous devices (e.g., a medium changer device)

3.2.106 Storage subsystem

integrated collection of storage controllers, storage devices, and necessary software, that provides storage services to one or more hosts

3.2.107 streamed Sequence

new sequence initiated by a Sequence Initiator in any class of service for an Exchange while it already has Sequences open for that Exchange (see FC-FS-6)

3.2.108 T10 Vendor ID

character string that uniquely identifies a vendor

Note 1 to entry: See 3.7.

3.2.109 TCP Port Number

identifier of a destination in Transmission Control Protocol

3.2.110 Telnet

protocol for communicating control of a character-oriented terminal over Transmission Control Protocol (see RFC 854)

**3.2.111 Transmission Control Protocol
TCP**

protocol communicating reliable flow-controlled byte streams over Internet Protocol allowing independent concurrent streams to multiple destinations at any IP Address (see RFC 793)

3.2.112 UDP Port Number

identifier of a destination in User Datagram Protocol (see RFC 768)

3.2.113 Upper Level

level above FC-2

**3.2.114 Upper Level Protocol
ULP**

protocol user of FC-4 (see FC-FS-6)

**3.2.115 User Datagram Protocol
UDP**

protocol communicating a packet stream with no incremental reliability over Internet Protocol allowing multiple independent concurrent destinations at any IP Address

Note 1 to entry: See RFC 768.

Note 1 to entry: See FC-BB-6.

3.2.116 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).

3.2.117 VFT Tagging PF_Port

PF_Port operating with a Multiplexer that has enabled processing of Virtual Fabric Tagging Headers (see FC-FS-6)

3.2.118 VFT Tagging PN_Port

PN_Port operating with a Multiplexer that has enabled processing of Virtual Fabric Tagging Headers (see FC-FS-6)

**3.2.119 Virtual E_Port
VE_Port**

instance of the FC-2V sublevel of Fibre Channel that communicates with another VE_Port

**3.2.120 Virtual Entity
VE**

virtualized resource

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

**3.2.121 Virtual Entities Manager
VEM**

entity managing Virtual Entities

Note 1 to entry: For example a Hypervisor.

3.2.122 Virtualization device

device that integrates one or more entities (either logical or physical), along with any additional functionality, for the purpose of providing a useful abstraction (e.g., a SCSI logical unit virtualization device)

3.2.123 VN_Port

instance of the FC-2V sublevel, synonymous with Nx_Port, that is used when it is desired to emphasize support for multiple Nx_Ports in a single Multiplexer (e.g., via a single PN_Port)

3.2.124 Wavelength division multiplexer

device that modulates/demodulates each of several data streams (e.g., Fibre Channel protocol data streams) to/from a different part of the light spectrum in an optical fiber

3.2.125 Well-known address

WKA

address identifier explicitly defined in this standard or other standards to access services (e.g., name server)

3.2.126 Worldwide_Name

WWN

Name_Identifier that is worldwide unique

Note 1 to entry: See FC-FS-6.

3.3 Editorial Conventions

In this standard, a number of conditions, mechanisms, sequences, parameters, events, states or other terms are printed with the first letter of each word in uppercase and the rest lowercase. This indicates that they have a special meaning in the context of this standard. The meaning is either described in the relevant text, in the glossary of this standard, or in a referenced standard (e.g., Exchange and Class). Any use of these terms in lowercase indicates that the words have the normal technical English meanings.

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

The ISO/British convention of decimal number representation is used in this standard. Numbers may be separated by single spaces into groups of three digits counting from the decimal position, and a period is used as the decimal marker. A comparison of the ISO/British, ISO/French, and American conventions is shown in table 1.

Table 1 – Comparison of numbering conventions

ISO/British	ISO/French	American
0.6	0,6	0.6
3.14159265	3,141 592 65	3.14159265
1 000	1 000	1,000
1 323 462.9	1 323 462,9	1,323,462.9

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 sections.

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 sections.

When the value of the bit or field is not relevant, x or xx appears in place of a specific value.

Unless stated otherwise: 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; and numbers or upper case letters immediately followed by lower-case h (xxh) are hexadecimal values.

3.4 Abbreviations and acronyms

Abbreviations and acronyms applicable to this standard are listed. Definitions of several of these items are included in 3.2.

ABTS	Abort Sequence
ACK	Acknowledgement
ADVC	Advise Credit
AE	Avionics Environment
AL_PA	Arbitrated Loop Physical Address
BA_ACC	Basic Accept
BB_Credit	buffer-to-buffer Credit
BB_Credit_CNT	buffer-to-buffer Credit_Count
BB_SCs	buffer-to-buffer State Change (SOF)
BB_SCr	buffer-to-buffer State Change (R_RDY)
BB_SC_N	buffer-to-buffer State Change Number
BSY	busy
Credit_CNT	Credit_Count
DF_CTL	Data_Field Control
D_ID	Destination_Identifier
DSCP	Differentiated Services Code Point
E_D_TOV	Error_Detect_Timeout value
EE_Credit	end-to-end Credit
EE_Credit_CNT	end-to-end Credit_Count
ELS	Extended Link Service
ELS_Command	Extended Link Service Command
ESB	Exchange Status Block
ESTC	Estimate Credit
ESTS	Establish Streaming
F_BSY	Fabric_Port_Busy
F_BSY(Df)	F_BSY response to a Data frame
F_BSY(LC)	F_BSY response to any Link_Control except P_BSY
FC	Fibre Channel
FC-2	FC-2 level (see FC-FS-6)
FC-4	FC-4 level (see FC-FS-6)
F_CTL	Frame Control (see FC-FS-6)
FFI	Fast Fabric Initialization
FLOGI	Fabric Login
F_RJT	Fabric Reject
HBA	Host Bus Adapter
hex	hexadecimal notation
HTTP	Hypertext Transfer Protocol

HTTPS	Secured Hypertext Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
LCF	Link Control Facility
LCR	Link Credit Reset
LESB	Link Error Status Block (see FC-FS-6)
LFA	Loop Fabric Address (see FC-AL-2)
LILP	Loop Initialization Loop Position (see FC-AL-2)
LIP	Loop Initialization Primitive (see FC-AL-2)
LISA	Loop Initialization Soft Assigned (see FC-AL-2)
LOGO	Logout
LR	Link Reset Primitive Sequence (see FC-FS-6)
LRR	Link Reset Response Primitive Sequence (see FC-FS-6)
LS_ACC	Link Service Accept
m	Meter
MB	MegaByte
ms	millisecond
μs	microsecond
N/A	not applicable
NAA	Network_Address_Authority
NAS	Network Attached Storage
NOP	No Operation
NOS	Not_Operational Primitive Sequence (see FC-FS-6)
NPIV	N_Port_ID Virtualization
ns	nanosecond
OLS	Offline Primitive Sequence (see FC-FS-6)
OX_ID	Originator Exchange_ID
P_BSY	N_Port_Busy
PDISC	Discover N_Port Service Parameters
PLOGI	N_Port Login
P_RJT	N_Port_Reject
PRLI	Process Login
PRLO	Process Logout
R_A_TOV	Resource_Allocation_Timeout value (see FC-FS-6)
R_CTL	Routing Control
RJT	reject
RO	relative offset
R_RDY	Receiver_Ready
R_T_TOV	Receiver_Transmitter_Timeout value
RTV	Read Timeout Value
RX_ID	Responder Exchange_ID
s	second
SBCCS	Single Byte Command Code Sets
SCR	State Change Registration
SEQ_CNT	Sequence Count
SEQ_ID	Sequence_ID
S_ID	Source_Identifier
SNMP	Simple Network Management Protocol
SOF	Start-of-Frame (see FC-FS-6)
SSB	Sequence Status Block
TCP	Transmission Control Protocol
TPLS	Test Process Login Status
TYPE	Data structure type
UDP	User Datagram Protocol

ULP	Upper Level Protocol
WWN	Worldwide_Name
X_ID	Exchange_Identifier

3.5 Symbols

Unless indicated otherwise, the following symbols have the listed meaning.

	concatenation
μ	micro (e.g., μm = micrometer)

3.6 Keywords

3.6.1 expected: A keyword used to describe the behavior of the hardware or software in the design models assumed by this standard. Other hardware and software design models may also be implemented.

3.6.2 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.

3.6.3 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.

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

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

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

3.6.7 meaningful: A control field or bit that shall be applicable and that shall be interpreted by the recipient.

3.6.8 not meaningful: A control field or bit that shall be ignored by the recipient.

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

3.6.10 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.

3.6.11 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.

3.6.12 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.

3.6.13 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.

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

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

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 most significant byte of the field, and successive characters of the T10 Vendor ID in successively less significant 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).

4 Extended Link Services

4.1 Introduction

An Extended Link Service (ELS) request solicits a destination Nx_Port to perform a function. An ELS reply shall be transmitted in response to an ELS request, unless otherwise specified. Each request or reply is composed of a single Sequence with the ELS_Command code being specified in the first word of the Payload of the first frame of the Sequence. If Zoning is active in the Fabric (see FC-GS-9), an ELS response from a well-known address (e.g., the Domain Controller) shall only include data relating to Nx_Ports that are in the same zone(s) as the requesting Nx_Port.

Each Sequence may be composed of one or more frames. Normal rules for Exchange and Sequence management apply to ELS frames, Sequences, and Exchanges. An Accept (LS_ACC) to an ELS shall terminate the Exchange by setting the Last Sequence bit to one on the last frame of the reply. An ELS request and the corresponding reply shall be performed within a single Exchange. Normal rules for Exchange and Sequence management as defined in FC-FS-6 shall apply.

The TYPE field for ELS frames shall be set to 01h.

The R_CTL field for ELS frames shall be set as specified in table 2.

Table 2 – Extended Link Services Routing Bits and Information Categories

R_CTL		Description
ROUTING	INFORMATION	
0010b	0001b	Solicited Data ^a
	0010b	Request
	0011b	Reply
	Others	Reserved
^a This value is only used by the Clock Synchronization Update (CSU) ELS.		

The first byte of the Payload (ELS_Command code) of the request or reply Sequence shall be as shown in table 9. The remainder of the Payload is ELS unique. Subsequent frames, if any, for a request or reply Sequence shall only contain additional Payload in the Payload field (i.e., the ELS_Command code is not repeated in each frame).

4.2 Link Service tag, length, value (TLV) descriptors

4.2.1 Overview

Some Link Service payloads are structured using tag, length, value (TLV) format descriptors. If a Link Service payload contains TLV format descriptors, then the TLV descriptor format shall be as specified in table 3.

Table 3 – Link Service TLV descriptor format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor tag			
1	Descriptor length			
2 - n	Descriptor value			

Descriptor tag: contains the tag value of the Link Service TLV descriptor (see table 6).

Descriptor length: contains the length in bytes of the descriptor value field. A descriptor length value of zero is valid (i.e., the tag has no descriptor value). The descriptor length value indicates the number of valid data bytes in the payload (i.e., any pad bytes are not included).

Descriptor value: contains the Link Service specific information for the descriptor tag. The descriptor value field shall be padded with zeros to an integral number of words.

The first word in a TLV format Link Service request payload or response payload shall contain an ELS command code (see table 9) followed by three zero bytes (i.e., XXh 00h 00h 00h).

The second word in a TLV format Link Service request payload or a LS_ACC response payload shall contain a length field that represents the entire length of the list of Link Service TLV descriptors that follow (see table 6), including pad bytes, if any.

4.2.2 Link Service request with Link Service TLV descriptor payload example

A Link Service request with a Link Service TLV descriptor payload example is specified in table 4.

Table 4 – Link Service request with Link Service TLV descriptor payload example

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Command code	00h	00h	00h
1	Descriptor list length = 28			
2	Descriptor tag = XXXX XXXXh			
3	Descriptor length = 4			
4	Descriptor value			
5	Descriptor tag = XXXX XXXXh			
6	Descriptor length = 8			
7	Descriptor value [first word]			
8	Descriptor value [last word]			

4.2.3 Link Service accept with Link Service TLV descriptor payload example

A Link Service accept with a TLV format link service descriptor payload example is specified in table 5.

Table 5 – Link Service accept with Link Service TLV descriptor payload example

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Descriptor list length = 40			
2	Descriptor tag = 0000 0001h			
3	Descriptor length = 4			
4	Descriptor value = request payload word 0 value			
5	Descriptor tag = XXXX XXXXh			
6	Descriptor length = 3			
7	Descriptor value			Pad byte
8	Descriptor tag = XXXX XXXXh			
9	Descriptor length = 8			
10	Descriptor value [first word]			
11	Descriptor value [last word]			

4.2.4 Link Service TLV descriptors

4.2.4.1 Overview

Link Service TLV descriptors are specified in table 6.

Table 6 – Link Service TLV descriptors

Tag value	Descriptor	Reference
0000 0000h	Reserved	
0000 0001h	Link Service Request Information	4.2.4.2
0000 0002h	Reserved	
0000 0003h	N_Port_ID descriptor	4.2.4.3
0001 0000h	SFP Diagnostics descriptor	4.3.49.5.4
0001 0001h	Port Speed descriptor	4.3.49.5.1
0001 0002h	Link Error Status Block descriptor	4.3.49.5.2
0001 0003h	Port Names descriptor	4.3.49.5.3
0001 0004h	QSFP Diagnostics descriptor	4.3.49.5.5
0001 0005h	FEC Status descriptor	4.3.49.5.6
0001 0006h	Buffer Credit descriptor	4.3.49.5.7
0001 0007h	Optical Element Data descriptor	4.3.49.5.8
0001 0008h	Optical Product Data descriptor	4.3.49.5.9
0001 0009h	Priority Range descriptor	4.3.50.4
0001 000Ah	VEM ID descriptor	4.3.51.3
0001 000Bh	Instantiated VE Mapping descriptor	4.3.51.3
0001 000Ch	Deinstantiated VE Mapping descriptor	4.3.51.3
0001 000Dh	Link Fault Capability descriptor	4.3.52.5.2
0001 000Eh	Elapsed Time descriptor	4.3.49.5.11
0001 000Fh	Congestion Signaling Capability descriptor	4.3.52.5.3

Table 6 – Link Service TLV descriptors (continued)

Tag value	Descriptor	Reference
0001 0010h	Frame Discard TOV descriptor	4.3.52.5.4
0001 FFFFh	Reserved for RDP	4.3.49.5.11
0002 0000h	Reserved	
0002 0001h	Link Integrity Notification descriptor	4.3.54.7.2
0002 0002h	Delivery Failure Notification descriptor	4.3.54.7.3
0002 0003h	Peer Congestion Notification descriptor	4.3.54.7.4
0002 0004h	Congestion Notification descriptor	4.3.54.7.5
0003 0000h	Reserved	
0003 0001h	FPIN Registration descriptor	4.3.53.5.2
All other values	Reserved	

4.2.4.2 Link Service Request Information descriptor

If Link Service TLV format requests are used, then the Link Service Request Information descriptor shall be the first Link Service TLV descriptor in a Link Service LS_ACC reply sequence.

The format of the Link Service Request Information descriptor is specified in table 7.

Table 7 – Link Service Request Information descriptor format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor tag = 0000 0001h			
1	Descriptor length = 4			
2	Request payload word 0 value			

Request payload word 0 value: contains the value of word 0 (i.e., the ELS word that contains the command code) specified in the associated Link Service request.

4.2.4.3 N_Port_ID descriptor

The format of the N_Port_ID descriptor is shown in table 8. The N_Port_ID descriptor contains a single N_Port ID.

Table 8 – N_Port_ID descriptor format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	N_Port_ID descriptor tag = 0000 0003h			
1	N_Port_ID descriptor Length (4)			
2	Reserved		N_Port_ID	

4.3 Extended Link Service requests

4.3.1 Introduction

A Sequence Initiator shall transmit an ELS Sequence in order to solicit the destination Nx_Port to perform a link-level function or service. Unless otherwise noted, Extended Link Service requests shall not be issued prior to completion of N_Port Login. Table 9 applies to ELSs sent to or received by all valid addresses, including well known addresses. FLOGI is required before any other ELS if a Fabric is present.

The LFA is used as the destination ID (D_ID) in the LINIT and LSTS ELS Request Sequences, and is used as the source ID (S_ID) in the Reply Sequences. No other Sequences shall be directed to a LFA.

Table 9 – ELS_Command codes

Value (Bits 31-24)	Description	Abbr.	Reference	N_Port Login Required	TLV Format
01h	Link Service Reject	LS_RJT	4.4.4	N/A	No
02h	Link Service Accept	LS_ACC	4.4.2	N/A	No
03h	N_Port Login	PLOGI	4.3.7	No	No
04h	F_Port Login	FLOGI	4.3.7	No	No
05h	Logout	LOGO	4.3.8	No	No
0Ah	Request Sequence Initiative	RSI	4.3.12	Yes	No
0Bh	Establish Streaming	ESTS	4.3.6	Yes	No
0Ch	Estimate Credit	ESTC	4.3.5	Yes	No
0Dh	Advise Credit	ADVC	4.3.3	Yes	No
0Eh	Read Timeout Value	RTV	4.3.10	Yes	No
0Fh	Read Link Error Status Block	RLS	4.3.9	Yes	No
10h	Echo	ECHO	4.3.4	No	No
11h	Test	TEST	4.3.13	Yes	No
12h	Reinstate Recovery Qualifier	RRQ	4.3.11	Yes	No
13h	Read Exchange Concise	REC	4.3.37	Yes	No
14h	Reserved for legacy implementations ^a				
16h	Fabric Performance Impact Notification	FPIN	4.3.54	No ^c	Yes
17h	Exchange Diagnostic Capabilities	EDC	4.3.52	- ^b	Yes
18h	Read Diagnostic Parameters	RDP	4.3.49	No ^c	Yes

^a Some early implementations of FCP-2 may have used the value 14h for SRR (Sequence Retransmission Request). This code is permanently reserved in this standard to avoid conflicts with such implementations. See FCP-5 for the standard implementation of SRR as an FC-4 Link Service.

^b N_Port Login required if the D_ID is the N_Port_ID of an N_Port. N_Port Login not required if the D_ID is the F_Port Controller.

^c If the S_ID or D_ID is set to the Fabric Controller well-known address, then FLOGI includes an implicit PLOGI for the purposes of these ELSs and extends the Receive Data_Field size specified in the FLOGI request and FLOGI LS_ACC to the implicit PLOGI.

Table 9 – ELS_Command codes (continued)

Value (Bits 31-24)	Description	Abbr.	Reference	N_Port Login Required	TLV Format
19h	Register Diagnostic Functions	RDF	4.3.53	No ^c	Yes
20h	Process Login	PRLI	4.3.19	Yes	No
21h	Process Logout	PRLO	4.3.20	Yes	No
23h	Test Process Login State	TPLS	4.3.21	Yes	No
24h	Third Party Process Logout	TPRLO	4.3.28	Yes	No
50h	Discover N_Port Service Parameters	PDISC	4.3.25	Yes	No
51h	Discover F_Port Service Parameters	FDISC	4.3.26	Yes	No
52h	Discover Address	ADISC	4.3.27	Yes	No
58h	Report Port Buffer Condition	RPBC	4.3.31	Yes	No
60h	Fabric Address Notification	FAN	4.3.14	No	No
61h	Registered State Change Notification	RSCN	4.3.17	No	No
62h	State Change Registration	SCR	4.3.18	No	No
63h	Report node FC-4 Types	RNFT	4.3.32	Yes	No
68h	Clock Synchronization Request	CSR	4.3.29	No	No
69h	Clock Synchronization Update	CSU	4.3.30	No	No
70h	Loop Initialize	LINIT	4.3.15	No	No
71h	Loop Port Control - obsolete	LPC	N/A	No	No
72h	Loop Status	LSTS	4.3.16	No	No
77h	Vendor Specific			N/A	

^a Some early implementations of FCP-2 may have used the value 14h for SRR (Sequence Retransmission Request). This code is permanently reserved in this standard to avoid conflicts with such implementations. See FCP-5 for the standard implementation of SRR as an FC-4 Link Service.

^b N_Port Login required if the D_ID is the N_Port_ID of an N_Port. N_Port Login not required if the D_ID is the F_Port Controller.

^c If the S_ID or D_ID is set to the Fabric Controller well-known address, then FLOGI includes an implicit PLOGI for the purposes of these ELSs and extends the Receive Data_Field size specified in the FLOGI request and FLOGI LS_ACC to the implicit PLOGI.

Table 9 – ELS_Command codes (continued)

Value (Bits 31-24)	Description	Abbr.	Reference	N_Port Login Required	TLV Format
78h	Request node Identification Data	RNID	4.3.22	No	No
79h	Registered Link Incident Report	RLIR	4.3.23	Yes	No
7Ah	Link Incident Record Registration	LIRR	4.3.24	Yes	No
7Bh	Scan Remote Loop	SRL	4.3.33	Yes	No
7Ch	Set Bit-error Reporting Parameters	SBRP	4.3.34	Yes	No
7Dh	Report Port Speed Capabilities	RPSC	4.3.36	Yes	No
7Eh	Query Security Attributes	QSA	4.3.35	No	No
7Fh	Exchange Virtual Fabrics Parameters	EVFP	4.3.38	N/A	No
80h	Link Keep Alive	LKA	4.3.39	No	No
81h	Link Cable Beacon	LCB	4.3.40	No	No
90h	Authentication ELS	AUTH_ELS	see FC-SP-2	see FC-SP-2	No
97h	Request Fabric Change Notification	RFCN	4.3.41	No	No

^a Some early implementations of FCP-2 may have used the value 14h for SRR (Sequence Retransmission Request). This code is permanently reserved in this standard to avoid conflicts with such implementations. See FCP-5 for the standard implementation of SRR as an FC-4 Link Service.

^b N_Port Login required if the D_ID is the N_Port_ID of an N_Port. N_Port Login not required if the D_ID is the F_Port Controller.

^c If the S_ID or D_ID is set to the Fabric Controller well-known address, then FLOGI includes an implicit PLOGI for the purposes of these ELSs and extends the Receive Data_Field size specified in the FLOGI request and FLOGI LS_ACC to the implicit PLOGI.

Table 9 – ELS_Command codes (continued)

Value (Bits 31-24)	Description	Abbr.	Reference	N_Port Login Required	TLV Format
A0h	Define FFI Domain Topology Map	FFI_DTM	4.3.42	Yes	No
A1h	Request FFI Domain Topology Map	FFI_RTM	4.3.43	Yes	No
A2h	FFI AE Principal Switch Selector	FFI_PSS	4.3.44	Yes	No
A3h	FFI Map Update Registration	FFI_MUR	4.3.45	Yes	No
A4h	FFI Registered Map Update Notification	FFI_RMUN	4.3.46	Yes	No
A5h	FFI Suspend Map Updates	FFI_SMU	4.3.47	Yes	No
A6h	FFI Resume Map Updates	FFI_RMU	4.3.48	Yes	No
B0h	Query Fabric Priority Allocation	QFPA	4.3.50	No	Yes
B1h	Update Virtual Entity Mappings	UVEM	4.3.51	^b	Yes
Others	Reserved				

^a Some early implementations of FCP-2 may have used the value 14h for SRR (Sequence Retransmission Request). This code is permanently reserved in this standard to avoid conflicts with such implementations. See FCP-5 for the standard implementation of SRR as an FC-4 Link Service.

^b N_Port Login required if the D_ID is the N_Port_ID of an N_Port. N_Port Login not required if the D_ID is the F_Port Controller.

^c If the S_ID or D_ID is set to the Fabric Controller well-known address, then FLOGI includes an implicit PLOGI for the purposes of these ELSs and extends the Receive Data_Field size specified in the FLOGI request and FLOGI LS_ACC to the implicit PLOGI.

An ELS Protocol is composed of an ELS request Sequence and an ELS Reply Sequence. The last Data frame of an ELS request Sequence shall transfer the Sequence Initiative to the Recipient in order to allow the reply to be transmitted (see FC-FS-6). If an ELS request Sequence is transmitted without the transfer of Sequence Initiative, the Recipient shall abort the Exchange and not perform the request.

The following ELS requests and their replies shall be supported by an Nx_Port (all others are optional for purposes of this standard, other standards or Technical Reports may require further ELS support):

- a) FLOGI;
- b) PLOGI; and
- c) LOGO.

An Nx_Port receiving an ELS request shall respond to it in accord with table 10, depending on its N_Port Login state with the Nx_Port sending the ELS request and the PLOGI requirement for the ELS specified in table 9.

Table 10 – Responses to Received ELSs

N_Port Login Required? (see table 9)	Logged in with Source N_Port?	
	Yes	No
Yes	Respond as appropriate for the ELS and the current state of the Nx_Port	If a reply sequence is defined for the ELS, originate a LOGO ELS Exchange to the sender of the received ELS or reply with an LS_RJT ELS Sequence with a reason code of “Unable to perform command request” and a reason code explanation of “N_Port Login required”. If a reply sequence is not defined for the ELS, it shall be discarded
No	Respond as appropriate for the ELS and the current state of the Nx_Port.	Respond as appropriate for the ELS and the current state of the Nx_Port.

An Nx_Port is not required to generate and send the PLOGI ELS request. However, if an Nx_Port receives a PLOGI ELS request, the Nx_Port shall respond with a LS_ACC Link Service Reply, or with LS_RJT. LS_RJT shall not be issued with a reason code of “Command not supported” in response to a PLOGI.

NOTE 1 – If an Nx_Port that does not generate PLOGI, is in a point-to-point topology, and has an N_Port_Name greater than the other Nx_Port's, the other Nx_Port may timeout, waiting to receive PLOGI.

One Nx_Port may transmit an RSI request to a second Nx_port at the same time that the second Nx_Port transmits an RSI request to the first Nx_Port for the same target Exchange. If such an instance occurs, the Originator Nx_Port of the target Exchange shall reject the RSI request sequence with an LS_RJT with a reason code of “command already in progress”. The Responder Nx_Port of the target Exchange shall process the RSI request Sequence normally.

4.3.2 Sequence and Exchange management

Extended Link Service communication shall observe all relevant rules of FC-FS-6.

All ELS requests, excluding ESTS, ESTC, and ADVC, and the corresponding replies shall be performed within a single Exchange, intended exclusively for the purpose of ELS processing (see FC-FS-6 for the procedure using ESTS, ESTC, and ADVC). The Advise Credit request may also be performed in a separate Exchange. Most ELS protocols are performed as a two Sequence Exchange. Each of these two Sequence Exchanges consist of a request Sequence by the Originator, transfer of Sequence Initiative, and a reply Sequence from the Responder that terminates the Exchange by setting the Last_Sequence bit (bit 20) in F_CTL.

Unless otherwise specified for a specific link service more than one frame may be used to form a request or reply Sequence.

The following rules regarding Sequence and Exchange management apply to ELSs in addition to the rules specified in FC-FS-6:

- a) Request and Reply Sequences may be sent in any class of service;
- b) Reply frames and Sequences shall be transmitted in the same class as the request;
- c) If Login has not been completed successfully, the default Login values shall be used;
- d) If Login has completed successfully, the Originator of the Exchange shall use the Discard multiple Sequences Error Policy for all ELS Exchanges (see FC-FS-6);
- e) The Originator of an ELS Exchange shall detect an Exchange error following Sequence Initiative transfer if the Reply Sequence is not received within a timeout interval of 2 X R_A_TOV;
- f) Unless otherwise specified (see 4.3.7, 4.3.8, and 4.3.26), if the Exchange Originator of an ELS Exchange detects an Exchange error, it shall abort the Exchange using ABTS-LS and retry the protocol of the aborted Exchange with a different Exchange;
- g) If the Sequence Initiator aborts a Sequence using ABTS due to receiving an ACK with the Abort Sequence bits set to 01b, the Sequence Initiator shall retry the Sequence after the BA_ACC is received for the aborted Sequence one time only. If the retry fails, the ELS Exchange shall be aborted using ABTS-LS;
- h) If the Sequence Initiator attempts to abort a Sequence using ABTS and it detects an E_D_TOV waiting for the ACK frame in response to the ABTS, it shall abort the Exchange using ABTS-LS, if conditions permit, and retry the original ELS with a different Exchange; and
- i) If the Sequence Initiator attempts to abort a Sequence using ABTS-LS and it detects an E_D_TOV waiting for the ACK frame in response to the ABTS-LS, it may retry the original ELS with a different Exchange.

4.3.3 Advise Credit (ADVC)

4.3.3.1 Description

The ADVC ELS is used to advise the destination Nx_Port of the estimated end-to-end Credit that the source Nx_Port requests to be allocated. The ADVC ELS request shall be a separate Sequence. It may also be requested in a separate Exchange. See FC-FS-6 for the usage of this ELS. The ADVC request may also be used independently from the Estimate Credit procedure (see FC-FS-6).

4.3.3.2 Protocol

- a) Advise Credit Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.3.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting Credit revision. The D_ID field designates the destination Nx_Port.

Payload: The format of the Payload is shown in table 11. The Payload shall contain the requested end-to-end Credit in the end-to-end Credit field of the appropriate Class Service Parameters (see 6.6.5) as indicated by the Class Validity bit. For each class in which a revised end-to-end Credit is

requested, the Class Validity bit shall be set to one. The recipient shall ignore the other Service Parameter fields.

Table 11 – ADVC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00	
0	ADVC (0Dh)	00h	00h	00h	
1	MSB	Common Service Parameters (16 bytes)			
..					
4					LSB
5	MSB				N_Port_Name (8 bytes)
6		LSB			
7	MSB	Node_Name (8 bytes)			
8					LSB
9	MSB	Obsolete (16 bytes)			
..					
12					LSB
13	MSB	Class 2 Service Parameters (16 bytes)			
..					
16					LSB
17	MSB	Class 3 Service Parameters (16 bytes)			
..					
20					LSB
21	MSB	Reserved (16 bytes)			
..					
24					LSB
25	MSB				Vendor Version Level (16 bytes)
..					
28		LSB			

4.3.3.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the ADVC command.

LS_ACC: LS_ACC signifies successful completion of the ADVC function and permanently replaces the end-to-end Credit in effect for the current N_Port Login.

The format of the LS_ACC Payload is shown in table 12. The Payload shall contain the revised end-to-end Credit allocated in the Credit field for the appropriate Class Service Parameters as indicated by the Class Validity bit. The revised end-to-end Credit shall replace the end-to-end Credit for the current Login for the Nx_Port transmitting the LS_ACC Sequence (see 6). For each class in which a revised end-to-end Credit is updated, the Class Validity bit shall be set to one. The recipient shall ignore the other Service Parameter fields. This revised end to-end Credit value is determined by the destination Nx_Port based on its buffering scheme, buffer management, buffer availability, and Nx_Port processing time (see FC-FS-6).

Table 12 – ADVC LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	MSB			
..	Common Service Parameters (16 bytes)			
4	LSB			
5	MSB			
6	N_Port_Name (8 bytes)			
7	LSB			
7	MSB			
8	Node_Name (8 bytes)			
9	LSB			
9	MSB			
..	Obsolete (16 bytes)			
12	LSB			
13	MSB			
..	Class 2 Service Parameters (16 bytes)			
16	LSB			
17	MSB			
..	Class 3 Service Parameters (16 bytes)			
20	LSB			
21	MSB			
..	Reserved (16 bytes)			
24	LSB			
25	MSB			
..	Vendor Version Level (16 bytes)			
28	LSB			

4.3.4 Echo (ECHO)

4.3.4.1 Description

The Echo ELS requests the Recipient to transmit the Payload contents that follows the ELS_Command back to the Initiator of the Echo command in the same order as received using the LS_ACC Reply Sequence. The Echo ELS Request provides a means to transmit a Data frame and have the Payload content returned for a simple loopback diagnostic function. The Echo command shall be transmitted as a one frame Sequence and the LS_ACC Reply Sequence is also a one frame Sequence.

4.3.4.2 Protocol

- a) Echo Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.4.3 Request Sequence

Addressing: The D_ID field designates the destination of the request while the S_ID field designates the source of the request.

Payload: The format of the Payload is shown in table 13. If a Login with the destination Nx_Port does not exist, the maximum size of the ECHO data field shall be 4 less than the default Receive Data_Field Size (see table 212). If a Login with the destination Nx_Port exists, the ECHO data field size is limited by 4 less than the smallest Receive Data_Field Size supported by the destination Nx_Port, the Fabric, and the source Nx_Port for the class of service being used. The ECHO data field in the LS_ACC frame shall be equal in size to the ECHO data field size in the ECHO Request Sequence.

Table 13 – ECHO Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	ECHO (10h)	00h	00h	00h
1	MSB			
..	ECHO data (up to max frame length - 4, any byte boundary)			
n	LSB			

4.3.4.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the ECHO command.

LS_ACC: LS_ACC signifies successful completion of the ECHO function. The format of the LS_ACC Payload is shown in table 14. The Payload shall contain the ECHO data contained in the Payload of the ECHO Request frame.

Table 14 – ECHO LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	MSB			ECHO data
..				(up to max frame length - 4, any byte boundary)
n				Excludes word 0 of ECHO Payload. LSB

4.3.5 Estimate Credit (ESTC)

4.3.5.1 Description

The ESTC ELS is used to estimate the minimum Credit required to achieve the maximum bandwidth for a given distance between an Nx_Port pair.

The class of the SOF of the ESTC Request identifies the class for which Credit is being estimated. The destination Nx_Port shall acknowledge Data frames as specified by its Login parameters. See FC-FS-6 for the usage of this frame.

4.3.5.2 Protocol

- a) Estimate Credit Request Sequence
- b) No Reply Sequence

4.3.5.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the Credit estimate. The D_ID field designates the destination Nx_Port specified in the Establish Streaming frame.

Payload: The format of the Payload is shown in table 15. The first word of the Payload of the first frame of the Sequence shall contain the ELS_Command code. The content of the Payload after the ELS_Command and for subsequent frames shall be valid data bytes.

Table 15 – ESTC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	ESTC (0Ch)	00h	00h	00h
1	MSB			Any data
..				(see FC-FS-6)
n				LSB

4.3.5.4 Reply Sequence

None.

4.3.6 Establish Streaming (ESTS)

4.3.6.1 Description

The ESTS ELS requests a temporary allocation of Credit known as Streaming Credit large enough to perform continuous streaming of Data frames. The SOF of the ESTS Request identifies the class for which Credit is being estimated. See FC-FS-6 for the usage of this frame.

4.3.6.2 Protocol

- a) Establish Streaming Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.6.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting Streaming. The D_ID field designates the destination Nx_Port addressed.

Payload: The format of the Payload is shown in table 16.

Table 16 – ESTS Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	ESTS (0Bh)	00h	00h	00h

4.3.6.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the ESTS command

LS_ACC: LS_ACC signifies successful completion of the ESTS function. The format of the LS_ACC Payload is shown in table 17. The Payload shall contain Streaming Credit allocated in the Nx_Port end-to-end Credit field of the appropriate Class Service Parameters (see 6.6.5). The Class Validity bit, if set to one, identifies the class that contains the Streaming Credit. The recipient shall ignore the other Service Parameter fields.

Table 17 – ESTS LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	MSB Common Service Parameters (16 bytes)			LSB
4				LSB
5	MSB N_Port_Name (8 bytes)			LSB
6				LSB
7	MSB Node Name (8 bytes)			LSB
8				LSB
9	MSB Obsolete (16 bytes)			LSB
..				LSB
12				LSB
13	MSB Class 2 Service Parameters (16 bytes)			LSB
..				LSB
16				LSB
17	MSB Class 3 Service Parameters (16 bytes)			LSB
..				LSB
20				LSB
21	MSB Reserved (16 bytes)			LSB
..				LSB
24				LSB
25	MSB Vendor Version Level (16 bytes)			LSB
..				LSB
28				LSB

4.3.7 Login (FLOGI/PLOGI)

4.3.7.1 Description

The FLOGI/PLOGI ELS shall transfer Service Parameters from the initiating Nx_Port to the Nx_Port associated with the D_ID. The FLOGI frame provides the means by which an Nx_Port may request Login with the Fabric (see 6.2). The PLOGI frame provides the means by which an Nx_Port may request Login with another Nx_Port prior to other Data frame transfers (see 6.3).

In order to Login with the Fabric and determine the Fabric operating characteristics, an Nx_Port shall specify the D_ID as the F_Port Controller Well-known address (i.e., FFFFEEh).

In order to direct the Login ELS frame to a Fibre Channel Service, an Nx_Port shall specify the N_Port_ID for the Nx_Port providing the service or the appropriate Well-known address (see FC-FS-6).

If an Nx_Port receives a Login from an Nx_Port, all open Sequences and Exchanges with the Nx_Port performing the Login shall be abnormally terminated.

If the Originator of a Login with the Fabric detects a time out, the FLOGI may be retried in a different Exchange without aborting the previous Exchange with ABTS-LS.

4.3.7.2 Protocol

- a) FLOGI or PLOGI Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.7.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting Login. If unidentified, as in Fabric Login, binary zeros are used. The D_ID field designates the destination Nx_Port or F_Port Controller of the Login.

Payload: The format of the Payload is shown in table 206. The Service Parameters are defined in 6.6.

4.3.7.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the FLOGI or PLOGI Request Sequence

LS_ACC: LS_ACC signifies successful completion of the FLOGI or PLOGI Request Sequence. The format of the LS_ACC Payload is shown in table 206. The Service Parameters are defined in 6.6.

4.3.8 Logout (LOGO)

4.3.8.1 Description

The LOGO ELS provides a method for explicitly removing service between two Nx_Ports or between an Nx_Port and a Fabric. Logout releases resources, identifiers, and relationships associated with maintaining service between an Nx_Port and a destination Nx_Port or Fabric.

If neither the S_ID nor D_ID of a LOGO ELS is the F_Port Controller (FFFFFFEh), the LOGO requests removal of service between the Nx_Port assigned to the N_Port_ID specified in the LOGO Payload and the Nx_Port assigned to the N_Port_ID specified in the D_ID. The N_Port_ID in the LOGO Payload may differ from the S_ID. This allows an Nx_Port to Logout its old Identifier using a new Identifier after its N_Port_ID has changed. Both the source Nx_Port and the destination Nx_Port of the Logout Request Sequence shall abnormally terminate all open Exchanges (see 6) that used the N_Port_ID indicated in the Payload of the Logout Request Sequence.

If either the S_ID or D_ID of a LOGO ELS is the F_Port Controller (FFFFFFEh), the LOGO ELS requests Fabric logout and release of the previously assigned N_Port_ID specified in the LOGO Payload. An Nx_Port that requests or accepts explicit logout from the Fabric shall implicitly log out the Nx_Port assigned to the N_Port_ID indicated in the Payload of the Logout Request Sequence from all other Nx_Ports and abnormally terminate all open Exchanges (see 6) that used the N_Port_ID indicated in the Payload of the Logout Request Sequence.

If the Originator of a LOGO detects a time out, the LOGO may be retried in a different Exchange without aborting the previous Exchange with ABTS-LS.

4.3.8.2 Protocol

- a) Logout Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.8.3 Request Sequence

Addressing: For an explicit Nx_Port Logout, the S_ID field designates the source N_Port_ID of the Nx_Port requesting Logout and the D_ID field designates the destination Nx_Port of the Logout Request. For an explicit Fabric Logout originated by an Nx_Port, the S_ID field shall be the N_Port_ID of the Nx_Port to be logged out (i.e., the same as the N_Port_ID in the Payload) and the D_ID shall be the F_Port Controller Well-known address (i.e., FFFFEh). For an explicit Fabric Logout originated by the Fabric, the S_ID field shall be the F_Port Controller Well-known address (i.e., FFFFEh) and the D_ID shall be the N_Port_ID of the Nx_Port to be logged out (i.e., the same as the N_Port_ID in the Payload).

Payload: The format of the Payload is shown in table 18.

Table 18 – LOGO Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	LOGO (05h)	00h	00h	00h
1	Reserved	N_Port_ID		
2	MSB	N_Port_Name		LSB
3		(8 bytes)		

4.3.8.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the LOGO command.

LS_ACC: LS_ACC signifies that service has been removed for the N_Port_ID indicated in the payload of the LOGO ELS. If the LOGO ELS requested Fabric Logout, LS_ACC signifies that the N_Port_ID specified in the Payload of the LOGO ELS has been logged out from the Fabric and released. The format of the LS_ACC Payload is shown in table 19.

Table 19 – LOGO LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.9 Read Link Error Status Block (RLS)

4.3.9.1 Description

The RLS ELS requests an FC_Port to return the identified Link Error Status Block (LESB) associated with the Port_ID specified in the Payload. This provides the Nx_Port transmitting the request with information regarding Link Errors detected within the designated FC_Port.

4.3.9.2 Protocol

- a) Read Link Error Status Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.9.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the LESB. The D_ID field shall be set as follows:

- a) to any Domain Controller well known address (FFFCxxh);
- b) to the F_Port Controller Well-known address (FFFFFFEh); or
- c) to any Nx_Port logged in with the S_ID address.

Payload: The format of the Payload is shown in table 20.

Table 20 – RLS Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RLS (0Fh)	00h	00h	00h
1	Reserved	N_Port_ID		

The requested LESB is identified as follows:

- a) If the D_ID is a Domain Controller well known address (FFFCxxh), the N_Port_ID field shall be set to an N_Port_ID within the associated domain. The LESB requested is for the F_Port that the N_Port_ID is logged in with;
- b) if the D_ID is the F_Port Controller Well-known address (FFFFFFEh), the N_Port_ID field is not meaningful and the LESB request is for the local Fx_Port that the S_ID is logged in with; or
- c) for all other D_IDs, the N_Port_ID field is not meaningful and the LESB requested is for the Nx_Port assigned to the D_ID.

4.3.9.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RLS command. The LS_RJT reason code and reason code explanation are set as follows:

- a) If an FC_Port does not support the LESB, it shall reply with an LS_RJT specifying a reason code of “Unable to perform command request” (09h) and should respond with a reason code explanation of “Request not supported” (2Ch); or
- b) if the N_Port_ID is not logged in with an F_Port within the domain, the Domain Controller should reply with an LS_RJT specifying a reason code of “Logical error” (03h) and reason code explanation “Invalid N_Port_ID” (1Fh).

LS_ACC: LS_ACC signifies that the FC_Port has transmitted the requested data. The format of the LS_ACC Payload is shown in table 21. The format of the Link Error Status Block is specified in FC-FS-6.

Table 21 – RLS LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	MSB			Link Error Status Block (see FC-FS-6) (24 bytes)
..				
6				LSB

4.3.10 Read Timeout Value (RTV)

4.3.10.1 Description

The RTV ELS requests an FC_Port to return the R_A_TOV and the E_D_TOV along with an indicator of the value of R_T_TOV. This provides the FC_Port transmitting the RTV ELS with information regarding these values from another FC_Port. Usage of R_A_TOV, E_D_TOV and R_T_TOV requirements are given in FC-FS-6.

4.3.10.2 Protocol

- a) Read Timeout Value (RTV) Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.10.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the timeout interval values. The D_ID field designates the destination FC_Port to which the request is being made.

Payload: The format of the Payload is shown in table 22.

Table 22 – RTV Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RTV (0Eh)	00h	00h	00h

4.3.10.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RTV command.

LS_ACC: LS_ACC returns the requested R_A_TOV and E_D_TOV values. The format of the LS_ACC Payload is shown in table 23. E_D_TOV Timeout values are specified as a count of either 1 ms or 1 ns increments, depending on the setting of the E_D_TOV Resolution (see FC-FS-6).

Table 23 – RTV LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Resource_Allocation_Timeout Value (R_A_TOV) (see FC-FS-6)			
2	Error_Detect_Timeout Value (E_D_TOV) (see FC-FS-6)			
3	Timeout Qualifier			

The Timeout Qualifier word is defined as follows:

a) Bits 31-27: Reserved;

b) Bit 26: E_D_TOV Resolution:

If the E_D_TOV Resolution bit is zero, the value specified in the E_D_TOV field shall indicate a count of 1 ms increments. If the E_D_TOV Resolution bit is one, the value specified in the E_D_TOV field shall indicate a count of 1 ns increments;

c) Bits 25-20: Reserved;

d) Bit 19: R_T_TOV Value;

If this bit is set to zero, the value of R_T_TOV shall be the default value of 100 milliseconds. If it is set to one, the value of R_T_TOV shall be 100 microseconds; and

e) Bits 18-0: Reserved.

4.3.11 Reinstate Recovery Qualifier (RRQ)

4.3.11.1 Description

The RRQ ELS shall be used to notify the destination Nx_Port that the Recovery_Qualifier shall be available for reuse. The Recovery_Qualifier (S_ID, D_ID, OX_ID, RX_ID, and low SEQ_CNT minus high SEQ_CNT) shall be associated with an Exchange in which the Abort Sequence or Abort Exchange was previously performed.

In the case of Abort Exchange (i.e., ABTS-LS), the ESB and Recovery_Qualifier are immediately available for reuse. In the case of Abort Sequence Protocol, the Recovery_Qualifier is purged.

A request to Reinstate the Recovery_Qualifier shall only be accepted if the Originator Nx_Port N_Port_ID or the Responder Nx_Port N_Port_ID of the target Exchange is the same as the N_Port_ID of the Nx_Port that makes the request. If the RRQ Request is not accepted, an LS_RJT

with reason code “Unable to perform command request” and reason code explanation “Invalid Originator S_ID” shall be returned.

A separate Exchange shall be used to reinstate the Recovery_Qualifier. The Payload shall contain the OX_ID and RX_ID for the Exchange Recovery_Qualifier, in addition to the S_ID of the Nx_Port that originated the Exchange being aborted. Resources associated with the OX_ID in the Originator, and with the RX_ID in the Responder, shall be released following transmission and reception of the LS_ACC reply Sequence if the Exchange had been aborted with ABTS-LS.

Both the Originator and Responder shall ensure that the OX_ID and RX_ID pair being terminated is currently associated with the OX_ID and RX_ID pair specified in the RRQ Request.

The Recovery_Qualifier range shall be timed out for an R_A_TOV timeout period (i.e., RRQ shall not be transmitted until an R_A_TOV timeout period after BA_ACC for ABTS has been received) by the Nx_Port that transmitted and successfully completed the ABTS frame.

4.3.11.2 Protocol

- a) Reinstatement Recovery Qualifier Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.11.3 Request Sequence

Addressing: The D_ID field designates the destination Nx_Port of the RRQ Request Sequence while the S_ID field designates the source Nx_Port that is requesting that the Recovery_Qualifier be reinstated.

Exchange: A separate Exchange is required.

SEQ_ID and SEQ_CNT: The SEQ_ID and the SEQ_CNT shall be appropriate for an open Sequence.

Payload: The format of the Payload is shown in table 24.

Table 24 – RRQ Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RRQ (12h)	00h	00h	00h
1	Reserved	Exchange Originator S_ID		
2	OX_ID		RX_ID	
3	MSB			
..	Obsolete (32 bytes)			
10	LSB			

4.3.11.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RRQ command.

LS_ACC: LS_ACC signifies that the destination Nx_Port reinstated the Recovery_Qualifier. The format of the LS_ACC Payload is shown in table 25.

Table 25 – RRQ LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.12 Request Sequence Initiative (RSI)

4.3.12.1 Description

The RSI ELS is used to request that Sequence Initiative be passed to the Sequence Recipient of an Exchange in progress. A request to pass Sequence Initiative shall only be accepted if the Originator Nx_Port or the Responder Nx_Port of the target Exchange makes the request. A separate Exchange shall be used to perform the Request Sequence Initiative. The Payload shall contain the OX_ID and RX_ID for the target Exchange, in addition to the S_ID of the Nx_Port that originated the Exchange. The LS_ACC Reply is sent subsequent to the transfer of Sequence Initiative on the target Exchange.

Transmission of RSI is allowed while the identified Exchange is open. Both the Originator and Responder shall ensure that the OX_ID and RX_ID pair for which Sequence Initiative is being passed are currently associated with the OX_ID and RX_ID pair specified in the RSI Request.

If there is a Sequence active for the target Exchange, the Sequence Initiator of the active Sequence of the target Exchange shall terminate them and transfer Sequence Initiative as follows:

- a) If there is an active Sequence for which the last Data frame has not been transmitted, the Sequence Initiator of the target Exchange shall terminate the Sequence by transmitting a Data frame with the End_Sequence and Sequence Initiative bits set to one in F_CTL; and
- b) If there are no Data frames to be sent for the active Sequence, the Sequence Initiator of the target Exchange shall transmit a NOP Basic Link Service frame (see FC-FS-6) with the End_Sequence and Sequence Initiative bits set to one in F_CTL.

If there is no Sequence active, the Sequence Initiator of the target Exchange shall transfer Sequence Initiative by initiating a new Sequence consisting of a single NOP Basic Link Service frame (a one frame Sequence) with the End_Sequence and Sequence Initiative bits set to 1 in F_CTL.

The LS_ACC to the Exchange requesting Sequence Initiative shall be transmitted after Sequence Initiative has been passed (see FC-FS-6) on the target Exchange.

4.3.12.2 Protocol

- a) Request Sequence Initiative Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.12.3 Request Sequence

Addressing: The D_ID field designates the destination Nx_Port of the Exchange for which Sequence Initiative is being requested and the S_ID field designates the source Nx_Port that is requesting Sequence Initiative.

Exchange: A separate Exchange is required other than the Exchange for which Sequence Initiative is being requested in order to properly track status.

SEQ_ID and SEQ_CNT: The SEQ_ID and the SEQ_CNT shall be appropriate for an open Sequence.

Payload: The format of the Payload is shown in table 26.

Table 26 – RSI Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RSI (0Ah)	00h	00h	00h
1	Reserved	Originator S_ID		
2	OX_ID		RX_ID	
3	MSB			
..	Obsolete (32 bytes)			
10	LSB			

4.3.12.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RSI command.

LS_ACC: LS_ACC signifies that the destination Nx_Port has transferred the Sequence Initiative for the target Exchange. The format of the LS_ACC Payload is shown in table 27.

Table 27 – RSI LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.13 Test (TEST)

4.3.13.1 Description

The TEST ELS shall consist of a single Sequence being transmitted from the Sequence Initiator to the Sequence Recipient. The Test Request may be used in diagnostic or testing procedures to provide system loading. There is no Reply Sequence. The Payload may consist of any frame size up to the maximum allowable for the class and other normal Sequence and frame limitations.

4.3.13.2 Protocol

- a) Test Request Sequence
- b) No Reply Sequence

4.3.13.3 Request Sequence

Addressing: The D_ID field designates the destination of the request while the S_ID field designates the source of the request.

Payload: The format of the Payload is shown in table 28.

Table 28 – TEST Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	TEST (11h)	00h	00h	00h
1	MSB	TEST data		
..		(up to max Data_Field size - 4, any byte boundary)		
n				LSB

The Payload size is limited by the smallest Data_Field size supported by the destination Nx_Port and the Fabric for the class being used.

4.3.13.4 Reply Sequence

none

4.3.14 Fabric Address Notification (FAN)

4.3.14.1 Description

The FAN ELS shall be sent by an F_Port Controller (FFFFFFEh) to all known previously logged in (via FLOGI) attached NL_Ports following an initialization event. This initialization event is typically Loop Initialization on an Arbitrated Loop, though other events that may cause a port to change its ID may also be considered. The F_Port Controller shall report the F_Port_Name and Fabric_Name as they were reported in the prior FLOGI, and shall report the current Loop Fabric Address. The F_Port Controller shall send this ELS using the default Login parameters (i.e., the parameters that are in effect prior to a FLOGI Request).

The attached ports may use this information to authenticate active Exchanges and operating parameters (e.g., Login BB_Credit).

The F_Port Controller shall report identical information to all attached NL_Ports. If the information changes in any way before the F_Port Controller is able to send the service to all attached ports, the F_Port Controller shall begin a new initialization event.

The attached NL_Ports shall not initiate a Reply Sequence to this ELS.

4.3.14.2 Protocol

- a) Fabric Address Notification Request Sequence
- b) No Reply Sequence

4.3.14.3 Request Sequence

Addressing: The S_ID is the F_Port Controller (FFFFFFEh) sending the FAN. The D_ID is the NL_Port receiving the FAN.

Payload: The format of the FAN Request Payload is shown in table 29.

Table 29 – FAN Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	FAN (60h)	00h	00h	00h
1	Reserved	Loop Fabric Address		
2	MSB F_Port_Name			LSB
3	(8 bytes)			
4	MSB Fabric_Name			LSB
5	(8 bytes)			

4.3.14.4 Reply Sequence

none

4.3.15 Loop Initialize (LINIT)

4.3.15.1 Description

The LINIT ELS requests the start of Loop Initialization on a designated loop. The FL_Port connected to the loop addressed by the Loop Fabric Address shall originate a LIP of the type specified in the Payload and begin the Initialization Process as described in FC-AL-2. The LS_ACC Reply Sequence shall not be initiated before the FL_Port has transitioned from the OPEN-INIT to MONITORING State, and the Fabric Controller has completed sending any FAN ELS (see FC-AL-2). If the LINIT ELS is received by a non-FL_Port, the port shall originate an LS_RJT with a reason code of “Command not supported”.

4.3.15.2 Protocol

- a) Loop Initialize Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.15.3 Request Sequence

Addressing: The S_ID designates the Nx_Port requesting Loop Initialization of the loop. The D_ID field shall be the Loop Fabric Address (LFA) of the loop to be initialized.

Payload: The format of the LINIT Request Payload is shown in table 30.

Table 30 – LINIT Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	LINIT (70h)	00h	00h	00h
1	Reserved	Initialization Function	LIP byte 3	LIP byte 4

The Initialization Function field defines modifications of the initialization to be performed. The format of the Initialization Function is shown in table 31.

Table 31 – Initialization Function

Function	Value
Normal Initialization - The Fabric determines the best method by which to complete the initialization.	0
Force Login - The L_bit shall be set in the Loop Initialization Soft Assigned (LISA) Sequence to force all affected L_Ports to perform FLOGI.	1
Reserved	2 - 255

The LIP Byte 3 field is the 3rd byte of the LIP to be originated. The LIP Byte 4 field is the 4th byte of the LIP Primitive Sequence to be originated. LIP Byte 3 and LIP Byte 4 should only be set to values permitted by FC-AL-2.

4.3.15.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the LINIT command

LS_ACC: LS_ACC signifies acceptance of the LINIT Request and completion of Loop Initialization. The format of the LS_ACC Payload is shown in table 32.

Table 32 – LINIT LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Reserved			Status

The format of the Status field is shown in table 33.

Table 33 – LINIT Status

State	Value
Reserved	0
Success - The requested function was completed.	1
Failure - The requested function could not be completed.	2
Reserved	3-255

4.3.16 Loop Status (LSTS)

4.3.16.1 Description

The LSTS ELS is used to request the Fabric Controller to report on the state of the specified loop.

4.3.16.2 Protocol

- a) Loop Status Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.16.3 Request Sequence

Addressing: The S_ID is the Nx_Port requesting status for the specified loop. The D_ID field shall be the Loop Fabric Address (LFA) of the loop for which status is being requested.

Payload: The format of the LSTS Payload is shown in table 34.

Table 34 – LSTS Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	LSTS (72h)	00h	00h	00h

4.3.16.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the LSTS command.

LS_ACC: LS_ACC supplies the requested status for the identified Loop. The format of the LS_ACC Payload is shown in table 35.

Table 35 – LSTS LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Reserved	Failed Receiver	FC-FLA Compliance Level - obsolete	Loop State
2	MSB	Current Public Loop Devices bit map (16 bytes)		
..				
5				LSB
6	MSB	Current Private Loop Devices bit map (16 bytes)		
..				
9				LSB
10	MSB	AL_PA Position Map (128 bytes)		
..				
41				LSB

The Payload fields are defined as follows:

- a) **Failed Receiver field:** The Failed Receiver field shall contain the AL_PA of the L_Port that detected the Loop Failure. This field is valid only if the Loop State indicates a Loop Failure. This field shall be set to 00h if the FL_Port detected the failure. This field shall be set to F7h if the Fabric Controller is unable to determine the port that detected the failure;
- b) **Compliance Level field:** obsolete;

c) **Loop State field:** The format of the Loop State field is shown in table 36;

Table 36 – Loop State

State	Value
Reserved	0
Online – The loop is not performing Loop Initialization, and no Failure has been detected.	1
Loop Failure - A Loop Failure has been detected. The AL_PA that detected the failure is reported in the Failed Receiver field.	2
Initialization Failure - The FL_Port has not been able to complete Loop Initialization.	3
Initializing - The loop is currently performing Loop Initialization.	4
Reserved	5 - 255

- d) **Current Public Loop Devices:** The format of the Current Public Loop Devices field follows the AL_PA bit mapped format defined in FC-AL-2. If a bit is set to one in this field, then a Loop Device that has performed FLOGI is present at the identified location. This field is valid only if the value in the Loop State field is Online;
- e) **Current Private Loop Devices:** The format of the Current Private Loop Devices field follows the AL_PA bit mapped format defined in FC-AL-2. If a bit is set to one in this field, then a Loop Device is present at the identified location. This field is valid only if the value in the Loop State field is Online; and
- f) **AL_PA Position Map:** The format of the AL_PA Position Map field follows the AL_PA position map format defined in FC-AL-2 for the Loop Initialization Loop Position (LILP) Sequence. If Byte 0 of the first word of the AL_PA position map is set to zero, then no AL_PA position map is available. This field is valid only if the value in the Loop State field is Online.

4.3.17 Registered State Change Notification (RSCN)

4.3.17.1 Introduction

A RSCN ELS shall be sent to registered Nx_Ports (see 4.3.18) when an event occurs that may have affected the state of one or more Nx_Ports, or the ULP state within the Nx_Port. The term state is used here to refer to any condition of an Nx_Port that is considered important enough to notify other Nx_Ports of a change in that condition. The definition of important is specific to an Nx_Port implementation, but should include the Login state or Link state. The RSCN additionally provides an indication of the change of state that is being reported.

RSCN is intended to provide a timely indication of changes in nodes to avoid the considerable traffic that polling may generate. RSCN may be used to indicate a failed node, allowing the release of resources tied up by the failed node. RSCN may also be used to notify interested nodes of new devices coming online, and of changes within an online node that affect the operation of the system

(e.g., more storage has become available). The sender of the RSCN Request may coalesce several events into a single report.

A RSCN Request may be sent by the Fabric to notify registered Nx_Ports of changes detected by the Fabric. The Fabric Controller (FFFFFDh) issues the RSCN Request to the registered Nx_Ports. An Nx_Port may also issue a RSCN Request to the Fabric Controller or another Nx_Port to indicate changes of state within the Nx_Port that are not otherwise detectable by the Fabric.

The Payload of a RSCN Request includes a list containing the addresses of the affected Nx_Ports. The RSCN includes a summary indication of the type of state change being reported to assist in analyzing the change. The sending of RSCN between Nx_Ports, neither of which are the Fabric Controller, is permitted.

If an Nx_Port registers for Fabric Detected and Nx_Port Detected events, then the RSCN Request shall include all affected N_Port_ID pages.

4.3.17.2 RSCNs issued by the Fabric Controller

The Fabric Controller shall issue an RSCN Request to all registered Nx_Ports for an affected Nx_Port when the Fabric detects an event. The Fabric Controller shall ensure that any Fabric-provided resources (e.g., the Name Service) have received updates to reflect changes caused by the event, prior to issuing the RSCN for the event. An event may include any of the following:

- a) an implicit Fabric Logout of the affected Nx_Port, including Loss-of-Signal, NOS, and OLS, or if the Fabric receives a FLOGI that contains new or different information from a port that had already completed FLOGI;
- b) a loop initialization of the affected L_Port, and the L_bit was set in the LISA Sequence;
- c) a Fabric Login from an affected Nx_Port not previously logged in;
- d) the Fabric path between the affected Nx_Port and any other Nx_Port has changed (e.g., a change to the Fabric routing tables that affects the ability of the Fabric to deliver frames in order, or an E_Port initialization or failure);
- e) any other Fabric-detected state change of the affected Nx_Port;
- f) an affected Nx_Port issues an RSCN Request to the Fabric Controller; or
- g) a Peer Zone (see FC-GS-9) is created, modified or deleted.

A registered Nx_Port that receives an RSCN Request may perform any operation to determine the nature of the state change. These operations include the PDISC ELS, the ADISC ELS, a query to the Name Service, or a ULP query. The Fabric may accumulate affected Nx_Port addresses for subsequent delivery to reduce the volume of RSCN traffic.

If the event was modification, deletion or addition of a Peer Zone, then the RSCN sent shall have an event qualifier of Changed Peer Zone (i.e., 0111b) in an Affected Port_ID page containing the N_Port_ID of the Principal member.

4.3.17.3 RSCN issued by the affected Nx_Port

An Nx_Port shall issue an RSCN Request to the Fabric Controller or the Nx_Port that has registered for receipt of RSCN Requests or both when an event is detected by an Nx_Port. Fabric Controllers

are implicitly registered to receive RSCNs after a successful implicit or explicit FLOGI. Unless it is providing a service associated with a well known address an Nx_Port shall only issue RSCNs with an event qualifier of 2h or 0h. An event may include any of the following:

- a) a failure within the affected Nx_Port; or
- b) any other important state change of the affected Nx_Port (e.g., an event defined by a FC-4).

4.3.17.4 RSCN initiative

An affected Nx_Port shall issue one RSCN Request for all state changes that occur prior to the initiation of the RSCN Request; multiple RSCN Requests shall not be queued for initiation. An RSCN Request shall be considered initiated when the SOF of the first frame of the RSCN Request Sequence has been transmitted.

4.3.17.5 RSCN registration

An Nx_Port, including a Well-known Address or Fabric Controller, shall only issue RSCN requests to Nx_Ports that have registered either explicitly or implicitly with the Nx_Port to be notified of state changes. These registrations shall be performed using the State Change Registration (SCR) ELS.

4.3.17.6 Protocol

- a) RSCN Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.17.7 Request Sequence

Addressing: If the Fabric is using RSCN to notify a registered Nx_Port of a state change (see 4.3.17.2), the S_ID is the Fabric Controller, (FFFFFDh) and the D_ID is the address of the registered Nx_Port destination. If an affected Nx_Port is using RSCN to notify the Fabric or a registered Nx_Port of a state change (see 4.3.17.3), the S_ID designates either the Nx_Port indicating a state change to the Fabric Controller or the registered Nx_Port and the D_ID is either the Fabric Controller, (FFFFFDh) or the address of the registered Nx_Port destination.

Payload: The format of the RSCN Request Payload is shown in table 37.

Table 37 – RSCN Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RSCN (61h)	Page Length (04h)	Payload Length	
1	affected Port_ID pages (1 to 255 pages, 4 bytes each)			
..				
N				

The RSCN Payload fields are defined as follows:

- a) **Page Length:** This field is the length in bytes of an affected Port_ID page. This value is fixed at 04h;

- b) **Payload Length:** This field is the length in bytes of the entire Payload, inclusive of the word 0. This value shall be a multiple of 4 bytes. The minimum value of this field is 8 bytes. The maximum value of this field is 1024 bytes; and
- c) **Affected Port_ID Pages:** Each affected Port_ID page contains the ID of the Nx_Port, Fabric Controller, domain, or area for which the event was detected. The RSCN Payload shall contain one or more of these pages. The generic format of the affected Port_ID page is shown in table 38.

Table 38 – Generic affected Port_ID page

Bit in Byte Byte	7	6	5	4	3	2	1	0
0 (Bits 31 - 24)	reserved		RSCN Event Qualifier				Address Format	
1 (Bits 23 - 16)	affected Port_ID byte 1 (Domain)							
2 (Bits 15 - 08)	affected Port_ID byte 2 (Area)							
3 (Bits 07 - 00)	affected Port_ID byte 3 (Port)							

A) **RSCN Event Qualifier:** The RSCN Event Qualifier values are shown in table 39; and

Table 39 – RSCN Event Qualifier values

RSCN event Qualifier	Value			
	Bit 5	Bit 4	Bit 3	Bit 2
Event is not specified	0	0	0	0
CHANGED NAME SERVER OBJECT - An object maintained by the Name Server has changed state for the port, area or domain indicated by the affected Port_ID.	0	0	0	1
CHANGED PORT ATTRIBUTE - An internal state of the port specified by the affected Port_ID has changed. The change of state is identified in a protocol specific manner.	0	0	1	0
CHANGED SERVICE OBJECT - An object maintained by the service identified by the well-known address contained in affected Port_ID has changed state. This Event Qualifier value shall not be used by services accessed through N_Port_ID that are not well-known addresses.	0	0	1	1
CHANGED SWITCH CONFIGURATION - Switch configuration has changed for the area or domain specified by the affected Port_ID.	0	1	0	0
REMOVED OBJECT - The port, area or domain indicated by the affected Port_ID is no longer accessible on the Fabric.	0	1	0	1
CHANGED FABRIC NAME - The Fabric_Name has changed for the Fabric. The Address Format shall be set to 3 (i.e., Fabric Address Group).	0	1	1	0
CHANGED PEER ZONE - The Peer Zone that contains the Principal member identified by the affected Port_ID has changed.	0	1	1	1
Reserved	All Other Values			

B) **Address Format:** The format of the Address Format field is shown in table 40.

Table 40 – Address Format

Format	Value
Port Address - Bytes 1, 2, and 3 of the affected Port_ID are valid, and indicate a single Nx_Port or service with a well-known address.	0
Area Address Group - Bytes 1 and 2 of the affected Port_ID are valid, and indicates a group of addresses that encompass an Area of Nx_Port address. Byte 3 shall be zero. Any links and ports within the area may be affected.	1
Domain Address Group - Byte 1 of the affected Port_ID is valid, and indicates a group of addresses that encompass a Domain. Bytes 2 and 3 shall be zero. Any links and ports within the domain may be affected.	2
Fabric Address Group - This format indicates a group of addresses that encompass the entire Fabric of Nx_Port addresses. Bytes 1, 2 and 3 shall be zero. Any links and ports within the area may be affected.	3

4.3.17.8 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the RSCN command

LS_ACC: LS_ACC signifies acceptance of the RSCN Request. The format of the LS_ACC Payload is shown in table 41.

Table 41 – RSCN LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.17.9 RSCN Recovery

If the RSCN sender receives an ABTS in response to the RSCN, then the RSCN sender shall terminate the Exchange with BA_ACC to the ABTS, and retry the original RSCN with a different Exchange at least once.

NOTE 2 – If the retry is sent immediately, the conditions that caused the ABTS to be sent may still be present, therefore a delay of up to E_D_TOV is recommended before each retry.

4.3.18 State Change Registration (SCR)

4.3.18.1 Description

The SCR ELS requests the Fabric Controller or Nx_Port to add the Nx_Port that is sending the SCR Request to the list of Nx_Ports registered to receive the RSCN ELS.

4.3.18.2 Protocol

- a) State Change Registration Request Sequence

b) LS_ACC or LS_RJT Reply Sequence

4.3.18.3 Request Sequence

Addressing: The S_ID designates the Nx_Port requesting registration for State Change Notification. The D_ID designates either the Fabric Controller, FFFFFDh, or the address of the Nx_Port from which RSCNs are desired.

Payload: The format of the SCR Request Payload is shown in table 42.

Table 42 – SCR Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	SCR (62h)	00h	00h	00h
	Reserved			Registration Function Bitmap

Registration Function Bitmap: The format of the Registration Function Bitmap field is shown in table 43.

Table 43 – Registration Function Bitmap

Function	Bit	Field Value
Reserved		0
Obsolete		FFh
Fabric Detected registration - Register to receive all RSCN Requests issued by the Fabric Controller for events detected by the Fabric.	0	01h
Nx_Port Detected registration - Register to receive all RSCN Requests issued for events detected by the affected Nx_Port.	1	02h
Fabric Name Change registration - Register to receive Fabric Name Change RSCN.	2	04h
Peer Zone Change registration - Register to receive Changed Peer Zone RSCN.	3	08h
Reserved	4	10h
Reserved	5	20h
Reserved	6	40h
-	7 ^a	N/A

^a This bit shall be set to zero.

4.3.18.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the SCR command.

LS_ACC: LS_ACC signifies acceptance of the SCR Request and registration for RSCN. The format of the LS_ACC Payload is shown in table 44.

Table 44 – SCR LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.19 Process Login (PRLI)

4.3.19.1 Introduction

The PRLI ELS is used to establish the operating environment between a group of related processes at the originating Nx_Port and a group of related processes at the responding Nx_Port (see 7.1).

4.3.19.2 Protocol

- a) Process Login Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.19.3 Request Sequence

Addressing: The S_ID field designates the Nx_Port requesting Process Login. The D_ID field designates the destination Nx_Port of the Process Login.

Payload: The format of the Payload is shown in table 45.

Table 45 – PRLI Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	PRLI (20h)	Page Length	Payload Length	
1	Service Parameter page			
..				
N				

The Payload Field are defined as follows:

- a) **Page length:** Byte 1 of word 0 contains an 8-bit value that specifies the length of the Service Parameter page. The right-most two bits shall be zeros. The minimum Page Length value is 12 (see table 46);
- b) **Payload length:** Bytes 2-3 of word 0 contain a 16-bit value that specifies the length of the PRLI Payload. The right-most two bits shall be zeros. The value specified shall be greater than or equal to 16, and less than or equal to 256 (see table 46); and
- c) **Service parameter page:** Words 1:N of the PRLI Payload contain the Service Parameter page. The Service Parameter page contains Service Parameters for a single image pair and is

associated with either a single FC-4 TYPE or is common to all FC-4 TYPE codes for the specified image pair. The format of PRLI Service Parameter pages is specified in table 46.

Table 46 – PRLI service parameter page format

Item	Word	Bit
TYPE Code or Common Service Parameters	0	31-24
TYPE Code Extension	0	23-16
Obsolete	0	15
Obsolete	0	14
Establish Image Pair	0	13
Reserved	0	12-0
Obsolete	1	31-0
Obsolete	2	31-0
Service Parameters (optional per FC-4)	3 - 63 as specified by each FC-4	31-0

The PRLI service parameter page format fields are defined as follows:

- a) **TYPE code or common service parameters:** Identifies the protocol associated with this Service Parameter page. If byte 0 of the first word of a Service Parameter page contains the value 00h, the Service Parameter page contains Service Parameters common to all FC-4 Types at that image pair or Nx_Port pair. If byte 0 of the first word of a Service Parameter page contains the value other than 00h, the Service Parameter page contains Service Parameters for the FC-4 TYPE indicated;
- b) **TYPE code extension:** Reserved for future use;
- c) **Establish Image Pair:**
 - 0 = Exchange Service Parameters only,
 - 1 = Establish image pair and exchange Service Parameters;
- d) **Service Parameters:** No Common Service Parameters are currently specified. The length of the optional Service Parameters may be from 0 to 60 words as specified in the respective FC-4 standard (see 7.1.1). The first three words of each PRLI service parameter page shall be as specified in table 46, words 0 to 2.

4.3.19.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the PRLI Request

LS_ACC: LS_ACC signifies successful completion of the PRLI Request. The format of the LS_ACC Payload is shown in table 47.

Table 47 – PRLI LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Page Length	Payload Length	
1				
..	Service Parameter Response page			
N				

The LS_ACC Payload fields are defined as follows:

- a) **Page length:** Byte 1 of word 0 contains an 8-bit value that specifies the length of the Service Parameter Response page. The right-most two bits shall be zeros. The value shall be the same value as in the Page Length field of the PRLI Request;
- b) **Payload length:** Bytes 2-3 of word 0 contain a 16-bit value that specifies the length of the PRLI LS_ACC Payload. The right-most two bits shall be zeros. The value specified shall be greater than or equal to 16, and less than or equal to 65 532; and
- c) **Service parameter response page:** Words 1:N of the PRLI LS_ACC Payload contain the Service Parameter Response page. The Service Parameter Response page contains Service Parameter responses for a single image pair or Nx_Port pair and is associated with a single FC-4 TYPE or common to all FC-4 Types at that image pair or Nx_Port pair.

The format of PRLI LS_ACC Service Parameter Response pages is described in table 48.

Table 48 – PRLI LS_ACC service parameter response page format

Item	Word	Bit
TYPE Code or Common Service Parameters	0	31-24
TYPE Code Extension	0	23-16
Obsolete	0	15
Obsolete	0	14
Image Pair Established	0	13
Reserved	0	12
Response Code (see table 49)	0	11-8
Reserved	0	7-0
Obsolete	1	31-0
Obsolete	2	31-0
Service Parameter Response (optional per FC-4)	3 - 63 as specified by each FC-4	31-0

The PRLI LS_ACC service parameter response page format fields are defined as follows:

- a) **TYPE code or common service parameters:** Identifies the protocol associated with this Service Parameter Response page. If byte 0 of the first word of a Service Parameter Response page contains the value 00h, the Service Parameter page contains Service Parameters common to all FC-4 Types at that image pair or Nx_Port pair. If byte 0 of the first word of a Service Parameter Response page contains a value other than 00h, the Service Parameter page contains Service Parameters for the FC-4 TYPE indicated;
- b) **TYPE code extension:** Reserved for future use;
- c) **Image Pair Established:** Image Pair Established is valid only if bit 13 was set to one on the corresponding Service Parameter page of the PRLI Request.

 0 = Image pair not established, see response code for additional information,
 1 = Image pair established;
- d) **Response code:** The response code field contains an encoded binary value indicating the result of the PRLI Request. The meanings of the encoded response code values are shown in table 49; and

- e) **Service parameter response:** Provides feedback to the Originator as to the resultant state of the Service Parameters as returned by the Responder. The length of the PRLI LS_ACC service parameter response may be 0 to 60 words long as specified in the respective FC-4 standard (see 7.1.1). The first three words of each PRLI LS_ACC service parameter response page shall be as specified in table 48, words 0 to 2.

Table 49 – PRLI accept response code

Encoded Value Word 0, Bits 11-8	Description
0000b	Reserved
0001b	Request executed
0010b	The Exchange recipient has no resources available for establishing image pairs between the specified source and destination Nx_Ports. The PRLI Request may be retried.
0011b	Initialization is not complete for the Exchange recipient. The PRLI Request may be retried.
0100b	Obsolete
0101b	The Exchange recipient has a predefined configuration that precludes establishing this image pair. The PRLI Request shall not be retried.
0110b	Request executed conditionally. Some Service Parameters were not able to be set to their requested state (see table 46)
0111b	Obsolete
1000b	Service Parameters are invalid
1001b to 1111b	Reserved

4.3.20 Process logout (PRLO)

4.3.20.1 Description

The PRLO ELS is used to request invalidation of the operating environment between an image at the initiating Nx_Port and an image at the recipient Nx_Port. PRLO frees resources committed by a previous PRLI function (see 7.2).

4.3.20.2 Protocol

- a) Process Logout Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.20.3 Request Sequence

Addressing: The S_ID field designates the Nx_Port requesting Process Logout. The D_ID field designates the destination Nx_Port of the Process Logout.

Payload: The format of the Payload is shown in table 50.

Table 50 – PRLO Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	PRLO (21h)	Obsolete (10h) ^a	Payload Length	
1	Logout Parameter page (4 words)			
..				
4				
^a This field is obsolete, but shall be set to 10h for compatibility.				

The PRLO Payload fields shall be defined as follows:

- a) **Payload length:** The Payload length shall be 20; and
- b) **Logout parameter page:** Words 1:4 of the PRLO Payload contain the Logout Parameter page. The Logout Parameter page contains logout parameters for a single image pair and is associated with either a single FC-4 TYPE or is common to all FC-4 TYPE codes for the specified image pair.

The format of PRLO Logout Parameter pages is described in table 51.

Table 51 – PRLO logout parameter page format

Item	Word	Bit
TYPE Code or Common Logout Parameters ^a	0	31-24
TYPE Code Extension	0	23-16
Obsolete	0	15
Obsolete	0	14
Reserved	0	13-0
Obsolete	1	31-0
Obsolete	2	31-0
Logout Service Parameters (Optional per FC-4)	3	31-0
^a If byte 0 of the first word of a Logout Parameter page is set to the value 00h, the Logout Parameter page is common to all FC-4 Types at that image pair.		

The PRLO logout parameter page format fields shall be defined as follows:

- a) **TYPE code or common logout parameters:** Identifies the protocol associated with this Logout Parameter page. If byte 0 of the first word of a Logout Parameter page contains the

value 00h, the Logout Parameter page contains Logout Parameters common to all FC-4 Types at that image pair or Nx_Port pair. If byte 0 of the first word of a Logout Parameter page contains the value other than 00h, the Logout Parameter page contains Logout Parameters for the FC-4 TYPE indicated;

- b) **TYPE code extension:** Reserved for future use; and
- c) **Logout Service Parameters:** Word 3 of the PRLO Payload contain the Logout Service Parameters. The Logout Service Parameters contain Service Parameters for a single image pair and is associated with either a single FC-4 TYPE or is common to all FC-4 TYPE codes for the specified image pair. No common Logout Service Parameters are currently specified.

4.3.20.4 Reply sequence

LS_RJT: LS_RJT signifies rejection of the PRLO Request

LS_ACC: LS_ACC signifies successful completion of the PRLO Request

LS_ACC Payload: The format of the LS_ACC Payload is shown in table 52.

Table 52 – PRLO LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Obsolete (10h) ^a	Payload Length	
1	Logout Parameter Response page (4 words)			
..				
N				
^a This field is obsolete, but shall be set to 10h for compatibility.				

The PRLO LS_ACC Payload fields shall be defined as follows:

- a) **Payload length:** The Payload length shall be 20; and
- b) **Logout parameter response page:** Words 1:N of the PRLO LS_ACC Payload contain the Logout Parameter Response page. The Logout Parameter Response page contains a logout parameter response for a single image pair and is associated with a single FC-4 TYPE or common to all FC-4 Types at that image pair or Nx_Port pair. See table 53.

The format of PRLO LS_ACC Logout Parameter Response pages is described in table 53.

Table 53 – PRLO LS_ACC logout parameter response page format

Item	Word	Bit
TYPE Code or Common Logout Parameters ^a	0	31-24
TYPE Code Extension	0	23-16
Obsolete	0	15
Obsolete	0	14
Reserved	0	13-12
Response Code (see table 54)	0	11-8
Reserved	0	7-0
Obsolete	1	31-0
Obsolete	2	31-0
Reserved	3	31-0
^a If byte 0 of the first word of a Logout Parameter page is set to the value 00h, the Logout Parameter page is common to all FC-4 Types at that image pair.		

The PRLO LS_ACC logout parameter response page format fields shall be defined as follows:

- a) **TYPE code or common logout parameters:** Identifies the protocol associated with this Logout Parameter Response page. If byte 0 of the first word of a Logout Parameter Response page contains the value 00h, the Logout Parameter page contains Logout Parameters common to all FC-4 Types at that image pair or Nx_Port pair. If byte 0 of the first word of a Logout Parameter Response page contains a value other than 00h, the Logout Parameter page contains Logout Parameters for the FC-4 TYPE indicated;
- b) **TYPE code extension:** Reserved for future use; and

- c) **Response code:** The Response code field contains an encoded binary value indicating the result of the PRLO Request and the status of the image pair. The meanings of the encoded Response code values are shown in table 54.

Table 54 – PRLO accept response code

Encoded Value Word 0, Bits 11-8	Description
0000b	Reserved
0001b	Request executed
0010b - 0011b	Reserved
0100b	Obsolete
0101b - 0110b	Reserved
0111b	Obsolete
1000b	Obsolete
1001b - 1111b	Reserved

4.3.21 Test Process Login State (TPLS)

4.3.21.1 Description

The TPLS ELS is used to determine whether image pairs are established for the image pairs specified by the combination of the S_ID || D_ID. Upon receiving a TPLS Request, the receiving Nx_Port checks whether it has an image pair established for each specified image.

TPLS verifies the Login state for the source Nx_Port at the destination Nx_Port.

The LS_ACC Reply Sequence confirms the successful completion of the TPLS function and indicates whether or not an image pair is established for the source specified by the S_ID and image pair(s) contained in the Payload. The Payload of the LS_ACC Reply indicates the state of the image pair.

4.3.21.2 Protocol

- a) Test Process Login State Request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.21.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port associated with the image pair. The D_ID field designates the destination Nx_Port associated with the image pair.

Payload: The format of the Payload is shown in table 55.

Table 55 – TPLS Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	TPLS (23h)	Obsolete (10h) ^a	Payload Length	
1	Image Pair ID page (4 words)			
..				
N				
^a This field is obsolete, but shall be set to 10h for compatibility.				

- a) **Payload length:** The Payload length shall be 20; and
- b) **Image pair ID page:** Words 1:N of the TPLS Payload contain the Image Pair ID pages. The Image Pair ID page contains parameters required to identify a single image pair. The format of Image Pair ID pages is described in table 56.

Table 56 – TPLS image pair ID page format

Item	Word	Bit
Reserved	0	31-16
Obsolete	0	15
Obsolete	0	14
Reserved	0	13-0
Obsolete	1	31-0
Obsolete	2	31-0
Reserved	3	31-0

4.3.21.4 Reply sequence

LS_RJT: LS_RJT signifies rejection of the TPLS Request.

LS_ACC: LS_ACC signifies successful completion of the TPLS Request. The format of the LS_ACC Payload is shown in table 57.

Table 57 – TPLS LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	(02h)	Obsolete (10h) ^a	Payload Length	
1	TPLS Response page (4 words)			
..				
N				
^a This field is obsolete, but shall be set to 10h for compatibility.				

- a) **Payload length:** Bytes 2-3 of word 0 contain a 16-bit value that specifies the length of the TPLS LS_ACC Payload. The value shall be 20; and
- b) **TPLS response page:** Words 1:N of the TPLS Payload contain the Image Pair ID page. The TPLS Response page contains TPLS response information associated with a single image pair.

The format of TPLS Response pages is described in table 58.

Table 58 – TPLS response page format

Item	Word	Bits
Reserved	0	31-16
Obsolete	0	15
Obsolete	0	14
Reserved	0	13-12
Response Code (see table 59)	0	11-8
Reserved	0	7-0
Obsolete	1	31-0
Obsolete	2	31-0
Reserved	3	31-1
Image Pair State	3	0

- a) **Response code:** The Response code field contains an encoded binary value indicating the result of the PRLI Request and the status of the image pair. The meanings of the encoded Response code values are shown in table 59; and

Table 59 – TPLS accept response code

Encoded Value Word 0, Bits 11-8	Description
0000b	Reserved
0001b	Request executed
0010b - 0110b	Reserved
0111b	Obsolete
1000b - 1111b	Reserved

- b) **Image pair state:**

1 = image pair established,
0 = image pair not established.

4.3.22 Request Node Identification Data (RNID)

4.3.22.1 Introduction

The RNID ELS is an ELS for acquiring Node Identification Data. The normal response shall be an Accept (i.e., LS_ACC) ELS Sequence with Node Identification Data in its payload. If the recipient Nx_Port or Fx_Port does not support the RNID ELS, it shall reply with an LS_RJT ELS Sequence with a reason code of “Command not supported”. If the recipient Nx_Port or Fx_Port does not support the requested Node Identification Data format, it shall either reply with an LS_RJT ELS Sequence with a reason code of “Unable to perform command request” and reason code explanation of “Unable to supply requested data” or reply with an LS_ACC ELS Payload containing only the Common Identification Data.

If an Nx_Port sends an RNID ELS and receives an LS_RJT Sequence with a reason code of “Unable to perform command request” and reason code explanation of “Unable to supply requested data”, the ULP may cause the Nx_Port to retry the RNID ELS with a different Node Identification Data Format request if an Nx_Port sends an RNID ELS and receives a reply with an LS_ACC ELS containing only Common Identification Data (see Table 65) in its payload, the ULP usage of the Common Identification Data and any further discovery processing is beyond the scope of this standard.

The RNID ELS may be sent to any Nx_Port. If the RNID ELS is sent to the Fabric Controller (i.e., FFFFDh), then the reply shall represent the Switch that contains the Fx_Port to which the requesting FC_Port is attached.

4.3.22.2 Protocol

- a) Request Node-Identification (RNID) request Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.22.3 Request Sequence

Addressing: The S_ID field designates the requesting source port or the Fabric Controller (FFFFFFDh). The D_ID field designates the Nx_Port or F_Port Controller receiving the RNID request or the Fabric Controller (FFFFFFDh).

Payload: The format of the RNID Payload is shown in table 60.

Table 60 – RNID Payload

Bits Word	31 ... 24 Byte 0	23 ... 16 Byte 1	15 ... 08 Byte 2	07 ... 00 Byte 3
0	RNID (78h)	00h	00h	00h
1	Node Identification Data Format	Reserved		

Node Identification Data Format: The format of the Node Identification Data field is shown in table 61.

Table 61 – Node Identification Data Format

Value	Description
00h	Shall be set if the requesting Nx_Port is requesting Common Identification Data only (see table 65).
01h – DEh	For Specific Identification Data corresponding to a specific ULP (e.g., FC-SB-6), shall be set to the FC-4 TYPE (see FC-FS-6) of that ULP.
DFh	Shall be used if the General Topology Discovery format (see 4.3.22.5) is to be returned in the RNID Accept Payload.
E0h – FFh	Shall be used to indicate that Specific Node Identification Data in a vendor specific format is to be returned.

4.3.22.4 Reply Sequence

LS_RJT: Signifies the rejection of the RNID request

LS_ACC: Signifies acceptance of the RNID request and returns Node Identification Data. The format of the Accept Payload is shown in table 62.

Table 62 – RNID Accept Payload

Bits Word	31 ... 24 Byte 0	23 ... 16 Byte 1	15 ... 08 Byte 2	07 ... 00 Byte 3
0	02h	00h	00h	00h
1	Node Identification Data Format	Common Identification Data Length (0 or 16)	VN_Port Phy Type	Specific Identification Data Length
2	Common Identification Data (0 bytes or 16 bytes)			
...				
5	Specific Identification Data (m) (0 - max bytes)			
6				
...				
m				

- a) **Node Identification Data Format:** The value of the Node Identification Data Format field is shown in table 63;

Table 63 – Node Identification Data Format

Value	Description
00h	The RNID Accept Payload only contains the Common Identification Data (see Table 65).
01h – DEh	The RNID Accept Payload may contain the Common Identification Data and shall contain the Specific Identification Data for the ULP that is assigned an FC-4 frame type (see FC-FS-6) equal to the value of the Node Identification Data Format from the RNID Payload (see Table 61).
DFh	The RNID Accept Payload shall contain the Common Identification Data and General Topology Discovery format Specific Identification Data.
E0h – FFh	The RNID Accept Payload may contain the Common Identification Data and shall contain vendor specific Specific Identification Data.

- b) **Common Identification Data Length:** If Common Node Identification Data is included in the RNID Accept payload, the Common Identification Data length shall specify 16 bytes as the size of the Common Identification Data field. If the ULP indicates in the Node Identification Data Format that no (see table 63) Common Node Identification Data is to be included in the RNID Accept payload, the Common Identification Data length shall be set to zero;

- c) **VN_Port Phy Type:** Word 1 bits 15-14 identify the type of physical interface for the PN_Port or PF_Port through which the RNID request was received (see table 64). Bits 8 - 13 of this field are reserved.

Table 64 – VN_Port Phy Type

Encoded Value Word 1, Bits 15 - 14	Description
00b	No Information about Phy Type Provided.
01b	The sending VN_Port uses an FC-FS-6 PN_Port or PF_Port.
10b	The sending VN_Port uses a lossless Ethernet MAC.
11b	Reserved

- d) **Specific Identification Data Length:** The Specific Identification Data Length field indicates the number of bytes in the Specific Identification Data field. The length is determined by the Node Identification Data Format requested for the respective ULP. This value shall be a multiple of 4. The minimum value of this field is zero bytes. The maximum value of this field is 252 bytes;
- e) **Common Identification-Data:** The format of the Common Identification Data field is shown in table 65; and

Table 65 – Common Identification Data

Bits Word	31 ... 24 Byte 0	23 ... 16 Byte 1	15 ... 08 Byte 2	07 ... 00 Byte 3
0	MSB		N_Port_Name	
1			(8 bytes) LSB	
2	MSB		Node_Name	
3			(8 bytes) LSB	

- A) **N_Port_Name:** The N_Port_Name field provides the Name_Identifier (see FC-FS-6) of the Nx_Port or Fx_Port to which the RNID ELS was directed; and
- B) **Node_Name:** The Node_Name field provides the Name_Identifier (see FC-FS-6) of the node associated with the Nx_Port or Fx_Port indicated in the N_Port_Name field.
- f) **Specific Identification Data:** The format of the Specific Identification Data field shall be dependent on the Node Identification Data Format field (see Table 61). If the value of the Node Identification Data Format field is set to DFh, the format is General Topology Discovery (see 4.3.22.5). For all other values of the Node Identification Data Format field the standard for the related ULP defines the format, fields used, and size of the parameters.

NOTE 3 – The information in the Specific Identification Data should only be used to determine the product identity of a node. The information in the Specific Identification Data should not be used to determine the functional characteristics or the service capabilities of a node.

4.3.22.5 General Topology Discovery format:

If the Node Identification Data Format field in the RNID LS_ACC Payload (see table 63) indicates the General Topology Discovery format (DFh) the RNID LS_ACC payload shall contain the following:

- a) Node Identification Data Format field set to DFh;
- b) Common Identification-Data length set to 16;
- c) Specific Identification-Data length set to 52;
- d) Common Node Identification Data (see table 65); and
- e) Specific Identification Data (see table 66).

Table 66 – General Topology Specific Identification Data

Bits Word	31 ... 24 Byte 0	23 ... 16 Byte 1	15 ... 08 Byte 2	07 ... 00 Byte 3
0	MSB			
...	Vendor Unique (16 bytes)			
3	LSB			
4	Associated Type			
5	Physical Port Number			
6	Number of Attached Devices			
7	Node Management	IP Version	UDP/TCP Port Number	
8	MSB			
...	IP Address (16 bytes)			
11	LSB			
12	Reserved		Vendor Specific	

- a) **Associated Type:** The Associated Type (see table 67) is the type of Fibre Channel functionality associated with the node of the Nx_Port or Fx_Port receiving the RNID request (e.g., switch, hub, storage device);
- b) **Physical Port Number:** A vendor unique value that identifies the physical port that has a Fibre Channel link attached;
- c) **Number of Attached Devices:** The number of LCFs attached to the LCF of the Nx_Port or Fx_Port receiving the RNID request. The minimum value is one. The maximum value is 126. For any Nx_Port receiving the RNID, the Number of Attached Devices shall be set to one. For any FL_Port receiving the RNID, the Number of Attached Devices shall be set to the number of nonzero ALPAs known to that FL_Port (i.e., the number of loop devices other than the FL_Port);
- d) **Node Management:** The Node Management field contains the protocol by which a Node Management Entity for the responding node may be accessed (see table 69);

Table 67 – Associated Type

Value (hex)	Type
00 00 00 00	Reserved
00 00 00 01	Unknown
00 00 00 02	Other (none of the following)
00 00 00 03	Hub
00 00 00 04	Switch
00 00 00 05	Gateway
00 00 00 06	obsolete
00 00 00 07	obsolete
00 00 00 09	Storage device (i.e., disk drive, CD-ROM drive, tape drive).
00 00 00 0A	Host
00 00 00 0B	Storage subsystem (e.g., raid, library)
00 00 00 0E	Storage Access Device (e.g., Media changer)
00 00 00 11	NAS server
00 00 00 12	Bridge
00 00 00 13	Virtualization device
xx xx xx FF	Multi-function device (see Table 68 for bit values for xx xx xx)
All Others	Reserved

Table 68 – Multi-function device bit definitions

Bit position	Function
31	Hub
30	Switch
29	Gateway
28	Storage device
27	Host
26	Storage subsystem
25	Storage access device
24	Wavelength division multiplexer
23	NAS server
22	Bridge
21	Virtualization device
20-8	Reserved

Table 69 – Node Management

Value	Protocol	Reference
00h	IP/UDP/SNMP	See RFC 791, RFC 2460, RFC 768, RFC 1157, and RFC 1901
01h	IP/TCP/Telnet	See RFC 791, RFC 2460, RFC 793 and RFC 854
02h	IP/TCP/HTTP	See RFC 791, RFC 2460, RFC 793, and RFC 2616
03h	IP/TCP/HTTPS	See RFC 791, RFC 2460, RFC 793, and RFC 2818
04h	IP/TCP/HTTP/XML	See RFC 791, RFC 2460, RFC 793, RFC 2616, CIM ^a , XML ^b , CIM-XML ^c , and CIM-HTTP ^d
05h - FFh	Reserved	

^a Distributed Management Task Force, *Common Information Model (CIM) in XML*, Version 2.2, June 14, 1999

^b World Wide Web consortium, *Extensible Markup Language (XML) 1.0*, Second Edition, October, 6, 2000

^c Distributed Management Task Force, *Specification for the Representation of CIM in XML (CIM-XML)*, Version 2.0, July 20, 1999

^d Distributed Management Task Force, *Specification for CIM Operations over HTTP (CIM-HTTP)*, Version 1.0, August 11, 1999

e) **IP Version:** The IP versions are shown in table 70;

Table 70 – IP Version

Value	Version	Description
00h	None	Indicates that this port is not providing IP support, and the Node Management, UDP Port, and IP Address fields shall be ignored.
01h	IPv4 (IP version 4)	See RFC 791
02h	IPv6 (IP version 6)	See RFC 2460
03h - FFh	Reserved	

f) **UDP/TCP Port Number:** The numerical value used in UDP (see RFC 768) or TCP (see RFC 793) to distinguish among multiple destinations at the same IP address; and

g) **IP Address:** The IP address by which a Node Management Entity for the responding node may be reached. If the IP version field is set to one, the IPv4 address shall be stored in the least significant word of the IP field and the remainder shall be set to zero.

4.3.23 Registered Link Incident Report (RLIR)

4.3.23.1 Description

The following terms are used to reference the ports, nodes, and records related to link incident reporting:

- a) incident port: A port using the LCF on which the link incident occurred and that detected the link incident. More than one VN_Port associated with a PN_Port may act as the incident port for link incidents at the PN_Port;
- b) connected port or connected node: A port or node using the LCF connected directly to the LCF of the incident port by a Fibre Channel link;
- c) reporting port: The port that sends the Link Incident Record;
- d) registered port: A port that is registered with the reporting port to receive link incident reports; and
- e) reporting node: The node that creates the Link Incident Record describing the incident on the incident port.

The RLIR ELS shall provide a method for a reporting Nx_Port to send a Link Incident Record to a registered Nx_Port. The normal response to an RLIR ELS Sequence shall be a LS_ACC ELS Sequence with no Payload. If the recipient Nx_Port does not support the RLIR ELS, it shall reply with an LS_RJT ELS Sequence with a reason code of "Command not supported". If the recipient Nx_Port is unable to accept the specified Link Incident Record type, it shall reply with an LS_RJT ELS Sequence with a reason code of "Unable to perform command request".

A port shall recognize a link incident if a condition is detected for which an incident code (see table 76) is defined. If a link incident is recognized and a recipient of Link Incident Records is registered for the incident port, a Link Incident Record that contains information related to the link incident shall be created. If a Link Incident Record is generated, the reporting port shall use the link incident reporting procedure to pass the Link Incident Record to a registered port. The link incident reporting mechanism, if any, for a port that does not use this link incident reporting procedure is beyond the scope of this standard.

After a Link Failure is detected, the reporting node may delay recognizing and considering this link failure as a reportable link incident condition, for a specific period of time. Some FC-4 protocols may be capable of transparently recovering from a Link Failure condition, and may not generate a Link Incident Record. If the Link Failure condition persists for longer than the FC-4 specified time out period, then a Link Incident Record is generated and reported.

4.3.23.2 Link Incident reporting procedure

The incident port's reporting node shall generate one Link Incident Record for each link incident. The number of Link Incident Records a reporting node may hold is vendor specific. If the incident port's reporting node attempts to generate a Link Incident Record but does not have resources to hold the record, the oldest Link Incident Record shall be discarded, and the new Link Incident Record shall be retained.

The port selected as the reporting port by the reporting node shall present the Link Incident Records by sending RLIR ELSs to ports registered as valid recipients of Link Incident Records. After sending

all required RLIR ELSs containing the Link Incident Record, the port may discard the Link Incident Record.

The reporting port shall select ports from the registrant list for each Link Incident Record Format. The RLIR shall first be sent to all registered ports that have registered to always receive RLIRs. If no registered ports are registered to always receive RLIR or if RLIR was not successfully transmitted to at least one of those ports registered to always receive RLIRs, then the reporting port shall send the RLIR to a registered port selected from among those that have registered to conditionally receive RLIR. If RLIR is not successfully transmitted to the selected registered port registered to conditionally receive RLIR, the sending port shall select another port, if any, from the established registration list. The RLIR ELS is considered successfully transmitted if a LS_ACC is received for the RLIR within 2xR_A_TOV.

If a reporting port is not able to successfully deliver a RLIR ELS to a registered port, the registered port's registration for the incident port may be discarded. If a reporting node is not able to successfully deliver an RLIR ELS through a particular reporting port, it may select a different reporting port and try again. If no registration is established, it is vendor specific as to whether the incident port's reporting node generates and retains or discards the Link Incident Record.

4.3.23.3 Protocol

- a) RLIR Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.23.4 Request Sequence

Addressing: The S_ID designates the reporting port that is providing the Link Incident Record. The D_ID field designates the registered port.

Payload: The format of the RLIR Payload is shown in table 71.

Table 71 – RLIR Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RLIR (79h)	00h	00h	00h
1	Link Incident Record Format	Common Link Incident Record Length	Common Link Incident-descriptor Length	Specific Link Incident Record Length
2	MSB			LSB
..	Common Link Incident Record (m) (m=4 or 16)			
m+1				
m+2	Common Link Incident descriptor			
	IQ	IC	Obsolete	
m+3	Specific Link Incident Record (0-max bytes)			
..				
n				

- a) **Link Incident Record Format:** The format of the Link Incident Record-Format field is shown in table 72.

Table 72 – Link Incident Record Format

Description	Value
Common Link Incident Record	00h
Specific-coded value	01h - FFh

If the length of the Specific Link Incident Record is zero, the Link Incident Record Format field shall be set to zero. If the length of the Specific Link Incident Record is not zero, the Link Incident Record Format field shall be set to the FC-4 TYPE (see FC-FS-6) associated with the format of the Specific Link Incident Record. A value of zero in the Specific-coded value field indicates that only the Common Link Incident Record and the Common Link Incident descriptor are being reported and that the Specific Link Incident Record Length shall be zero;

- b) **Common Link Incident Record Length:** If the incident port's Name_Identifier (i.e., the N_Port_Name and the Node_Name) only are included in the RLIR ELS Payload, the Common Link Incident Record length shall specify 16 bytes, indicating the size of the Common Link Incident Record field. If the additional Common Link Incident Record fields are included in the Payload, the Common Link Incident Record Length shall specify 64 bytes. The link identified by the N_Port_Name and the Node_Name shall be one end of the link for which the report is being made. The report may be passed through that link or by any other path to the node that has registered to receive RLIR ELSs;
- c) **Common Link Incident descriptor Length:** The Common Link Incident descriptor length shall be set to 4h, specifying the number of bytes in the Link Incident descriptor field;
- d) **Specific Link Incident Record Length:** The Specific Link Incident Record Length shall be specified according to the Link Incident Record Format indicated for the respective ULP. The length shall specify the number of bytes in the Specific Link Incident Record field. This value shall be a multiple of 4. The minimum value of this field is 0. The maximum value of this field is 252;
- e) **Common Link Incident Record:** If the Common Link Incident Length field value is 16 (i.e., m=4), then the Common Link Incident Record shall contain only the first four words (i.e., the Incident port N_Port_Name and Incident port Node Name fields) in table 73. If the Common Link Incident Length field value is 64 (i.e., m=16), then the Common Link Incident Record shall contain all the words specified in table 73. Optional fields (i.e., words 4 through 15) that have unknown values shall be set to zero. The Common Link Incident Record format is specified in table 73;

Table 73 – Common Link Incident Record Data

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	MSB _____ Incident port N_Port_Name _____			
1				LSB
2	MSB _____ Incident port Node_Name _____			
3				LSB
4	Incident Port Type	Obsolete		
5	MSB _____ Connected port N_Port_Name _____			
6				LSB
7	MSB _____ Connected port Node_Name _____			
8				LSB
9	MSB _____ Fabric_Switch_Name _____			
10				LSB
11	Incident Port Number			
12	Transaction ID			
13	reserved			Time Stamp Format
14	MSB _____ Time Stamp _____			
15				LSB

- A) **Incident port N_Port Name:** The Name_Identifier that is the incident port’s N_Port_Name;
- B) **Incident port Node_Name:** The Name_Identifier that is the incident port’s Node_Name;
- C) **Incident Port Type:** The value in the Incident Port Type field specifies the type of the incident port. Port Type values are defined in FC-GS-3. A value of zero in the field indicates the Incident Port Type is unknown or unspecified;
- D) **Connected port N_Port_Name:** The Name_Identifier that is the connected port’s N_Port_Name;
 - aa) If the incident port is an L_Port in a private loop, the connected port shall be the port associated with the link incident, if known;
 - bb) If the incident port is an L_Port in a public loop, the connected port shall be the FL_Port;
 - cc) If the incident port is an Nx_Port connected to a Fabric, the connected port shall be the F_Port;
 - dd) If the incident port is an FL_Port, the connected port shall be the port associated with the link incident, if known; and
 - ee) If the connected port’s N_Port_Name is unknown or unspecified, the Connected Port N_Port_Name may be zero.

- E) **Connected Port Node_Name:** The Name_Identifier that is the Node_Name of the port described by the Connected Port N_Port_Name. If the connected port is an Fx_Port, the Node_Name may be the Fabric_Name, the Node_Name of the local Fabric element, or unspecified. If the Node_Name is unknown or unspecified, the corresponding field shall be zero;
- F) **Fabric_Switch_Name:** The Fabric_Name or the locally-attached Switch_Name for the incident port. For Nx_Port devices this is the value obtained in the FLOGI LS_ACC Payload. Private loop devices do not have a Fabric_Switch_Name and shall report zeros;
- G) **Incident Port Number:** The vendor specific identification of the LCF on which the link incident occurred within the unit having Fibre Channel ports. If this field is presented in the Common Link Incident Record, the Incident Port Number shall be valid;
- H) **Transaction ID:** A 32-bit value that starts at one and is incremented by one for each link incident record generated by the reporting node. Applications receiving link incident records may use this value to eliminate duplicates or detect missing records. The value of zero shall be used only to indicate that the Transaction ID is unknown or unspecified;
- I) **Time Stamp Format:** This field specifies the format of the Time Stamp field as shown in table 74; and
- J) **Time Stamp:** The time stamp value in the format specified by the Time Stamp Format field. If the Time Stamp Format value is 0, the Time Stamp field is unknown or unspecified and may have any value.

Table 74 – Time Stamp Format values

Value	Meaning
00h	The Time Stamp field is unknown or unspecified.
01h	Time Server: The 64 bit time stamp is reported in units of seconds and fractions of a second. The time stamp uses the value obtained from the Time Server (see FC-GS-4).
02h	Clock synchronization format: The 64-bit time stamp is reported as defined for the Clock Synchronization Update (CSU) ELS.
03h to FFh	reserved

f) **Common Link Incident descriptor:**

- A) **Incident Qualifier (IQ):** This field (byte 0 of the Common Link Incident descriptor) qualifies the manner in which the contents of the Link Incident Record shall be interpreted. The meaning is defined in table 75;

Table 75 – Incident Qualifier

Bits	Meaning
31	Reserved.

Table 75 – Incident Qualifier(continued)

Bits	Meaning										
30	Reserved.										
29	Switch: If set to one, indicates that the incident port is a port on a switch. If set to zero, indicates that the incident port is not a port on a switch.										
28	Expansion Port: If set to one, indicates that the switch port is an Inter-Switch-Link Expansion port (E_Port). If zero, bit 28 indicates that the switch port is not an Inter-Switch-Link Expansion port.										
27-26	Severity Indication: Bits 27-26 constitute a two-bit code that identifies the severity indication for the link incident. The codes and their meanings are as follows:										
	<table border="1"> <thead> <tr> <th data-bbox="428 722 509 772">Code</th> <th data-bbox="509 722 1419 772">Meaning</th> </tr> </thead> <tbody> <tr> <td data-bbox="428 772 509 863">0</td> <td data-bbox="509 772 1419 863">Informational report: Indicates link incident notification of an informational purpose.</td> </tr> <tr> <td data-bbox="428 863 509 984">1</td> <td data-bbox="509 863 1419 984">Link degraded but operational: Indicates if the link associated with the incident port is not in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.</td> </tr> <tr> <td data-bbox="428 984 509 1106">2</td> <td data-bbox="509 984 1419 1106">Link not operational: Indicates if the link associated with the incident port is in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.</td> </tr> <tr> <td data-bbox="428 1106 509 1163">3</td> <td data-bbox="509 1106 1419 1163">Reserved.</td> </tr> </tbody> </table>	Code	Meaning	0	Informational report: Indicates link incident notification of an informational purpose.	1	Link degraded but operational: Indicates if the link associated with the incident port is not in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.	2	Link not operational: Indicates if the link associated with the incident port is in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.	3	Reserved.
	Code	Meaning									
	0	Informational report: Indicates link incident notification of an informational purpose.									
	1	Link degraded but operational: Indicates if the link associated with the incident port is not in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.									
2	Link not operational: Indicates if the link associated with the incident port is in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.										
3	Reserved.										
0	Informational report: Indicates link incident notification of an informational purpose.										
1	Link degraded but operational: Indicates if the link associated with the incident port is not in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.										
2	Link not operational: Indicates if the link associated with the incident port is in a Link-Failure or Offline State as a result of the event that generated the Link Incident Record.										
3	Reserved.										
25	Subassembly type: If set to one, specifies that the type of subassembly used for the port that is the subject of this Link Incident Record is a laser. If set to zero, specifies that the type of subassembly used for the port that is the subject of this Link Incident Record is not a laser.										
24	FRU identification: If set to one, specifies that the Specific-Link Incident Record Data is in a format that provides field-replaceable-unit (FRU) identification. If set to zero, specifies that the Specific-Link Incident Record Data is not in a format that provides field-replaceable-unit (FRU) identification.										

- B) Incident Code (IC): This field (byte 1 of the Common Link Incident descriptor) contains the incident code that describes the incident that was observed by the incident node; and
 - aa) Bit 23 is reserved; and
 - bb) Bits 22-16 is the value that specifies the type of incident that was observed. The values specified and their meanings are shown in table 76.

Table 76 – Incident Code values

Value	Meaning
00h	Reserved.
01h	Implicit incident: A condition, caused by an event known to have occurred within the LCF of the incident port, has been recognized by the incident port. The condition affects the attached link in such a way that it may cause a link incident to be recognized by the LCF of the connected port.
02h	Bit-error-rate threshold exceeded: The LCF of the incident port has detected that the Error Interval Count equals the Error Threshold (see FC-FS-6).
03h	Link Failure - Loss-of-Signal or synchronization: The LCF of the incident port has recognized a Loss-of-Synchronization condition, and it persisted for more than the R_T_TOV timeout period (see FC-FS-6).
04h	Link Failure - NOS recognized: The NOS has been recognized by the LCF of the incident port (see FC-FS-6)
05h	Link Failure - Primitive Sequence timeout: The LCF of the incident port has recognized either a Link-Reset-Protocol timeout (see FC-FS-6), or a timeout when timing for the appropriate response while in the LF1 State and after NOS is no longer recognized (see FC-FS-6).
06h	Link Failure - Invalid Primitive Sequence for port state: The LCF of the incident port recognized either a LR or LRR Primitive Sequence while in the OL3 State (see FC-FS-6).
07h	Link Failure - Loop Initialization time out: The LCF of the incident port failed to complete loop initialization within the normal loop time out period (see FC-AL-2).
08h	Link Failure – receiving LIP(F8): The LCF of the incident port is receiving LIP(F8) indicating some other LCF on the loop is experiencing a Loss-of-Signal condition.
09h - FFh	Reserved.

C) **Specific Link Incident Record:** The format of the Specific Link Incident Record Field is dependent on the Link Incident Record Format. Specific Link Incident Record formats are defined in the specification for the specific FC-4 (e.g., FC-SB-6) indicated by the Link Incident Record Format.

4.3.23.5 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the RLIR Request

LS_ACC: LS_ACC signifies acceptance of the RLIR Request and its Link Incident Record. The format of the LS_ACC Payload is shown in table 77.

Table 77 – RLIR LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.24 Link Incident Record Registration (LIRR)

4.3.24.1 Description

See 4.3.23.1 for terms related to link incident reporting.

The LIRR ELS requests the recipient to add or remove this source Nx_Port to or from the list of Nx_Ports registered to receive the Registered Link Incident Report (RLIR) ELS. The normal response to a LIRR ELS Sequence shall be a LS_ACC ELS Sequence with no Payload. If the recipient Nx_Port does not support the LIRR ELS, it shall reply with a LS_RJT ELS Sequence with a reason code of “Command not supported”. If the recipient Nx_Port is unable to perform the requested LIRR registration function or the specified format is not supported, it shall reply with a LS_RJT ELS Sequence with a reason code of “Unable to perform command request”.

If an Nx_Port is registering or de-registering with the Fabric for receipt of Link Incident Records from the Fabric, the LIRR ELS shall be sent to the Management Server (FFFFFFAh).

4.3.24.2 Registration for Link Incident Records

To obtain and process Link Incident Records, a ULP shall register its port with other ports of interest. The term “ports of interest” refers to all ports that the registered port may address and for which the port requires their Link Incident Records.

An Nx_Port shall attempt to register as a valid recipient of subsequent Link Incident Records by sending a LIRR ELS as soon as the requesting Nx_Port has determined its address identifier and has determined the address identifiers for the ports of interest. The ULP shall indicate the format of Link Incident Records of interest that are being registered. The recipient Nx_Port shall respond with a LS_ACC ELS with no Payload.

If a registered ULP is no longer interested in receiving Link Incident Records from any port or ports of interest, it shall send a LIRR ELS to de-register their port as a valid recipient of Link Incident Records. The ULP shall indicate the format of Link Incident Records that are being de-registered.

If a LIRR ELS requesting a de-registration function for a port that is not registered is received, the recipient port accepts the request by sending a LS_ACC ELS and performs no further action. If a LIRR ELS requesting a registration for a port that is already registered for the specified format, the recipient port shall accept by sending a LS_ACC ELS request and perform no further action.

The recipient of a LIRR ELS shall maintain a separate established registration list of valid Link Incident Record recipients for each registration format specified. The size of the established registration list maintained by each port for each specified format is vendor specific. If the established registration list is full for an add registration function request, the port shall reply with a LS_RJT ELS Sequence with a reason code of "Unable to perform command request".

A Logout occurring between a reporting port and a registered port shall cause de-registration of the valid-recipient port in the established registration list maintained by the port.

4.3.24.3 Responsibilities of Valid-Registered Recipients

A valid-registered recipient shall be capable of receiving and processing Link Incident Records that are generated according to the Link Incident reporting procedure (see 4.3.23.2). The usage and coordination of Link Incident Records by ULPs is beyond the scope of this standard.

If a valid registered recipient receives a Link Incident Record for a format that does not match its registration format, the recipient port shall reject the RLIR ELS with a reason code of “Unable to perform command request”.

4.3.24.4 Protocol

- a) LIRR Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.24.5 Request Sequence

Addressing: The S_ID designates the source port requesting registration for Link Incident Records. The D_ID field designates the port receiving the registration request.

Payload: The format of the LIRR Payload is shown in table 78.

Table 78 – LIRR Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	LIRR (7Ah)	00h	00h	00h
1	Registration Function	Link Incident Record-Registration Format	Reserved	

The LIRR Payload field shall be defined as follows:

- a) **Registration Function:** The Registration function shall specify the mode of registration (e.g., whether the recipient node adds or removes the sending Nx_Port to or from its list of registered

nodes, hereafter referred to as the established registration list). The format of the Registration Function field is shown in table 79; and

Table 79 – Registration Function

value	Function
00h	Reserved
01h	Set registration – conditionally receive: The source port is registered as a valid recipient of subsequent RLIR ELSs for the format specified. The port is added to the appropriate format specific established registration list. This source port is chosen as the recipient of a link incident record only if no other recipients from this established registration list have been chosen.
02h	Set registration – always receive: The source port is registered as a valid recipient of subsequent RLIR ELSs for the format specified. The port is added to the appropriate format specific established registration list. This source port is always chosen as a recipient of a link incident record.
03h - FEh	Reserved
FFh	Clear registration: The source port is de-registered as a valid recipient of subsequent RLIR ELSs for the format specified (i.e., remove from the established registration list).

b) **Link Incident Record-Registration Type:** The type of the requested Link Incident Record Registration is shown in table 80.

Table 80 – Link Incident Record-Registration Format

Value	Description
00h	Common Format
01h - FFh	Specific-coded value

A Link Incident Record Registration Format of 00h specifies that only the Common Link Incident Record and the Link Incident descriptor is reported. If the Link Incident Record Registration format is non-zero, it shall be set to the Specific-coded value (FC-4 Device Type) of the registration format being requested for a specific client ULP (e.g., see FC-SB-6). The Specific-coded value uses the TYPE codes defined in FC-4 data structures of this standard. These codes allow clients to define and register for Link Incident Record formats for their own particular usage according to the unique FC-4 protocol.

4.3.24.6 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the LIRR Request

LS_ACC: LS_ACC Signifies acceptance of the LIRR Request and the registration for RLIR ELSs. The format of the LS_ACC Payload is shown in table 81.

Table 81 – LIRR LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.25 Discover N_Port/Service Parameters (PDISC)

4.3.25.1 Description

The PDISC ELS shall transfer Service Parameters from the initiating Nx_Port to the Nx_Port associated with the D_ID without affecting the operating environment between the two ports. This provides the means for exchange of Service Parameters without terminating open Sequences or open Exchanges.

4.3.25.2 Protocol

- a) PDISC Sequence
- b) LS_ACC or LS_RJT Reply Sequence

4.3.25.3 Request Sequence

The Payload in PDISC and the LS_ACC are, except for the ELS_Command, identical to the Payloads in PLOGI (see 4.3.7) and the corresponding LS_ACC (see 4.4.2). The interchange of PDISC information shall not modify the operating environment or Service Parameters between the two ports. Service Parameters contained in PDISC shall be ignored and no error condition shall be reported.

4.3.25.4 Reply Sequence

A response shall not be sent to a PDISC if a LR or loop initialization occurs before the LS_ACC is sent. An LS_ACC for a PDISC shall be ignored if a LR or loop initialization occurred between the PDISC and the LS_ACC.

4.3.26 Discover F_Port Service Parameters (FDISC)

4.3.26.1 Description

The FDISC ELS shall transfer Service Parameters from the initiating Nx_Port to the Fx_Port at well-known F_Port_ID (i.e., FFFFFFFh). This provides the means for the exchange of Service Parameters and the assignment of additional N_Port_IDs without changing service parameters.

4.3.26.2 Protocol

- a) FDISC Sequence
- b) LS_ACC, LS_RJT Reply Sequence

4.3.26.3 Request Sequence

The Payload in FDISC and the corresponding LS_ACC are, except for the ELS_Command, identical to the Payloads in FLOGI (see 4.3.7) and the LS_ACC (see 4.4.2). If the S_ID of the FDISC ELS is equal to zero, the FDISC ELS shall transfer an N_Port_Name and Node_Name, and request the assignment and Login of an additional N_Port_ID (see 6.3). The FDISC ELS with an S_ID equal to zero shall only be sent by a PN_Port with at least one N_Port_ID that is currently logged in with the Fabric, and the assignment of an additional N_Port_ID shall not modify the Service Parameters between the two FC_Ports.

If the S_ID of the FDISC ELS is set to a previously-assigned N_Port_ID, the FDISC ELS shall contain the N_Port_Name and Node_Name corresponding to that N_Port_ID. The interchange of FDISC information shall not modify the operating environment or Service Parameters between the two FC_Ports. Service Parameters contained in FDISC shall be ignored and no error condition shall be reported.

If the Originator of a FDISC with S_ID equal to zero detects a time out, the FDISC may be retried in a different Exchange without aborting the previous Exchange with ABTS-LS.

4.3.26.4 Reply Sequence

If the S_ID of the FDISC ELS is zero, the D_ID field of the LS_ACC shall be set to the additional N_Port_ID being assigned. After the F_Port Controller sends the LS_ACC, it shall consider Fabric Login to have occurred for the Nx_Port to which the additional N_Port_ID was assigned.

If the S_ID of the FDISC ELS is not zero and is currently logged in, the D_ID of the LS_ACC shall be set to the S_ID of the FDISC ELS.

If an FDISC ELS with an S_ID set to zero is received by a F_Port Controller and no N_Port_ID is logged in, or if the S_ID is set to a non-zero N_Port_ID that is not currently logged, then an F_RJT with the reason code set to 'Login required' shall be returned to the S_ID of the FDISC ELS if it was issued in Class 2, and the FDISC ELS shall be discarded if it was issued in Class 3.

4.3.27 Discover Address (ADISC)

4.3.27.1 Description

The ADISC ELS shall exchange addresses and identifiers of communicating Nx_Ports.

4.3.27.2 Protocol

- a) ADISC Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.27.3 Request Sequence

The ADISC Request payload format is shown in table 82.

Table 82 – ADISC Request payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	ADISC (52h)	00h	00h	00h
1	Reserved	Hard Address of Originator		
2	MSB	N_Port_Name of Originator		LSB
3		(8 bytes)		
4	MSB	Node_Name of Originator		LSB
5		(8 bytes)		
6	Reserved	N_Port_ID of Originator		

- a) **Hard Address of Originator:** A 24-bit L_Port_ID that consists of:
- A) the 8-bit Domain address in the most significant byte;
 - B) the 8-bit Area address in the next most significant byte; and
 - C) the 8-bit AL_PA that an L_Port attempts to acquire in the LIHA sequence during loop initialization in the least significant byte (see FC-AL-2).

If an L_Port does not have a hard address, or if a port does not have FC-AL-2 capability, it shall report zeroes in this field.

- b) **N_Port_Name of Originator:** The 8-byte N_Port_Name of the Originator;
- c) **Node_Name of Originator:** The 8-byte Node_Name of the Originator; and
- d) **N_Port_ID of Originator:** The 24-bit S_ID used in the header of the ADISC Request frame.

4.3.27.4 Reply Sequence

The LS_ACC Payload format is shown in table 83.

Table 83 – ADISC LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Reserved	Hard Address of Responder		
2	MSB N_Port_Name of Responder			LSB
3	(8 bytes)			
4	MSB Node_Name of Responder			LSB
5	(8 bytes)			
6	Reserved	N_Port_ID of Responder		

- a) **Hard Address:** A 24-bit L_Port_ID that consists of:
 - A) the 8-bit Domain address in the most significant byte;
 - B) the 8-bit Area address in the next most significant byte; and
 - C) the 8-bit AL_PA that an L_Port attempts to acquire in the LIHA sequence during loop initialization in the least significant byte (see FC-AL-2).

If an L_Port does not have a hard address, or if a port does not have FC-AL-2 capability, it shall report zeroes in this field.

- b) **N_Port_Name of Responder:** The 8-byte N_Port_Name of the Responder;
- c) **Node_Name of Responder:** The 8-byte Node_Name of the Responder; and
- d) **N_Port_ID of Responder:** The 24-bit S_ID used in the header of the ADISC LS_ACC frame.

A response shall not be sent to an ADISC if an LR or loop initialization occurs before the LS_ACC is sent. An LS_ACC for a ADISC shall be ignored if an LR or loop initialization occurred between the ADISC and the LS_ACC.

Table 84 summarizes the responses to FDISC, PDISC and ADISC.

Table 84 – Response summary to FDISC/PDISC

ELS command	Responding Port Status	Responding Nx_Port	Responding F_Port Controller	
		Class 2, or 3	Class 2	Class 3
FDISC	Logged in	LS_RJT ^c	LS_ACC	LS_ACC
	Not Logged in	LS_RJT ^c	F_RJT ^b	Discard
PDISC	Logged in	LS_ACC	LS_RJT ^c	Discard
	Not Logged in	LS_RJT ^a	F_RJT ^b	Discard
ADISC	Logged in	LS_ACC	LS_RJT ^c	Discard
	Not Logged in	LS_RJT ^a	F_RJT ^b	Discard

^a A LOGO ELS sequence or an LS_RJT ELS Sequence with the reason code set to “Unable to perform command request” and the reason code explanation set to “N_Port Login required” shall be returned.

^b An F_RJT with the Reject reason code set to “Login required” shall be returned.

^c A LOGO ELS Sequence or an LS_RJT ELS Sequence with the reason code set to “Command not supported” and the reason code explanation set to “Request not supported” shall be returned.

4.3.28 Third Party Process Logout (TPRLO)

4.3.28.1 Description

The TPRLO ELS is used to invalidate the operating environments (i.e., remove image pairs and associated resources) at the recipient Nx_Port for the specified TYPE.

TPRLO has the same effect on the recipient Nx_Port as if all Nx_Ports that have performed Process Login with the recipient Nx_Port for the specified TYPE, performed PRLO with the recipient Nx_Port.

4.3.28.2 Protocol

- a) TPRLO Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.28.3 Request Sequence

Addressing: The S_ID field designates the Nx_Port initiating TPRLO. The D_ID field designates the destination Nx_Port of the TPRLO.

Payload: The TPRLO format is shown in table 85.

Table 85 – TPRLO Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	TPRLO (24h)	Obsolete (10h) ^a	Payload Length	
1	Logout Parameter page (4 words)			
..				
n				
^a This field is obsolete, but shall be set to 10h for compatibility.				

- a) **Payload length:** The payload length shall be 20; and
- b) **Logout parameter page:** Words 1:N of the TPRLO Payload contain the Logout Parameter page. The Logout Parameter page contains logout parameters for a single image pair and is associated with either a single FC-4 TYPE or is common to all FC-4 TYPE codes for the specified image pair.

The TPRLO Logout parameter page is shown in table 86.

Table 86 – TPRLO logout parameter page

Item	Word	Bit
TYPE Code or Common Service Parameters	0	31-24
TYPE Code Extension	0	23-16
Third Party Originator Process_Associator Validity - obsolete	0	15
Responder Process_Associator Validity - obsolete	0	14
Third Party Originator N_Port_ID Validity - obsolete	0	13
Global Process Logout	0	12
Reserved	0	11-0
Third Party Originator Process_Associator - obsolete	1	31-0
Responder Process_Associator - obsolete	2	31-0
Reserved	3	31-24
Third Party Originator N_Port_ID - obsolete	3	23-0

- a) **TYPE Code or Common Logout Parameters:** Identifies the protocol associated with this TPRLO Logout Parameter page. If byte 0 of the first word of a TPRLO Logout Parameter page contains the value 00h, the TPRLO Logout Parameter page contains logout parameters

common to all FC-4 Types at that image pair or Nx_Port pair. If byte 0 of the first word of a TPRLO Logout Parameter page contains the value other than 00h, the TPRLO Logout Parameter page contains logout parameters for the FC-4 TYPE indicated; and

- b) **TYPE code extension:** Reserved for future use;
- c) **Global Process Logout:** The Global Process Logout bit shall be set to one, and specifies:
 - A) only the TYPE code and TYPE code extension fields shall have meaning, along with the Global Process Logout bit itself;
 - B) all image pairs for all Nx_Ports with which Process Login has been performed shall be removed from the recipient Nx_Port for the specified TYPE; and
 - C) all resources associated with the establishment of all image pairs of the specified TYPE at the recipient Nx_Port shall be released.

The TPRLO LS_ACC Payload format is shown table 87.

Table 87 – TPRLO LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Obsolete (10h) ^a	Payload Length	
1	Logout Parameter page (4 words)			
..				
n				
^a This field is obsolete, but shall be set to 10h for compatibility.				

- a) **Payload Length:** The Payload Length field specifies the length of the TPRLO LS_ACC payload. The Payload Length value shall be 20; and
- b) **Logout Parameter page:** The TPRLO LS_ACC Logout Parameter page format is the same as the TPRLO Request Logout Parameter page format (see table 86) and the field values shall be set to the values received in the TPRLO Request Logout Parameter page.

4.3.29 Clock Synchronization Request (CSR)

4.3.29.1 Description

The CSR ELS is used to request the Clock Synchronization Server to either send or to quit sending periodic Clock Synchronization Update (CSU) ELS frames or Clock Synchronization primitives, depending on the method implemented (see FC-FS-6).

4.3.29.2 Protocol

- a) Clock Synchronization Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.29.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting Clock Synchronization updates. The D_ID field designates either the Clock Synchronization Server (FFFFF6h) or the Fabric Controller (FFFFFDh).

Payload: The format of the Payload is shown in table 88.

Table 88 – CSR Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	CSR (68h)	00h	00h	00h
1	Clock Sync Mode	CS_Accuracy	CS_Implement- ed_MSB	CS_Implement- ed_LSB
3	CS_Update_Period			

- a) **Clock Sync Mode:** The meaning of the Clock Sync Mode byte in the CSR Payload is defined in table 89;

Table 89 – CSR Clock Sync Mode Meaning

Value	S_ID set to the Clock Sync Server	S_ID set to the Fabric Controller
00h	Enable Clock Synchronization service to this client. The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields are not meaningful in the CSR Request.	Return Quality of Service parameters. The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields are not meaningful in the CSR Request.
01h	Enable Clock Synchronization service to this client. The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields contain the requested Quality of Service parameters.	Return Quality of Service parameters. The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields contain the requested Quality of Service parameters.
02h – FEh	Reserved	Reserved
FFh	Disable Clock Synchronization service to this client	Reserved

- b) **CS_Accuracy (Mantissa and Exponent):** This field contains the CS_Accuracy_Mantissa (Bits 23-21) and CS_Accuracy_Exponent (Bits 20-16). These bits indicate the requested accuracy of the Clock Synchronization value as it leaves the server port. Specifically, the request is that the Clock Count value is always within the range of:

$$T_{reference} \pm (0.5 + CS_Accuracy_Mantissa * 2^{-4}) * 2^{(CS_Accuracy_Exponent-30)},$$

where

- A) T_reference is the clock reference value internal to the server;
- B) CS_Accuracy_Mantissa is a value from 000b to 111b; and
- C) CS_Accuracy_Exponent is a value from 00000b to 11111b;

Example #1, if CS_Accuracy Mantissa and Exponent = 001b and 01011b, respectively, the Clock Synchronization value as it exits the server is requested to be within the range of:

$$T_reference \pm 1.073 \mu\text{sec}$$

Example #2, if CS_Accuracy Mantissa and Exponent = 111b and 11000b, respectively, the Clock Synchronization value as it exits the server is requested to be within the range of:

$$T_reference \pm 14.65 \text{ msec}$$

- c) **CS_Implemented_MSB:** This field is a value that is constrained to the range of 0 to 63. These bits indicate the requested most significant bit position within the 64-bit Clock Count field (e.g., a value of '110111b' indicates that the client requests that the highest bit that contains meaningful information be the MSB of byte 1 of the Clock Count field);
- d) **CS_Implemented_LSB:** This field is a value that is constrained to the range of 0 to 63. These bits indicate the requested least significant bit position within the 64-bit Clock Count field (e.g., a value of 001000b indicates that the client requests that the lowest bit that contains meaningful information be the LSB of byte 6 of the Clock Count field); and
- e) **CS_Update_Period:** This field is a value. It represents the requested time, in microseconds, between consecutive updates from the Clock Synchronization server.

4.3.29.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the CSR command.

LS_ACC: LS_ACC signifies that the Clock Synchronization Server agrees to perform the action requested in the CSR Payload. The format of the LS_ACC Payload is shown in table 90. The Clock Sync Mode item in the CSR LS_ACC Payload shall contain the value that was received in the Clock Sync Mode item of the CSR Payload.

Table 90 – CSR LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Clock Sync Mode	CS_Accuracy	CS_Implement- ed_MSB	CS_Implement- ed_LSB
2	CS_Update_Period			

- a) **Clock Sync Mode:** The meaning of the Clock Sync Mode byte in the CSR Payload is defined in table 91;

Table 91 – CSU Clock Sync Mode Meaning

Hex Value	Meaning
00h	Clock synchronization service enabled to this client
01h - FEh	Reserved
FFh	Clock synchronization service disabled to this client

- b) **CS_Accuracy (Mantissa and Exponent):** This field consists of two values, CS_Accuracy_Mantissa (Bits 23-21) and CS_Accuracy_Exponent (Bits 20-16). These bits indicate the accuracy of the Clock Synchronization value as it leaves the server port. Specifically, the server shall supply a CS_Accuracy value such that the Clock Count value is always within the range of:

$$T_{\text{reference}} \pm (0.5 + \text{CS_Accuracy_Mantissa} * 2^{-4}) * 2^{(\text{CS_Accuracy_Exponent}-30)},$$

where

- A) T_reference is the clock reference value internal to the server;
- B) CS_Accuracy_Mantissa is a value from 000b to 111b; and
- C) CS_Accuracy_Exponent is a value from 00000b to 11111b;

Example #1, if CS_Accuracy Mantissa and Exponent = 001b and 01011b, respectively, the Clock Synchronization value as it exits the server shall always be within the range of:

$$T_{\text{reference}} \pm 1.073 \mu\text{sec}$$

Example #2, if CS_Accuracy Mantissa and Exponent = 111b and 11000b, respectively, the Clock Synchronization value as it exits the server shall always be within the range of:

$$T_{\text{reference}} \pm 14.65 \text{ msec}$$

- c) **CS_Implemented_MSB:** This field is a value that is constrained to the range of 0 to 63. It represents the most significant bit position within the Clock Count field that shall contain meaningful information (e.g., a value of 110111b indicates that the MSB of byte 1 of the Clock Count field is the highest bit that contains meaningful information);
- d) **CS_Implemented_LSB:** This field is a value that is constrained to the range of 0 to 63. It represents the least significant bit position within the Clock Count field that shall contain meaningful information (e.g., a value of 001000b indicates that the LSB of byte 6 of the Clock Count field is the lowest bit that contains meaningful information); and
- e) **CS_Update_Period:** This field is a value. It represents the time, in microseconds, between consecutive updates from the Clock Synchronization server.

4.3.30 Clock Synchronization Update (CSU)

4.3.30.1 Description

The CSU ELS is used by the Clock Synchronization Server to send its current clock value to its clients (see FC-FS-6).

4.3.30.2 Protocol

- a) Clock Synchronization Update Request Sequence; and
- b) No Reply Sequence.

4.3.30.3 Request Sequence

Information Category: The Information Category value (Header word 0, bits 24-27) shall indicate Solicited Data (0001b).

Addressing: The S_ID field designates the Clock Synchronization Server well-known address (FFFFF6h). The D_ID field designates the Nx_Port(s) that is/are to receive the clock information.

Payload: The format of the Payload is shown in table 92.

Table 92 – CSU Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	CSU (69h)	Reserved		
1	Clock Count (8 bytes)			
2				

The meaning of the Clock Count field is given in table 93.

Table 93 – Clock Count Field Meaning

Byte Number	Meaning
0	Counter value, byte 0 (MSB)
1	Counter value, byte 1
2	Counter value, byte 2
3	Counter value, byte 3
4	Counter value, byte 4
5	Counter value, byte 5
6	Counter value, byte 6
7	Counter value, byte 7 (LSB)

The bit values are derived from clock frequencies that are used in 1Gbits/s Fibre Channel and shall be defined as follows. The value of the Bit 7 in Word 2 shall be equal to 1/106.25MHz, roughly 9.4 ns.

Every other bit value is a binary multiple of this value. The next most significant bit is 2x that value, or 18.8ns. The next least significant value is ½ that value, or 4.7ns. The overall least significant bit is 73.5ps. The overall range that may be represented is 1.36 x 10⁹ sec, approximately equal to 43 years.

The Clock Count value shall represent the time at which the most significant bit was placed on the link by the CSU ELS originator.

Any bits outside the range of CS_Implemented_MSB to CS_Implemented_LSB shall be set to zero. This applies to both the Clock Sync Server and to the Fabric.

4.3.30.4 Reply Sequence

none.

4.3.31 Report Port Buffer Conditions (RPBC)

4.3.31.1 Description

The RPBC ELS shall provide a method for a Port to report its buffer conditions. The normal response to an RPBC ELS Sequence shall be a LS_ACC ELS Sequence. If the recipient Port does not support the RPBC ELS, it shall reply with an LS_RJT ELS Sequence with a reason code of “Command not supported”.

4.3.31.2 Protocol

- a) Report Port Buffer Conditions (RPBC) Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.31.3 Request Sequence

Addressing: The S_ID designates the source Nx_port that is requesting port buffer conditions. The D_ID field designates the destination Nx_Port or F_Port Controller (FFFFFFEh).

Payload: The format of the RPBC Payload is shown in table 94.

Table 94 – RPBC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RPBC (58h)	Reserved		
1	ELS Buffer Parameters			
2	Reserved	Originator S_ID - obsolete		

The ELS Buffer Parameter field contained in word 1 of the RPBC Payload relates to the ELS buffer conditions of the sender of the RPBC ELS and the contents are shown in table 95.

Table 95 – ELS Buffer Parameters Field

Bits	Field Name
31	Multi-frame ELS sequence supported
30-12	Reserved.
11-00	ELS Receive Data_Field Size

The Multi-frame ELS sequence supported bit indicates that the port either supports multi-frame ELS's (bit 31=1) or does not support multi-frame ELS's (bit 31=0).

The ELS Receive Data_Field Size field (word 1, bits 11-0) specifies the largest ELS frame that may be received by the FC_Port providing the Buffers Parameter Field. The ELS Receive Data_Field Size field (word 1, bits 11-0) specifies the largest ELS frame that may be received by the FC_Port responding with the RPBC LS_ACC. Values shall be a multiple of four bytes. Values less than 256 or greater than 2 112 are invalid. For each class of service, the maximum usable ELS Receive Data_Field Size is the lessor of the reported ELS Receive Data_Field Size or the Buffer-to-Buffer Receive Data_Field Size reported during Login.

When provided in the Request Sequence, the ELS Receive Data_Field Size specifies the largest ELS frame that may be received by the Originator FC_Port.

4.3.31.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RPBC command

LS_ACC: LS_ACC signifies acceptance of the RPBC Request, and the included payload provides associated values from the sender. The format of the RPBC LS_ACC Payload is shown in table 96.

Table 96 – RPBC LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Reserved		
1	ELS Buffer Parameters			

The ELS Buffer Parameter field contained in word 1 of the RPBC LS_ACC Payload relates to the ELS buffer conditions of the sender of the RPBC LS_ACC ELS and the contents are shown in table 95.

When provided in the Reply Sequence, the ELS Receive Data_Field Size specifies the largest ELS frame that may be received by the Responder FC_Port.

4.3.32 Report node FC-4 Types (RNFT)

4.3.32.1 Description

The RNFT ELS provides for the exchange of supported FC-4 protocol lists. It may be used any time after N_Port Login to verify that the remote Nx_Port supports a given FC-4 protocol.

The RNFT requests a list of the FC-4 protocols supported by the responder, and the RNFT LS_ACC provides the requested list to the source of the RNFT.

A specific FC-4 may require that its Nx_Ports support RNFT, and therefore may conclude that a remote Nx_Port that returns LS_RJT with reason code of “Command not supported” does not support that FC-4.

4.3.32.2 Protocol

- a) Report node FC-4 Types (RNFT) Request sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.32.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the FC-4 Types information. The D_ID field designates the destination Nx_Port to which the request is being made.

Payload: The format of the request Payload is shown in table 97.

Table 97 – RNFT Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RNFT (63h)	Reserved	Maximum Size	
1	Reserved			Index

- a) **Maximum Size:** Bytes 2-3 of Word 0 contain a 16-bit value that specifies the maximum length, in bytes, of the RNFT LS_ACC that the originator is able to accept. The value zero implies the RNFT LS_ACC may be any size;
- b) **Index:** Byte 3 of Word 1 contains an 8-bit value that specifies the index of the first FC-4 Entry to be returned in the RNFT Reply; and

Each FC-4 protocol supported by the responder has an index in the range from zero to (List Length – 1) that should be used to specify a subset of the entries if the entire list does not fit into one reply. NOTE - The index of the entry for a particular FC-4 TYPE may not be consistent between subsequent RNFT Requests (e.g., due to additions or deletions of supported FC-4 TYPEs).

4.3.32.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RNFT command.

LS_ACC: LS_ACC Signifies that the destination Nx_Port has transmitted the requested data. The format of the LS_ACC Payload is shown in table 98.

Table 98 – RNFT LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Reserved	Payload Length (M)	
1	Reserved	List Length	Reserved	Index
2	FC-4 Entry 1			
..	...			
N+1	FC-4 Entry N			

- a) **Payload Length:** Bytes 2-3 of Word 0 contain a 16-bit value that specifies the length M of the RNFT LS_ACC Payload in bytes:

$$M = 8 + N \times 4$$

where

N is the number of FC-4 Entries contained in the Payload;

- b) **List Length:** Byte 1 of Word 1 contains an 8-bit value that specifies the total number of FC-4 protocols supported by the responder.

If List Length exceeds Index+N then the originator may request additional records with another RNFT in which Index is increased by N;

- c) **Index:** Byte 3 of Word 1 contains an 8-bit value that specifies the index of the first FC-4 Entry returned in the RNFT reply, and should reflect the value indicated in the request; and
- d) **FC-4 Entry:** The FC-4 Entry record contains a FC-4 Entry and is shown in table 99.

Table 99 – RNFT FC-4 Entry

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
1	FC-4 Type	FC-4 Qualifier		

- A) **FC-4 Type:** The FC-4 TYPE code of a FC-4 protocol that is supported by the sending Nx_Port. The values are defined in FC-FS-6; and
- B) **FC-4 Qualifier:** The FC-4 Qualifier may be used to distinguish between two protocols that use the same FC-4 TYPE code.

For FC-4 type codes that are reserved or assigned for specific use in this standard (00h - DFh), the value of the FC-4 Qualifier shall be zero.

For Vendor specific FC-4 TYPE codes (E0h through FFh), the FC-4 Qualifier shall be selected from one of the 24-bit Company_ID values assigned by the IEEE Registration Authority to the organization that defines the Vendor specific FC-4 protocol, and that Company_ID shall be

used to qualify that FC-4 TYPE in all implementations. It is up to the organization that defines the Vendor specific FC-4 protocol to assure that the protocol has a unique qualified FC-4 Type.

4.3.33 Scan Remote Loop (SRL)

4.3.33.1 Description

The SRL ELS shall require a switch to scan attached loops to determine if any L_Ports have been disabled or removed. If the switch determines any L_Ports that are currently logged in with the Fabric have been removed or disabled it shall update the name server and send an RSCN to all registered Nx_Ports.

The SRL Payload indicates whether the switch shall scan all attached loops or a single loop. If a single loop is to be scanned the Payload shall contain the FL_Port_ID of the loop to be scanned.

4.3.33.2 Protocol

- a) Scan Remote Loop Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.33.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting Scan Remote Loop. The D_ID field designates the destination of the address identifier for the Domain Controller of the switch for which loops are being scanned. The format of the Domain Controller address is FFFCh || Domain_ID. Domain_ID is the Domain_ID of the switch being queried.

Payload: The format of the Payload is shown in table 100.

Table 100 – SRL Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	SRL (7Bh)	Reserved		
1	Flag	Flag Parameter		

- a) **Flag:** Byte 0 of word 1 indicates if the FL_Port shall be scanned, if all FL_Ports within the domain shall be scanned, or the scan period that all FL_Ports within the domain shall be scanned. The meaning of bits 0-7 is given in Table 101; and

Table 101 – Flag field definitions

Value (hex)	Meaning	Flag Parameter
00	All the FL_Ports within the domain shall be scanned.	Ignored
01	Only the loop attached to the FL_Port addressed in the address identifier of the FL_Port field shall be scanned.	Address identifier of the FL_Port
02	Enable periodic scanning for all FL_ports.	Scan period ^a
03	Disable periodic scanning for all FL_ports.	Ignored
All Others	Reserved	

^a Scan period in seconds. If the scan period is set to zero the scan period is vendor specific. If the switch does not support this option it shall reject the SRL ELS with a reason code of “Unable to perform command request” and a reason code explanation of “Periodic Scanning not supported”. If the switch does not support the selected value it shall reject the SRL ELS with a reason code of “Unable to perform command request” and a reason code explanation of “Periodic Scan Value not allowed”.

- b) **Flag Parameter:** See Table 101.

4.3.33.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the SRL Request

LS_ACC: LS_ACC signifies acceptance of the SRL Request. If the period scanning is enabled then the switch shall return the value of the periodic scanning period in the LS_ACC payloads Scan Period field, If the periodic scanning period is disabled then the switch shall set the LS_ACC payload Scan Period field to zero. The format of the LS_ACC Payload is shown in table 102.

Table 102 – SRL LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	Scan Period		

4.3.34 Set Bit-error Reporting Parameters (SBRP)

4.3.34.1 Description

Set SBRP ELS is used to communicate a set of bit error reporting parameters to a Port or to all Ports in a particular Domain in a Fabric. There are 3 parameters, Error Interval, Error Window, and Error Threshold. Error Interval is the time period over which bit error bursts are integrated to produce a single reported error. An Error Window is composed of one or more Error Intervals. The Error Interval Count is the number of Error Intervals occurring in an Error Window. If the Error Interval Count is greater than or equal to the Error Threshold, a Registered Link Incident Report (RLIR) is generated

with an Incident Code specifying Bit-error-rate threshold exceeded (see FC-FS-6). At the end of the Error Window, the count is set to zero and the process is repeated. See figure 1 for illustration of the parameters.

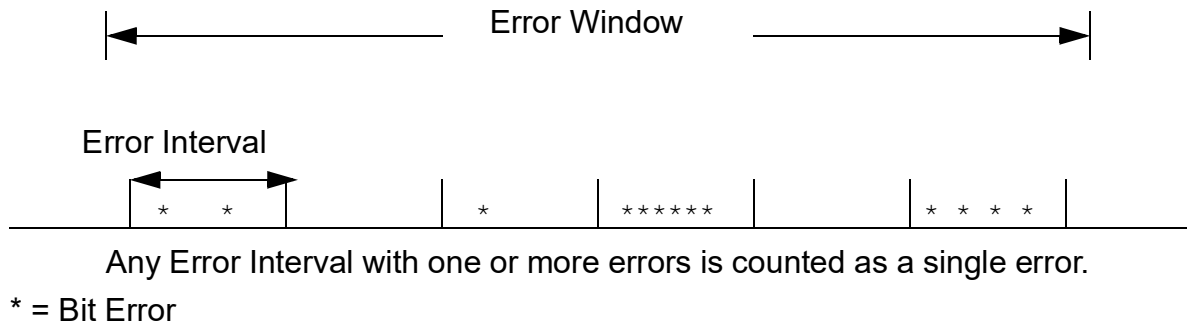


Figure 1 – Illustration of parameters

SBRP may be used to determine an acceptable set or the current set of bit error reporting parameters in the destination Port or Domain, without changing the settings.

The setting of parameters in a particular Port or Domain is done by a controlling entity that is outside the scope of this standard.

4.3.34.2 Protocol

- a) Set Bit-error Reporting Parameters (SBRP) Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.34.3 Request Sequence

Addressing: The S_ID designates the source Nx_Port requesting the bit error rate reporting parameters. The D_ID designates the destination Nx_Port or the Domain Controller (FFFCh || <Domain_ID>). Domain_ID is the Domain_ID of the recipient switch) to process the SBRP ELS.

Payload: The format of the SBRP Payload is shown in table 103.

Table 103 – SBRP Payload

Bits Word	31 .. 28	27 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 03	02 .. 00
0	SBRP (7Ch)		00h	00h		00h	
1	Error Flags						
2	Error Window exponent	Error Window value			Error Interval exponent	Error Interval value	
3	Error Threshold						

- a) **Error Flags:** The following bits are mutually exclusive and only one bit shall be set for each instance of the SBRP ELS:

Bit 0 – Set Error Reporting Parameters: The bit is set to request that the destination set the Error Window, the Error Interval and the Error Threshold parameters. If the destination is the Domain Controller then the request is for the switch to set all ports to the requested values,

Bit 1 – Report Error Reporting Parameters: The bit is set to request that the destination return the active parameters, currently being enforced,

Bits 2-31: Reserved;

- b) **Error Window exponent and Error Window value:** The Error Window is a time duration described by a 16-bit value. Twelve bits, Bits 28-16, are used for the base value. Four bits, Bits 31-28 are encoded to form the exponent. The product of the base value and the encoded exponent yields the time duration in seconds. The encoded exponent is defined as follows:

0h represents 10^0

1h represents 10^{-1}

2h represents 10^{-2}

·
·
·

Fh represents 10^{-15}

(e.g., a base value of 300h multiplied by an encoded exponent of 0h would yield a value of 1 times 10^0 or 300h seconds). The tolerance for the Error Window is -0 to + 1 Error Interval;

- c) **Error Interval exponent and Error Interval value:** The Error Interval is a time duration in seconds and has the same definitions for exponent and base value as Error Window (e.g., a base value of 15 multiplied by an exponent of 1h would yield a value of 15 times 10^{-1} or 1.5 seconds). The Error Interval has a tolerance of – 50% to + 50%; and
- d) **Error Threshold:** The Error Threshold fields specifies the basis for a comparison value with the Error Interval Count.

4.3.34.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the SBRP Request.

LS_ACC: LS_ACC signifies acceptance of the SBRP Request and presents SBRP data. The format of the LS_ACC Payload is shown in table 104.

The LS_ACC Payload conveys either:

- a) the current Bit Error Rate Reporting Parameters (i.e., the requestor is querying for currently set parameters);
- b) the accepted Bit Error Rate Reporting Parameters (i.e., a request to set the parameters was accepted); or

- c) acceptable Bit Error Rate Report Parameters (i.e., a request to set the parameters was rejected because the parameters requested are not within the range supported by the port or switch).

Table 104 – SBRP LS_ACC Payload

Bits Word	31	30 .. 28	27 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 00
0	02h			00h	00h		00h
1	SBRP Request Accepted	Reserved					
2	Error Window exponent		Error Window value		Error Interval exponent	Error Interval value	
3	Error Threshold						

- a) **SBRP Request Accepted:** If a request to set the reporting parameters is indicated (i.e., request Payload word 1 bit 0 is set to one) and the recipient accepts the requested parameters, it shall set the SBRP Request Accepted bit to zero (i.e., Accepted) and echo the requested Error Window, Interval and Threshold parameters in words 2 and 3.

If a request to set the reporting parameters is indicated (i.e., request Payload word 1 bit 0 is set to one) and the recipient does not accept the parameters as requested the recipient shall set the SBRP Request Accepted bit to one and echo in words 2 and 3 those parameters that are accepted. Values not accepted by the recipient shall be set to zero.

If a request to report the parameters is indicated (i.e., request Payload word 1 bit 1 is set to one) the recipient shall set the SBRP Request Accepted bit to zero (i.e., Accepted) report the current Bit Error Rate reporting parameters being enforced. If no Bit Error Rate reporting is being enforced, words 2 and 3 shall be set to zero.

4.3.35 Query Security Attributes (QSA) Version 2

4.3.35.1 Overview

The Query Security Attributes (QSA) ELS allows an N_Port_ID to obtain a summary of selected Fabric-wide policies and register for notifications of relevant policy changes. N_Port_IDs may use the returned policy summary to determine if the Fabric provides the level of security required in their particular environment (see FC-SP-2).

4.3.35.2 QSA Version 2 Request Sequence

Protocol: Query Security Attributes (QSA) Request Sequence

Addressing: The S_ID designates the source N_Port_ID that is requesting policy summary. The D_ID field designates the Fabric Controller (FFFFFDh).

Payload: The format of the QSA request payload is shown in table 105.

Table 105 – QSA Request Payload

Item	Size (Bytes)
7E00 0000h	4
Revision	4
Enforced Security Attribute Registration Mask	4
Extended Security Attribute Registration Mask	4

Revision: Shall be set to 0000 0002h.

Enforced Security Attribute Registration Mask: A bit field that registers the N_Port_ID for change notification for the enforced Fabric Security Attributes. The defined bits are shown in table 106.

Table 106 – Enforced Security Attribute Registration Mask

Bit	Description
31 .. 2	Reserved
1	Insistent Domain_ID Change Notification: If set to one indicates that change notification is required for the Insistent Domain_ID Enforced Security Attribute. If set to zero indicates that change notification for the Insistent Domain_ID Enforced Security Attribute shall not be performed.
0	Fabric Binding Change Notification: If set to one indicates that change notification is required for the Fabric Binding Enforced Security Attribute. If set to zero indicates that change notification for the Fabric Binding Enforced Security Attribute shall not be performed.

Extended Security Attribute Registration Mask: A bit field that registers the N_Port ID for change notification for the extended Fabric Security Attributes. The defined bits are shown in table 107.

Table 107 – Extended Security Attribute Registration Mask

Bit	Description
31 .. 1	Reserved
0	Vendor Specific

4.3.35.3 QSA Version 2 Reply Sequence

LS_RJT: LS_RJT shall be sent as a reply to signify the rejection of the QSA Request Sequence.

LS_ACC: LS_ACC shall be sent as a reply to signify the acceptance of the QSA Request Sequence for processing. The format of the QSA LS_ACC payload is shown in table 108.

Table 108 – QSA LS_ACC Payload

Item	Size (Bytes)
0200 0000h	4
Revision	4
Enforced Security Attribute Field	4
Extended Security Attribute Field	4

Revision: Shall be set to 0000 0002h.

Enforced Security Attribute Field: A bit field describing the security attributes currently enforced by the Fabric. The defined bits are shown in table 109.

Table 109 – Enforced Security Attribute Field

Bit	Description
31 .. 2	Reserved
1	Insistent Domain_ID: If set to one indicates that the Fabric is currently enforcing a policy that requires all Switches to exploit the Insistent Domain_ID behavior, as defined in FC-SP-2. This may be done by the Fabric enforcing a non-NULL Switch Membership List not containing a wildcard entry, and with all Switches listed in the Switch Membership List having the Insistent Domain_ID flag set to one and the Domain_ID attribute not set to zero (see FC-SP-2).
0	Fabric Binding: If set to one indicates that the Fabric is currently enforcing a policy that allows only certain administratively specified Switches to be part of the Fabric itself. This may be done by the Fabric enforcing a non-NULL Switch Membership List not containing only a wildcard entry or by vendor specific methods (see FC-SP-2).

Extended Security Attributes Field: A bit field describing the extended security attributes currently enforced by the Fabric. The defined bits are shown in table 110.

Table 110 – Extended Security Attribute Field

Bit	Description
31 .. 1	Reserved
0	Vendor Specific: Enabled if the bit is set to one, disabled if the bit is set to zero.

4.3.36 Report Port Speed Capabilities (RPSC)

4.3.36.1 Description

The RPSC ELS provides a method for a FC_Port to report its current and potential link operating speeds. The normal response to an RPSC ELS Sequence shall be a LS_ACC ELS Sequence. If the recipient Port does not support the RPSC ELS, it shall reply with an LS_RJT ELS Sequence with a reason code of "Command not supported".

4.3.36.2 Protocol

- a) RPSC Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.36.3 Request Sequence

Addressing: The S_ID designates the source Nx_Port that is requesting port speeds. The D_ID field designates the destination Nx_Port or F_Port Controller (FFFFFFEh). If the D_ID designates the F_Port Controller (FFFFFFEh), Port Speed Capabilities are reported back for the F_Port to which the Nx_Port represented in the S_ID is directly attached.

Payload: The format of the RPSC Payload is shown in table 111.

Table 111 – RPSC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RPSC (7Dh)	00h	00h	00h

4.3.36.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the RPSC Request.

LS_ACC: LS_ACC signifies acceptance of the RPSC Request. The format of the LS_ACC Payload is shown in table 112.

Table 112 – RPSC LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	02h	Flags	Number of entries	
1	Port Speed Capabilities		Port Operating Speed	

Flags: The following bits are defined:

Bit 23 .. 17: Reserved.

Bit 16: Phy Type. Shall be set to zero to indicate a FC-FS-6 PN_Port or PF_Port Phy Type. Shall be set to one to indicate a Lossless Ethernet MAC Phy Type.

Number of Entries: Shall be set to 1.

Port Speed Capabilities: Identifies the speed capabilities of the port, as shown in table 113.

Table 113 – Port Speed Capabilities

Bit	FC-FS-6 PN_Port or PF_Port Phy Type	Lossless Ethernet MAC Phy Type
31	1GFC capable	1GE capable
30	2GFC capable	10GE capable
29	4GFC capable	40GE capable
28	10GFC capable	100GE capable
27	8GFC capable	25GE capable
26	16GFC capable	50GE capable
25	32GFC capable	400GE capable
24	64GFC capable	Reserved
23	128GFC capable	Reserved
22	256GFC capable	Reserved
21 .. 18	Reserved	Reserved
17	Configured: this bit shall be set to one to indicate the Port Operating Speed is administratively configured to one of the supported Port Speed Capabilities	
16	Unknown	

Port Operating Speed: Identifies the current operating speed if set, as shown in table 114.

Table 114 – Port Operating Speed

Bit	FC-FS-6 PN_Port or PF_Port Phy Type	Lossless Ethernet MAC Phy Type
15	1GFC operation	1GE operation
14	2GFC operation	10GE operation
13	4GFC operation	40GE operation
12	10GFC operation	100GE operation
11	8GFC operation	25GE operation
10	16GFC operation	50GE operation
9	32GFC operation	400GE operation
8	64GFC operation	Reserved
7	128GFC operation	Reserved
6	256GFC operation	Reserved
5 .. 2	Reserved	Reserved
1	Unknown	
0 ^a	Speed not established	

^a Applicable only when Port Operating Speed is returned in the Port Speed descriptor of an RDP ELS response that is the response from the Domain Controller for the indicated PF_Port.

If Bit 17 of Port Speed Capabilities is set to one, the Port Operating Speed bit corresponding to the configured port speed shall be set to one. If the port is currently not in an active state, bit 0 shall also be set to one.

4.3.37 Read Exchange Concise (REC)

4.3.37.1 Description

This ELS shall be used only for purposes specific to an FC-4. The REC (Read Exchange Concise) Extended Link Service requests an Nx_Port to return Exchange information for the RX_ID and OX_ID originated by the S_ID specified in the Payload of the request Sequence.

A Read Exchange Concise Request shall only be accepted if the Originator Nx_Port N_Port_ID or the Responder Nx_Port N_Port_ID of the target Exchange is the same as the N_Port_ID of the Nx_Port that makes the request. If the REC Request is not accepted, an LS_RJT with reason code "Unable to perform command request" and reason code explanation "Invalid Originator S_ID" shall be returned.

The specification of OX_ID and RX_ID shall be provided for the destination Nx_Port to locate the status information requested. A Responder destination Nx_Port shall use the RX_ID and verify that the OX_ID is consistent, unless the RX_ID is unassigned (i.e., RX_ID = FFFFh). If the RX_ID is unassigned in the request, the Responder shall identify the Exchange by means of the S_ID specified in the Payload of the request Sequence and OX_ID. An Originator Nx_Port shall use the OX_ID and verify that the RX_ID is consistent.

If the destination Nx_Port of the REC request determines that the Originator S_ID, OX_ID, or RX_ID are inconsistent, then it shall reply with an LS_RJT Sequence with a reason code of "Unable to perform command request" and a reason code explanation of "Invalid OX_ID-RX_ID combination".

The value of the Parameter field in the Frame_Header of an REC ELS and an LS_ACC in response to an REC ELS shall be specified by the FC-4 that sends the frame. The Relative offset present bit in the Frame_Header of an REC ELS or an LS_ACC in response to an REC ELS shall be set to zero.

4.3.37.2 Protocol

- a) Read Exchange Concise (REC) Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.37.3 Request Sequence

Addressing: The S_ID designates the source port that is requesting exchange information. The D_ID field designates the recipient that is requested to provide exchange information.

Payload: The format of the REC Payload is shown in table 115.

Table 115 – REC Payload

Bits Word	31 .. 24 Byte 0	23 .. 16 Byte 1	15 .. 08 Byte 2	07 .. 00 Byte 3
0	REC (13h)	00h	00h	00h
1	Reserved	Exchange Originator S_ID		
2	OX_ID		RX_ID	

Exchange Originator S_ID: shall be set to the address identifier of the target Exchange Originator.

OX_ID: shall be set to the Originator Exchange_ID value of the target Exchange.

RX_ID: shall be set to the Responder Exchange_ID value of the target Exchange.

4.3.37.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the REC Request

LS_ACC: LS_ACC signifies acceptance of the REC Request and returns the requested exchange information. The format of the LS_ACC Payload is shown in table 116.

Table 116 – REC LS_ACC Payload

Bits Word	31 .. 24 Byte 0	23 .. 16 Byte 1	15 .. 08 Byte 2	07 .. 00 Byte 3
0	LS_ACC (02h)	00h	00h	00h
1	OX_ID		RX_ID	
2	Reserved	Originator Address Identifier		
3	Reserved	Responder Address Identifier		
4	FC4VALUE			
5	E_STAT			

The Originator Address Identifier field shall be set to the address identifier of the Originator of the exchange about which information was requested.

The Responder Address Identifier field shall be set the address identifier of the Responder of the exchange about which information was requested.

The value of the FC4VALUE field shall be specified by the FC-4 that sends the LS_ACC to a REC ELS.

E_STAT shall be as defined in FC-FS-6 for the E_STAT field in the Exchange Status Block for the exchange about which information was requested. The bits specifying whether the Exchange is

complete (i.e., Bit 29) and whether the responder holds Sequence Initiative (i.e., Bit 30) shall be valid. The setting of other bits may not be valid.

4.3.38 Exchange Virtual Fabrics Parameters (EVFP)

4.3.38.1 EVFP Messages Structure

4.3.38.1.1 EVFP Request Sequence

Protocol: Exchange Virtual Fabrics Parameters (EVFP) Request Sequence

Addressing: If an EVFP ELS is transmitted from a PN_Port, the S_ID field shall be set to FFFFF0h, indicating the N_Port Controller of the originating PN_Port. The D_ID field shall be set to FFFFFEh, indicating the F_Port Controller of the destination PF_Port. If an EVFP ELS is transmitted from a PF_Port, the S_ID field shall be set to FFFFFEh, indicating the F_Port Controller of the originating PF_Port. The D_ID field shall be set to FFFFF0h, indicating the N_Port Controller of the destination PN_Port.

Payload: Two types of EVFP messages are defined. All EVFP Request messages share the same message structure, shown in table 117.

Table 117 – EVFP Request Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	EVFP (7Fh)	00h	00h	00h
1	Protocol Version	Message Code	Transaction Identifier	
2	MSB		Core N_Port_Name /	
3			Core Switch_Name LSB	
4	Reserved		Message Payload Length	
5	MSB			
..			Message Payload	
N			LSB	

Protocol Version: Shall be set to one.

EVFP Message Code: Specifies the EVFP message that is to be transmitted from the source to the destination. The defined EVFP message codes are shown in table 118.

Table 118 – EVFP Message Codes

Value	Description	Reference
01h	EVFP_SYNC	4.3.38.2
02h	EVFP_COMMIT	4.3.38.3
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 4 – The usage of the Transaction Identifier is very similar to the usage of an OX_ID if an Exchange Originator is enforcing uniqueness via the OX_ID mechanism (see FC-FS-6), but it is not related in any way to the OX_ID present in the Fibre Channel frames carrying the EVFP messages.

Core N_Port_Name / Core Switch_Name: If the EVFP ELS is transmitted from a PN_Port, this field shall be set to its Core N_Port_Name. If the EVFP ELS is transmitted from a PF_Port, this field shall be set to the Core Switch_Name of the Switch it belongs to.

Payload Length: Shall be set to the total length in bytes of the EVFP Payload (i.e., 20 + the Message Payload length).

4.3.38.1.2 EVFP Reply Sequence

Accept (LS_ACC)
Signifies acceptance of the EVFP request.

Accept Payload: All EVFP Accept messages share the same message structure, shown in table 119.

Table 119 – EVFP Accept Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Protocol Version	Message Code	Transaction Identifier	
2	MSB		Core N_Port_Name /	
3			Core Switch_Name	
			LSB	
4	Reserved		Message Payload Length	
5	MSB			
..			Message Payload	
N			LSB	

The fields in table 119 are the same as defined in table 117.

Service Reject (LS_RJT)
Signifies the rejection of the EVFP request

Table 120 shows the use of reason codes and reason code explanations under some error conditions.

Table 120 – LS_RJT Reason Codes for EVFP

Error Condition	Reason Code	Reason Code Explanation
EVFP ELS not supported	Command not supported	No additional explanation
EVFP collision	Command already in progress	No additional explanation
Protocol Version not supported	Protocol error	No additional explanation
EVFP_COMMIT before EVFP_SYNC	Logical error	No additional explanation
Insufficient Resources	Unable to perform command request	No additional explanation
Invalid Payload Message	Protocol error	No additional explanation

4.3.38.2 EVFP_SYNC Message Payload

4.3.38.2.1 Overview

The EVFP_SYNC Message Payload carries a list of descriptors. Each descriptor is self-identifying (see table 122). The format of the EVFP_SYNC Message Payload is shown in table 121. This Message Payload is used in both EVFP_SYNC Request and EVFP_SYNC Accept.

Table 121 – EVFP_SYNC Message Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor #1 = Tagging Administrative Status (see 4.3.38.2.2) ^a			
1				
2	Descriptor #2 = Port VF_ID (see 4.3.38.2.3) ^b			
3				
4	Descriptor #3 = Locally-Enabled VF_ID List (see 4.3.38.2.4) ^c			
..				
132				
...	...			
H	Descriptor #m			
...				
K				

^a Descriptor #1 is required to be present in EVFP_SYNC request.
^b Descriptor #2 is required to be present in EVFP_SYNC request.
^c Descriptor #3 is required to be present in EVFP_SYNC request.

All descriptors share the same format, shown in table 122.

Table 122 – Descriptor Format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor Control	Descriptor Type	Descriptor Length	
1	MSB			
..	Descriptor Value			
M	LSB			

Descriptor Control: Specifies the behavior of the receiving entity if the descriptor is unsupported. The defined codes are shown in table 123.

Table 123 – Descriptor Control Codes

Value	Description
01h	Critical. Abort the EVFP transaction if the descriptor is unsupported. ^a
02h	Non critical. Skip the descriptor if unsupported and continue the EVFP transaction. ^a
all others	Reserved
^a 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 124.

Table 124 – Descriptor Types

Value	Description	Reference
01h	Tagging Administrative Status descriptor	4.3.38.2.2
02h	Port VF_ID descriptor	4.3.38.2.3
03h	Locally-Enabled VF_ID List descriptor	4.3.38.2.4
F0h .. FEh	Vendor Specific descriptor	4.3.38.2.5
all others	Reserved	

Descriptor Length: Specifies the length in bytes of the Descriptor Value.

4.3.38.2.2 Tagging Administrative Status descriptor

The format of the Tagging Administrative Status descriptor is shown in table 125.

Table 125 – Tagging Administrative Status descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor Control = 01h	Descriptor Type = 01h	Descriptor Length = 0004h	
1	Administrative Tagging Mode			

The defined Administrative Tagging Modes are shown in table 126.

Table 126 – Administrative Tagging Modes

Value	Notation	Description
0000 0001h	OFF	The FC_Port shall not perform VFT Tagging.
0000 0002h	ON	The FC_Port may perform VFT Tagging if the peer does not prohibit it.
0000 0003h	AUTO	The FC_Port may perform VFT Tagging if the peer requests it.
all others	Reserved	

In absence of any explicit configuration, the default Administrative Tagging Mode of a VF capable N_Port or F_Port should be AUTO.

Table 127 shows how VFT tagging is negotiated between peer FC_Ports.

Table 127 – 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

4.3.38.2.3 Port VF_ID descriptor

The format of the Port VF_ID descriptor is shown in table 128.

Table 128 – Port VF_ID descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor Control = 01h	Descriptor Type = 02h	Descriptor Length = 0004h	
1	Port Flags		Port VF_ID	

Port Flags: Reserved. Shall be set to zero.

Port VF_ID: The 12 least significant bits of this field shall be set to the Port VF_ID. The four most significant bits shall be set to zero. In the absence of any explicit configuration, the value 001h should be used as Port VF_ID.

4.3.38.2.4 Locally-Enabled VF_ID List descriptor

The format of the Locally-Enabled VF_ID List descriptor is shown in table 129.

Table 129 – Locally-Enabled VF_ID List descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor Control = 01h	Descriptor Type = 03h	Descriptor Length = 0200h	
1	MSB			
..	VF_ID Bitmap			
128	LSB			

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-6) 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.

4.3.38.2.5 Vendor Specific descriptor

The format of the Vendor Specific descriptor is shown in table 130.

Table 130 – Vendor Specific descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Descriptor Control	Descriptor Type	Descriptor Length	
1	MSB	T10 Vendor ID		LSB
2				
3	MSB	Vendor Specific		
..				
N				LSB

T10 Vendor ID: Shall be set to the Vendor’s T10 Vendor ID.

4.3.38.3 EVFP_COMMIT Message Payload

Both EVFP_COMMIT Request and EVFP_COMMIT Accept have no Message Payload.

4.3.39 Link Keep Alive (LKA)

4.3.39.1 Overview

The LKA ELS is used for traffic generation. It provides a means to generate traffic in order to confirm that the link is still intact and/or to ensure the link is not terminated due to lack of traffic. The LKA ELS was specifically designed to keep Fibre Channel backbone links alive (e.g., some TCP implementations will disconnect connections that are not used for some time period).

The LKA ELS is sent by a VE_Port or B_Access portal (see FC-BB-6) to a remote peer in order to determine the health of a link between them, or simply to generate traffic to keep a link from being terminated. If a link is comprised of more than one physical or virtual connection, the LKA may be transmitted on each of the connections. If a connection is configured to handle only specific class(es) of traffic, the LKA shall be sent on a class of service the connection is configured for.

The LKA ELS request Sequence shall consist of a single frame requesting the recipient to reply using the ACC reply Sequence consisting of a single frame. The LKA ELS request frame shall indicate End_Sequence and Sequence Initiative transfer as well as other appropriate F_CTL bits as defined in FC-FS-2. The LKA ELS shall be transmitted as a single frame Sequence and the ACC reply Sequence is also a single frame Sequence. The LKA ELS shall be transmitted as an Exchange that is separate from any other Exchange. The LKA ELS is applicable to Class F, 2, and 3.

The LKA ELS may be sent at any time. The LKA ELS should be sent at least every K_A_TOV if no traffic has been sent and/or received on the connection. The default value for K_A_TOV shall be 1/2 E_D_TOV.

If an accept is not received within E_D_TOV, a new LKA ELS may be transmitted in a new Exchange. The Exchange used for the previous LKA request shall be aborted.

Upon discovering an error (e.g. due to service reject or failure to receive a timely accept in response to one or more LKA ELS requests), the initiator shall initiate appropriate exception handling. The definition of appropriate exception handling is topology-specific.

4.3.39.2 Protocol

- a) Link Keep Alive request Sequence; and
- b) LS_ACC or LS_RJT reply Sequence.

4.3.39.3 Request Sequence

Addressing: The S_ID field shall be set to FFFFDh, indicating the Fabric Controller of the VE_Port or B_Access portal originating the request. The D_ID field shall be set to FFFFDh, indicating the Fabric Controller of the remote peer.

Payload: The format of the LKA Request Payload is shown in table 131.

Table 131 – LKA Payload

Bits Word	31... 24	23... 16	15... 08	07... 00
0	80h	00h	00h	00h

4.3.39.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the LKA request.

LS_ACC: LS_ACC signifies that the connection is intact. The format of the LS_ACC payload is shown in table 132.

Table 132 – LKA LS_ACC Payload

Bits Word	31... 24	23... 16	15... 08	07... 00
0	02h	00h	00h	00h

4.3.40 Link Cable Beacon (LCB)

The LCB ELS is used to turn on and turn off a beacon (e.g., a LED) that is associated with the connector of the receiving port. It provides a means to visually locate the corresponding end of a specific link and, therefore, may not be affected by the state of the attached link or transitions in the state of the attached link.

The LCB ELS is transmitted by:

- a) an N_Port to the F_Port Controller of the attached F_Port;
- b) an F_Port to an attached N_Port; or
- c) an N_Port to an N_Port in a point-to-point topology.

4.3.40.1 Protocol

- a) Link Cable Beacon request Sequence; and
- b) LS_ACC or LS_RJT reply Sequence.

4.3.40.2 Request Sequence

Addressing: For an N_Port, the S_ID field shall be set to the address identifier of the originating N_Port. The D_ID field shall be set to FFFFEh, indicating the F_Port Controller of the destination F_Port, or to the address identifier of the destination N_Port. For an F_Port, the S_ID field shall be set to FFFFEh, indicating the F_Port Controller of the originating F_Port. The D_ID field shall be set to the address identifier of the destination N_Port.

Payload: The format of the LCB Request Payload is shown in table 133.

Table 133 – LCB Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	81h	00h	00h	00h
1	Subcommand	Reserved		Capability
2	Status	Frequency	Duration	

Subcommand: The subcommand field functions are specified in table 134.

Table 134 – Subcommand field functions

Functions	Value
Reserved	0
Beacon on	1
Beacon off	2
Reserved	3 - 255

The status, capability, frequency, and duration fields are only applicable for the Beacon on subcommand and the response to the Beacon on subcommand.

Capability: The capability field functions are specified in table 135. The capability field indicates the functions supported by the requesting N_Port or F_Port.

Table 135 – Capability field functions

Bits	Description
2 - 7	Reserved
1	Frequency Supported: 0 = Support not indicated. 1 = N_Port or F_Port supports specific frequency selections.
0	Duration Supported: 0 = Support not indicated. 1 = N_Port or F_Port supports specific duration selections.

Status: The status field functions are specified in table 136.

Table 136 – Status field functions

Functions	Value
Reserved	0
Normal ^a	1
Warning ^a	2
Critical ^a	3
Reserved	4 - 255
^a Meaning of conditions associated with the function are implementation dependant.	

Frequency: The frequency field specifies the number of blinks in a 10 second interval. A frequency field value of zero specifies the beacon is always on.

Duration: The duration field specifies the duration, in seconds, that the subcommand function shall operate (e.g., beaconing is on at the specified frequency). A duration field value of zero specifies the subcommand function shall operate (e.g., continue beaconing at the specified frequency) until a beacon off subcommand is received.

4.3.40.3 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the LCB request.

LS_ACC: LS_ACC signifies acceptance of the LCB request. The format of the LS_ACC Payload is found in table 137.

Table 137 – LCB LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h
1	Subcommand	Reserved		Capability
2	Status	Frequency	Duration	

The values of the fields of the response payload shall be the same as those in the request payload unless otherwise indicated by the Capability field.

The Capability field indicates the functions supported by the receiving N_Port or F_Port as follows:

- a) If frequency is supported then the Frequency Supported flag shall be set to one and the Frequency field shall be set to the value requested or the nominally supported value (i.e. an implementation is not required to support the exact Frequency value provided in the request);
- b) if frequency is not supported then the Frequency Supported flag shall be set to zero and the Frequency field should be set to the value in the request;
- c) if duration is supported then the Duration Supported flag shall be set to one and the Duration field shall be set to the value requested or the nominally supported value (i.e. an implementation is not required to support the exact Duration value provided in the request); or
- d) if duration is not supported then the Duration Supported flag shall be set to zero and the Duration field should be set to the value in the request.

4.3.41 Registered Fabric Change Notification (RFCN)

4.3.41.1 Overview

The Registered Fabric Change Notification (RFCN) ELS allows the Fabric Controller to notify an N_Port_ID of relevant changes in selected Fabric-wide policies (see FC-SP-2). RFCN is a unidirectional ELS (i.e., it has no reply Sequence).

If a Fabric detects an upcoming change in a security attribute for which change notification has been registered, the Fabric Controller shall send the RFCN ELS to the N_Port_IDs registered for change notification for that security attribute. The RFCN ELS shall not be sent to an N_Port_ID unless the N_Port_ID has registered for change notification of a security attribute, and the attribute is changing.

4.3.41.2 RFCN Request Sequence

Protocol: Registered Fabric Change Notification (RFCN) Request Sequence

Addressing: The S_ID designates the Fabric Controller (FFFFFFDh). The D_ID field designates the recipient N_Port_ID that is being notified of the change in the Fabric.

Payload: The format of the RFCN request payload is shown in table 138.

Table 138 – RFCN Request Payload

Item	Size (Bytes)
9700 0000h	4
Change Flags	4

Change Flags: The defined Change Flags are shown in table 139.

Table 139 – Change Flags

Bit	Description
31 .. 1	Reserved
0	Fabric Security Attributes: If set to one, a change is occurring in a Fabric Security Attribute (e.g., the Fabric Binding or the Insistent Domain_ID Attributes) for which change notification has been registered (see FC-SP-2). If set to zero, the Fabric Security Attributes for which change notification is registered are not changing.

4.3.41.3 RFCN Reply Sequence

None.

4.3.42 Define FFI Domain Topology Map (FFI_DTM)

4.3.42.1 Description

The FFI_DTM ELS Request shall transfer a complete initial or replacement Domain Topology Map to the Domain Controller of the AE Principal Switch.

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal Switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_DTM request currently is not the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of "Logical error".

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of "Command not supported".

4.3.42.2 Protocol

- a) FFI Domain Topology Map Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.42.3 Request Sequence

Addressing: The S_ID designates the Nx_Port sending an FFI Domain Topology Map. The D_ID designates the Domain Controller of the AE Principal Switch in the form FFFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_DTM Request Payload is shown in table 140.

Table 140 – FFI_DTM Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A0h	00h	Payload Length	
1	FFI Incarnation Number			
2	Reserved	Reserved	Number of FFI Link State Records	
3	MSB			
..	FFI Link State Records			
n	LSB			

Payload Length: This field is the length in bytes of the entire Payload, inclusive of the length of word 0. This value shall be a multiple of 4. The minimum value of this field is 28. The maximum value of this field is 65 532.

FFI Incarnation Number: This field contains the new incarnation of the FFI Domain Topology Map.

Number of FFI Link State Records: This field shall specify the number of FFI Link State Records that follow this field. The minimum value is 2.

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 FC-SW-4.

4.3.42.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_DTM Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of the FFI_DTM Request. The format of the LS_ACC Payload for FFI_DTM is shown in table 141.

Table 141 – FFI_DTM LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.43 Request FFI Domain Topology Map (FFI_RTM)

4.3.43.1 Description

The FFI_RTM ELS Request shall request the Domain Controller of the AE Principal Switch to return the current Domain Topology Map in the LS_ACC Reply Sequence.

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal Switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_RTM request currently is not the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of "Logical error".

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of "Command not supported."

4.3.43.2 Protocol

- a) FFI Domain Topology Map Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.43.3 Request Sequence

Addressing: The S_ID field value identifies the Nx_Port requesting the Domain Topology Map from the Domain Controller of the AE Principal Switch. The D_ID designates the Domain Controller of the AE Principal Switch in the form FFFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_RTM Request Payload is shown in table 142.

Table 142 – FFI_RTM Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A1h	00h	00h	00h

4.3.43.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_RTM Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of a valid FFI_RTM Request and that the AE Principal Switch has transmitted the requested data. The format of the LS_ACC Payload for FFI_RTM is shown in table 143.

Table 143 – FFI_RTM LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	Payload Length	
1	FFI Incarnation Number			
2	Reserved	Reserved	Number of FFI Link State Records	
3	MSB			
..	FFI Link State Records			
n	LSB			

Payload Length: This field is the length in bytes of the entire Payload, inclusive of the length of word 0. This value shall be a multiple of 4 bytes. The minimum value of this field is 28 bytes. The maximum value of this field is 65 532.

FFI Incarnation Number: This field contains the current incarnation of the FFI Domain Topology Map.

Number of FFI Link State Records: This field shall specify the number of FFI Link State Records that follow this field. The minimum value is 2.

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 FC-SW-4.

4.3.44 FFI AE Principal Switch Selector (FFI_PSS)

4.3.44.1 Description

The FFI_PSS ELS Request shall be sent to an AE Switch that is not currently the AE Principal Switch in order to command the recipient to become the AE Principal Switch.

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal Switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_PSS request currently is the AE Principal Switch or is not capable of becoming the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of "Logical error".

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of "Command not supported".

4.3.44.2 Protocol

- a) FFI AE Principal Switch Selector Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.44.3 Request Sequence

Addressing: The S_ID field value identifies the Nx_Port requesting a change of AE Principal Switch. The D_ID field value identifies the Domain Controller of an AE Switch that is not the current AE Principal Switch but is capable of becoming an AE Principal Switch in the form FFFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_PSS Request Payload is shown in table 144.

Table 144 – FFI_PSS Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A2h	00h	00h	00h

4.3.44.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_PSS Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of the FFI_PSS Request. The format of the LS_ACC Payload for FFI_PSS is shown in table 145.

Table 145 – FFI_PSS LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.45 FFI Map Update Registration (FFI_MUR)

4.3.45.1 Description

The FFI_MUR ELS Request shall request the Domain Controller of the AE Principal Switch to add or remove the Nx_Port that is sending the FFI_MUR Request (S_ID value) to/from the list of Nx_Ports registered to receive the FFI_RMUN ELS Requests (see 4.3.46).

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_MUR request currently is not the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of “Logical error”.

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of “Command not supported”.

4.3.45.2 Protocol

- a) FFI Map Update Registration Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.45.3 Request Sequence

Addressing: The S_ID designates the Nx_Port requesting registration for FFI Map Update Notification. The D_ID designates the Domain Controller of the AE Principal Switch in the form FFFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_MUR Request Payload is shown in table 146.

Table 146 – FFI_MUR Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A3h	00h	00h	00h
1	00h	00h	00h	Registration Function

Registration Function: The format of the Registration Function field is shown in table 147.

Table 147 – Registration Function

Function	Value
Reserved	0
Full registration - Register to receive FFI_RMUN Requests. If the requesting Nx_Port is already registered, this request is treated as a NOP function.	3
Reserved	4 - 254
Clear registration – Remove the current FFI_RMUN registration, if any. If the requesting Nx_Port is not registered, this request is treated as a NOP function.	255

4.3.45.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_MUR Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of the FFI_MUR Request and registration to receive FFI_RMUN Requests. The format of the LS_ACC Payload for FFI_MUR is shown in table 148.

Table 148 – FFI_MUR LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.46 FFI Registered Map Update Notification (FFI_RMUN)

4.3.46.1 Description

A unidirectional FFI_RMUN Request shall be sent to registered Nx_Ports (see 4.3.45) after the AE Principal Switch sends a Fast Fabric Initialization (FFI) Request Sequence initiated by the AE Principal Switch.

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal Switch. Support for this ELS is optional for Nx_Ports. However, if the Nx_Port invokes the FFI_MUR ELS, it shall support receipt of the FFI_RMUN Request.

FFI_RMUN Request is intended to provide a timely indication of Map Updates to avoid the considerable time to discover any such failure.

Sending FFI_RMUN Requests between Nx_Ports is prohibited. Sending FFI_RMUN Requests by a Domain Controller that is not an AE Principal Switch is prohibited.

4.3.46.2 Registration to Receive FFI_RMUN Requests

See 4.3.45.

4.3.46.3 Protocol

- a) FFI_RMUN Request Sequence.

4.3.46.4 Request Sequence

Addressing: The S_ID is the Domain Controller of the AE Principal Switch (FFFCxxh) and the D_ID is the address of the Registered Nx_Port destination.

Payload: The format of the FFI_RMUN Request Payload is shown in table 149.

Table 149 – FFI_RMUN Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A4h	00h	Payload Length	
1	FFI Incarnation Number			
2	Reserved	Reserved	Number of FFI Link State Records	
3	MSB			
..	FFI Link State Records			
n	LSB			

Payload Length: This field is the length in bytes of the entire Payload, inclusive of the length of word 0. This value shall be a multiple of 4 bytes. The minimum value of this field is 28 bytes. The maximum value of this field is 65 532.

FFI Incarnation Number: This field contains the current incarnation of the FFI Domain Topology Map.

Number of FFI Link State Records: This field shall specify the number of FFI Link State Records that follow this field. The minimum value is 2.

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 FC-SW-4.

4.3.46.5 Reply Sequence

LS_RJT: none

LS_ACC: none

4.3.47 FFI Suspend Map Updates (FFI_SMU)

4.3.47.1 Description

The FFI_SMU ELS Request shall request the Domain Controller of the AE Principal Switch, to suspend Domain Topology Map updates as specified in FC-SW-4. If Domain Topology Map updates are currently suspended, this request is treated as a NOP function. Successful execution of this ELS shall also suspend Registered Map Updates Notifications, if any (see 4.3.46).

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_SMU request currently is not the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of "Logical error".

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of "Command not supported".

4.3.47.2 Protocol

- a) FFI Suspend Map Updates Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.47.3 Request Sequence

Addressing: The S_ID field value identifies the Nx_Port requesting suspension of Domain Topology Map updates by the Domain Controller of the AE Principal Switch. The D_ID field value identifies the Domain Controller of the AE Principal Switch in the form FFFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_SMU Request Payload is shown in table 150.

Table 150 – FFI_SMU Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A5h	00h	00h	00h

4.3.47.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_SMU Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of a valid FFI_SMU Request and indicates that the Domain Controller of the AE Principal Switch has suspended Domain Topology Map updates. The format of the LS_ACC Payload for FFI_SMU is shown in table 151.

Table 151 – FFI_SMU LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.48 FFI Resume Map Updates (FFI_RMU)

4.3.48.1 Description

The FFI_RMU ELS Request shall request the Domain Controller of the AE Principal Switch, to resume Domain Topology Map updates as specified in FC-SW-4. If Domain Topology Map updates are currently not suspended, this request is treated as a NOP function. Successful execution of this ELS shall also resume Registered Map Updates Notifications, if any (see 4.3.46).

Support for this ELS is mandatory for a Domain Controller in an AE Principal Switch or an AE Switch that is capable of becoming an AE Principal Switch. Support for this ELS is optional for Nx_Ports.

If the Domain Controller that receives an FFI_RMU request currently is not the AE Principal Switch, it shall respond with an LS_RJT reply with a reason code of “Logical error”.

If a destination receives this ELS Request that is not a Domain Controller, it shall respond with an LS_RJT reply with a reason code of “Command not supported”.

4.3.48.2 Protocol

- a) FFI Resume Map Updates Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.48.3 Request Sequence

Addressing: The S_ID field value identifies the Nx_Port requesting resumption of Domain Topology Map updates by the Domain Controller of the AE Principal Switch. The D_ID field value identifies the Domain Controller of the AE Principal Switch in the form FFCxxh, where xx is the one-byte value assigned to the Domain Controller.

Payload: The format of the FFI_RMU Request Payload is shown in table 152.

Table 152 – FFI_RMU Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	A6h	00h	00h	00h

4.3.48.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the FFI_RMU Request. The LS_RJT reply contains an appropriate reject reason code.

LS_ACC: LS_ACC signifies acceptance of a valid FFI_RMU Request and indicates that the Domain Controller of the AE Principal Switch has resumed Domain Topology Map updates. The format of the LS_ACC Payload for FFI_RMU is shown in table 153.

Table 153 – FFI_RMU LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	02h	00h	00h	00h

4.3.49 Read Diagnostic Parameters (RDP)

4.3.49.1 Description

The RDP ELS requests an FC_Port to return the diagnostic parameters associated with the N_Port_ID specified in the Payload. This provides the Nx_Port transmitting the request with information that may be used for diagnosis of link or port related errors, or degraded conditions associated with the designated FC_Port.

4.3.49.2 Protocol

- a) Read Diagnostic Parameters Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence

4.3.49.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the diagnostic parameters. The D_ID field shall be set as follows:

- a) to any Domain Controller (FFFCxxh);
- b) to the F_Port Controller (FFFFFEh); or
- c) to any Nx_Port N_Port_ID

Payload: The format of the RDP Request Payload is shown in table 154.

Table 154 – RDP Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	RDP (18h)	00h	00h	00h
1	Descriptor List Length = 12			
2	MSB			
3	N_Port_ID descriptor			
4	LSB			

N_Port_ID descriptor: See 4.2.4.3. The requested diagnostic parameters are described as follows:

- a) if the D_ID is a Domain Controller well known address (FFFCxxh), the N_Port_ID in the N_Port_ID descriptor shall be set to an N_Port_ID within the associated domain. The diagnostic parameters requested are for the Fx_Port through which the N_Port_ID is logged in;
- b) if the D_ID is the F_Port Controller Well-known address (FFFFFEh), the N_Port_ID in the N_Port_ID descriptor shall be set to the S_ID of the request. The diagnostic parameters requested are for the Fx_Port through which the N_Port_ID is logged in; or
- c) for all other D_IDs, the N_Port_ID in the N_Port_ID descriptor shall be set to the D_ID of the request. The diagnostic parameters requested are for Nx_Port to which the D_ID is assigned.

4.3.49.4 Reply Sequence

If the Nx_Port receiving an RDP ELS request is not logged in with the source N_Port, it shall respond with LS_RJT with a reason code of “Login Required” or respond with a single frame sequence as appropriate for the ELS and the current state of the Nx_Port.”

LS_RJT: LS_RJT signifies rejection of the RDP command. The LS_RJT reason code and reason code explanation are set as follows:

- a) If an FC_Port does not support the RDP request, it shall reply with an LS_RJT specifying a reason code of “Command not supported” (i.e., 0Bh) and a reason code explanation of “Request not supported” (i.e., 2Ch);
- b) if the N_Port_ID in the N_Port_ID descriptor is not logged in with an F_Port within the domain, the Domain Controller shall reply with an LS_RJT specifying a reason code of “Logical error” (i.e., 03h) and reason code explanation “Invalid N_Port_ID” (i.e., 1Fh); or
- c) if the FC_Port receiving an RDP ELS request requires a login with the source FC_Port and is not logged in, it shall reply with an LS_RJT specifying a reason code of “Unable to perform command request” (i.e., 09h) and a reason code explanation of “N_Port Login required” (i.e., 1Eh).

LS_ACC: LS_ACC signifies that the FC_Port has transmitted the requested data.

An FC_Port sending an RDP ELS prior to login implies the originator has a Receive Data_Field size capable of receiving the maximum length of the response. If the FC_Port receiving an RDP ELS request does not require a login and is not logged in with the source FC_Port it shall:

- 1) respond with a single frame Sequence with a payload that fits within the implied Receive Data_Field size; or
- 2) respond with a multi-frame Sequence using the default login value for Receive Data_Field size.

NOTE 5 – This could be incompatible with N_Ports that cannot receive multi-frame ELS sequences.

The format of the LS_ACC Payload is shown in table 155.

Table 155 – RDP LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	LS_ACC (02h)	00	00	00
1	Descriptor list length ((n-1)*4) bytes)			
2	MSB			
3	Link Service Request Information descriptor (see 4.2.4.2)			
4	LSB			
5 - n	Diagnostic parameter descriptors			

The descriptor list length shall not be greater than 2040-bytes (i.e. n is less than or equal to 511).

The diagnostic parameter descriptors (see 4.3.49.5) in an RDP LS_ACC shall include;

- a) a Link Service Request Information descriptor;
- b) a Port Speed descriptor;
- c) a Link Error Status Block descriptor; and
- d) a Port Names descriptor for the Fx_Port or Nx_Port for which diagnostic parameters are being provided followed by a Port Names descriptor for the directly attached Fx_Port or Nx_Port, if any.

The diagnostic parameter descriptors in an RDP LS_ACC shall also include one of the following descriptors;

- a) an SFP Diagnostics descriptor; or
- b) a QSFP Diagnostics descriptor.

The diagnostic parameter descriptors in an RDP LS_ACC may include one or more of the following descriptors;

- a) a FEC Status descriptor;
- b) a Buffer Credit descriptor;
- c) one or more Optical Element Data descriptors;
- d) an Optical Product Data descriptor;

- e) a Port Congestion descriptor; or
- f) an Elapsed Time descriptor.

4.3.49.5 Diagnostic parameter descriptors

4.3.49.5.1 Port Speed descriptor

The Port Speed descriptor is shown in table 156.

Table 156 – Port Speed descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Port Speed descriptor tag = 0001 0001h			
1	Port Speed descriptor Length (4 bytes)			
2	Port Speed Capabilities		Port Operating Speed	

Port Speed Capabilities and Port Operating Speed are as defined for the RPSC ELS (see 4.3.36). The Phy Type is specified in the Link Error Status Block descriptor (see 4.3.49.5.2).

4.3.49.5.2 Link Error Status Block descriptor

The Link Error Status Block descriptor is shown in table 157.

Table 157 – Link Error Status Block descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Link Error Status Block descriptor tag = 0001 0002h			
1	Link Error Status Block descriptor Length (28 bytes)			
2	MSB			
3				
4	Link Error Status Block (see FC-FS-6 or FC-			
5	BB-6)			
6				
7	LSB			
8	Px_Port Phy Type	Reserved		

The content of the Link Error Status Block is specified in FC-FS-6 or FC-BB-6, depending on the Px_Port Phy Type.

Px_Port Phy Type: Word 8bits 31 - 30 identify the type of physical interface for the PN_Port or PF_Port through which the RDP request was received (see table 158). Bits 29 - 24 of this field are reserved.

Table 158 – Px_Port Phy Type

Encoded Value Word 8, Bits 31 - 30	Description
00b	No Information about Phy Type provided.
01b	The sending Vx_Port uses an FC-FS-6 Px_Port or PF_Port.
10b	The sending Vx_Port uses a lossless Ethernet MAC.
11b	Reserved

4.3.49.5.3 Port Names descriptor

The Port Names descriptor contains the Node WWN and the FC_Port WWN for the FC_Port as shown in table 159. If describing a VN_Port, the Node WWN and FC_Port WWN fields contain the names associated with the VN_Port.

Table 159 – Port Names descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Port Names descriptor tag = 0001 0003h			
1	Port Names descriptor Length (16 bytes)			
2	MSB			
3	Node WWN			LSB
4	MSB			
5	Port WWN			LSB

4.3.49.5.4 SFP Diagnostic Parameters descriptor

The SFP Diagnostic Parameters descriptor is shown in table 160.

Table 160 – SFP Diagnostic Parameters descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	SFP Diagnostics descriptor tag = 0001 0000h			
1	SFP Diagnostics descriptor Length (12 bytes)			
2	Temperature		Vcc	
3	Tx Bias		Tx Power	
4	Rx Power		SFP Flags	

SFP diagnostic parameters are as provided in Diagnostic Monitoring Interface for Optical Transceivers (see SFF-8472). Parameters provided are;

- a) **Temperature:** The temperature field contains the internally measured transceiver temperature in units of 1/256 C, range -128C to + 128C;
- b) **Vcc:** The Vcc field contains internally measured supply voltage in units of 100uV, range 0 - 6.55V;
- c) **Tx Bias:** The Tx Bias field contains the measured transmitter laser bias current in units of 2 uA, Range 0 - 131mA;
- d) **Tx Power:** The Tx Power field contains the measured coupled TX output power in units of 0.1 uW, range 0 - 6.5mW;
- e) **Rx Power:** The Rx Power field contains the measured received optical power in units of 0.1 uW, range 0 - 6.5mW; and

- f) **SFP Flags:** The defined SFP Flags are shown in table 161.

Table 161 – SFP Flags

Bit	Description
15	16GFC FEC Active: 0b = 16G FEC status is inactive or unknown 1b = 16G FEC is active
14 - 8	Reserved
7 - 6	Connector Type: 00b = Other or unknown 01b = SFP+ 10b - 11b = reserved
5	SFP Diagnostic Parameters not valid: If set to 1, the response does not include valid values for; a) Temperature; b) Vcc; c) Tx Bias; d) Tx Power; and e) Rx Power.
4	Optical Port: If set to 1, the FC_Port is an optical port. If set to 0, the FC_Port is not an optical port.
3 - 0	Port Tx Type: 0000b = Other or unknown 0001b = Short Wave Laser 0010b = Long Wave Laser LC 1310nm 0011b = Long Wave Laser LL 1550nm 0100b - 1111b = reserved

SFP diagnostic parameters may be periodically obtained from an SFP by the FC_Port and parameters returned shall be the most recent values obtained. If the recipient of the RDP request is not able to provide diagnostic parameters, the values in these parameters shall be set to 0 and the SFP Diagnostic Parameters not Valid bit 5 shall be set to 1.

4.3.49.5.5 QSFP Diagnostic Parameters descriptor

The QSFP Diagnostic Parameters descriptor is shown in table 162.

Table 162 – QSFP Diagnostic Parameters descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	QSFP Diagnostics descriptor tag = 0001 0004h			
1	QSFP Diagnostics descriptor Length (40 bytes)			
2	Temperature		Vcc	
3	Reserved		QSFP Flags	
4	Lane 0 – Tx Bias		Lane 0 – Tx Power	
5	Lane 0 – Rx Power		Reserved	
6	Lane 1 – Tx Bias		Lane 1 – Tx Power	
7	Lane 1 – Rx Power		Reserved	
8	Lane 2 – Tx Bias		Lane 2 – Tx Power	
9	Lane 2 – Rx Power		Reserved	
10	Lane 3 – Tx Bias		Lane 3 – Tx Power	
11	Lane 3 – Rx Power		Reserved	

QSFP diagnostic parameters are as provided in Common Management Interface (see SFF-8636). Parameters provided are;

- a) Temperature: The temperature field contains the internally measured transceiver temperature in units of 1/256 C, range -128C to + 128C;
- b) Vcc: The Vcc field contains internally measured supply voltage in units of 100uV, range 0 - 6.55V;
- c) Tx Bias: The Tx Bias fields contains the measured transmitter laser bias current in units of 2 uA, Range 0 - 131mA for each corresponding QSFP lane;
- d) Tx Power: The Tx Power fields contains the measured coupled TX output power in units of 0.1 uW, range 0 - 6.5mW for each corresponding QSFP lane;
- e) Rx Power: The Rx Power fields contains the measured received optical power in units of 0.1 uW, range 0 - 6.5mW for each corresponding QSFP lane; and

f) QSFP Flags: The defined QSFP Flags are shown in table 163.

Table 163 – QSFP Flags

Bits	Description
15	16GFC FEC Active: 0b = 16G FEC status is inactive or unknown 1b = 16G FEC is active
14 – 11	Reserved
10 – 9	Lane Identifier: 00b = Lane 0 01b = Lane 1 10b = Lane 2 11b = Lane 3
8	Lane Identifier Valid: If set to 1, the response is for the single lane indicated by the Lane Identifier. If set to 0, the Lane Identifier is not meaningful.
7 – 6	Connector Type: 00b = Other or unknown 01b = QSFP 10b - 11b = reserved
5	QSFP Diagnostic Parameters not valid: If set to 1, the response does not include valid values for; a) Temperature; b) Vcc; c) Tx Bias for any lane; d) Tx Power for any lane; and e) Rx Power for any lane.
4	Optical Port: If set to 1, the FC_Port is an optical port. If set to 0, the FC_Port is not an optical port.
3 – 0	Port Tx Type: 0000b = Other or unknown 0001b = Short Wave Laser 0010b = Long Wave Laser 0100b - 1111b = reserved

QSFP diagnostic parameters may be periodically obtained from an QSFP by the FC_Port and parameters returned shall be the most recent values obtained. If the recipient of the RDP request is not able to provide diagnostic parameters, the values in these parameters shall be set to 0 and the QSFP Diagnostic Parameters not Valid bit 5 shall be set to 1.

4.3.49.5.6 FEC Status descriptor

The FEC Status descriptor is shown in table 164.

Table 164 – FEC Status descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	FEC Status descriptor tag = 0001 0005h			
1	FEC Status descriptor Length (8 bytes)			
2	Corrected blocks			
3	Uncorrectable blocks			

Corrected blocks: Number of FEC encoding blocks reported as corrected by hardware.

Uncorrectable blocks: Number of FEC encoding blocks reported as uncorrectable by hardware.

4.3.49.5.7 Buffer Credit descriptor

The Buffer Credit descriptor contains the buffer-to-buffer credit (see 6.6.2.3). values shared during login between N_Ports and F_Ports as shown in table 165.

Table 165 – Buffer Credit descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Buffer Credit descriptor tag = 0001 0006h			
1	Buffer Credit descriptor Length (12 bytes)			
2	FC_Port buffer-to-buffer credit			
3	Attached FC_Port buffer-to-buffer credit			
4	Nominal FC_Port RTT			

FC_Port buffer-to-buffer credit: The advertised buffer-to-buffer credit for the Fx_Port or Nx_Port for which diagnostic parameters are being provided.

Attached FC_Port buffer-to-buffer credit: The advertised buffer-to-buffer credit received from the directly attached Fx_Port or Nx_Port, if any. The Attached FC_Port buffer-to-buffer credit shall be zero if there is no directly attached Fx_Port or Nx_Port.

Nominal FC_Port RTT: The nominal round trip time (i.e., RTT) of the link between the FC_Port and the attached FC_Port as determined by the FC_Port, where a value of zero indicates the RTT is unknown and a non-zero value indicates the RTT of the link expressed in nanoseconds.

4.3.49.5.8 Optical Element Data descriptor

The Optical Element Data descriptor contains the optical element data of the SFP or QSFP as shown in table 166.

Table 166 – Optical Element Data descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Optical Element Data descriptor tag = 0001 0007h			
1	Optical Element Data descriptor Length (12 bytes)			
2	High Alarm		Low Alarm	
3	High Warning		Low Warning	
4	Optical Element Function Flags			

Optical Element Data descriptor parameters are as provided in “Diagnostic Monitoring Interface for Optical Transceivers” for SFPs (see SFF-8472) and in “Common Management Interface” for QSFPs (see SFF-8636). The parameters provided are:

- a) **High Alarm:** The high alarm value for the operating range of the optical element.
- b) **Low Alarm:** The low alarm value for the operating range of the optical element.
- c) **High Warning:** The high warning value for the operating range of the optical element.
- d) **Low Warning:** The low warning value for the operating range of the optical element.
- e) **Optical Element Function Flags:** The defined Optical Element Function Flags are shown in table 167.

Table 167 – Optical Element Function Flags

Bits	Description
31 – 28	Optical Element Type: 0 = Reserved 1 = Module Temperature 2 = Module Voltage 3 = Tx Bias 4 = Tx Power 5 = Rx Power 6-15 = reserved
27 – 16	Reserved
15 – 12	Lane 3 Optical Element Type Transgression Flags
11 – 8	Lane 2 Optical Element Type Transgression Flags
7 – 4	Lane 1 Optical Element Type Transgression Flags
3 – 0	Lane 0 Optical Element Type Transgression Flags

If the Optical Element Type field is set to Module Temperature or Module Voltage then the Optical Element Type Transgression flags are in bits 3-0 and bits 15-4 shall be set to 0.

If the Optical Element Function Flags are for an SFP then only Optical Element Type Transgression Flags bits 3-0 are valid and bits 15-4 shall be set to 0.

If the Optical Element Function Flags are for a QSFP then bits 15-0 are valid. All lanes are reported regardless of the operating mode of the QSFP.

The defined Optical Element Type Transgression Flags are shown in table 168.

Table 168 – Optical Element Type Transgression Flags

Bits	Description
3	High Alarm - Set to 1 if the operating value has transgressed the high alarm level.
2	Low Alarm - Set to 1 if the operating value has transgressed the low alarm level.
1	High Warning - Set to 1 if the operating value has transgressed the high warning level.
0	Low Warning - Set to 1 if the operating value has transgressed the low warning level.

4.3.49.5.9 Optical Product Data descriptor

The Optical Product descriptor contains the product data of the SFP or QSFP as shown in table 169.

Table 169 – Optical Product Data descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Optical Product Data descriptor tag = 0001 0008h			
1	Optical Product Data descriptor Length (60 bytes)			
2	MSB			
3	Vendor name			
4				
5				

Table 169 – Optical Product Data descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
6	MSB			
7				
8	Part number			
9	LSB			
10	MSB			
11				
12	Serial number			
13	LSB			
14	Revision			
15				
16	Date			

Optical Product Data descriptor parameters are as provided in “Diagnostic Monitoring Interface for Optical Transceivers” for SFPs (see SFF-8472) and in “Common Management Interface” for QSFPs (see SFF-8636). The parameters provided are:

- a) **Vendor name:** The vendor name is the full name of the corporation, a commonly accepted abbreviation of the name of the corporation, the SCSI company code for the corporation, or the stock exchange code for the corporation (ASCII).
- b) **Part number:** The vendor part number is a field that defines the vendor part number or product name (ASCII).
- c) **Serial number:** The vendor serial number is a field that defines the vendor's serial number (ASCII).
- d) **Revision:** The vendor revision is a field that defines the vendor's revision level for the part number provided (ASCII). SFF vendor revision values consisting of 2-bytes shall be stored in bits 31-16 of the Revision field.
- e) **Date:** The vendor date code is a field that defines the vendor's manufacturing date code (encoded ASCII).

4.3.49.5.10 Port Congestion descriptor

The Port Congestion descriptor is shown in table 170.

Table 170 – Port Congestion descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Port Congestion descriptor tag = 0001 0009h			
1	Port Congestion descriptor Length (56 bytes)			
2	Counter Validity Mask			
3	Tx Zero Credit Transition to Zero Count			
4	Rx Zero Credit Transition to Zero Count			
5	Tx Zero Credit Count			
6	Rx Zero Credit Count			
7	Credit Interval			
8	Tx Discard Count			
9	Tx Discard Interval			
10	Active State Tx LR Count			
11	Active State Rx LR Count			
12	Vendor Specific Counter 1			
13	Vendor Specific Counter 2			
14	Vendor Specific Counter 3			
15	Vendor Specific Counter 4			

Counter Validity Mask: A field containing a bit mask that represents each of the possible counters as shown in table 171.

Table 171 – Counter Validity Bit Mask Values

Bits	Description
31	Credit Sampled 0b – Credit counters for Tx Zero Credit Count and Tx Credit Count fields are not sampled 1b - Credit counters for Tx Zero Credit Count and Tx Credit Count fields are sampled
12	Vendor Specific Counter 4 Validity Bit 0b – Counter is invalid 1b – Counter is valid
11	Vendor Specific Counter 3 Validity Bit 0b – Counter is invalid 1b – Counter is valid

Table 171 – Counter Validity Bit Mask Values (continued)

Bits	Description
10	Vendor Specific Counter 2 Validity Bit 0b – Counter is invalid 1b – Counter is valid
9	Vendor Specific Counter 1 Validity Bit 0b – Counter is invalid 1b – Counter is valid
8	Active State Rx LR Count Validity Bit 0b – Counter is invalid 1b – Counter is valid
7	Active State Tx LR Count Validity Bit 0b – Counter is invalid 1b – Counter is valid
6	Tx Discard Interval Validity Bit 0b – Counter is invalid 1b – Counter is valid
5	Tx Discard Count Validity Bit 0b – Counter is invalid 1b – Counter is valid is bit 6 is also set to one
4	Credit Interval Validity Bit 0b – Counter is invalid 1b – Counter is valid
3	Rx Zero Credit Count Validity Bit 0b – Counter is invalid 1b – Counter is valid if bit 4 is also set to one
2	Tx Zero Credit Count Validity Bit 0b – Counter is invalid 1b – Counter is valid if bit 4 is also set to one
1	Rx Zero Credit Transition to Zero Count Validity Bit 0b – Counter is invalid 1b – Counter is valid
0	Tx Zero Credit Transition to Zero Count Validity Bit 0b – Counter is invalid 1b – Counter is valid
All others	Reserved

Tx Zero Credit Transition to Zero Count: Number of times that transmit buffer-to-buffer Credit has transitioned to zero. Validity of this counter is indicated with bit 0 in the Counter Validity Mask.

Rx Zero Credit Transition to Zero Count: Number of times that receive buffer-to-buffer Credit has transitioned to zero. Validity of this counter is indicated with bit 1 in the Counter Validity Mask.

Tx Zero Credit Count: If Counter Validity Mask bit 31 (i.e., Credit Sampled) is set to 1, this count contains the number of times that transmit buffer-to-buffer Credit is zero when sampled. Sampling is performed at the rate indicated by the Credit Interval field.

If Counter Validity Mask bit 31 is set to 0, this count contains the number of times that transmit buffer-to-buffer Credit has remained at zero for an entire Interval indicated by the Credit Interval field.

Validity of this counter is indicated with bit 2 and bit 4 in the Counter Validity Mask.

Rx Zero Credit Count: If Counter Validity Mask bit 31 (i.e., Credit Sampled) is set to 1, this count contains the number of times that receive buffer-to-buffer Credit is zero when sampled. Sampling is performed at the rate indicated by the Credit Interval field.

If Counter Validity Mask bit 31 is set to 0, this count contains the number of times that the receive buffer-to-buffer Credit has remained at zero for an entire Interval indicated by the Credit Interval field.

Validity of this counter is indicated with bit 3 and bit 4 in the Counter Validity Mask.

Credit Interval: The Credit Interval represents the time interval associated with the Tx Zero Credit Count and Rx Zero Credit Count fields. Validity of this field is indicated with bit 4 in the Counter Validity Mask. A value of zero (0) for this counter with an accompanying validity bit shall indicate that the interval is unknown.

If Counter Validity Mask bit 31 (i.e., Credit Sampled) is set to 1, this field contains the approximate sample rate, in milliseconds, at which the Rx Credit or Tx Credit counters are sampled.

If Counter Validity Mask bit 31 is set to 0, this field contains the value, in nanoseconds, that the Rx Credit counter or Tx Credit counter remains at zero before counting. The upper limit for this value is E_D_TOV.

Tx Discard Count: Number of times that a frame has been discarded by the fabric. Validity of this counter is indicated with bit 5 and bit 6 in the Counter Validity Mask.

Tx Discard Interval: Length of time in microseconds that a frame is held for transmission before it is discarded. Validity of this counter is indicated with bit 6 in the Counter Validity Mask.

Active State Tx LR Count: Number of times that the interface has transmitted an LR (Link Reset) while in the Active State (see FC-FS-6). Validity of this counter is indicated with bit 7 in the Counter Validity Mask.

Active State Rx LR Count: Number of times that the interface has received an LR (Link Reset) while in the Active State (see FC-FS-6). Validity of this counter is indicated with bit 8 in the Counter Validity Mask.

Vendor Specific Counter 1: A counter used to indicate a vendor unique counter related to port link congestion. Validity of this counter is indicated by bit 9 in the Counter Validity Mask.

Vendor Specific Counter 2: A counter used to indicate a vendor unique counter related to port link congestion. Validity of this counter is indicated by bit 10 in the Counter Validity Mask.

Vendor Specific Counter 3: A counter used to indicate a vendor unique counter related to port link congestion. Validity of this counter is indicated by bit 11 in the Counter Validity Mask.

Vendor Specific Counter 4: A counter used to indicate a vendor unique counter related to port link congestion. Validity of this counter is indicated by bit 12 in the Counter Validity Mask.

No means are provided to reset a counter in the Port Congestion descriptor, however, on overflow it shall be set to zero and then continue counting.

4.3.49.5.11 Elapsed Time descriptor

The Elapsed Time descriptor (see table 172) descriptor contains a list of RDP descriptor tags and associated time values representing when the counters or values contained in that RDP descriptor were last reset or read.

Table 172 – Elapsed Time descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Elapsed Time descriptor tag = 0001 000Eh			
1	Elapsed Time descriptor Length (in bytes)			
2	RDP Descriptor Tag #1			
3	Elapsed Time (for RDP Descriptor Tag #1)			
4	RDP Descriptor Tag #2			
5	Elapsed Time (for RDP Descriptor Tag #2)			
	..			
	RDP Descriptor Tag #n			
	Elapsed Time (for RDP Descriptor Tag #n)			

RDP Descriptor Tag: A descriptor tag for a single RDP descriptor that is included in the RDP reply LS_ACC payload. The tag value of 0001 FFFFh may be used to indicate that the Elapsed Time field pertains to any of the following RDP descriptor tags included in the RDP Reply LS_ACC payload:

- a) Link Error Status Block descriptor;
- b) FEC Status descriptor; or
- c) Port Congestion descriptor.

Elapsed Time: Time in seconds since the counters or values for the associated RDP descriptor were last reset or read. Contains a 32 bit value of the amount of time, in nearest whole seconds, as follows:

Indicates the amount of time since the values were last reset for the following RDP descriptors:

- a) Link Error Status Block descriptor;
- b) FEC Status descriptor; or
- c) Port Congestion descriptor.

Indicates the amount of time since the values were last read from the transceiver for the following RDP descriptors:

- a) SFP Diagnostic Parameters descriptor; or
- b) QSFP Diagnostic Parameters descriptor.

If the counters or values have not been reset or read since the device initialized, then if this RDP descriptor is included it shall indicate the time since the device initialized.

Inclusion of an elapsed time entry for each descriptor is independent and optional.

4.3.50 Query Fabric Priority Allocation (QFPA)

4.3.50.1 Overview

The QFPA ELS is used by an N_Port supporting Priority Tagging (see 9) to query the attached Switch about the Priority field values useable as local VE IDs and the associated QoS priority processing the Fabric provides for those values.

4.3.50.2 Protocol

- a) Query Fabric Priority Allocation Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.50.3 Request Sequence

Addressing: The S_ID field shall be set to the N_Port_ID of the N_Port originating the QFPA Request. The D_ID shall be set to the F_Port Controller (FFFFFFEh).

Payload: The format of the QFPA Request Payload is shown in table 173.

Table 173 – QFPA Request Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	QFPA (B0h)	00h	00h	00h

4.3.50.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the QFPA Request.

LS_ACC: LS_ACC supplies the requested Fabric Priority allocation. The format of the QFPA LS_ACC Payload is shown in table 174.

Table 174 – QFPA LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	02h	00h	00h	00h
1	Descriptor list length = Number of Priority Range descriptors (n) x 12			
2 .. 4	Priority Range descriptor #1			
5 .. 7	Priority Range descriptor #2			
..	..			
n x 3 + 1	Priority Range descriptor #n			

The format of the Priority Range descriptor is shown in table 175.

Table 175 – Priority Range descriptor Format

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Descriptor Tag = 0001 0009h			
1	Descriptor Length = 4			
2	Low Range Value	High Range Value	QoS Priority	Local VE ID Use

Low Range Value: contains, in the lowest significant seven bits, the lowest value of the described priority range, inclusive. The most significant bit is reserved.

High Range Value: contains, in the lowest significant seven bits, the highest value of the described priority range, inclusive. The most significant bit is reserved.

QoS Priority: indicates if the described Priority range is part of the subset of Priorities for which the Fabric provides QoS priority processing. The following values are defined for this field:

00h: No QoS priority provided (i.e., best effort).

01h: QoS priority provided.

02h .. FFh: Reserved.

NOTE 6 – As defined in FC-FS-6, an implementation may provide QoS priority processing only for a contiguous subset of priority values. Priority values within this subset are grouped in priority ranges representing increasing QoS priorities.

Local VE ID Use: indicates if the values of the Priority field associated with the described Priority range are able to be used as local VE IDs. This field is a 8-bit bitmask defined as follow:

Bit 0: If set to one then Priority field values with the seven most significant bits belonging to the described Priority range and the Tagging Extension bit set to zero may be used as local VE IDs. If set to zero then Priority field values with the seven most significant bits belonging to the described Priority range and the Tagging Extension bit set to zero shall not be used as local VE IDs.

Bit 1: If set to one then Priority field values with the seven most significant bits belonging to the described Priority range and the Tagging Extension bit set to one may be used as local VE IDs. If set to zero then Priority field values with the seven most significant bits belonging to the described Priority range and the Tagging Extension bit set to one shall not be used as local VE IDs.

Bit 8 .. 2: Reserved.

NOTE 7 – As defined in FC-FS-6, the Priority field is an 8-bit field of which the seven most significant bits are the priority and the least significant bit is the Tagging Extension bit.

4.3.51 Update VE Mappings (UVEM)

4.3.51.1 Overview

The UVEM ELS is used by an N_Port that supports Priority Tagging (see 9) to update the VE Identification Server with changes in the global VE ID to Fabric VE ID mappings when one or more VEs are instantiated or deinstantiated. If the N_Port is logged in to a device that supports Priority Tagging, the UVEM ELS is used to notify the device about relevant changes in the global VE ID to Fabric VE ID mappings when one or more VEs are instantiated or deinstantiated.

4.3.51.2 Protocol

- a) Update VE Mappings Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence.

4.3.51.3 Request Sequence

Addressing: The S_ID field shall be set to the N_Port_ID of the N_Port originating the UVEM Request. The D_ID shall be set to:

- a) the F_Port Controller (FFFFFEh) if updating the VE Identification Server;
- b) the N_Port_ID of the destination N_Port if notifying another N_Port.

Payload: The format of the UVEM Request Payload is shown in table 176.

Table 176 – UVEM Request Payload (Sheet 1 of 2)

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	UVEM (B1h)	00h	00h	00h
1	Descriptor list length = $24 + m*28 + q*28$			
2 .. 7	VEM ID descriptor			
8 .. 14	Instantiated VE Mapping descriptor #1			
15 .. 21	Instantiated VE Mapping descriptor #2			
	..			
$7*(m+0)+1 .. 7*(m+1)$	Instantiated VE Mapping descriptor #m			
$7*(m+1)+1 .. 7*(m+2)$	Deinstantiated VE Mapping descriptor #1			
$7*(m+2)+1 .. 7*(m+3)$	Deinstantiated VE Mapping descriptor #2			

Table 176 – UVEM Request Payload (Sheet 2 of 2)

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
	..			
7*(m+q)+1 .. 7*(m+q+1)	Deinstantiated VE Mapping descriptor #q			

The format of the VEM ID descriptor is shown in table 177.

Table 177 – VEM ID descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Descriptor Tag = 0001 000Ah			
1	Descriptor Length = 16			
2	MSB			
3	VEM ID			
4				
5				

VEM ID: contains the VEM ID of the originating VEM.

The format of the Instantiated VE Mapping descriptor is shown in table 178.

Table 178 – Instantiated VE Mapping descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Descriptor Tag = 0001 000Bh			
1	Descriptor Length = 20			
2	MSB			
3	Global VE ID			
4				
5				
6	N_Port_ID			Local VE ID

The format of the Deinstantiated VE Mapping descriptor is shown in table 179.

Table 179 – Deinstantiated VE Mapping descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Descriptor Tag = 0001 000Ch			
1	Descriptor Length = 20			
2	MSB			
3	Global VE ID			
4				
5				
6	N_Port_ID			Local VE ID

Global VE ID: the identifier used to globally identify a VE.

N_Port_ID: the N_Port_ID of the originating N_Port.

Local VE ID: the value selected to locally identify a VE.

Deinstantiating a VE mapping not registered in the VE Identification Server shall not be treated as an error.

A Deinstantiated VE Mapping descriptor 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 descriptor and the VEM ID specified in the VEM ID descriptor. An UVME request payload containing a Deinstantiated VE Mapping descriptor with a NULL Global VE ID and a NULL Local VE ID shall contain only a VEM ID descriptor and the Deinstantiated VE Mapping descriptor with a NULL Global VE ID and a NULL Local VE ID.

4.3.51.4 Reply Sequence

LS_RJT: LS_RJT signifies the rejection of the UVEM Request.

LS_ACC: LS_ACC signifies the success of the UVEM processing. The format of the UVEM LS_ACC Payload is shown in table 180.

Table 180 – UVEM LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	02h	00h	00h	00h

4.3.52 Exchange Diagnostic Capabilities (EDC)

4.3.52.1 Description

The EDC ELS is used to exchange diagnostic capabilities parameters between requesting and responding FC_ports.

4.3.52.2 Protocol

- a) Exchange Diagnostic Capabilities Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence

4.3.52.3 Request Sequence

Addressing: The S_ID field designates the source Nx_Port requesting the exchange of diagnostic capabilities. The D_ID field shall be set as follows:

- a) to the F_Port Controller (FFFFFFEh); or
- b) to any Nx_Port N_Port_ID.

Payload: The format of the EDC Request Payload is shown in table 181.

Table 181 – EDC Request Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	EDC (17h)	00h	00h	00h
1	Descriptor list length = ((n-1)*4) bytes			
2 - n	Diagnostic capability descriptors			

An EDC request shall include zero or more Diagnostic Capability descriptors (see 4.3.52.5).

4.3.52.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the EDC command.

LS_ACC: LS_ACC signifies acceptance of the request and provides the requested data. The format of the LS_ACC Payload is shown in table 182.

Table 182 – EDC LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	02h	00h	00h	00h
1	Descriptor list length ((n-1)*4) bytes			
2	MSB			
3	Link Service Request Information descriptor (see 4.2.4.2)			
4	LSB			
5 - n	Diagnostic capability descriptors			

An EDC LS_ACC shall include zero or more Diagnostic Capability descriptors (see 4.3.52.5).

4.3.52.5 Diagnostic capability descriptors

4.3.52.5.1 Overview

The diagnostic capability descriptors included in an EDC request correspond to functions supported by the requesting FC_Port.

The diagnostic capability descriptors included in an EDC response correspond to functions supported by the responding FC_Port.

4.3.52.5.2 Link Fault Capability descriptor

Link Degrade Signaling is supported by monitoring errors in the Reed-Solomon FEC logic (see FC-FS-6). The Link Fault Capability descriptor provides a mechanism to exchange the Link Degrade Signaling function parameters between two ports. This allows each port to determine the error rate associated with the Link Degrade Signal when it is received.

The Link Fault Capability descriptor is shown in table 183.

Table 183 – Link Fault Capability descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Link Fault Capability descriptor tag = 0001 000Dh			
1	Descriptor Length = 12			
2	Degrade Activate Threshold			
3	Degrade Deactivate Threshold			
4	FEC Degrade Interval			

Degrade Activate Threshold: The value of the Degrade_Activate_Threshold register (see FC-FS-6).

Degrade Deactivate Threshold: The value of the Degrade_Deactivate_Threshold register (see FC-FS-6).

FEC Degrade Interval: The value of the FEC_Degrade_interval register (see FC-FS-6).

4.3.52.5.3 Congestion Signaling Capability descriptor

Congestion Signaling is supported by sending and/or receiving the Warning Congestion Signal and Alarm Congestion Signal (see FC-FS-6). The Congestion Signaling Capability descriptor provides a mechanism to exchange the Congestion Signaling Capability parameters between two ports. The format of the Congestion Signaling Capability descriptor is shown in Table 184.

Table 184 – Congestion Signaling Capability descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Congestion Signaling Capability Descriptor tag = 0001 000Fh			
1	Descriptor Length = 16			
2	Transmit Signal Capability			
3	Transmit Signal Frequency			
4	Receive Signal Capability			
5	Receive Signal Frequency			

Transmit Signal Capability: Defines the ability of the FC_Port to transmit the primitives used for Congestion Signals (see FC-FS-6). An FC_Port that is capable of transmitting only one signal value shall indicate support of the Warning Congestion Signal. The definition of the Transmit Signal Capability methods is listed in Table 185.

Table 185 – Signal Capability description

Bits	Description
31 - 4	Reserved
3 - 0	Signal Capability values: 0 = The Warning Congestion Signal and the Alarm Congestion Signal are not supported 1 = The Warning Congestion Signal is supported 2 = The Warning Congestion Signal and the Alarm Congestion Signal are supported All other values are reserved

Transmit Signal Frequency: The Transmit Signal Frequency describes the maximum frequency the transmitter is capable of transmitting Congestion Signals expressed as a minimum time between signals. The Transmit Signal Frequency fields are defined in Table 186.

Table 186 – Signal Frequency description

Bits	Description
31 - 26	Reserved
25 - 16	Signal Frequency Count: 0 = Reserved 1 to 999 = The time between signals in the units indicated in the Signal Frequency Units field. All other values are reserved
15 - 4	Reserved
3 - 0	Signal Frequency Units: 0 = Reserved 1 = Seconds 2 = Milliseconds All other values are reserved

Receive Signal Capability: Defines the ability of the FC_Port to receive the primitives used for Congestion Signals (see FC-FS-6). An FC_Port that is capable of receiving only one signal value shall indicate support of the Warning Congestion Signal. The definition of the Receive Signal Capability methods is listed in Table 185.

Receive Signal Frequency: The Receive Signal Frequency describes the maximum frequency the receiver is capable of receiving Congestion Signals expressed as a minimum time between signals. The Receive Signal Frequency fields are defined in Table 186.

At the completion of the EDC ELS exchange, each end-point of the link adopts the least capable settings for Signal Capability and Signal Frequency of the two end-points. If a port indicates it is not able to receive a signal then the connected port shall not transmit that signal. A port's frequency of transmission shall be based on the larger of the value in that port's Transmission Signal Frequency field and the value in the connected port's Receive Signal Frequency field. A sender of Congestion Signals shall not transmit signals more frequently than the Receive Signal Frequency of the connected port.

Example #1, the EDC sender is capable of transmitting and receiving the Warning Congestion Signal and the Alarm Congestion Signal at a frequency of one signal per 10 milliseconds. The EDC responder is capable of transmitting and receiving only the Warning Congestion Signal at a frequency of one signal per 100 milliseconds. At the conclusion of the EDC exchange, both end-points adopt transmission and reception of the Warning Congestion Signal at a frequency of one signal per 100 milliseconds and no Alarm Congestion Signal is sent.

Example #2, the EDC sender is capable of transmitting and receiving the Warning Congestion Signal and the Alarm Congestion Signal at a frequency of one signal per 10 milliseconds. The EDC responder is capable of transmitting the Warning Congestion Signal at a frequency of one signal per 100 milliseconds and is capable of receiving the Warning Congestion Signal and the Alarm Congestion Signal at a frequency of one signal per 10 milliseconds. At the conclusion of the EDC exchange, the EDC sender adopts

transmitting the Warning Congestion Signal and the Alarm Congestion Signal at a frequency of one signal per 10 milliseconds and adopts receiving the Warning Congestion Signal at a frequency of one signal per 100 milliseconds. The EDC responder adopts receiving the Warning Congestion Signal and the Alarm Congestion Signal at a frequency of one signal per 10 milliseconds and adopts transmitting the Warning Congestion Signal at a frequency of one signal per 100 milliseconds.

4.3.52.5.4 Frame Discard TOV descriptor

The Frame Discard Timeout provides a mechanism to free a buffer if processing of a frame extends beyond the timeout value (see FC-FS-6). The Frame Discard TOV descriptor provides a mechanism to exchange the Frame Discard TOV value between two FC_Ports. The format of the Frame Discard TOV descriptor is shown in table 187.

Table 187 – Frame Discard TOV descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Frame Discard TOV Descriptor tag = 0001 0010h			
1	Descriptor Length = 8			
2	F_D_TOV value			
3	Operational Flags			

F_D_TOV value: Defines the value used by the FC_Port for the Frame Discard Timeout (i.e., F_D_TOV) processing (see FC-FS-6).

Operational Flags: The Operational Flags field describes the characteristics of the F_D_TOV value. The Operational Flags field is described in table 188.

Table 188 – Operational Flags field description

Bit	Description
31 – 1	Reserved
0	This bit provides the units for the value provided in the F_D_TOV value field: 0 = milliseconds 1 = microseconds

4.3.53 Register Diagnostic Functions (RDF)

4.3.53.1 Description

The RDF command provides a method for Nx_Ports to indicate to the Fabric Controller the specific diagnostic functions supported by the Nx_Port and the range of diagnostic capabilities associated with each of the supported functions.

An RDF ELS request sent by an Nx_Port to the Fabric Controller requests the Fabric Controller to add the Nx_Port that is sending the RDF ELS to the list of Nx_Ports registered to receive the ELSs specified by the registered diagnostic function descriptors (see 4.3.53.5).

An RDF ELS request sent by the Fabric Controller to a registered Nx_Port indicates there is a change to the diagnostic functions supported by the Fabric. An Nx_Port receiving an RDF ELS request from the Fabric Controller should reregister with the Fabric Controller by sending a new RDF ELS request to the Fabric Controller.

Any subsequent RDF ELS request sent by an Nx_Port to the Fabric Controller indicates there is a change to the diagnostic functions supported by the Nx_Port . The Fabric Controller removes all previously registered diagnostic functions for that Nx_Port then adds the Nx_Port that is sending the RDF ELS to the list of Nx_Ports registered to receive the ELSs specified by the registered diagnostic function descriptors (see 4.3.53.5).

4.3.53.2 Protocol

- a) Register Diagnostic Functions Request Sequence; and
- b) LS_ACC or LS_RJT Reply Sequence

4.3.53.3 Request Sequence

Addressing: If an Nx_Port sends an RDF ELS to register supported diagnostic functions, the S_ID field designates the source Nx_Port requesting registration of the diagnostic functions and the D_ID designates the Fabric Controller (FFFFFFDh). If the Fabric Controller sends an RDF ELS to indicate a change to the diagnostic functions supported by the Fabric, the S_ID designates the Fabric Controller (FFFFFFDh) and the D_ID designates the address of the registered Nx_Port destination.

Payload: The format of the RDF ELS request Payload is shown in table 189.

Table 189 – RDF Request payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	RDF (19h)	00h	00h	00h
1	Descriptor list length = ((n-1)*4) bytes)			
2 - n	Diagnostic Function descriptor list			

Descriptor list length: The descriptor list length shall not be greater than 2040 bytes (i.e., n is less than or equal to 511). The descriptor list length in an RDF ELS request sent by the Fabric Controller shall be set to zero.

Diagnostic Function descriptor list: The list of registered diagnostic function descriptors (see 4.3.53.5) identifying the registered diagnostic functions supported by the requesting Nx_Port.

4.3.53.4 Reply Sequence

LS_RJT: LS_RJT signifies rejection of the RDF command.

LS_ACC: LS_ACC sent by the Fabric Controller signifies acceptance of the RDF request and returns the list of successfully registered diagnostic functions. The format of the LS_ACC Payload is shown in table 190.

LS_ACC sent by an Nx_Port signifies acceptance of the RDF request. The Descriptor list length shall be set to 12 and only the Link Service Request Information descriptor shall be returned.

Table 190 – RDF LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	LS_ACC (02h)	00h	00h	00h
1	Descriptor list length = ((n-1)*4) bytes)			
2	MSB			
3	Link Service Request Information descriptor			
4	(see 4.2.4.2) LSB			
5 - n	Diagnostic Function descriptor list (if any)			

Descriptor list length: The descriptor list length shall not be greater than 2040 bytes (i.e., n is less than or equal to 511).

Diagnostic Function descriptor list: The list of registered diagnostic function registration descriptors (see 4.3.53.5) describing the subset of diagnostic functions supported and successfully registered by the Fabric Controller.

If a registered diagnostic function (see table 191) is not registered by an Nx_Port then the Fabric Controller shall not transmit the associated ELS to that Nx_Port. A registered diagnostic function ELS transmitted to the Nx_Port shall contain only the descriptors successfully registered in the RDF ELS information exchange.

If none of the registered diagnostic functions registered by an Nx_Port are supported by the Fabric, then the Fabric Controller shall reply with an RDF ELS response with an empty Diagnostic Function descriptor list (i.e., a payload containing only the Link Service Request Information descriptor) to indicate the Fabric does not support any of the diagnostic functions in the Diagnostic Function descriptor list of the RDF ELS request.

RDF and registered diagnostic function ELS requests and responses are sent using a single frame Sequence.

An FC_Port sending an RDF ELS request or response to the Fabric Controller implies the originator has a Receive Data_Field size equal to the Receive Data_Field size provided in the FLOGI request for the class of service being used (see table 9 footnote c).

4.3.53.5 Registered Diagnostic function descriptors

4.3.53.5.1 Overview

The registered diagnostic function descriptors included in an RDF ELS request correspond to functions supported by the requesting Nx_Port and to be registered with the Fabric Controller. The list of registered diagnostic functions is shown in table 191.

Table 191 – Registered diagnostic functions

Function	Reference	Registration Descriptor	Reference
FPIN ELS	4.3.54	FPIN Registration descriptor	4.3.53.5.2

The registered diagnostic function descriptors included in an RDF ELS LS_ACC correspond to functions supported by the Fabric Controller and the requesting Nx_Port.

4.3.53.5.2 FPIN Registration descriptor

The FPIN Registration descriptor registers an Nx_Port with the Fabric Controller to receive FPIN ELS requests. The format of the FPIN Registration descriptor is shown in table 192.

Table 192 – FPIN Registration descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	FPIN Registration descriptor tag = 0003 0001h			
1	FPIN Registration descriptor length ((n-1)*4) bytes			
2	Descriptor Tag count			
3	Descriptor Tag list			

Descriptor Tag count: Number of descriptor tags indicating the FPIN functions supported (see 4.3.54).

Descriptor Tag list: List of FPIN descriptor tags supported. The Descriptor Tag list in an RDF ELS request (see 4.3.54) indicates the functions supported by the requesting Nx_Port. The Descriptor Tag list in an RDF ELS LS_ACC indicates the functions supported by both the requesting Nx_Port and the Fabric Controller, which may be a subset of the functions listed in the RDF ELS request.

4.3.54 Fabric Performance Impact Notifications (FPIN)

4.3.54.1 Description

The FPIN ELS is used to notify registered Nx_Ports of events detected by an FC_Port (see Annex A for examples). The FPIN ELS contains one or more notification descriptors for each detected event. Simultaneously detected events may be coalesced into a single FPIN ELS. Each descriptor contains a description of the event.

4.3.54.2 FPIN events sent by the Fabric Controller

If the Fabric detects an event associated with the descriptors listed in this subclause, the Fabric Controller may issue an FPIN ELS to the registered Nx_Ports that includes one or more of the following descriptors:

- a) Link Integrity Notification descriptor (see 4.3.54.7.2);
- b) Delivery Failure notification descriptor (see 4.3.54.7.3);
- c) Peer Congestion Notification descriptor (see 4.3.54.7.4); or
- d) Congestion Notification descriptor (see 4.3.54.7.5).

4.3.54.3 FPIN events sent by the Nx_Port

If a registered Nx_Port detects an event associated with the descriptors listed in this subclause, then that Nx_Port may issue an FPIN ELS to the Fabric Controller. FPIN ELSs received by the Fabric Controller from an Nx_Port are forwarded by the Fabric Controller according to the definition of the associated Fabric notification descriptor (see 4.3.54.7). The FPIN ELS includes one or more of the following:

- a) Link Integrity Notification descriptor (see 4.3.54.7.2);
- b) Delivery Failure notification descriptor (see 4.3.54.7.3);
- c) Peer Congestion Notification descriptor (see 4.3.54.7.4); or
- d) Congestion Notification descriptor (see 4.3.54.7.5).

4.3.54.4 Protocol

- a) Fabric Performance Impact Notification Request Sequence
- b) No Reply Sequence

4.3.54.5 Request Sequence

Addressing: If the Fabric is using FPIN to notify a registered Nx_Port of a detected event (see 4.3.54.2), the S_ID is the Fabric Controller (FFFFFFDh) and the D_ID is the address of the registered Nx_Port destination. If an Nx_Port is using FPIN to notify the Fabric of a detected event (see 4.3.54.3), the S_ID designates the Nx_Port indicating the event and the D_ID is the Fabric Controller (FFFFFFDh).

Payload: The format of the FPIN Request Payload is shown in table 193.

Table 193 – FPIN ELS format

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	FPIN (16h)	00h	00h	00h
1	Descriptor list length = ((n-1)*4) bytes)			
2 - n	Descriptor list			

Descriptor list length: The descriptor list length shall not be greater than 2040 bytes (i.e., n is less than or equal to 511).

Descriptor list: The list of notification descriptors identifying the event detected by the Fabric Controller or an Nx_Port (see 4.3.54.7).

4.3.54.6 Reply Sequence

None.

4.3.54.7 Fabric notification descriptors

4.3.54.7.1 Overview

The Fabric notification descriptors included in an FPIN request from the Fabric Controller correspond to events detected by the Fabric (see 4.3.54.2).

The Fabric notification descriptors included in an FPIN request from an Nx_Port correspond to events detected by the Nx_Port (see 4.3.54.3).

4.3.54.7.2 Link Integrity notification descriptor

The Link Integrity notification descriptor indicates an error threshold has been exceeded. The format of the Link Integrity notification descriptor is shown in table 194.

Table 194 – Link Integrity notification descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Link Integrity notification descriptor tag = 0002 0001h			
1	Descriptor list length = ((n-1)*4) bytes)			
2	MSB			
3	Detecting Port Name			LSB
4	MSB			
5	Attached Port Name			LSB
6	Event Type		Event Modifier	
7	Event Threshold			
8	Event Count			
9	Port Name Count			
10 - n	Port Name List			

Detecting Port Name: The Name_Identifier of the FC_Port detecting the notification event.

Attached Port Name: The Name_Identifier of the FC_Port connected to the detecting FC_Port.

Event Type: The type of event reported in the Link Integrity notification descriptor. The Event Type values are shown in table 195.

Table 195 – Link Integrity notification Event Type values

Value (hex)	Description
0000	Unknown
0001	Link Failure
0002	Loss-of-Synchronization
0003	Loss-of-Signal
0004	Primitive Sequence Protocol Error
0005	Invalid Transmission Word
0006	Invalid CRC
0007	Uncorrectable FEC Error
000F	Device Specific
All other values are reserved	

Event Modifier: A value further describing the Event Type (e.g., Device Specific Event Type). The Event Modifier values for the Device Specific Event Type are shown in table 196.

Table 196 – Event Modifier values for the Device Specific Event Type

Value (hex)	Description
0000	None/Unknown
0001	Resource Contention
0002 - 000F	Reserved
All other values are implementation specific	

Event Threshold: The duration in milliseconds of the Link Integrity event detection cycle.

Event Count: The minimum number of event occurrences during the Event Threshold to cause generation of a Link Integrity event.

NOTE 8 – In cases where a single anomaly could cause multiple occurrence of an event, an Event Count for link integrity should take into account Error Intervals (see FC-FS-6) so that thresholds are not exceeded prematurely.

Port Name Count: Number of Port Names in the Port Name List. If the Port Name Count is zero then the Port Name List is empty. If the Port Name List is empty then:

- a) If the event is detected by an F_Port, the Attached Port Name identifies the only Nx_Port affected by the event; or
- b) if the event is detected by an Nx_Port, then the Detecting Port Name identifies the only Nx_Port affected by the event.

Port Name List: The list of N_Port_Names of the Nx_Ports accessible through the link connecting the Attached Port Name and Detecting Port Name (e.g., the Nx_Ports affected by the event). An N_Port_Name is the Name_Identifier associated with an Nx_Port (see FC-FS-6).

Link Integrity events detected by an F_Port are distributed by the Fabric Controller to all registered Nx_Ports in the zone membership list(s) associated with the PN_Port attached to the detecting F_Port. The Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

Link Integrity events detected by an Nx_Port are sent to the Fabric Controller. The Fabric Controller distributes the Link Integrity events to all registered Nx_Ports in the zone membership list(s) associated with the PN_Port of the detecting Nx_Port and to all registered Nx_Ports in the zone membership list(s) associated with the N_Port_Names in the Port Name List. For events distributed by the Fabric Controller, the Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

Link Integrity events detected by an E_Port are distributed by the Fabric Controller to all registered Nx_Ports affected by the event. The Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

4.3.54.7.3 Delivery Failure notification descriptor

The Delivery Failure notification descriptor indicates a delivery event has been detected (i.e., a frame has been discarded by the Fabric or a device). The format of the Delivery Failure notification descriptor is shown in table 197.

Table 197 – Delivery Failure notification descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Delivery Failure notification descriptor tag = 0002 0002h			
1	Descriptor list length = (44 bytes)			
2	MSB			
3	Detecting Port Name			LSB
4	MSB			
5	Attached Port Name			LSB
6	Reason Code			
7 - 12	Event Data (discarded frame header)			

Detecting Port Name: The Name_Identifier of the FC_Port detecting the notification event.

Attached Port Name: The Name_Identifier of the FC_Port connected to the detecting FC_Port.

Reason Code: The reason for the Delivery Failure notification event. The Reason Code values are shown in table 198.

Table 198 – Delivery Reason Code values

Value (hex)	Description
0000 0000	Not specified
0000 0001	Timeout (e.g., protocol, ULP)
0000 0002	Unable to route (e.g., change in Fabric topology)
0000 0003	Frame Discard Timeout (see FC-FS-6)
0000 0004 to 0000 000E	Reserved
0000 000F	Device specific
0000 0010 to 0000 001F	Vendor specific
All other values are reserved	

Event Data: The Event Data for a Delivery Failure notification event contains the Frame_Header of the affected frame (see FC-FS-6).

Delivery Failure notification events detected by an F_Port or an E_Port are distributed by the Fabric Controller. If the FC_Port identified by the S_ID of the affected frame is in the zone membership list(s) associated with the FC_Port identified by the D_ID of the affected frame, then the Fabric Controller distributes the Delivery Failure notification event to the FC_Port identified by the S_ID of the affected frame.

Delivery Failure notification events detected by an Nx_Port are sent to the Fabric Controller. If the Nx_Port is in the zone membership list(s) associated with the FC_Port identified by the S_ID in the Delivery Event Data field, then the Fabric Controller distributes the Delivery Failure notification event to the FC_Port identified by the S_ID in the Delivery Event Data field.

The detection of delivery notification events that are sent by the Fabric Controller or by the Nx_Port are implementation dependent and may differ based on the upper layer protocol (see Annex A).

4.3.54.7.4 Peer Congestion notification descriptor

The Peer Congestion notification descriptor indicates a congestion condition has been detected that affects ports in the same zone as the congested FC_Port (i.e., the peers of the congested FC_Port). The format of the Peer Congestion notification descriptor is shown in table 199.

Table 199 – Peer Congestion notification descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Peer Congestion notification descriptor tag = 0002 0003h			
1	Descriptor list length = ((n-1)*4) bytes			

Table 199 – Peer Congestion notification descriptor (continued)

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
2	MSB			
3	Detecting Port Name			LSB
4	MSB			
5	Attached Port Name			LSB
6	Event Type		Event Modifier	
7	Event Threshold			
8	Port Name Count			
9 - n	Port Name List			

Detecting Port Name: The Name_Identifier of the FC_Port detecting the notification event.

Attached Port Name: The Name_Identifier of the FC_Port connected to the detecting FC_Port.

Event Type: The type of congestion event detected (see Annex A). The Congestion Event Type values are shown in table 200.

Table 200 – Congestion Event Type values

Value (hex)	Description
0000	Congestion cleared
0001	Lost credit
0002	Credit stall
0003	Oversubscription
0003 to 000E	Reserved
000F	Device specific
0010 to 001F	Vendor specific
All other values are reserved	

Event Modifier: A value further describing the Event Type (e.g., Device Specific Event Type). The Event Modifier values for the Device Specific Event Type are shown in table 196.

Event Threshold: The duration in milliseconds of the congestion event detection cycle.

Port Name Count: Number of Port Names in the Port Name List. If the Port Name Count is zero then the Port Name List is empty. If the Port Name List is empty then:

- a) If the event is detected by an F_Port, the Attached Port Name identifies the only Nx_Port affected by the event; or

- b) if the event is detected by an Nx_Port, then the Detecting Port Name identifies the only Nx_Port affected by the event.

Port Name List: The list of N_Port_Names of the Nx_Ports accessible through the link connecting the Attached Port Name and Detecting Port Name (e.g., the Nx_Ports affected by the event). An N_Port_Name is the Name_Identifier associated with an Nx_Port (see FC-FS-6).

If an F_Port detects a congestion condition, then the Fabric Controller distributes Peer Congestion events to all registered Nx_Ports in the zone membership list(s) associated with the PN_Port attached to the detecting F_Port. The Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

If an Nx_Port detects a congestion condition, then the Nx_Port sends Peer Congestion events to the Fabric Controller. The Fabric Controller distributes the Peer Congestion events to all registered Nx_Ports in the zone membership list(s) associated with the PN_Port of the detecting Nx_Port and to registered the Nx_Ports in the zone membership list(s) associated with the N_Port_Names in the Port Name List. For events distributed by the Fabric Controller, the Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

If an E_Port detects a congestion condition, then the Fabric Controller distributes Peer Congestion events to all registered Nx_Ports affected by the event. The Port Name List shall contain the N_Port_Names of the Nx_Ports affected by the event and that are in the zone membership list(s) associated with the Nx_Port(s) receiving the notification.

While the congestion condition persists, Peer Congestion events are sent by the detecting FC_Port at the interval indicated by the value in the Event Threshold field.

The status of the Peer Congestion condition is cleared when:

- a) Peer Congestion events cease for a period of two-times the Event Threshold;
- b) a Peer Congestion event with a Congestion Event Type value of “Congestion Cleared” is received; or
- c) the Attached Port Name is no longer in the Fabric (e.g., an Nx_Port identified by the Attached Port Name or the Detecting Port Name in the Peer Congestion Event is no longer reachable).

4.3.54.7.5 Congestion notification descriptor

The Congestion notification descriptor indicates a congestion condition has been detected at the FC_Port. Note, Congestion notifications are distinct from Congestion Signals (see Annex A). The format of the Congestion notification descriptor is shown in table 201.

Table 201 – Congestion notification descriptor

Bits Word	31 .. 24	23 .. 16	15 .. 8	7 .. 0
0	Congestion notification descriptor tag = 0002 0004h			
1	Descriptor list length = 12			
2	Event Type		Event Modifier	
3	Event Threshold			
4	Severity	Reserved		

Event Type: The type of congestion event detected (see Annex A). The Congestion Event Type values are shown in table 200.

Event Modifier: A value further describing the Event Type (e.g., Device Specific Event Type). The Event Modifier values for the Device Specific Event Type are shown in table 196.

Event Threshold: The duration in milliseconds of the congestion event detection cycle.

Severity: A value indicating the distress level of the detected congestion event. The Severity values are shown in table 202.

Table 202 – Congestion notification Severity values

Value (hex)	Description
F1	Warning
F7	Alarm
All other values are reserved	

If an F_Port detects a congestion condition, then the Fabric Controller distributes Congestion events to the attached FC_Port.

If an Nx_Port detects a congestion condition, then the Nx_Port sends Congestion events to the Fabric Controller.

If an E_Port detects a congestion condition, then the Fabric Controller distributes Congestion events to the attached FC_Port.

While the congestion condition persists, Congestion events are sent by the detecting FC_Port to the attached FC_Port at the interval indicated by the value in the Event Threshold field. The Congestion notification descriptor is the only descriptor included in the FPIN ELS payload.

The status of the Congestion condition is cleared when:

- a) Congestion events cease for a period of two-times the Event Threshold,
- b) a Congestion event with a Congestion Event Type value of “Congestion Cleared” is received,
or
- c) the attached FC_Port is no longer reachable.

4.4 Extended Link Service Reply Sequences

4.4.1 Overview

An ELS Reply Sequence shall signify that the ELS request Sequence is completed. The reply Sequence may contain data in the Payload following the ELS_Command code word. The format and meaning of the Payload is specified in the request ELS definition.

4.4.2 LS_ACC

The Link Service Accept (LS_ACC) ELS reply Sequence shall notify the originator of an ELS request that the ELS request Sequence has been completed. The Responder shall terminate the Exchange by setting the Last Sequence bit (Bit 20) in F_CTL on the last Data frame of the reply Sequence. The first byte of the Payload shall contain 02h. The remainder of the Payload is unique to the ELS request.

Protocol: LS_ACC is the reply Sequence for several ELSs as indicated in the applicable clause.

Addressing: The D_ID field designates the source of the ELS Sequence being accepted while the S_ID field designates the destination of the request Sequence being accepted.

Payload: The Payload content following the ELS_Command code (02XXXXXXh) is defined within individual ELS requests.

4.4.3 Reply Sequence

none

4.4.4 Link Service Reject (LS_RJT)

4.4.4.1 Description

The Link Service Reject (LS_RJT) shall notify the transmitter of a Link Service request that the Link Service request Sequence has been rejected. A four-byte reason code shall be contained in the Data_Field. Link Service Reject may be transmitted for a variety of conditions that may be unique to a specific Link Service request (e.g., if the Service Parameters specified in a Login frame were logically inconsistent or in error, a P_RJT frame would not be transmitted in response, but rather a Link Service Reject).

4.4.4.2 Payload

Addressing: The D_ID field designates the source of the ELS request being rejected while the S_ID field designates the destination of the request Sequence being rejected.

Payload: The first word of the Payload shall contain the ELS_Command code (01000000h). The next four bytes of this field shall indicate the reason for rejecting the request (see figure 2 and tables 203 and 204). The first error condition encountered shall be the error reported.

NOTE 9 – The applicable ELSs column in table 204 is not necessarily complete (i.e., a given reply may contain a reason code explanation other than what is indicated for it in the table).

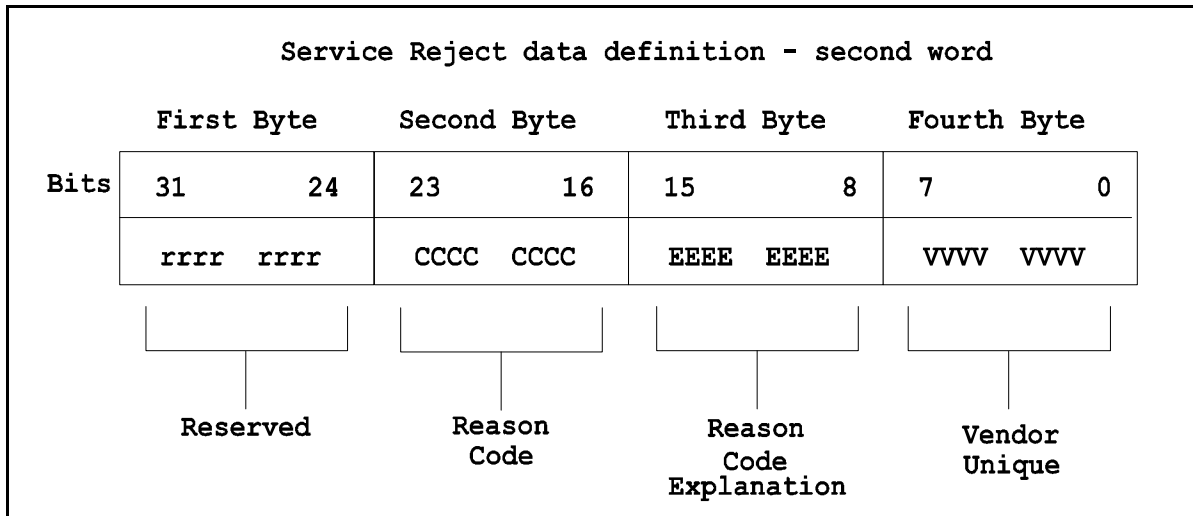


Figure 2 – LS_RJT format

Table 203 – LS_RJT Reason Codes

Encoded Value (Bits 23-16)	Description	Explanation
01h	Invalid ELS_Command code	The ELS_Command code in the Sequence being rejected is invalid.
03h	Logical error	The request identified by the ELS_Command code and Payload content is invalid or logically inconsistent for the conditions present.
05h	Logical busy	The Link Service is logically busy and unable to process the request at this time.
07h	Protocol error	This indicates that an error has been detected that violates the rules of the ELS Protocol that are not specified by other error codes.
09h	Unable to perform command request	The Recipient of a Link Service command is unable to perform the request at this time.
0Bh	Command not supported	The Recipient of a Link Service command does not support the command requested.
0Eh	Command already in progress	
20h	FIP Error	See FC-BB-6.
FFh	Vendor specific error (See bits 7-0)	The Vendor specific error bits may be used by Vendors to specify additional reason codes.
Others	Reserved	

4.4.4.3 Reply Sequence

none

Table 204 – LS_RJT Reason Code Explanations

Encoded Value (Bits 15-8)	Description
00h	No additional explanation
01h	Service Parm error - Options
03h	Service Parm error - Initiator Ctl
05h	Service Parm error - Recipient Ctl
07h	Service Parm error - Receive Data_Field Size
09h	Service Parm error - Concurrent Seq
0Bh	Service Parm error - Credit
0Dh	Invalid N_Port_Name/F_Port_Name
0Eh	Invalid Node_Name/Fabric Name
0Fh	Invalid Common Service Parameters
11h	Invalid Association_Header
13h	Association_Header required
15h	Invalid Originator S_ID
17h	Invalid OX_ID-RX_ID combination
19h	Command (request) already in progress
1Eh	N_Port Login required
1Fh	Invalid N_Port_ID
21h	Obsolete
23h	Obsolete
25h	Obsolete
27h	Obsolete
29h	Insufficient resources to support Login

Table 204 – LS_RJT Reason Code Explanations(continued)

Encoded Value (Bits 15-8)	Description
2Ah	Unable to supply requested data
2Ch	Request not supported
2Dh	Invalid Payload length
30h	Obsolete
31h	Obsolete
32h	Obsolete
33h	Obsolete
34h	Obsolete
35h	Obsolete
36h	Obsolete
37h	Obsolete
38h	Obsolete
40h	Obsolete
41h	Obsolete
42h	Obsolete

Table 204 – LS_RJT Reason Code Explanations(continued)

Encoded Value (Bits 15-8)	Description
44h	Invalid Port/Node_Name
46h	Login Extension not supported
48h	Authentication required (see FC-SP-2)
50h	Periodic Scan Value not allowed
51h	Periodic Scanning not supported
52h	No Resources Assigned
60h	MAC addressing mode not supported
61h	Proposed MAC address incorrectly formed
62h	VN2VN_Port not in Neighbor Set.
Others	Reserved

5 FC-4 Link Service

An FC-4 Link Service request solicits a destination Port (i.e., an Fx_Port or an Nx_Port) to perform a function or service in order to support an individual FC-4 Device_Data protocol. The Information Category for a request shall be specified as Unsolicited Control. A FC-4 Link Service reply may be transmitted in answer to a FC-4 Link Service request. The Information Category for a reply shall be specified as Solicited Control. Each request or reply shall be composed of a single Sequence. The format of the request or reply shall be specified by the individual FC-4 being supported and is beyond the scope of this standard. Each Sequence may be composed of one or more frames.

The protocols supported by the FC-4 Link Services shall be performed within a single Exchange, intended exclusively for the purpose. FC-4 Link Service protocols are performed using a two Sequence Exchange. The protocols consist of a request Sequence by the Originator (i.e., an Nx_Port), transfer of Sequence Initiative (see FC-FS-6), and a reply Sequence from the Responder (i.e., an Nx_Port or an Fx_Port). The execution of a FC-4 Link Service may perform sequence abort functions and modify sequence initiative of other exchanges in a protocol specific manner. The Sequence Initiator and Sequence Recipient shall follow the rules for Sequence management and Recovery_Qualifier reuse as specified in FC-FS-6. The following rules regarding Sequence and Exchange management apply to FC-4 Link Services in addition to the rules specified in FC-FS-6:

- a) FC-4 Link Services shall only be Exchanges originated following N_Port Login;
- b) the Originator of the Exchange shall use the Discard multiple Sequences Exchange Error Policy (see FC-FS-6) for all FC-4 Link Service Exchanges;
- c) the Originator of an FC-4 Link Service Exchange shall detect an Exchange error following Sequence Initiative transfer if the reply Sequence is not received within a timeout interval equal to twice the value of R_A_TOV;
- d) if the Exchange Originator of an FC-4 Link Service Exchange detects an Exchange error, it shall abort the Exchange using ABTS and retry the protocol of the aborted Exchange with a different Exchange; and
- e) if the Sequence Initiator aborts a Sequence using ABTS (Abort Sequence Protocol) due to receiving an ACK with the Abort Sequence bits (5-4) set to 01b, the Sequence Initiator shall retry the Sequence after the Basic Accept (see FC-FS-6) is received for the aborted Sequence one time only.

R_CTL bits 31-28 (Word 0) are set = to 0011b to indicate a FC-4 Link_Data frame. The TYPE field for each FC-4 Link Service frame shall match the FC-4 Device_Data TYPE field as specified in FC-FS-6.

6 Login and Service Parameters

6.1 Introduction

The Login procedure is a method by which an Nx_Port establishes its operating environment with a Fabric, if present, and other destination Nx_Ports with which it communicates. Fabric Login and N_Port Login are both accomplished with a similar procedure using different D_IDs and possibly different S_IDs.

Login between an Nx_Port and the Fabric or between two Nx_Ports is long-lived. The number of concurrent Nx_Ports with which an Nx_Port may be logged in with is a function of the Nx_Port facilities available.

Login between an Nx_Port and the Fabric or between two Nx_Ports may use an explicit or implicit method. When Login is referred to throughout other sections of this standard, either the explicit or implicit procedure may be used. Implicit Login is assumed to provide the same functionality as Explicit Login.

Explicit Login is accomplished using a Login (FLOGI or PLOGI) ELS (see 4.3.7) within a new Exchange to transfer the Service Parameters (contained in the Payload) of the Nx_Port initiating the Login Exchange. The LS_ACC contains the Service Parameters of the Responder (contained in the Payload).

Implicit Login is a method of defining and specifying the Service Parameters of destination Nx_Ports by means other than the explicit use of the Login ELS. Specific methods of implicit Login are not defined in this standard.

Implicit Fabric Login, unlike Explicit Fabric Login, does not require the support, within an Nx_Port, of the FLOGI Link Service request and reply. Implicit Fabric Login may be supported in a variety of ways, requiring varying levels of support within an Nx_Port.

An Nx_Port determines its own native N_Port_ID through explicit or implicit Login by

- a) the Fabric, if present;
- b) implicit definition; or
- c) assignment in the PLOGI Sequence transmitted to a destination Nx_Port attached in a point-to-point topology.

Nx_Ports may collect address identifiers from other potential destination Nx_Ports from:

- a) a name server function, if present;
- b) implicit definition; or
- c) an alternate initialization procedure.

6.2 Fabric Login

6.2.1 Introduction

Login with the Fabric is required for all Nx_Ports, regardless of the class supported. Communication with other Nx_Ports shall not be attempted until the Fabric Login procedure is complete.

Fabric Login accomplishes the following functions:

- a) It determines the presence or absence of a Fabric;
- b) If a Fabric is present, it provides the Nx_Port with the F_Port_Name, the Fabric_Name or locally-attached Switch_Name, and the specific set of operating characteristics associated with the entire Fabric;
- c) If a Fabric is present, it provides the Fabric with the specific set of operating characteristics, N_Port_Name and Node_Name of the Nx_Port with the F_Port_Name, the Fabric_Name or locally-attached Switch_Name, and the specific set of operating characteristics associated with the entire Fabric;
- d) If a Fabric is present, the Fabric shall optionally assign or shall confirm the N_Port_ID of the Nx_Port that initiated the Login;
- e) If a Fabric is present, it initializes the buffer-to-buffer Credit;
- f) If the Nx_Port and the Fabric support Authentication, it enables the Nx_Port to perform Fabric Authentication (see FC-SP-2); and
- g) If the Nx_Port and the Fabric support Virtual Fabrics, it enables the subsequent negotiation of Virtual Fabrics parameters (see 8.2).
- h) If a Fabric is present, and the Fabric supports FDISC with S_ID=0, the Fabric and the Nx_Port may exchange advisory information on the potential use and availability of N_PORT_ID resources.

6.2.2 Explicit Fabric Login Procedure

6.2.2.1 Introduction

The explicit Fabric Login procedure shall require an Nx_Port to transmit a Fabric Login (FLOGI) Link Service ELS (see 4.3.7).

Explicit Fabric Login replaces previous Service Parameters. The Login procedure shall follow the Exchange and Sequence management rules, the buffer-to-buffer flow control rules, and the end-to-end flow control rules as specified in FC-FS-6.

6.2.2.2 Explicit Fabric Login Request

The Nx_Port shall transmit the FLOGI in a new Exchange. The Payload of FLOGI contains the Service Parameters of the Nx_Port, a 64-bit N_Port_Name of the Nx_Port, and a 64-bit Node_Name. The Service Parameters are as specified for F_Port Login in 6.6. The applicability of the Service Parameters to Fabric Login are given in tables 207 and 212. The Nx_Port shall assign an OX_ID and set the D_ID to the F_Port Controller Well-known address (i.e., FFFFFFFh).

If the Nx_Port is unidentified, an Nx_Port shall set the S_ID in the FLOGI to 000000h or 0000h || YY. If the Nx_Port sets the S_ID to 000000h, the Nx_Port is requesting the Fabric assign all 24 bits of the N_Port_ID. If the Nx_Port sets the S_ID to 0000h || YY, the Nx_Port is requesting the Fabric assign the upper 16 bits, bits 23 to 8, and validate the lower 8 bits, bits 7 to 0, of the N_Port_ID. An example of the use of S_ID of 0000h || YY is FC-AL-2. The lower 8 bits of the N_Port_ID are the AL_PA.

6.2.2.3 Responses to Explicit Fabric Login

The following are possible responses the Nx_Port may receive after transmitting a FLOGI:

- a) LS_ACC reply Sequence with OX_ID equal to the OX_ID of the FLOGI, and the Common Service N_Port/F_Port bit set to one (i.e., an Fx_Port). LS_ACC is the normal response to a Fabric Login request. The D_ID of the LS_ACC shall be the N_Port_ID assigned by the Fabric. If the S_ID in the FLOGI was 000000h, the D_ID shall be XXXXXX. If the S_ID in the FLOGI was 0000h || YY, the D_ID shall be XXXXY. If the S_ID in the FLOGI was XXXXXX, the D_ID shall be same value of XXXXXX. The Payload shall include the Service Parameters for the entire Fabric, a F_Port_Name, and the current Fabric_Name or the locally-attached Switch_Name. The Service Parameters are as specified for F_Port Login in 6.6. The applicability of the Service Parameters to Fabric Login are given in tables 207 and 212. The Nx_Port may continue operation with other Nx_Ports if the N_Port_ID, F_Port_Name, and Node_Switch_Fabric_Name (see 6.6.4) are the same as in a previous Fabric Login or proceed to N_Port Login;
- b) LS_ACC reply Sequence with OX_ID equal to the OX_ID of the FLOGI, and the Common Service N_Port/F_Port bit set to zero (i.e., an Nx_Port). This indicates a point-to-point connection with another Nx_Port. The D_ID of the LS_ACC shall be the S_ID of the FLOGI. The Payload shall include the Service Parameters from the FLOGI with all classes marked invalid, a 64-bit N_Port_Name and 64-bit Node_Name of the connected Nx_Port. If the received N_Port_Name is less than its N_Port_Name, the Nx_Port proceeds to N_Port Login. If the received N_Port_Name is greater than its N_Port_Name, the Nx_Port waits for PLOGI from the attached N_Port.;
- c) F_BSY with OX_ID equal to the OX_ID of the FLOGI. The D_ID shall be the S_ID of the FLOGI. The Fabric is busy. The Nx_Port may retry the FLOGI again later;
- d) P_BSY Sequence with OX_ID equal to the OX_ID of the FLOGI. The D_ID shall be the S_ID of the FLOGI. This indicates a point-to-point connection with another Nx_Port that is currently busy. The Nx_Port may proceed to N_Port Login after a delay to allow the destination Nx_Port to become not busy;
- e) F_RJT Sequence with OX_ID equal to the OX_ID of the FLOGI. The D_ID shall be the S_ID of the FLOGI. The Fabric has rejected the FLOGI request. The reason code contained in the Payload determines the Nx_Port's action. If the reason code is "Class not supported", the Nx_Port may originate a FLOGI in a different class. If the reason code is "Invalid S_ID", the Nx_Port may originate a FLOGI with a different S_ID:
 - A) If the S_ID of the rejected FLOGI was 000000h or 0000h || YY, the Nx_Port may select a 24 bit value, XXXXXX, for its N_Port_ID by a method outside this standard and originate a FLOGI with this value in the S_ID; and
 - B) If the S_ID of the rejected FLOGI was XXXXXX, the Nx_Port may select a value of '00 00 00' or '00 00 yy', or a new value 'XX XX XX' for its N_Port_ID by a method outside this standard and originate a FLOGI with this value in the S_ID.
- f) P_RJT Sequence, with OX_ID equal to the OX_ID of the FLOGI. The D_ID shall be the S_ID of the FLOGI. This indicates a point-to-point connection with another Nx_Port. The reason code contained in the Payload determines the Nx_Port's action. If the reason code is "Class not supported", the Nx_Port may proceed to N_Port Login in a different class than used for the FLOGI. For other reason codes, the Nx_Port should respond accordingly;

- g) LS_RJT Sequence with OX_ID equal to the OX_ID of the FLOGI. The D_ID of the LS_RJT shall be the N_Port_ID assigned by the Fabric. If the S_ID in the FLOGI was 000000h, the D_ID shall be XXXXXX. If the S_ID in the FLOGI was 0000h || YY, the D_ID shall be XXXXYY. If the S_ID in the FLOGI was XXXXXX, the D_ID shall be XXXXXX. The reason code contained in the Payload determines the Nx_Port's action. The Nx_Port may alter the Service Parameters based on the reason code and originate a new FLOGI;
- h) No Response may indicate a delivery error, e.g., error on the physical transport. The Nx_Port shall perform ELS error recovery (see 4.3.2). The Nx_Port may originate a new FLOGI after recovery; and
- i) If the received N_Port_Name is equal to its N_Port_Name, then the Nx_Port is connected to itself and this case is outside the scope of this standard. The FLOGI is discarded.

6.2.2.4 Nx_Port response to FLOGI

If an Nx_Port receives a FLOGI and it is not already logged into a Fabric, then the Nx_Port shall respond to the received FLOGI with an LS_ACC reply Sequence with the OX_ID equal to the OX_ID of the received FLOGI and the Common Service Parameter N_Port/F_Port bit set to zero (i.e., an Nx_Port). This indicates a point-to-point connection with another Nx_Port. The D_ID of the LS_ACC shall be the S_ID of the received FLOGI. The Payload shall include the Service Parameters from the received FLOGI with all classes marked invalid, and the 64-bit N_Port_Name and 64-bit Node_Name of the connected Nx_Port. If the received N_Port_Name is less than its N_Port_Name, the Nx_Port proceeds to N_Port Login. If the received N_Port_Name is greater than its N_Port_Name, the Nx_Port waits for PLOGI from the attached N_Port. If an Nx_Port receives a FLOGI and it is logged into a Fabric, it shall reject the FLOGI (see 4.3.7.4).

6.2.2.5 Relogin with the Fabric

During a Login with the Fabric, if the Nx_Port was previously logged in with the Fabric and the N_Port_ID, F_Port_Name, and the Node_Switch_Fabric_Name (see 6.6.4) are the same as the previous Login, the Nx_Port may continue current communications with other Nx_Ports that it has established Logins with; if the Nx_Port detects that the N_Port_ID, F_Port_Name, and/or the Node_Switch_Fabric_Name (see 6.6.4) has changed since the last Fabric Login, the Nx_Port shall implicitly logout with all Nx_Ports and wait an R_A_TOV timeout period before initiating or accepting communication with other Nx_Ports. The timeout period shall start when the Nx_Port detects the change. After waiting the timeout period, new N_Port Logins are required before the Nx_Port may communicate with other Nx_Ports.

6.2.3 SOFs

Fabric Login shall only be performed in Class 2, or 3. Since the Fabric may not support both classes, the FLOGI Sequence may require retry in a different Class with the appropriate SOF.

Fabric Login is valid for all supported classes as indicated by the validity bits in the FLOGI LS_ACC Reply Sequence.

Selection of the SOF for the FLOGI Sequence is based on the Classes supported by the originating Nx_Port. The FLOGI Sequence is transmitted and the appropriate action is specified in 6.2.2.3. If an F_RJT with reason code "Class of service not supported by entity at FFFFFEh" is received, another supported SOF shall be attempted until the Login procedure is complete or until all supported SOF types have been attempted. If all supported SOF types have been attempted and the Fabric has rejected all or timed out, the Fabric and Nx_Port are incompatible and outside intervention is required.

6.2.4 Frequency

Login between an Nx_Port and the Fabric should be long-lived. If Implicit Logout with the Fabric has occurred, it is necessary to perform a new Login with the Fabric (see 6.4.4).

6.2.5 Fabric Login completion - Originator

The Originator of the FLOGI request considers Fabric Login to have ended if

- a) in Class 2, the Originator has transmitted the ACK (EOft) to the LS_ACC; or
- b) in Class 3, the Originator has received the LS_ACC.

If Login is ended, the values of buffer-to-buffer Credit are initialized.

6.2.6 Fabric Login completion - Responder

The Responder of the FLOGI request considers Fabric Login to have ended if

- a) in Class 2, the Responder has received the ACK (EOft) to the LS_ACC; or
- b) in Class 3, the Responder has transmitted the LS_ACC.

If Fabric Login has ended successfully, the values of buffer-to-buffer Credit are initialized.

6.3 N_Port Login

6.3.1 Introduction

N_Port Login follows the Fabric Login procedure. If a Fabric is present, as determined by performing the Fabric Login procedure, an Nx_Port proceeds with N_Port Login according to 6.3.2.2. If a Fabric is not present, as determined by performing the Fabric Login procedure, an Nx_Port proceeds with N_Port Login according to 6.3.2.4.

N_Port Login accomplishes the following functions:

- a) It provides each Nx_Port with the other Nx_Port's operating characteristics, N_Port_Name and Node_Name;
- b) If a Fabric is not present, it assigns the native N_Port_ID for both Nx_Ports; and
- c) It initializes the Nx_Port end-to-end Credit.

In point-to-point topology or between NL_Ports on the same loop, buffer-to-buffer Credit is initialized. N_Port Login between two Nx_Ports is complete after each Nx_Port has received the Service Parameters of the other Nx_Port. This may be accomplished by either implicit or explicit N_Port Login.

An Nx_Port is required to Login with each Nx_Port with which it intends to communicate. This includes reserved and well-known address identifiers since they are considered to be N_Ports (see FC-FS-6).

NOTE 10 – It is not required that an Nx_Port provide the same Login information with each destination Nx_Port or with the Fabric. However, an Nx_Port should avoid using contradictory or conflicting parameters with different Login destinations.

The N_Port Common Service Parameters during N_Port Login are specified in 6.6.2 (See table 207 for applicability). The N_Port Class Service Parameters during N_Port Login are specified in 6.6.5 (See table 212 for applicability). Both the Common Service Parameters and Class Service Parameters apply to each Nx_Port during N_Port Login.

NOTE 11 – If an Nx_Port (A) receives a PLOGI from another Nx_Port (B), Nx_Port (A) should verify that it is not already logged in with an Nx_Port (C) with the same N_Port_Name but different Nx_Port native address identifier and Node_Name. If so, it should consider the prior Login to be ended and all open Sequences that it originated with or received from the destination Nx_Port are terminated before accepting the new Login. Such a situation may arise if configuration changes have occurred.

N_Port Login provides each Nx_Port with the other Nx_Port's Service Parameters. Knowledge of a destination Nx_Port's receive and transmit characteristics is required for data exchanges. Service Parameters of destination Nx_Ports are saved and used to communicate with those Nx_Ports. The Service Parameters interchanged between two Nx_Ports may be asymmetrical. Saving the Service Parameters of destination Nx_Ports with which an Nx_Port communicates requires Nx_Port resources. These resources should be released using the destination N_Port Logout procedure (see 6.4).

Due to the resetting behavior of a PLOGI (e.g., termination of all open exchanges with the destination port), a port shall only send a PLOGI to a destination port if it is not logged in with the destination port. Examples of why a port is not logged in include:

- a) it has determined that a configuration change has occurred;
- b) it has lost knowledge of the Login parameters with the destination port;
- c) the destination port has responded with a frame that indicates an error condition (e.g., LOGO, P_RJT);
- d) the local port has logged out the destination port, either implicitly or explicitly, due to resource constraints; and
- e) the destination port failed to respond after 2 times R_A_TOV has expired.

A configuration change shall be determined by comparing the Port_Name, Node_Name, and Address_Identifier received in the ACC from an ADISC or PDISC with the values previously established during the previous Login process. A configuration change has occurred if either N_Port_Name or N_Port_ID do match and any of the three parameters do not match.

6.3.2 Explicit N_Port Login

6.3.2.1 Introduction

The explicit N_Port Login procedure shall require an Nx_Port to transmit a PLOGI Link request Sequence.

Explicit N_Port Login replaces previous Service Parameters. The Login procedure shall follow the Exchange and Sequence management rules, the buffer-to-buffer flow control rules, and the end-to-end flow control rules as specified in FC-FS-6.

A well-behaved Nx_Port shall Logout with another Nx_Port prior to initiating a new N_Port Login. However, if an Nx_Port receives or transmits a PLOGI request with another Nx_Port, it shall abnormally terminate open Sequences and respond to any new Sequences with that Nx_Port as though a Logout had been previously performed. During the N_Port Login procedure, other communication with the destination Nx_Port shall not be initiated or accepted. Once the N_Port Login procedure has been successfully completed, communication between the Nx_Ports may be initiated or accepted, (e.g., if Nx_Port(A) performs a PLOGI request with Nx_Port(B) and Nx_Port(B) transmits the LS_ACC reply, then either Nx_Port(A) or Nx_Port(B) may initiate communication for other protocols. Nx_Port(B) shall not be required to transmit a PLOGI request Sequence to Nx_Port(A) unless it wishes to invalidate or alter the existing Login parameters).

6.3.2.2 N_Port Login - Fabric present

The destination Nx_Port explicit Login procedure requires transmission of a N_Port Login (PLOGI) Link Service Sequence. The PLOGI is sent within an Exchange with an assigned OX_ID, the D_ID of the destination Nx_Port and a S_ID of originating Nx_Port. The Payload of this Sequence contains the Service Parameters, N_Port_Name, and Node_Name of the Nx_Port originating the PLOGI Sequence. The N_Port Service Parameters are as specified in 6.6. The applicability of the Service Parameters to N_Port Login are given in tables 207 and 212.

The normal reply Sequence to a PLOGI Link Service Sequence by an Nx_Port is a LS_ACC Link Service Reply Sequence within the Exchange identified by the OX_ID of the Login Sequence and the RX_ID assigned by the Responder with a D_ID of the originating Nx_Port (PLOGI Sequence) and a S_ID of the responding Nx_Port. The Payload of the LS_ACC contains the Service Parameters of the responding Nx_Port.

6.3.2.3 Responses to N_Port Login - Fabric present

The following are possible responses the Nx_Port may receive in response to transmitting a PLOGI with a Fabric present:

- a) LS_ACC reply Sequence with OX_ID equal to the OX_ID of the PLOGI, and the Common Service N_Port/F_Port bit = 0 (i.e., Responder is an Nx_Port) - This is the normal response to a N_Port Login request. The D_ID of the LS_ACC shall be the S_ID from the PLOGI. The S_ID of the LS_ACC shall be D_ID from the PLOGI. The Payload shall include the Service Parameters for the destination Nx_Port, a 64-bit N_Port_Name and a 64-bit Node_Name. The N_Port Service Parameters are as specified for in 6.6. The applicability of the Service Parameters to N_Port Login are given in tables 207 and 212. The Nx_Port may begin normal communication with the remote Nx_Port;
- b) F_BSY with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The Fabric is busy. The Nx_Port may retry the PLOGI again later;
- c) F_RJT Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The Fabric has rejected the PLOGI request. The reason code contained in the Payload determines the Nx_Port's action. If the reason code is "Invalid D_ID", N_Port Login is not possible with the addressed Nx_Port. The Nx_Port may attempt Login with other destination Nx_Ports. For other reason codes, the Nx_Port should respond according to the code;

- d) P_BSY Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The destination Nx_Port is busy. The Nx_Port may retry the PLOGI again later;
- e) P_RJT Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The reason code contained in the Payload determines the Nx_Port's action. If the reason code is "Class not supported", the Nx_Port may attempt PLOGI in a different class. For other reason codes, the Nx_Port should responded according to the code;
- f) PLOGI Sequence. The D_ID is the N_Port_ID of receiving Nx_Port. The S_ID is the N_Port_ID of the originating Nx_Port. The OX_ID is as assigned by the originating Nx_Port. The Payload shall include a 64-bit N_Port_Name and 64-bit Node_Name of the Nx_Port originating the PLOGI. This indicates a collision with N_Port Login from the destination Nx_Port. If the received N_Port_Name is less than the receiving Nx_Port's N_Port_Name, the Nx_Port sends LS_RJT to the originating Nx_Port with reason code "Command already in progress". If the received N_Port_Name is greater than its N_Port_Name, the Nx_Port processes the received PLOGI;
- g) LS_RJT Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID of the LS_ACC shall be the N_Port_ID of the destination Nx_Port. The reason code contained in the Payload determines the Nx_Port's action. The Nx_Port may alter the Service Parameters based on the reason code and originate a new PLOGI;
- h) No Response may indicate a delivery error, e.g., error on the physical transport. The Nx_Port shall perform ELS error recovery (see 4.3.2). The Nx_Port may originate a new PLOGI after recovery; and
- i) If the received N_Port_Name is equal to its N_Port_Name, then the Nx_Port is connected to itself and this case is outside the scope of this standard.

6.3.2.4 N_Port Login - No Fabric present

This procedure is based on the Nx_Port discovering the Fabric is not present during an attempted Fabric Login. (see 6.2.2.3) The destination N_Port explicit Login procedure requires transmission of a PLOGI Link Service Sequence in a new Exchange.

Only one Nx_Port in the point-to-point connection is required to transmit a PLOGI. If the N_Port_Names are exchanged during Fabric Login, the Nx_Port with the highest N_Port_Name shall transmit the PLOGI.

If either Nx_Port does not have access to the N_Port_Name of the connected Nx_Port, it may send a PLOGI. The processing requirements for responses received after transmitting a PLOGI resolves the condition of both Nx_Ports transmitting PLOGI.

An Nx_Port in a point-to-point configuration transmits PLOGI within a new Exchange. The S_ID shall be different than the D_ID. The Payload of this Sequence contains the Service Parameters, N_Port_Name, and Node_Name of the Nx_Port originating the PLOGI Sequence. The N_Port Service Parameters are as specified for in 6.6. The applicability of the Service Parameters to N_Port Login are given in tables 207 and 212.

6.3.2.5 Responses to N_Port Login - No Fabric present

The following are possible responses the Nx_Port may receive in response to transmitting a PLOGI in a point-to-point configuration:

- a) LS_ACC reply Sequence with OX_ID equal to the OX_ID of the PLOGI, and the Common Service N_Port/F_Port bit = 0 (i.e., the Responder is an Nx_Port). The D_ID of the LS_ACC shall be the S_ID from the PLOGI. The S_ID shall be the destination Nx_Port's N_Port_ID assigned by the D_ID in the PLOGI. The Payload shall include the Service Parameters for the destination Nx_Port, a 64-bit N_Port_Name and a 64-bit Node_Name. The N_Port Service Parameters are as specified in 6.6. The applicability of the Service Parameters to N_Port Login are given in tables 207 and 212. This is the normal response to a N_Port Login request. The Nx_Port may begin normal communication with the remote Nx_Port;
- b) P_BSY Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The destination Nx_Port is busy. The Nx_Port may retry the PLOGI again later;
- c) P_RJT Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID shall be the S_ID of the PLOGI. The reason code contained in the Payload determines the Nx_Port's action. If the reason code is "Class not supported", the Nx_Port may attempt N_Port Login in a different Class. For other reason codes, the Nx_Port should respond according to the code;
- d) LS_RJT Sequence with OX_ID equal to the OX_ID of the PLOGI. The D_ID of the LS_RJT shall be the N_Port_ID of the destination Nx_Port. The reason code contained in the Payload determines the Nx_Port's action. The Nx_Port may alter the Service Parameters based on the reason code and originate a new PLOGI;
- e) PLOGI Sequence. The D_ID is the address identifier of receiving Nx_Port. The S_ID is the N_Port_ID of the originating Nx_Port. The OX_ID is as assigned by the originating Nx_Port. The Payload shall include a 64-bit N_Port_Name and 64-bit Node_Name of the Nx_Port originating the PLOGI. This indicates a collision with N_Port Login from the destination Nx_Port. If the received N_Port_Name is less than the receiving Nx_Port's N_Port_Name, the Nx_Port sends LS_RJT to the originating Nx_Port with reason code "Command already in progress". If the received N_Port_Name is greater than its N_Port_Name, the Nx_Port processes the received PLOGI;
- f) No Response may indicate a delivery error, e.g., error on the physical transport. The Nx_Port shall perform ELS error recovery (see 4.3.2). The Nx_Port may originate a new PLOGI after recovery; and
- g) If the received N_Port_Name is equal to its N_Port_Name, then the Nx_Port is connected to itself and this case is outside the scope of this standard.

6.3.3 SOFs

N_Port Login is only supported in Class 2, and 3. Since the destination Nx_Port may not support all these classes for Login, the PLOGI Sequence may require retransmission in a different Class with the appropriate SOF in the same manner described for Fabric Login (see 6.2.3). Login is valid for all supported classes as indicated by the validity bits in the PLOGI LS_ACC Reply Sequence.

6.3.4 Frequency

The frequency of N_Port Login is installation dependent based on the frequency of configuration changes that may alter the N_Port_ID within an installation. Service Parameters of other Nx_Ports are retained until the next N_Port Login or until N_Port Logout (implicit or explicit) is performed.

6.3.5 N_Port Login completion - Originator

The Originator of the PLOGI request considers Login to have ended if

- a) in Class 2, the Originator has transmitted the ACK (EOFt) to the LS_ACC; or
- b) in Class 3, the Originator has received the LS_ACC.

If N_Port Login is ended with a Fabric present, the value of end-to-end Credit is initialized. If N_Port Login is ended in a point-to-point topology, the values of buffer-to-buffer and end-to-end Credit are initialized.

6.3.6 N_Port Login completion - Responder

The Responder of the PLOGI request considers Login to have ended if

- a) in Class 2, the Responder has received the ACK (EOFt) to the LS_ACC; or
- b) in Class 3, the Responder has transmitted the LS_ACC.

If N_Port Login is ended with a Fabric present, the value of end-to-end Credit is initialized. If N_Port Login is ended in a point-to-point topology, the values of buffer-to-buffer and end-to-end Credit are initialized.

6.4 Logout

6.4.1 Introduction

The destination Logout procedure provides a method for removing service between two N_Port_IDs or between an N_Port_ID and a Fabric. Logout releases resources, identifiers, and relationships associated with maintaining service between an N_Port_ID and a destination N_Port_ID or Fabric, including abnormal termination of all open Sequences and Exchanges between the two involved entities. Explicit Nx_Port Logout may be requested by an Nx_Port to another Nx_Port (see 6.4.2). Explicit Fabric Logout may be requested by an Nx_Port to a Fabric (e.g., to remove a virtual N_Port_ID previously assigned by the Fabric) or by a Fabric to an Nx_Port (see 6.4.3). Implicit Logout may occur between an Nx_Port and the Fabric (see 6.4.4).

6.4.2 Explicit N_Port Logout Procedure

The explicit N_Port Logout procedure shall require an Nx_Port to transmit a Logout (LOGO) Link Service request Sequence (see 4.3.8) to a destination Nx_Port. The Logout procedure is complete if the responding Nx_Port transmits a LS_ACC Link Service reply Sequence.

To explicitly Logout, the initiating Nx_Port shall terminate other open Sequences that it initiated with the destination Nx_Port prior to performing Logout, otherwise, the state of other open Sequences is unpredictable. If an Nx_Port receives a Logout request while another Sequence is open that was initiated from the requesting Nx_Port, it may reject the Logout request using an LS_RJT (Link Service Reject).

After an explicit Logout is performed with an Nx_Port, the default Login Service Parameters specified in table 207 and table 212 shall be functional if Login was explicit. After an explicit Logout is performed with an Nx_Port, the implicit Login Service Parameters shall be functional if Login was implicit.

6.4.3 Explicit Fabric Logout Procedure

The explicit Fabric Logout procedure shall require either:

- a) an Nx_Port to transmit a Logout (LOGO) Link Service request Sequence (see 4.3.8) to an F_Port Controller (i.e., Well-known address FFFFFEh); or
- b) an F_Port Controller to transmit a Logout (LOGO) Link Service request to an Nx_Port.

The explicit Fabric Logout procedure shall be complete for the responding Nx_Port, including an F_Port controller, if it transmits a LS_ACC Link Service reply Sequence. The explicit Fabric Logout procedure shall be complete for the requesting Nx_Port, including an F_Port controller, if it receives a LS_ACC Link Service reply Sequence.

If an Nx_Port, including an F_Port controller, that does not support explicit Fabric Logout receives a LOGO that requests explicit Fabric Logout, it shall reject the explicit Fabric Logout request using an LS_RJT with reason code "Command not supported" and reason code explanation "No additional explanation". If an Nx_Port, including an F_Port Controller, that sends a LOGO that requests explicit Fabric Logout receives in reply an LS_RJT with reason code "Command not supported" and reason code explanation "No additional explanation" the Nx_Port that sent the LOGO shall implicitly log out the other Nx_Port.

An F_Port Controller that supports N_Port_ID virtualization shall support explicit Fabric Logout originated by an Nx_Port.

An Nx_Port that has requested or accepted explicit Fabric Logout of an N_Port_ID shall implicitly log out the N_Port_ID with all other N_Port_IDs with which the N_Port_ID was logged in, and shall not originate frames from the N_Port_ID until after it has subsequently been reassigned the N_Port_ID by the Fabric. An F_Port Controller that has explicitly logged out an N_Port_ID shall neither originate nor route frames to the N_Port_ID until after it has completed a subsequent reassignment of the N_Port_ID.

6.4.4 Implicit Logout

If a PN_Port receives or transmits an NOS or OLS, all Nx_Ports at that PN_Port shall be implicitly logged out from the Fabric, if present, or attached Nx_Port in a point-to-point topology. Communication with other Nx_Ports shall not be accepted until the Fabric Login procedure is complete (implicit or explicit).

During Login with the Fabric, if the Nx_Port detects that the N_Port_ID, F_Port_Name and/or the Node_Switch_Fabric_Name (see 6.6.4) has changed since the last Fabric Login, and the Clean Address bit is zero, the Nx_Port shall implicitly logout with all Nx_Ports and wait an R_A_TOV timeout period before initiating or accepting communication with other Nx_Ports. The timeout period shall start if the Nx_Port detects the change. After waiting the timeout period, new N_Port Logins are required before the Nx_Port may communicate with other Nx_Ports.

During Login with the Fabric, if the Nx_Port detects that the N_Port_ID, F_Port_Name and/or the Node_Switch_Fabric_Name (see 6.6.4) has changed since the last Fabric Login, and the Clean Address bit is set to one, the Nx_Port shall implicitly logout with all Nx_Ports before initiating or accepting communication with other Nx_Ports.

NOTE 12 – If an Nx_Port receives OLS from the Fabric, the Fabric may be indicating configuration changes internal to the Fabric using the Online to Offline Protocol.

NOTE 13 – If an Nx_Port is concerned that a partial Fabric Login may be in process using its link immediately preceding its attempted Fabric Login, it may wait an R_A_TOV in order to ensure that the response it receives from the Fx_Port during Fabric Login is associated with its Login request.

6.4.5 The effects of FLOGI, FDISC, and LOGO

The effects of FLOGI, FDISC, and LOGO on pre-existing Fabric Logins is summarized in table 205.

Table 205 – Effects of FLOGI, FDISC, & LOGO on Permanent Port Name (PPN)

ELS Received (with D_ID FFFFEEh)	Condition of F_Port Controller		
	Condition 1: FLOGI not completed, or all VN_Ports logged out	Condition 2: FLOGI Completed, & at least one VN_Port logged in	Condition 3: FDISC(s) completed, & at least one VN_Port logged in
FLOGI S_ID=0^d	- LS_ACC, assign an N_Port_ID to first VN_Port. - Set PPN of new VN_Port to the F_Port_Name in the FLOGI LS_ACC.	-set BB_Credit to zero -implicit logout of all logged-in VN_Ports - disassociate the logged-out N_Port_IDs with PPN -re-assign one N_Port_ID to new VN_Port - Set PPN of new VN_Port to the F_Port_Name in the FLOGI LS_ACC.	- See Condition 2 (previous column).
FLOGI S_ID not = 0^d	-confirm or reject S_ID (see 6.2.2.3), “Response to Explicit Fabric Login.” - If confirmed, set PPN of FLOGI VN_Port (i.e., new N_Port_ID) to the F_Port_Name in the FLOGI LS_ACC.	-set BB_Credit to zero -Implicit logout of all logged-in VN_Ports - disassociate logged-out VN_Ports from PPN -confirm or reject S_ID (see 6.2.2.3), “Response to Explicit Fabric Login.” - If VN_Port logged-in, set PPN of VN_Port to the F_Port_Name in the FLOGI LS_ACC.	- See Condition 2 (previous column).
FDISC S_ID=0^d	- F_RJT (RC=Login required) for Class 2. - Discard for Class 3.	- If N_Port ID virtualization supported, LS_ACC (D_ID = assigned N_Port_ID), and set PPN of new VN_Port to the F_Port_Name in the FLOGI LS_ACC. ^{a, b} - If N_Port_ID virtualization not supported, LS_RJT (RC=command not supported, RCE=Request not supported.)	- LS_ACC (D_ID = assigned N_Port_ID), and set PPN of new VN_Port to the F_Port_Name in the FLOGI LS_ACC. ^{a, b}
<p>^a If another VN_Port is currently logged into the same F_Port with the same N_Port_Name as contained in the FDISC request payload, the VN_Port associated with the N_Port_Name in the FDISC request shall be implicitly logged out and then the FDISC request processed.</p> <p>^b If an FDISC with S_ID=0 is received and no more N_Port_IDs are available, the F_Port Controller shall respond with an LS_RJT with an LS_RJT Reason Code of Unable to perform command request and an LS_RJT Reason Code Explanation of insufficient resources.</p> <p>^c Any assigned N_Port_ID may be individually logged out and disassociated from the Permanent Port Name, including the FLOGI-assigned N_Port_ID; Name Server attributes for the N_Port_ID are cleared. Other logged-in VN_Ports continue to be associated with their current Permanent Port Names. See FC-GS-9 for a definition of Permanent Port Name.</p> <p>^d Whenever a new N_Port_ID is assigned by either FLOGI or FDISC, or logged out, the Name Server database is updated and the applicable RSCNs are sent.</p>			

Table 205 – Effects of FLOGI, FDISC, & LOGO on Permanent Port Name (PPN)

ELS Received (with D_ID FFFFh)	Condition of F_Port Controller		
	Condition 1: FLOGI not completed, or all VN_Ports logged out	Condition 2: FLOGI Completed, & at least one VN_Port logged in	Condition 3: FDISC(s) completed, & at least one VN_Port logged in
FDISC S_ID not = 0	- F_RJT (RC=Login required) for Class 2. - Discard for Class 3.	- If S_ID logged-in, LS_ACC - if S_ID not logged-in: F_RJT (RC=Login required) for Class 2. - Discard for Class 3.	- See Condition 2 (previous column).
OLS/NOS	- Perform Primitive Sequence Protocols (see FC-FS-6).	- implicit logout of all logged-in VN_Ports; Perform Primitive Sequence Protocols (see FC-FS-6). - disassociate the logged-out VN_Ports from PPN.	- See Condition 2 (previous column).
LOGO^d	- LS_ACC, no action	- If S_ID logged-in, log out the individual S_ID only, and disassociate only the logged-out VN_Port from the PPN. ^c - if S_ID not logged-in, LS_ACC no action.	- See Condition 2 (previous column).
<p>^a If another VN_Port is currently logged into the same F_Port with the same N_Port_Name as contained in the FDISC request payload, the VN_Port associated with the N_Port_Name in the FDISC request shall be implicitly logged out and then the FDISC request processed.</p> <p>^b If an FDISC with S_ID=0 is received and no more N_Port_IDs are available, the F_Port Controller shall respond with an LS_RJT with an LS_RJT Reason Code of Unable to perform command request and an LS_RJT Reason Code Explanation of insufficient resources.</p> <p>^c Any assigned N_Port_ID may be individually logged out and disassociated from the Permanent Port Name, including the FLOGI-assigned N_Port_ID; Name Server attributes for the N_Port_ID are cleared. Other logged-in VN_Ports continue to be associated with their current Permanent Port Names. See FC-GS-9 for a definition of Permanent Port Name.</p> <p>^d Whenever a new N_Port_ID is assigned by either FLOGI or FDISC, or logged out, the Name Server database is updated and the applicable RSCNs are sent.</p>			

6.5 Extended Login Processing

Support for an Extended Login request is specified by the Payload Bit (see 6.6.2.4.21) set to one in a PLOGI or FLOGI request and the associated LS_ACC. An Extended Login request provides the following information:

- a) Services Availability (see 6.6.8);
- b) Login Extension Data (see 6.6.9); and
- c) Clock Synchronization QoS (see 6.6.10).

If a Login of 256 bytes or more is required and the buffer conditions of the destination port are unknown, a timeout may be avoided with the following procedure:

- 1) issue a Login request with the Payload Bit (see 6.6.2.4.21) set to zero;
- 2) if the Query Buffer Conditions bit (see 6.6.2.4.16) is set to one in the LS_ACC, issue an RPBC ELS (see 4.3.31).; and
- 3) if both the ELS Receive Data_Field Size field in the RPBC LS_ACC Payload and the Buffer-to-Buffer Receive Data_Field Size field in the Login LS_ACC Payload are at least 256 bytes, issue a Login request with the Payload Bit set to one.

6.6 Service Parameters

6.6.1 ELS and LS_ACC Payload

Table 206 defines the Payload format for the FLOGI, PLOGI, and FDISC ELSs and the LS_ACCs. The definitions of the parameters are applicable to PLOGI, FLOGI, FDISC, PLOGI LS_ACC, FLOGI LS_ACC, and FDISC LS_ACC unless stated otherwise.

NOTE 14 – The Link Service may further limit values supplied during Login as specified by individual Upper Level Protocols.

Table 206 – FLOGI, PLOGI, FDISC or LS_ACC Payload

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	ELS_Command code			
1	MSB			
..	Common Service Parameters			
4	(16 bytes)			
				LSB
5	MSB			
6	Port_Name			
				LSB
7	MSB			
8	Node_Switch_Fabric_Name			
				LSB
9	MSB			
..	Obsolete			
12	(16 bytes)			
				LSB
13	MSB			
..	Class 2 Service Parameters			
16	(16 bytes)			
				LSB
17	MSB			
..	Class 3 Service Parameters			
20	(16 bytes)			
				LSB
21	MSB			
..	Auxiliary Parameter Data			
24	(16 bytes)			
				LSB
25	MSB			
..	Vendor Version Level			
28	(16 bytes)			
				LSB
29	MSB			
30	Services Availability ^a			
				LSB
31	Login Extension Data Length ^a			
32				
..	Reserved			
61				
62	Clock Synchronization QoS ^a			
63	(8 bytes)			
64 to n	Login Extension Data (if any)			

^a These fields are only present if the Payload Bit (see 6.6.2.4.21) is set to one. If the Payload bit is set to zero, these fields are not present in the Payload (i.e., the Payload is 116 bytes long).

6.6.2 Common Service Parameters

6.6.2.1 Applicability

Table 207 defines the applicability, by class as well as by PLOGI, FLOGI, PLOGI LS_ACC and FLOGI LS_ACC, of the Common Service Parameters to N_Port Login and Fabric Login. The Default Login Value column (see table 207) refers to the Login values to be used prior to a successful Login. These are words 1-4 in the Payload (see table 207).

Table 207 – Common Service Parameter applicability

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
FC-PH Version - obsolete	0	31-16	2020h ^g	n	n	n	n	n	n
Buffer-to-buffer Credit	0	15-0	0 or 1 ^d	y	y	y	y	y	y
Common Features	1	31-16							
Continuously increasing relative offset	1	31	0	y	y	n	n	n	n
Clean Address	1	31	0	n	n	n	n	y	y
Multiple N_Port_ID Support	1	31	0	n	n	y	y	n	n
Random relative offset	1	30	0	y	y	n	n	n	n
Virtual Fabrics bit	1	30	0	n	n	y	y	y	y
<p>Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning); “v” indicate the definition is vendor specific.</p> <p>^a E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology or while doing PLOGI with an NL_Port on the same loop. ^b The Common Service Parameter applicability is specified in FC-SP-2. ^c This field shall be set to 00h. ^d Default buffer-to-buffer credit = 1 for all ports but an L_Port, and Buffer-to-buffer credit=0 for an L_Port. ^e N_Port/F_Port=0 for an N_Port, and N_Port/F_Port=1 for an F_Port. ^f BB_Credit Management=0 for an N_Port or F_Port, BB_Credit_Management=1 for an L_Port ^g Legacy implementations may check this field. Implementations that are compliant with this standard shall not check this field.</p>									

Table 207 – Common Service Parameter applicability(continued)

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
Valid Vendor Version Level	1	29	0	y	y	y	y	n	n
Multiple N_Port_ID Assignment	1	29	0	n	n	n	n	y	y
N_Port/F_Port	1	28	0 or 1 ^e	y	y	y	y	y	y
BB_Credit Management	1	27	0 or 1 ^f	y	y	y	y	n	n
Name Server Session Started	1	27	0	n	n	n	n	y	y
Name Server Session Begin	1	26	0	n	n	y	y	n	n
E_D_TOV Resolution	1	26	0	y ^a	y ^a	n	n	y	y
Energy Efficient LPI Mode Supported	1	25	0	y	y	y	y	y	y
Application_Header Support	1	24	0	y	y	n	n	n	n
Broadcast supported by Fabric	1	24	0	n	n	n	n	y	y
Priority Tagging Supported	1	23	0	y	y	y	y	y	y
Query Data Buffer conditions	1	22	0	y	y	y	y	y	y
<p>Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning); “v” indicate the definition is vendor specific.</p> <p>^a E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology or while doing PLOGI with an NL_Port on the same loop. ^b The Common Service Parameter applicability is specified in FC-SP-2. ^c This field shall be set to 00h. ^d Default buffer-to-buffer credit = 1 for all ports but an L_Port, and Buffer-to-buffer credit=0 for an L_Port. ^e N_Port/F_Port=0 for an N_Port, and N_Port/F_Port=1 for an F_Port. ^f BB_Credit Management=0 for an N_Port or F_Port, BB_Credit_Management=1 for an L_Port ^g Legacy implementations may check this field. Implementations that are compliant with this standard shall not check this field.</p>									

Table 207 – Common Service Parameter applicability(continued)

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
Security bit (see FC-SP-2)	1	21	0	_b	_b	_b	_b	_b	_b
Clock Synchronization Primitive Capable	1	20	0	y	y	y	y	y	y
R_T_TOV Value	1	19	0	y	y	y	y	y	y
Dynamic Half Duplex Supported	1	18	0	y	y	y	y	y	y
SEQ_CNT/Vendor Specific	1	17	0	y	y	v	v	v	v
Payload Bit	1	16	0	y	y	y	y	y	y
BB_SC_N	1	15-12	0	y	y	y	y	y	y
Buffer-to-Buffer Receive Data_Field Size	1	11-0	256	y	y	y	y	y	y
Set to 00h ^c	2	31-24	0	-	-	-	-	-	-
Nx_Port Total Concurrent Sequences	2	23-16	1	y	y	n	n	n	n
Relative offset by Info Category	2	15-0	0	y	y	n	n	n	n
<p>Legend: "y" indicates yes, applicable (i.e., has meaning); "n" indicates no, not applicable (i.e., has no meaning); "v" indicate the definition is vendor specific.</p> <p>^a E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology or while doing PLOGI with an NL_Port on the same loop. ^b The Common Service Parameter applicability is specified in FC-SP-2. ^c This field shall be set to 00h. ^d Default buffer-to-buffer credit = 1 for all ports but an L_Port, and Buffer-to-buffer credit=0 for an L_Port. ^e N_Port/F_Port=0 for an N_Port, and N_Port/F_Port=1 for an F_Port. ^f BB_Credit Management=0 for an N_Port or F_Port, BB_Credit Management=1 for an L_Port ^g Legacy implementations may check this field. Implementations that are compliant with this standard shall not check this field.</p>									

Table 207 – Common Service Parameter applicability(continued)

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
R_A_TOV	2	31-0	10 000	n	n	n	n	y	y
E_D_TOV Value	3	31-0	2 000	y ^a	y ^a	n	n	y	y
<p>Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning); “v” indicate the definition is vendor specific.</p> <p>^a E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology or while doing PLOGI with an NL_Port on the same loop. ^b The Common Service Parameter applicability is specified in FC-SP-2. ^c This field shall be set to 00h. ^d Default buffer-to-buffer credit = 1 for all ports but an L_Port, and Buffer-to-buffer credit=0 for an L_Port. ^e N_Port/F_Port=0 for an N_Port, and N_Port/F_Port=1 for an F_Port. ^f BB_Credit Management=0 for an N_Port or F_Port, BB_Credit_Management=1 for an L_Port ^g Legacy implementations may check this field. Implementations that are compliant with this standard shall not check this field.</p>									

6.6.2.2 Payload

The Common Service Parameters Payload for FLOGI is shown in table 208.

Table 208 – Common Service Parameters - FLOGI

Bits Word	31 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 00
0	FC-PH Version - obsolete		Buffer-to-buffer Credit		
1	Common Features (see table 207)		BB_SC_N	Buffer-to-Buffer Receive Data_Field Size	
2	Reserved		Reserved		
3	Reserved		Reserved		

The Common Service Parameters Payload for PLOGI and PLOGI LS_ACC is shown in table 209.

Table 209 – Common Service Parameters - PLOGI and PLOGI LS_ACC

Bits Word	31 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 00
0	FC-PH Version - obsolete		Buffer-to-buffer Credit		
1	Common Features (see table 207)		BB_SC_N	Buffer-to-Buffer Receive Data_Field size	
2	00h	Total Concurrent Sequences	Relative offset by Information Category		
3	E_D_TOV				

The Common Service Parameters Payload for FLOGI LS_ACC is shown in table 210.

Table 210 – Common Service Parameters - FLOGI LS_ACC

Bits Word	31 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 00
0	FC-PH Version - obsolete		Buffer-to-buffer Credit (Fx_Port)		
1	Common Features (see table 207)		BB_SC_N	Buffer-to-Buffer Receive Data_Field size	
2	R_A_TOV				
3	E_D_TOV				

6.6.2.3 Buffer-to-buffer Credit

The buffer-to-buffer Credit field (word 0, bits 15-0) defines the number of buffers available for holding Class 2, or Class 3 frames received. An FC_Port tracks Buffer-to-buffer Credit as a single entity for all frames subject to buffer-to-buffer flow control (see FC-FS-6). Values in the Buffer-to-buffer Credit field are 1 to 32 767. The value 0 is reserved.

For N_Port Login, this field shall only be meaningful for an Nx_Port in a point-to-point topology and between two NL_Ports on the same loop.

6.6.2.4 Common Features

6.6.2.4.1 Continuously increasing relative offset

- 0 = not supported
- 1 = supported

If the continuously increasing relative offset bit (word 1, bit 31) is set to one, the Nx_Port supplying this parameter shall be capable of supporting continuously increasing relative offset, if present (F_CTL bit 3), within a Sequence on a frame by frame SEQ_CNT basis. This bit shall only be applicable to those Information Categories in which an Nx_Port supports relative offset (i.e., word 2, bits 15-0). See FC-FS-6 for the use of continuously increasing relative offset.

This bit shall be applicable to a Sequence Initiator in addition to a Sequence Recipient for all Classes of Service supported by the Nx_Port.

6.6.2.4.2 Clean Address

0 = No information
 1 = Clean Address

The Clean Address bit (word 1, bit 31) provides an indication to an Nx_Port as to whether the address it was assigned by the Fabric had been previously used by another device within R_A_TOV. If this bit is set to zero, the assigned address may or may not have been used by a previous device within R_A_TOV. If this bit is set to one, the assigned address has not been used by any other device within R_A_TOV, or has been assigned to the current device for a previous FLOGI and not been changed within R_A_TOV. This bit is only meaningful in the FLOGI LS_ACC, it is not meaningful in the FLOGI request.

6.6.2.4.3 Multiple N_Port_ID Support

0 = not supported
 1 = supported

The Multiple N_Port_ID Support bit (word 1, bit 31) shall be set to one to indicate that the PN_Port supplying this parameter is capable of requesting multiple N_Port_IDs using the FDISC ELS. The N_Port_ID Support bit shall be set to zero to indicate that the PN_Port supplying this parameter is not capable of requesting additional N_Port_IDs. This bit is only meaningful in the FLOGI request, it is not meaningful in the FLOGI LS_ACC.

6.6.2.4.4 Random relative offset

0 = not supported
 1 = supported

The random relative offset bit (word 1, bit 30) indicates that the Nx_Port supplying this parameter shall be capable of supporting random relative offset values, if present (F_CTL bit 3). Random values may increase, decrease, or otherwise fluctuate within a Sequence. This bit shall only be applicable to those Information Categories in which an Nx_Port supports relative offset (i.e., word 3, bits 15-0). See FC-FS-6 for the use of random relative offset.

This bit shall be applicable to a Sequence Initiator in addition to a Sequence Recipient for all Classes of Service supported by the Nx_Port.

6.6.2.4.5 Virtual Fabrics bit

0 = not supported
 1 = supported

The Virtual Fabrics bit (word 1, bit 30) indicates support for Virtual fabrics (see clause 8).

6.6.2.4.6 Valid Vendor Version Level

0 = not valid
 1 = Valid

In PLOGI, PLOGI LS_ACC, and FLOGI, if the Valid Vendor Version Level bit (word 1, bit 29) is set to one, the Vendor Version Level (words 25 through 28 in table 206) contains valid information. If it is set to zero, the Vendor Version Level field is not meaningful.

6.6.2.4.7 Multiple N_Port_ID Assignment

0 = not supported
1 = supported

If the Multiple N_Port_ID Support bit (word 1, bit 31) in the FLOGI request is one, the Multiple N_Port_ID Assignment bit (word 1, bit 29) shall be set to one if the F_Port supplying this parameter is capable of assigning multiple N_Port_IDs to the attached PN_Port using the FDISC ELS. The Multiple N_Port_ID Assignment bit shall be set to zero if the Multiple N_Port_ID Support bit in the FLOGI request is zero or to indicate that the F_Port is not capable of assigning multiple N_Port_IDs to the attached PN_Port if the Multiple N_Port_ID Support bit in the FLOGI request is one. This bit is only meaningful in the FLOGI LS_ACC, it is not meaningful in the FLOGI request.

NOTE 15 – The definition above has been modified from previous revisions of the standard.

6.6.2.4.8 N_Port/F_Port

0 = Nx_Port
1 = Fx_Port

An Nx_Port shall set the N_Port/F_Port bit (word 1, bit 28) to zero for PLOGI, PLOGI LS_ACC and FLOGI. If an Nx_Port is connected in a Fabric topology, the Fx_Port shall set the N_Port/F_Port bit to one in the FLOGI LS_ACC. If an Nx_Port is connected in a point-to-point topology, the Nx_Port shall set the N_Port/F_Port bit to zero in the FLOGI LS_ACC.

6.6.2.4.9 BB_Credit Management

0 = BB_Credit management specified in FC-FS-6 shall be used.
1 = Alternate BB_Credit management specified in FC-AL-2 shall be used.

The BB_Credit Management bit (word 1, bit 27) specifies the type of BB_Credit Management to be used.

6.6.2.4.10 Name Server Session Begin

0 = not requested
1 = requested

The Name Server Session Begin bit (i.e., word 1 bit 26) provides an indication whether a Server Session (see FC-GS-9) is requested to be established with the Name Server by emulating the acceptance of a Server Session Begin (i.e., SSB) CT_IU as part of processing the FLOGI.

The Name Server Session Started bit (see 6.6.2.4.11) provides an indication whether a Server Session has been established with the Name Server.

6.6.2.4.11 Name Server Session Started

0 = not established
1 = established

The Name Server Session Started bit (i.e., word 1 bit 27) provides an indication whether a Server Session (see FC-GS-9) has been established with the Name Server as described in 6.6.2.4.10. If the Name Server Session Begin bit (see 6.6.2.4.10) is set to zero in the FLOGI request, then the Name Server Session Started bit shall be set to zero in the corresponding FLOGI LS_ACC.

6.6.2.4.12 E_D_TOV Resolution

0 = 1 millisecond
 1 = 1 nanosecond

The E_D_TOV resolution bit (word 1, bit 26) indicates the resolution of the E_D_TOV timer. If the bit is set to zero, the timer shall be in increments of 1 millisecond. If the bit is set to one, the timer shall be in increments of 1 nanosecond. See FC-FS-6 for the definition of E_D_TOV.

6.6.2.4.13 Application_Header Support

0 = no support
 1 = supported

The Application Header Support bit (word 1, bit 24) indicates support for receiving an Application_Header (see FC-FS-6). If the Application Header Support bit is set to zero, then the FC_Port does not support the reception of an Application_Header. If the Application Header Support bit is set to one, then the FC_Port supports the reception of an Application_Header.

6.6.2.4.14 Broadcast Supported by Fabric

0 = Broadcast not supported by the Fabric
 1 = Broadcast supported by the Fabric

The Broadcast bit (word 1, bit 24) indicates whether broadcast is supported by the Fabric or not. It is only meaningful in the FLOGI LS_ACC. For broadcast requirements see FC-FS-6.

6.6.2.4.15 Priority Tagging Supported

0 = Priority Tagging not supported
 1 = Priority Tagging supported

The Priority Tagging Supported bit (word 1, bit 23) indicates whether Priority Tagging (see 9) is supported by the Fabric or by a VN_Port. The Priority Tagging Supported bit may be set to one only if the Priority bit (word 0, bit 23) of the Class Service Parameters is set to one.

N_Port and Fabric should indicate at FLOGI time support for Priority Tagging via the Priority Tagging Supported bit. N_Ports should indicate at PLOGI time support for Priority Tagging via the Priority Tagging Supported bit. N_Ports that support Priority Tagging shall preserve the value of the Priority field across all frames of all Sequences of an Exchange (i.e., the Exchange Originator shall use the same value in the Priority field of the frames of an Exchange it originates and the Exchange Responder shall use the same value as set by the Exchange Originator in the Priority field of the frames that it originates for that Exchange).

6.6.2.4.16 Query Buffer Conditions

0 = No buffer conditions to report.
 1 = Buffer conditions to report.

The Query Buffer Conditions bit (word 1, bit 22) indicates whether or not the FC_Port has unusual buffer conditions to report. A Report Port Buffer Conditions (RPBC) ELS may be issued to gather the buffer conditions. An FC_Port shall only set the Query Buffer Conditions to 1 if the FC_Port supports the RPBC ELS, and any of the following conditions are true:

- a) The ELS Receive Data_Field Size field is different that the Buffer-to-Buffer Receive Data_Field Size field in the common service parameters; or
- b) multi-frame ELSs are not supported.

6.6.2.4.17 Clock Synchronization Primitive Capable

0 = not Clock Synchronization Primitive Capable
 1 = Clock Synchronization Primitive Capable

The Clock Synchronization Primitive Capable bit (word 1, bit 20) indicates support for the Primitive method of Clock Synchronization (see FC-FS-6). If the bit is set to zero, the Primitive method of Clock Synchronization is not supported. If the bit is set to one, the meaning is as defined in table 211.

Table 211 – Clock Synchronization Applicability

Type	Meaning if bit set to one
Clock Synchronization Server (N_Port Login only)	The Clock Synchronization Server is capable of generating clock synchronization Primitive Signals (see FC-FS-6).
Other FC_Ports (N_Port Login only)	The FC_Port is capable of receiving the clock synchronization Primitive Signals (see FC-FS-6) and acting upon them.
Fabric (FLOGI LS_ACC only)	The Fabric is capable of receiving the clock synchronization Primitive Signals (see FC-FS-6) and acting upon them.

6.6.2.4.18 R_T_TOV value

0 = Default value of 100 milliseconds.
 1 = Short value of 100 microseconds.

The R_T_TOV value (word 1, bit 19) indicates support for the short value of R_T_TOV. If this bit is set to zero, the default value of 100 milliseconds is specified. If it is set to one, the value of 100 microseconds is specified.

6.6.2.4.19 Dynamic Half Duplex Supported

0 = DHD not supported
 1 = DHD supported

The Dynamic Half Duplex Supported bit (word 1, bit 18) indicates support for Dynamic Half Duplex. If it is set to zero, Dynamic Half Duplex is not supported. If it is set to one, the FC_Port supports the reception of the DHD primitive.

NOTE 16 – DHD is applicable to FC-AL topologies. See FC-AL-2 for a description of DHD behavior.

6.6.2.4.20 SEQ_CNT

Word 1, bit 17 has the following meaning in a PLOGI request and PLOGI response:

0 = Normal rules apply to SEQ_CNT
 1 = Continuously Increasing SEQ_CNT shall be used.

SEQ_CNT (word 1, bit 17) indicates the requirement on SEQ_CNT. If the bit is set to zero, normal rules shall apply to SEQ_CNT usage. If the bit is set to one, the Nx_Port is guaranteeing that it shall

transmit all frames within an Exchange using a continuously increasing SEQ_CNT. Each Exchange shall start with SEQ_CNT set to zero in the first frame, and every frame transmitted after that shall increment the previous SEQ_CNT by one, even across transfers of Sequence Initiative. Any frames received from the other Nx_Port in the Exchange shall have no effect on the transmitted SEQ_CNT (see FC-FS-6).

The definition of word 1, bit 17 is vendor specific in a FLOGI request and FLOGI response.

6.6.2.4.21 Payload Bit

0 = Payload length is 116 bytes

1 = Payload length in bytes is 256 plus four times the value of the Login Extension Length field

The Payload Bit (word 1 bit 16) indicates the length of the FLOGI or PLOGI Payload. If it is set to zero, the Payload length shall be 116 bytes. If it is set to one, the Payload length in bytes shall be 256 plus four times the value of the Login Extension Length field.

If Payload Bit is 0 in PLOGI or FLOGI, then LS_ACC sent in response shall have the Payload Bit set to 0.

6.6.2.4.22 BB_SC_N

The Buffer-to-buffer State Change Number (BB_SC_N) field (word 1, bits 15-12) specifies the Buffer-to-buffer State Change Number. It indicates that the sender of the PLOGI or FLOGI frame is requesting $2^{BB_SC_N}$ number of frames to be sent between two consecutive BB_SCs primitives, and $2^{BB_SC_N}$ number of R_RDY primitives to be sent between two consecutive BB_SCr primitives. See FC-FS-6 for a description of the BB_Credit recovery process.

6.6.2.4.23 Buffer-to-Buffer Receive Data_Field Size

The Buffer-to-Buffer Receive Data_Field Size field (word 1, bits 11-0) specifies the largest frame Data_Field size that may be received by the Nx_Port supplying the Service Parameters as a Sequence Recipient for:

- a) a Class 2 Data frame; or
- b) a Class 3 Data frame.

The value shall be a multiple of four bytes. Values less than 256 or greater than 2 112 are invalid. An Fx_Port shall support a Data_Field size of at least 256 bytes.

6.6.2.4.24 Energy Efficient LPI Mode Supported

0b = Nx_Port does not support Energy Efficient LPI Mode Operation (see FC-FS-6)

1b = Nx_Port supports Energy Efficient LPI Mode Operation (see FC-FS-6)

The Energy Efficient Operation Support bits (Word 1, Bit 25), indicates support for Energy Efficient LPI Mode Operation (see FC-FS-6).

6.6.2.4.25 Total Concurrent Sequences

Total Concurrent Sequences field (word 2, bits 23 - 16) specifies the total number of Concurrent Sequences for all classes that the Nx_Port is capable of supporting as a Recipient.

The Total Concurrent Sequences specified by an Nx_Port shall be less than or equal to the sum of the Concurrent Sequences supported on a Class by Class basis (e.g., an Nx_Port may specify that it is capable of supporting ten Concurrent Sequences in Class 2 and ten Concurrent Sequences in Class 3. However, the total number of Concurrent Sequences while both Class 2 and 3 are open may be fifteen).

6.6.2.4.26 Relative offset by category

The relative offset by category field (word 2, bits 15 - 0) shall indicate on a bit-position basis, whether or not relative offset shall be supported for the corresponding Information Category (e.g., if bit 14 = 1 and bit 2 = 1 and the others are set to zero, Information Category 1110b and 0010b frames shall be capable of using relative offset as a Sequence Recipient or a Sequence Initiator). See FC-FS-6 for definition of the Information Category field.

6.6.2.4.27 R_A_TOV

The R_A_TOV value shall be specified as a count of 1 ms increments. Therefore, a value of 0000000Ah specifies a time period of 10 milliseconds.

6.6.2.4.28 E_D_TOV

If the E_D_TOV Resolution bit (word 1, bit 26) is set to zero, the E_D_TOV value shall be specified as a count of 1 millisecond increments. If the E_D_TOV Resolution bit is set to one, the E_D_TOV value shall be specified as a count of 1 nanosecond increments (e.g., based on the setting of the E_D_TOV Resolution bit, a value of 0000000Ah specifies a time period of either 10 milliseconds or 10 nanoseconds).

For PLOGI, the E_D_TOV value in the LS_ACC to the PLOGI shall be greater than or equal to the value in the PLOGI. The E_D_TOV value in the LS_ACC shall be the value used by each Nx_Port. See FC-FS-6 for definition of E_D_TOV.

6.6.3 Port_Name

The Port_Name is an eight-byte field (words 5-6) that identifies an Nx_Port. Each Nx_Port, including Nx_Ports that have Well-known addresses, shall provide a Name_Identifier. Nx_Ports that are not assigned to Well-known addresses shall provide a Name_Identifier that is unique within the Fibre Channel interaction space of the Nx_Port. Bits 63-60 specify the format of the Name_Identifier. The formats are defined in FC-FS-6.

6.6.4 Node_Switch_Fabric_Name

The Node_Switch_Fabric_Name is an eight-byte field (words 7-8) that labels a Node, Fabric, or locally-attached Switch for identification purposes, such as diagnostics. Each name shall be unique within the Fibre Channel interaction space. Bits 63-60 specify the format of the name. The formats are defined in FC-FS-6.

In PLOGI, PLOGI LS_ACC, and FLOGI link service commands the Node_Switch_Fabric_Name provides the name identifier of the node (i.e., Node_Name) which the port performing the login is associated with. In a FLOGI LS_ACC link service command the Node_Switch_Fabric_Name identifies either the name of the Fabric (i.e., Fabric_Name) or the locally-attached Switch (i.e., Switch_Name, see FC-SW-8) with which the port has successfully logged in.

6.6.5 Class Service Parameters

6.6.5.1 Applicability

Table 212 defines the applicability, by class as well as by PLOGI, FLOGI, PLOGI LS_ACC and FLOGI LS_ACC, of the Class Service Parameters to N_Port Login and Fabric Login. The Class 2 Service Parameters are given in words 13 - 16. The Class 3 Service Parameters are given in words 17 - 20. The words given in the second column and in the following subclauses are relative to the start of the specific class service parameters field (see table 212). The Default Login Value column (see table 212) refers to the Login values to be used prior to a successful Login.

Table 212 – Class Service Parameters Applicability

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
Class Validity	0	31	0	y	y	y	y	y	y
Service Options	0	30-16							
Obsolete	0	30							
Obsolete	0	29-28							
Obsolete	0	27	0	n	n	y ^a	y ^a	y ^b	y ^b
Simplex dedicated connection - obsolete	0	26	0	n	n	n	n	n	n
Camp-On - obsolete	0	25	0	n	n	n	n	n	n
Buffered Class 1 - obsolete	0	24	0	n	n	n	n	n	n
Priority	0	23	0	y	y	y	y	y	y
Preference	0	22	0	y	y	y	y	y	y
DiffServ QoS	0	21	0	y	y	y	y	y	y
Reserved	0	20-16	0	n	n	n	n	n	n
Initiator Control	0	15-0							
X_ID Reassignment - obsolete	0	15-14	0	n	n	n	n	n	n
Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning)									
^a This bit has no meaning. ^b This bit shall be set to one.									

Table 212 – Class Service Parameters Applicability (continued)

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
Obsolete	0	13-12							
ACK_0 capable	0	11	0	y	n	n	n	n	n
ACK_N Capable - obsolete	0	10	0	n	n	n	n	n	n
ACK generation assistance	0	9	0	y	n	n	n	n	n
Data compression capable - obsolete	0	8	0	n	n	n	n	n	n
Data compression history buffer size - obsolete	0	7-6	0	n	n	n	n	n	n
Data Encryption Capable - obsolete	0	5	0	n	n	n	n	n	n
Clock Synchronization ELS capable	0	4	0	y	y	y	y	y	y
Reserved	0	3-0	0	n	n	n	n	n	n
Recipient Control	1	31-16							
ACK_0 Capable	1	31	0	y	n	n	n	n	n
ACK_N Capable - obsolete	1	30	0	n	n	n	n	n	n
X_ID interlock	1	29	1	y	n	n	n	n	n
Error policy support	1	28-27	0	y	y	n	n	n	n
Reserved	1	26	0	n	n	n	n	n	n
Categories per Sequence	1	25-24	1	y	y	n	n	n	n
Data compression capable - obsolete	1	23	0	n	n	n	n	n	n
Data compression history buffer size - obsolete	1	22-21	0	n	n	n	n	n	n
Data decryption capable – obsolete	1	20	0	n	n	n	n	n	n
Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning)									
^a This bit has no meaning. ^b This bit shall be set to one.									

Table 212 – Class Service Parameters Applicability (continued)

Service Parameter	Word	Bits	Default Login Value	PLOGI and PLOGI LS_ACC Parameter applicability		FLOGI Parameter applicability		FLOGI LS_ACC Parameter applicability	
				Class		Class		Class	
				2	3	2	3	2	3
Clock Synchronization ELS capable	1	19	0	y	y	y	y	y	y
SEQ_ID unique per Exchange	1	18	0	y	y	n	n	n	n
Reserved	1	17-16	0	n	n	n	n	n	n
Reserved	1	15-12	0	n	n	n	n	n	n
Receive Data_Field Size	1	11-0	256	y	y	n	n	n	n
Reserved	2	31-24	0	n	n	n	n	n	n
Concurrent Sequences	2	23-16	1	y	y	n	n	n	n
Nx_Port end-to-end Credit	2	15-0	1	y	n	n	n	n	n
Reserved	3	31-24	0	n	n	n	n	n	n
Open Sequences per Exchange	3	23-16	1	y	y	n	n	n	n
Reserved	3	15-0	0	n	n	n	n	n	n
Obsolete	3	31-0							
Legend: “y” indicates yes, applicable (i.e., has meaning); “n” indicates no, not applicable (i.e., has no meaning)									
^a This bit has no meaning. ^b This bit shall be set to one.									

6.6.5.2 Payload

The Payload Class Service Parameters using FLOGI is shown in table 213.

Table 213 – Class Service Parameters - FLOGI

Bits Word	31 .. 24	23 .. 16	15 .. 12	11 .. 08	07 .. 00
0	Service Options		Initiator Control		
1	Recipient Control		Reserved		
2	Reserved	Concurrent Sequences	Nx_Port End-to-end Credit		
3	Reserved	Open Sequences per Exchange	Reserved		

The Payload Service Parameters using PLOGI and PLOGI LS_ACC is shown in table 214.

Table 214 – Class Service Parameters - PLOGI and PLOGI LS_ACC

Bits Word	31 .. 24	23 .. 16	15 .. 00
0	Service Options		Initiator Control
1	Recipient Control		Reserved Receive Data_Field Size
2	Reserved	Concurrent Sequences	Nx_Port End-to-end Credit
3	Reserved	Open Sequences per Exchange	Reserved

The Payload Parameters using FLOGI LS_ACC is shown in table 215.

Table 215 – Class Service Parameters - FLOGI LS_ACC

Bits Word	31 .. 16	15 .. 00
0	Service Options	Reserved
1	Recipient Control	Reserved
2	Reserved	Reserved
3	Obsolete	

6.6.5.3 Class validity

- 0 = Invalid - Class not supported
- 1 = Valid - Class supported

The Class validity bit (word 0, bit 31) shall indicate whether this Class is supported or not. If the Class validity bit is set to zero, this set of sixteen bytes shall be ignored. If the Class validity bit is one, this Class shall be supported.

The Class 2 Service Parameters are given in words 13 - 16. The Class 3 Service Parameters are given in words 17 - 20.

6.6.5.4 Service options

6.6.5.4.1 Introduction

The service options shall specify optional features of a class of service supported by the port supplying the service parameters.

6.6.5.4.2 Priority

6.6.5.4.2.1 Nx_Port

- 0 = non-zero Priority may be tolerated
- 1 = non-zero Priority shall be tolerated

The Priority bit (word 0, bit 23) has meaning in Class 2, and 3.

If an Nx_Port performs Login with another Nx_Port, it sets the Priority bit (word 0, bit 23) to one within the Class of Service to indicate tolerance for non-zero values in the Priority field (see FC-FS-6) in the frame header if CS_CTL/Priority Enable bit (F_CTL bit 17) is set to one. The other Nx_Port indicates tolerance for non-zero values in the Priority field by setting this bit to one in the LS_ACC. An Nx_Port that indicates tolerance of non-zero values in the Priority field shall not reject or otherwise deprecate a frame solely because the Priority field is non-zero.

NOTE 17 – Even if an Nx_Port never intends to set the Priority Field to any value other than zero, the Nx_Port may still have reason to set the Login Priority bit to one.

An Nx_Port should not set a non-zero value in the Priority field in a frame sent within a class of service if Login Priority is not indicated by both the destination Nx_Port and the Fabric. If an Nx_Port sets a non-zero value in the Priority field in a frame sent within a class of service and Login Priority is not indicated by both the destination Nx_Port and the Fabric, then the results are unpredictable.

6.6.5.4.2.2 Fx_Port

- 0 = Priority is not supported
- 1 = Priority is supported

The Priority bit (word 0, bit 23) has meaning in Class 2, and 3.

If an Nx_Port performs Login with a Fabric, it requests support for use of the Priority field (See FC-FS-6) by setting the Priority bit (word 0, bit 23) to one. If Priority is set to one in the Login and the LS_ACC, then both the Nx_Port and Fabric have agreed that Priority is available for use.

The set of values specified in table 216 give the meaning of the combination of the Priority bit.

Table 216 – Fabric Login Priority and Preemption Support

Nx_Port	F_Port Controller	Meaning
0	0	Neither supports Priority
0	1	Fabric is capable of supporting Priority
1	0	Nx_Port support requested, Fabric does not support Priority
1	1	Nx_Port requested, Fabric is capable of supporting Priority, available for use

6.6.5.4.3 Preference

6.6.5.4.3.1 Nx_Port

- 0 = non-zero CS_CTL may be tolerated
- 1 = non-zero CS_CTL shall be tolerated

If an Nx_Port performs Login with another Nx_Port, it shall indicate tolerance for non-zero CS_CTL within the Class of Service by setting the Preference bit (word 0, bit 22) to one. The other Nx_Port indicates tolerance for non-zero CS_CTL by setting this bit to one in the LS_ACC. An Nx_Port that

tolerates a non-zero CS_CTL shall not reject or otherwise deprecate a frame solely because the CS_CTL field is non-zero.

NOTE 18 – Even if an Nx_Port never intends to set the PREF bit to any value other than zero, the Nx_Port may still have reason to set the Login Preference bit to one. Setting the bit to one indicates to the other Nx_Port that the Nx_Port shall accept frames with a non-zero CS_CTL field value.

In Class 2 and 3, if this bit is set to one, the Nx_Port shall tolerate the PREF field in the CS_CTL field of the Frame_Header. Tolerance for CS_CTL as a Sequence Initiator means that the PREF field may specify Preference to the Fabric. Tolerance for CS_CTL as a Sequence Recipient means that the Nx_Port shall ignore the PREF field (see FC-FS-6).

6.6.5.4.3.2 Fx_Port

0 = Normal delivery
 1 = Preferred delivery functional

If the Preference bit (word 0, bit 22) is set to one by an Nx_Port, then it is requested that all frames transmitted by the Nx_Port requesting this function be delivered according to the setting of the PREF field in the CS_CTL field of the Frame_Header. If this bit is set to one by the F_Port Controller, the F_Port Controller is indicating that it shall deliver Class 2 and 3 frames transmitted by the requesting Nx_Port according to the setting of the PREF field.

NOTE 19 – An F_Port Controller that responds with bit 22 set to zero may not itself support Preferred delivery, but other Fabric Elements in the path to the destination may support it. An Nx_Port may attempt Preferred delivery even if the F_Port Controller does not indicate support.

If this bit is set to one, the Fabric shall deliver both Data and Link_Control (class 2 only) frames according to the setting of the PREF field in the CS_CTL field of the Frame_Header (see FC-FS-6).

Table 217 summarizes the function of the PREF bit for both Class 2 and Class 3.

Table 217 – Class 2 and 3 Preference Bit Function

Nx_Port Word 0, Bit 22	F_Port Controller Word 0, Bit 22	Meaning
0	0	Fabric may be capable of providing Preferred delivery
0	1	Fabric is capable of providing Preferred delivery
1	0	Nx_Port support requested, Fabric may not support Preferred delivery
1	1	Nx_Port requested, Fabric is capable, Preferred delivery is available for use

6.6.5.4.4 DiffServ QoS

6.6.5.4.4.1 N_Port Login

0 = DiffServ QoS not supported
 1 = DiffServ QoS supported

If an Nx_Port performs Login with another Nx_Port, it shall indicate support for Differentiated Services QoS by setting the DiffServ QoS bit (word 0, bit 21) to one. The other Nx_Port indicates support for Differentiated Services QoS by setting this bit to one in the LS_ACC. Support of Differentiated Services QoS as an Exchange Originator means that the CS_CTL/Priority Enable bit (F_CTL bit 17) is set to zero to indicate that the DSCP field of the CS_CTL field in the Frame_Header (word 1, bits 29-24) specifies the Differentiated Services QoS policy. Support of Differentiated Services QoS as an Exchange Responder means that the CS_CTL/Priority Enable bit (F_CTL bit 17) is set to zero, indicates that the DSCP field of the CS_CTL field in the Frame_Header (word 1, bits 29-24) specifies the Differentiated Services QoS policy (see FC-FS-6).

The DiffServ QoS bit has meaning for both Classes 2 and 3.

6.6.5.4.4.2 Fabric Login

- 0 = DiffServ QoS not Supported
- 1 = DiffServ QoS Supported

While doing Fabric Login, an Nx_Port shall request support for Differentiated Services QoS by setting the DiffServ QoS bit (word 0, bit 21) to one. If the LS_ACC reply from the F_Port Controller has this bit set to one, both the Nx_Port and F_Port Controller have agreed that Differentiated Services QoS is available for use.

The set of values specified in table 218 give the meaning of the combination of Word 0, bit 21 between the requesting Nx_Port and the responding F_Port Controller (see FC-FS-6).

The DiffServ QoS bit has meaning for both Classes 2 and 3.

Table 218 – DiffServ QoS bit definition

Nx_Port	F_Port Controller	Description
0	0	Neither supports Differentiated Services QoS
0	1	Fx_Port is capable of supporting Differentiated Services QoS
1	0	Nx_Port support requested, Fx_Port does not support Differentiated Services QoS
1	1	Nx_Port requested, Fx_Port is capable of supporting Differentiated Services QoS, Differentiated Services QoS is available for use

6.6.5.5 Initiator control

6.6.5.5.1 Introduction

The Initiator Control Flags shall specify which protocols, policies or functions the Sequence Initiator function in the Nx_Port supplying the Service Parameters requests of the recipient or is capable of as a Sequence initiator.

6.6.5.5.2 ACK_0 capability

- 0 = ACK_0 incapable
- 1 = ACK_0 capable

The ACK_0 capability bit (word 0, bit 11) specifies if the Nx_Port supplying these Class Service Parameters is capable of support for ACK_0 as a Sequence Initiator for acknowledgement of an entire Sequence in either Discard or Process Exchange Error Policies. As a Sequence Initiator an Nx_Port receives and processes ACK frames in response to Data frame transmission. ACK_0 support is applicable to acknowledged class of service Sequences (see FC-FC-2).

The conditions under which ACK_0 is supported are defined in table 219 and described in the following text.

Table 219 – ACK_0 Support Conditions (Initiator Control)

Nx_Port A Word 0, Bit 11	Nx_Port B Word 1, Bit 31	Nx_Port A as Sequence Initiator
0	0	ACK_0 not supported
0	1	ACK_0 not supported
1	0	ACK_0 not supported
1	1	ACK_0 supported

If one Nx_Port (e.g., Nx_Port A) is capable of receiving ACK_0 as a Sequence Initiator (word 0, bit 11 set to one) and the other Nx_Port (e.g., Nx_Port B) is capable of transmitting ACK_0 as a Sequence Recipient (word 1, bit 31 set to one), ACK_0 is supported while Nx_Port A is the Sequence Initiator and Nx_Port B is the Sequence Recipient. Otherwise, ACK_0 shall not be supported while Nx_Port A is the Sequence Initiator and Nx_Port B is the Sequence Recipient. ACK_0 usage shall take precedence over ACK_1.

ACK_0 capability may be asymmetrical for a single Nx_Port (i.e., an Nx_Port may be capable processing ACK_0 as a Sequence Initiator, but not be capable of ACK_0 transmission as a Sequence Recipient). Similarly, an Nx_Port may be capable of generating ACK_0 as a Sequence Recipient, but not be capable of ACK_0 reception as a Sequence Initiator.

6.6.5.5.3 ACK generation assistance

- 0 = No ACK generation assistance is provided to Sequence Recipient.
- 1 = ACK generation assistance is provided to Sequence Recipient.

Usage of the ACK generation assistance bit (word 0, bit 9) is specified in FC-FS-6.

6.6.5.5.4 Clock synchronization ELS capable

- 0 = Initiator does not have clock synchronization ELS capability
- 1 = Initiator has clock synchronization ELS capability

The Clock synchronization ELS capable bit (word 0, bit 4) is only meaningful from the Clock Synchronization Server well-known address (i.e., FFFFF6h). This bit indicates support for the ELS method of Clock Synchronization. If this bit is set to zero, the Nx_Port does not support the ELS method of Clock Synchronization. If this bit is set to one, the Nx_Port is capable of generating the CSU ELS frames. See FC-FS-6.

6.6.5.6 Recipient control

6.6.5.6.1 Introduction

The Recipient Control Flags shall specify which protocols, policies or functions are supported by the Recipient Initiator function in the Nx_Port supplying the Service Parameters while acting as a recipient of Data frames.

6.6.5.6.2 ACK_0 capability

- 0 = ACK_0 incapable
- 1 = ACK_0 capable

The ACK_0 capability bit (word 1, bit 31) specifies that the Nx_Port supplying these Class Service Parameters may or may not be capable of support for ACK_0 as a Sequence Recipient for acknowledgement of an entire Sequence in either Discard or Process Exchange Error Policies. As a Sequence Recipient an Nx_Port shall support infinite buffering and be capable of transmitting ACK_0 frames in response to Data frame transmission. ACK_0 support is applicable to acknowledged class of service Sequences (see FC-FS-6).

The conditions under which ACK_0 is supported are defined in table 220 and described in the following text.

Table 220 – ACK_0 Support Conditions (Recipient Control)

Nx_Port A Word 0, Bit 11	Nx_Port B Word 1, Bit 31	Nx_Port A as Sequence Recipient
0	0	ACK_0 not supported
0	1	ACK_0 not supported
1	0	ACK_0 not supported
1	1	ACK_0 supported

If one Nx_Port (e.g., Nx_Port A) is capable of receiving ACK_0 as a Sequence Initiator (Word 0, Bit 11 set to one) and the other Nx_Port (e.g., Nx_Port B) is capable of transmitting ACK_0 as a Sequence Recipient (Word 1, Bit 31 set to one), then ACK_0 may be used while Nx_Port A is the Sequence Initiator and Nx_Port B is the Sequence Recipient. Otherwise, ACK_0 shall not be supported while Nx_Port A is the Sequence Initiator and Nx_Port B is the Sequence Recipient.

ACK_0 capability may be asymmetrical for a single Nx_Port (i.e., an Nx_Port may be capable processing ACK_0 as a Sequence Initiator, but not be capable of ACK_0 transmission as a Sequence Recipient. Similarly, an Nx_Port may be capable of generating ACK_0 as a Sequence Recipient, but not be capable of ACK_0 reception as a Sequence Initiator). If an Nx_Port sets both Word 0, bit 11 and Word 1, bit 31 to one, then it is capable of ACK_0 support as either a Sequence Initiator or a Sequence Recipient.

6.6.5.6.3 X_ID interlock

- 0 = X_ID interlock not required
- 1 = X_ID interlock required

X_ID interlock (word 1, bit 29) only applies to Class 2. This bit indicates that the Nx_Port supplying this parameter requires that an interlock be used during X_ID assignment in Class 2. 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 (see FC-FS-6).

6.6.5.6.4 Error policy supported

The definition of the Error policy supported bits (word 1, bits 28-27) is shown in table 221.

Table 221 – Error Policy Bits Definition

Word 1, bits 28-27	Meaning
00b	Only discard policy supported
01b	Reserved
10b	Both discard and process policies supported
11b	Reserved

These bits are set to specify the types of support possible for missing frame conditions processed by the destination Nx_Port. The policy used for a given Exchange shall be specified as discard or process by the Exchange Originator (see FC-FS-6).

6.6.5.6.5 Categories per Sequence

The definition of the Categories per Sequence bits (word 1, bits 25-24) is shown in table 222.

Table 222 – Categories per Sequence Bits Definition

Word 1, bits 25-24	Meaning
00b	1 Category/Sequence
01b	2 Categories/Sequence
10b	Reserved
11b	More than 2 Categories/Sequence

The setting of these bits shall specify that the Recipient is capable of processing one, two, or more than two Information Categories (R_CTL bits 27-24 in the Frame_Header) in a single Sequence. Bits 25-24 are applicable to each Class of Service since support for an individual Class may offer different capabilities in the same Nx_Port.

While an Nx_Port is acting as a Sequence Initiator, it shall restrict the number of Information Categories per Sequence based on the Sequence Recipient's capability as specified during N_Port

Login. An Nx_Port's capability for processing Information Categories in a single Sequence may prohibit that Nx_Port from communicating in certain FC-4 protocols.

Each FC-4 should allow the ability to communicate using only one Information Category per Sequence but always provide the ability to communicate using multiple Information Categories per Sequence where possible, and when performance may be enhanced.

6.6.5.6.6 Clock synchronization ELS capable

0 = Recipient does not have clock synchronization ELS capability

1 = Recipient has clock synchronization ELS capability

The Clock synchronization capable bit (word 1, bit 19) indicates support for the ELS method of Clock Synchronization. If this bit is set to zero, the Nx_Port or Fabric does not support the ELS method of Clock Synchronization. If this bit is set to one, the Nx_Port or Fabric is capable of receiving CSU ELS frames and acting upon them (see FC-FS-6).

6.6.5.6.7 SEQ_ID unique per Exchange

0 = SEQ_ID is required to be unique per S_ID/D_ID, independent of X_ID

1 = SEQ_ID is only required to be unique per Exchange (see FC-FS-6)

SEQ_ID unique per Exchange (word 1, bit 18) specifies whether the Sequence Recipient requires SEQ_ID uniqueness per S_ID/D_ID or allows uniqueness per X_ID. If the SEQ_ID unique per Exchange bit is set to zero in the PLOGI request or PLOGI LS_ACC, then the SEQ_ID shall have a value that is unique among all concurrently open Sequences between the Sequence Initiator and the Sequence Recipient, independent of the X_ID. If the SEQ_ID unique per Exchange bit is set to one in the PLOGI request and PLOGI LS_ACC, then the SEQ_ID is only required to be unique per Exchange (see FC-FS-6).

6.6.5.7 Receive Data_Field Size

The Receive Data_Field Size is a value (word 1, bits 11-0) that specifies the largest Data_Field size for a frame (see FC-FS-6) that may be received by the Nx_Port supplying the Service Parameters as a Sequence Recipient. Values shall be a multiple of four bytes. Values less than 256 or greater than 2 112 are invalid. An Nx_Port shall support a Data_Field size of at least 256 bytes.

The Receive Data_Field Size for Class 2, and Class 3 shall be equal to or less than the Buffer-to-Buffer Receive Data_Field Size specified in the Common Service Parameters.

6.6.5.8 Concurrent Sequences

Concurrent Sequences (word 2, bits 23-16) shall specify the number of Sequence Status Blocks available in the Nx_Port supplying the Service Parameters for tracking the progress of a Sequence as a Sequence Recipient. The maximum number of Concurrent Sequences that may be specified is 255 per Nx_Port as a Recipient that may be allocated across all classes. The total number of Concurrent Recipient Sequences that may be open across all classes by a single Nx_Port is specified in the Common Service Parameter field (see 6.6.2.4.25). Allowable values for this field is 01h - FFh. The value 00 is reserved.

NOTE 20 – The maximum number of open Sequences between two Nx_Ports is the sum of the Concurrent Sequences fields reported in the PLOGI and LS_ACC of the PLOGI.

In Class 2 and 3, the SEQ_ID values shall range from 0 to 255. In Class 2 an Nx_Port may respond with a P_BSY to a frame initiating a new Sequence if Nx_Port resources are not available.

6.6.5.9 end-to-end Credit

The Nx_Port End-to-end Credit field (word 2, bits 14-0) is the maximum number of Data frames that may be transmitted by an Nx_Port without receipt of an accompanying ACK or Link_Response frames. The minimum value of the Nx_Port End-to-end Credit field is one. The Nx_Port End-to-end Credit field specified is associated with the number of buffers available for holding the Data_Field of a frame and processing the contents of that Data_Field by the Nx_Port supplying the Service Parameters. The Nx_Port End-to-end Credit field is not applicable to Class 3 since ACK frames are not used.

In order to ensure frame identification integrity, the Nx_Port end-to-end Credit field is defined as a 15-bit field while SEQ_CNT is a 16-bit field. This ensures that end-to-end Credit never exceeds one-half of the maximum SEQ_CNT. Bit 15 shall be set to zero.

Valid values for the Nx_Port End-to-end Credit field are 1 to 32 767. The value 0 is reserved.

6.6.5.10 Open Sequences per Exchange

The value of open Sequences per Exchange field (word 3, bits 23 - 16) shall specify the maximum number of Sequences that may be open at the Recipient at one time between a pair of Nx_Ports for one Exchange. The value of X+2 specifies 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 of X limits the link facility resources required for error detection and recovery. The value of X is specified in bits 23-16 (see FC-FS-6).

NOTE 21 – The number of SSBs specified at X+2 to be retained in the ESB ensures that if Sequence streaming rules are followed, the ESB shall contain at least one "good" Sequence that ended normally. Another SSB position was allocated in order to allow for any race or timing conditions that might impact that "good" Sequence.

The open Sequences per Exchange field is valid for PLOGI and PLOGI LS_ACC only.

6.6.6 Auxiliary Parameter Data

Auxiliary Parameter Data is valid for all classes.

The format of the Auxiliary Parameter Data is as shown in table 223.

Table 223 – Auxiliary Parameter Data

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Flags	Reserved		
1	MSB			
2	Platform Name_Identifier			LSB
3	Reserved		NPIV_CNT	

The Flags field in FLOGI, FDISC, FLOGI LS_ACC, and FDISC LS_ACC, is defined as follows:

- a) word 0 bit 31 shall be zero;
- b) word 0 bit 30 shall be one if the Auxiliary Parameter data is valid, and shall be zero if the Auxiliary Parameter Data is not valid;
- c) word 0 bit 29 shall be one if the Platform Name_Identifier field is valid, and shall be zero if the Platform Name_Identifier field is not valid; and
- d) word 0 bits 29..24 reserved.

The Platform Name_Identifier field is valid for FLOGI, FLOGI LS_ACC, PLOGI, and PLOGI LS_ACC. The Platform Name_Identifier field contains a Name_Identifier (see FC-FS-6) and identifies the platform associated with the FC_Port that sent the FLOGI, FLOGI LS_ACC, PLOGI, or PLOGI LS_ACC.

The NPIV_CNT field in the FDISC is defined as follows:

- a) word 3 bits 0..15 reserved.

The NPIV_CNT field in the FLOGI is defined as follows:

- b) word 3 bits 0..15 advises the fabric of the number of FDISC with S_ID=0 operations that the PN_Port intends to originate. This value shall not limit the use of FDISC with S_ID=0 operations by either the originator or the responder.

The NPIV_CNT field in the LS_ACC response to FLOGI is defined as follows:

- c) word 3 bits 0..15 contain the number of FDISC with S_ID=0 operations that are likely to be successful, as defined in table 224. A value of FFFFh indicates this field is not specified.

NOTE 22 – The number of FDISCS with S_ID=0 that are likely to be successful is dependent on the number of other FDISC with S_ID=0 requests that are received and other implementation specific factors.

The NPIV_CNT field in the LS_ACC response to FDISC is defined as follows:

- d) word 3 bits 0..15 contain the number of additional FDISC with S_ID=0 operations that are likely to be successful subsequent to this one, as defined in table 224.

Table 224 – NPIV_CNT field in LS_ACC

Value	Meaning
0 - FFFDh	The number of FDISC with S_ID=0 operations that are likely to be successful.
FFFEh	Greater than FFFDh FDISC with S_ID=0 operations are likely to be successful.
FFFFh	It is not known how many FDISC with S_ID=0 operations are likely to be successful.

6.6.7 Vendor Version Level

Vendor Version Level field (words 25-28) specifies vendor-specific information. If the Valid Version Level bit in the Common Service Parameters field (word 1, bit 29) is set to one, the Vendor Version Level field contains valid information. If the Valid Version Level bit is set to zero, the Vendor Version Level field is not meaningful.

6.6.8 Services Availability

6.6.8.1 Introduction

This field returns information regarding the Fabric's ability to route to the defined well-known addresses. It is meaningful only for FLOGI LS_ACC. Only bits 10 - 3 of word 30 are meaningful. Word 29 and bits 31 - 11 and 2 - 0 of word 30 are reserved (see FC-FS-6).

6.6.8.2 Clock Synchronization Server

If set to one, the Clock Synchronization Server bit (word 30, bit 9) shall indicate that the Fabric supports routing to the well-known Clock Synchronization Server address identifier (FFFFF6h). If set to zero, this bit shall indicate that the Fabric does not support routing to the well-known Clock Synchronization Server address identifier. See FC-FS-6 for the specification of this server.

6.6.8.3 Security Key Distribution Server

If set to one, the Security Key Distribution Server bit (word 30, bit 8) shall indicate that the Fabric supports routing to the well-known Key Distribution Server address identifier (FFFFF7h). If set to zero, this bit shall indicate that the Fabric does not support routing to the well-known Key Distribution Server address identifier.

6.6.8.4 Management Server

If set to one, the Management Server bit (word 30, bit 5) shall indicate that the Fabric supports routing to the well-known Management Server address identifier (i.e., FFFFFAh). If set to zero, this bit shall indicate that the Fabric does not support routing to the well-known Management Server address identifier. See FC-GS-9 for the specification of this server.

6.6.8.5 Time Server

If set to one, the Time Server bit (word 30, bit 4) shall indicate that the Fabric supports routing to the well-known Time Server address identifier (i.e., FFFFFBh). If set to zero, this bit shall indicate that the Fabric does not support routing to the well-known Time Server address identifier.

6.6.8.6 Directory Server

If set to one, the Directory Server bit (word 30, bit 3) shall indicate that the Fabric supports routing to the well-known Directory Server address identifier (i.e., FFFFFCh). If set to zero, this bit shall indicate that the Fabric does not support routing to the well-known Directory Server address identifier. See FC-GS-9 for the specification of this server.

6.6.9 Login Extension

6.6.9.1 General

If a port does not support a Login request with the Payload Bit set to one, the port shall reject the Login request with an LS_RJT. The reason code shall be set to “Logical error” with a reason code explanation of “Invalid Payload length”.

If a port does not support a Login request with a non-zero Login Extension Data Length field, the port shall reject the Login request with an LS_RJT. The reason code shall be set to “Logical error” with a reason code explanation of “Login Extension not supported”.

If a port receives a Login request containing a page code that it does not support, the port shall reject the Login request with an LS_RJT. The reason code shall be set to “Logical error” with a reason code explanation of “Login Extension not supported”.

If a port does not support a LS_ACC reply Sequence with the Payload Bit set to one, the port shall perform an explicit Logout with the port that sent the reply.

6.6.9.2 Login Extension Data Length

If the Login Extension Data Length field (word 31) is non-zero, a Login Extension follows the normal payload. The Login Extension Data Length field indicates the length of the Login Extension field in words. The Payload Bit (see 6.6.2.4.21) shall be set to one if this field is non-zero.

6.6.9.3 Login Extension format

The Login Extension field is a sequence of zero or more Login Extension Pages. The length in words of the sequence of Login Extension Pages shall be equal to the value of the Login Extension Length field. The format of a Login Extension Page is given in Table 225. Word 0 of the first Login Extension Page, if present, is word 64 of the FLOGI/PLOGI/ACC payload.

Table 225 – Login Extension Page format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Page Length (n)		Reserved	Page Code
1	Page Code Specific Data			
...				
n - 1				

Page Length Field: The length in words of the following page, including the word containing the page length and page code fields.

Page Code Field: The Page Code field specifies the type of page as shown in table 226.

Table 226 – Page Code Definitions

Page Code	Meaning
00h - EFh	Reserved
F0h	Vendor Specific
F1h - FFh	Reserved

Vendor Specific Page: The format of the Vendor Specific Page is shown in table 227.

Table 227 – Vendor Specific Page format

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
0	Page Length (n)		Reserved	Page Code (F0h)
1	(MSB) _____ Vendor Identification Data _____ (LSB)			
2	Vendor Specific Data			
3				
...				
n - 1				

Vendor Identification field: The value of the Vendor Identification field shall be a T10 Vendor ID. The format and interpretation of the Vendor Specific Data field is vendor specific to the vendor identified by the value of the Vendor Identification field.

Vendor Specific Data Field: The Vendor Specific Data field contains vendor specific data and shall be padded to a word boundary.

6.6.10 Clock Synchronization Quality of Service

6.6.10.1 N_Port Login

6.6.10.1.1 Applicability

The Clock Synchronization Quality of Service (QoS) field in PLOGI ELS or LS_ACC (words 62-63) is only meaningful if sent to or received from the Clock Synchronization Server (i.e., FFFFF6h). This field contains meaningful information only if either the Clock Synchronization Primitive Capable bit of the Common features field (word 1, bit 20) is set to one; or if the Clock Synchronization ELS capable bit of one of the N_Port Class Service Parameter Recipient control fields (word 1, bit 19) is set to one. If this field does not contain meaningful information, it shall be set to zero (see FC-FS-6).

The Clock Synchronization QoS field is defined in table 228.

Table 228 – N_Port Clock Synchronization QoS

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
62	CS_QoS_Request	CS_Accuracy	CS_Implemented_MSB	CS_Implemented_LSB
63	CS_Update_Period			

6.6.10.1.2 CS_QoS_Request

For PLOGI and FLOGI Request, the meaning is defined in table 229. This field is not meaningful for PLOGI LS_ACC and shall be set to zero.

Table 229 – FLOGI/PLOGI CS_QoS_Request

Word 62, Bits 31-24	Meaning
00h	The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields are not meaningful in the FLOGI/PLOGI Request.
01h	The CS_Accuracy, CS_Implemented_MSB, CS_Implemented_LSB, and CS_Update_Period fields contain the requested Quality of Service Parameters.
02h – FFh	Reserved

6.6.10.1.3 CS_Accuracy (Mantissa and Exponent)

This field contains the CS_Accuracy_Mantissa (word 62, bits 23-21) and CS_Accuracy_Exponent (bits word 62, 20-16).

If sent to the Fabric during FLOGI, these bits indicate the accuracy that the Fabric is requested to maintain in passing along to the clients the clock synchronization value it receives from the Clock Synchronization Server (FFFFFF6h).

If sent to the Clock Synchronization Server (FFFFFF6h), these bits indicate the requested accuracy of the clock synchronization value as it leaves the server port.

If received from the Clock Synchronization Server (i.e., FFFFFFF6h), these bits indicate the accuracy of the clock synchronization value as it leaves the server port. Specifically, the server shall supply a CS_Accuracy value such that the Clock Synchronization value is always within the range of:

$$T_{reference} \pm (0.5 + CS_Accuracy_Mantissa * 2^{-4}) * 2^{(CS_Accuracy_Exponent-30)},$$

where

- a) T_reference is the clock reference value internal to the server;
- b) CS_Accuracy_Mantissa is a value from 000b to 111b'; and

c) CS_Accuracy_Exponent is a value from 00000b to 11111b'.

Example #1, if CS_Accuracy Mantissa and Exponent are set to 001b and 01011b, respectively, the Clock Synchronization value as it exits the server shall always be within the range of:

$$T_{\text{reference}} \pm 1.073 \mu\text{sec}$$

Example #2, if CS_Accuracy Mantissa and Exponent are set to 111b and 11000b, respectively, the Clock Synchronization value as it exits the server shall always be within the range of:

$$T_{\text{reference}} \pm 14.65 \text{ msec}$$

6.6.10.1.4 Clock Synchronization Implemented MSB

The Clock Synchronization Implemented MSB field (word 62, bits 13 - 8) is a 6-bit value. Word 62, bits 15-14 are reserved and shall be set to zero.

If sent to the Clock Synchronization Server (i.e., FFFFF6h) during PLOGI, these bits indicate the requested most significant bit position within the 64-bit Clock Count field in the CSU ELS Payload.

If received from the Clock Synchronization Server (i.e., FFFFF6h) this field represents the most significant bit position within the 64-bit Clock Count field that contains meaningful information.

NOTE 23 – The value in the Clock Count field shall wrap around to zero if an overflow occurs from the Clock Synchronization Implemented MSB.

(e.g., a value of 110111b indicates that the MSB of byte 1 of the Clock Count field is the highest bit that contains meaningful information).

6.6.10.1.5 Clock Synchronization Implemented LSB

The Clock Synchronization Implemented LSB field (word 62, bits 13 - 8) field is a 6-bit value. Word 62, bits 7-6 are reserved and shall be set to zero.

If sent to the Clock Synchronization Server (i.e., FFFFF6h) during PLOGI, these bits indicate the requested least significant bit position within the 64-bit Clock Count field in the CSU ELS Payload.

If received from the Clock Synchronization Server (i.e., FFFFF6h), this field represents the least significant bit position within the 64-bit Clock Count field that contains meaningful information.

(e.g., a value of 001000b indicates that the LSB of byte 6 of the Clock Count field is the lowest bit that contains meaningful information).

6.6.10.1.6 Clock Synchronization Update Period

If sent to the Clock Synchronization Server (i.e., FFFFF6h), the Clock Synchronization Update Period field (word 63) indicates the requested time, in microseconds, between consecutive updates from the Clock Synchronization server.

If received from the Clock Synchronization Server (i.e., FFFFF6h), it represents the time, in microseconds, between consecutive updates from the Clock Synchronization server.

This field is not meaningful for FLOGI and shall be set to zero.

6.6.10.2 Fabric Login

6.6.10.2.1 Applicability

The Clock Synchronization Quality of Service field contains meaningful information only if either Word 1, bit 20 - Clock Synchronization Primitive Capable of the Common features field is set to one, or Word 1, bit 19 - Clock Synchronization ELS capable of the Recipient control field is set to one. If this field does not contain meaningful information, it shall be set to zero (see FC-FS-6).

The Fx_Port Clock Synchronization Quality of Service field is illustrated in table 230.

Table 230 – Fx_Port Clock Synchronization QoS

Bits Word	31 .. 24	23 .. 16	15 .. 08	07 .. 00
62	Reserved	CS_Transfer_Accuracy	CS_Implemented_MSB	CS_Implemented_LSB
63	Reserved			

6.6.10.2.2 CS_Transfer_Accuracy

The CS_Transfer_Accuracy field contains CS_Transfer_Accuracy_Mantissa (word 62, bits 23-21) and CS_Transfer_Accuracy_Exponent (word 62, bits 20-16).

These bits indicate the accuracy that the Fabric maintains in passing along to the clients the clock synchronization value it receives from the Clock Synchronization Server. Specifically, the Fabric shall supply a CS_Transfer_Accuracy value such that the Clock Synchronization value supplied to the clients is always within the range of:

$$(T_{server} + T_{fabric_delay}) \pm (0.5 + CS_Accuracy_Mantissa * 2^{-4}) * 2^{(CS_Accuracy_Exponent-30)}$$

where:

- a) T_server is the value received from the Clock Synchronization Server;
- b) T_fabric_delay is the time period between a given value received from the server until the corresponding value is delivered to the client;
- c) CS_Accuracy_Mantissa is a value from 000b to 111b'; and
- d) CS_Accuracy_Exponent is a value from 00000b to 11111b'.

Example #1, if CS_Transfer_Accuracy Mantissa and Exponent are set to 001b and 01011b, respectively, the Clock Synchronization value supplied to the clients shall always be within the range of:

$$(T_{server} + T_{fabric_delay}) \pm 1.073 \mu\text{sec}$$

Example #2, if CS_Transfer_Accuracy Mantissa and Exponent are set to 111b and 11000b, respectively, the Clock Synchronization value supplied to the clients shall always be within the range of:

$$(T_{server} + T_{fabric_delay}) \pm 14.65 \text{ msec}$$

6.6.10.2.3 Clock Synchronization Implemented MSB

The Clock Synchronization Implemented MSB field (word 62, bits 13 - 8) is a 6-bit value that represents the most significant bit position within the 64-bit Clock Count field that contains meaningful information (e.g., a value of 110111b indicates that the MSB of byte 1 of the Clock Count field is the highest bit that contains meaningful information). Word 62, bits 15-14 are reserved and shall be set to zero.

This field refers to the capabilities of the Fabric in transferring the clock synchronization value that was received from the Clock Synchronization Server (i.e., FFFFF6h) to the clients. It does not refer to the capabilities of the Clock Synchronization Server itself.

6.6.10.2.4 Word 0, Bits 7-0 Clock Synchronization Implemented LSB

This Clock Synchronization Implemented LSB field (word 62, bits 5 - 0) is a 6-bit value that represents the least significant bit position within the 64-bit Clock Count field that contains meaningful information (e.g a value of 001000b indicates that the LSB of byte 6 of the Clock Count field is the lowest bit that contains meaningful information). Word 62, bits 7-6 are reserved and shall be set to zero.

This field refers to the capabilities of the Fabric in transferring the clock synchronization value that was received from the Clock Synchronization Server (i.e., FFFFF6h) to the clients. It does not refer to the capabilities of the Clock Synchronization Server itself.

7 Process Login/Logout

7.1 Process Login

7.1.1 Introduction

The Process Login (PRLI) ELS request shall be used to establish the operating environment between processes at the originating Nx_Port and processes at the responding Nx_Port.

Establishing the operating environment may include the establishment of image pairs and the exchange of Service Parameters. The establishment of image pairs is FC-4 independent and is system structure dependent. The exchange of Service Parameters is FC-4 dependent, and if required by a particular FC-4, shall be specified in the corresponding FC-4 standard.

A Process Login remains in effect until a process logout occurs. The number of concurrent Process Logins in effect at an Nx_Port is a function of the Nx_Port facilities available. Process Login is separate from N_Port Login. Process Login may be either implicit or explicit.

There are 2 types of Process Login:

- a) Implicit Process Login is a method of establishing an operating environment by means other than the explicit use of the PRLI Exchange. Specific methods of Implicit Process Login are not defined in this standard; and
- b) Explicit Process Login is accomplished by using the PRLI ELS Sequence within a separate Exchange to establish an operating environment.

The combination of the D_ID, S_ID, and TYPE identify an image pair (see figure 3). Either a single group or multiple groups of related processes may exist behind an Nx_Port.




Figure 3 – Image pairs

PRLI, if required, is performed after N_Port Login is successful and prior to other FC-4 transfers. Examples of use of the Process Login function may include image initialization, image re-configuration, or if the Nx_Port receives an indication that the image pair no longer exists. PRLI allows each image behind an Nx_Port to separately manage its resources.

PRLI may be used to establish an operating environment between any of the following combinations of Nx_Port facilities:

- a) two Nx_Ports;
- b) one Nx_Port and one Nx_Port image; and
- c) two Nx_Port images.

An image pair may be established or modified with a PRLI request and LS_ACC reply Sequence set. Failure to establish a particular image pair does not affect existing image pairs or the ability to establish other image pair.

PRLI may also be used to exchange Service Parameters without establishing image pairs. However, if an image pair is currently established, a subsequent PRLI request targeted to the same Nx_Port pair shall identify an image pair in order to modify Service Parameter settings for that image pair.

If a PRLI request is received for an established image pair, the established image pair is unaffected and the PRLI request is processed normally. This allows the exchange of Service Parameters for a FC-4 not specified if the original image pair was established. PRLO shall be used to remove an established image pair.

It shall be the responsibility of the ULPs to ensure that all active operations over an image pair have been properly terminated prior to issuing a PRLI request that replaces Service Parameters. If the replacement of Service Parameters affects any active operations, all open Sequences and Exchanges shall be terminated by invoking Abort Sequence (ABTS). Whether or not the replacement of Service Parameters affects an active operation shall be specified for each Service Parameter by the associated FC-4.

The Nx_Port originating the PRLI request shall not consider the image pair to be established until it has taken the necessary action to establish the image pair, and has received an LS_ACC reply Sequence indicating that the image pair has been established. The Nx_Port responding to the PRLI request does not consider the image pair to be established until the necessary action is taken at the Nx_Port to establish the image pair, and an LS_ACC reply Sequence is sent.

If a link error is detected while a PRLI request is received, the appropriate response, if any, is made, and the image pair is not established. If the recipient Nx_Port is not logged in with the requesting port it shall reply with a LOGO ELS Sequence, or with an LS_RJT ELS Sequence with a reason code of "Unable to perform command request" and a reason code explanation of "N_Port Login required". If an LS_RJT is sent in response to a PRLI request for an image pair that is already established, the existing image pair is unaffected. If an LS_RJT, P_BSY, F_BSY, P_RJT, or F_RJT response is received to a PRLI request, the PRLI request may be retried until the image pair is established. The number of retries is system dependent. In the case of LS_RJT, whether or not the PRLI is retried depends on the LS_RJT reason code.

In the event that there is an error in the response to establish an image pair, the originating Nx_Port shall not assume that the requested action has or has not taken place. If the Nx_Port that originates a PRLI request receives no valid response, the Nx_Port should retry the request. The number of retries is system dependent.

7.1.2 Mode of operation

7.1.2.1 Informative mode

Service Parameter information is exchanged enabling subsequent negotiation for image pair establishment.

7.1.2.2 Binding mode

Information is exchanged that explicitly establishes a relationship between processes in the communicating Nx_Ports. The relationship does not allow any communication types or paths other than those established by the PRLI.

The use of a Binding PRLI page requires that the Originator have precise and detailed knowledge of the capabilities available in the Responder. That information may be obtained from Directory Services, implicitly from configuration information obtained outside the scope of FC, or by performing an Informative PRLI.

Binding or Informative mode is determined by the setting of the Establish Image Pair bit in the PRLI request page.

The Service Parameters included in a page may be either requirements or capabilities. Capabilities indicate those FC-4 properties that describe the role and state of the node in the FC-4 (e.g., channel or device for FC-SB-6, initiator or target for FCP-5, and similar values). Requirements indicate those FC-4 properties that shall be agreed upon by both nodes for operation with a FC-4.

7.2 Process Logout

The Process Logout (PRLO) ELS request shall be used to request invalidation of the operating environment between an image at the initiating Nx_Port and an image at the recipient Nx_Port. PRLO frees resources committed by a previous PRLI function. ULP behavior following successful execution of the PRLO function is specified in the corresponding FC-4 standard.

Examples of PRLO usage include image re-configuration and TYPE-specific reset of Process Login Service Parameters.

TYPE-specific Service Parameter settings may be reset or image pairs removed with a PRLO request and LS_ACC reply Sequence set. Other TYPE-specific Service Parameter settings or image pairs associated with the same or different image pairs or Nx_Ports shall be unaffected. After TYPE-specific Service Parameter settings are reset or image pair is removed, Information Units may not be sent or received for the specified FC-4 TYPE using that image pair and Nx_Port combination specified in the PRLO request.

If a PRLO request is received for an image pair or FC-4 TYPE that does not exist, the request is accepted, provided that no link errors are detected, and the LS_ACC response is sent.

The Nx_Port originating the PRLO request shall not consider an image pair to be removed until it receives an LS_ACC reply Sequence. The Nx_Port responding to the PRLO request shall not consider an image pair to be removed or TYPE-specific Service Parameter settings to be reset, as appropriate, until the LS_ACC reply Sequence is sent. An Nx_Port that receives a P_BSY, F_BSY, P_RJT, or F_RJT reply in response to a PRLO request may retry the PRLO request. The number of retries is system dependent.

Unless the requesting Nx_Port receives a valid response to a PRLO request, that Nx_Port shall not assume that the requested action has or has not taken place. If the Nx_Port that originates a PRLO request receives no valid response, the Nx_Port should retry the request. The number of retries is vendor unique.

A PRLO page identifies a particular image pair to logout. Only that image pair is logged out. No further communication for the affected FC-4(s) is possible between these two images. It is the responsibility of the ULPs to ensure that all active operations over an image pair have been orderly

and properly terminated prior to issuing a PRLO request. Following PRLO execution, all active Sequences and Exchanges shall be terminated by invoking Abort Sequence (ABTS). On-going operations and states for other image pairs are not affected.

If a ULP attempts to communicate over an image pair that has not been established or has been abnormally terminated, the communication shall be acknowledged in the normal manner. The Originator may then perform a PRLO operation for the affected image pair in order to properly terminate the operating environment at both the Originator and Responder.

8 Virtual Fabrics PN_Port Support

8.1 Overview

The Virtual Fabric Tagging Header (VFT_Header, see FC-FS-6) 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. The VFT_Header may be supported by Multiplexers for PN_Ports, PF_Ports, and PE_Ports (see FC-FS-6).

8.2 Enabling VFT Tagging on PN_Ports

The use of Virtual Fabrics is negotiated if FLOGI is performed via the Virtual Fabrics bit of the FLOGI Common Service Parameters (see table 207).

If set to one in the FLOGI request, the Virtual Fabrics bit indicates that the sending PN_Port may negotiate Virtual Fabrics parameters. If set to one in the FLOGI LS_ACC reply, the Virtual Fabrics bit indicates that the Fabric requires the PN_Port to initiate an EVFP transaction (see 4.3.38). Table 231 shows the usage of the Virtual Fabrics bit.

Table 231 – Virtual Fabrics Bit Usage

Requesting PN_Port	PF_Port Configuration	Behavior
Virtual Fabrics Bit = 0b	Virtual Fabrics Not Allowed	LS_ACC with Virtual Fabrics Bit = 0b
Virtual Fabrics Bit = 1b	Virtual Fabrics Not Allowed	LS_ACC with Virtual Fabrics Bit = 0b
Virtual Fabrics Bit = 0b	Virtual Fabrics Allowed	LS_ACC with Virtual Fabrics Bit = 0b
Virtual Fabrics Bit = 1b	Virtual Fabrics Allowed	LS_ACC with Virtual Fabrics Bit = 1b

If the PF_Port requests the PN_Port to initiate an EVFP transaction, the D_ID of the FLOGI LS_ACC shall be set to 000000h. In this manner the link flow control parameters (e.g., Buffer-to-buffer Credit and Buffer-to-Buffer Receive Data_Field Size) are negotiated, but no N_Port_ID is assigned to the PN_Port. The PN_Port may continue the initialization process by using the N_Port Controller Well Known Address (i.e., FFFF0h, see FC-FS-6) as S_ID to perform an Authentication transaction or an EVFP transaction. Figure 4 shows the initialization process to enable Virtual Fabrics at a PN_Port.

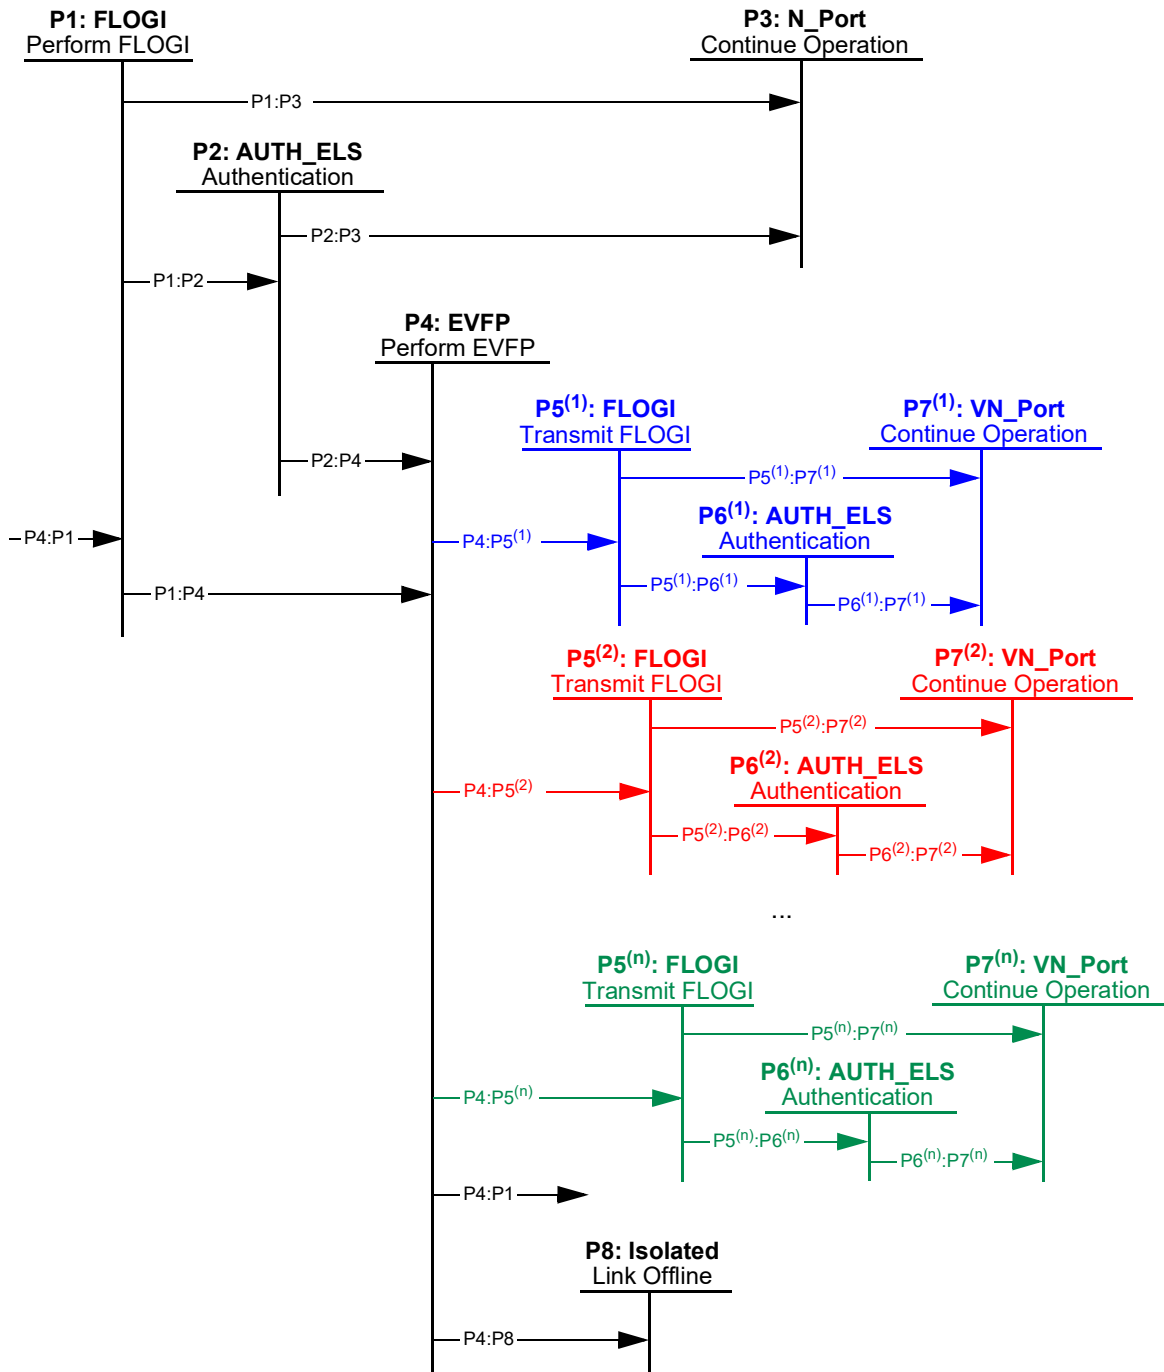


Figure 4 – PN_Port Initialization of Virtual Fabrics

State P1: FLOGI. The PN_Port performs FLOGI. The N_Port_Name of the VN_Port to be associated with the Port VF_ID shall be used by the PN_Port in the FLOGI request; the Fabric_Name of the Virtual Fabric or Virtual Switch_Name of the locally-attached Switch associated with the Port VF_ID shall be used by the PF_Port in the FLOGI LS_ACC.

State P2: AUTH_ELS. In this state an Authentication transaction shall be performed (see FC-SP-2). The Virtual Fabrics bit may be set to one or to zero; the resulting behavior is as specified in table 231. If the FLOGI LS_ACC had the Virtual Fabrics bit set to one and the D_ID set to 000000h, then the S_ID of the AUTH_ELS shall be the N_Port Controller WKA, otherwise the S_ID of the AUTH_ELS shall be the assigned Address Identifier.

State P3: N_Port. The N_Port continues its operations. The PN_Port may acquire additional N_Port_IDs by following the Additional N_Port_ID procedure.

Transition P1:P2. Occurs if the FLOGI LS_ACC has the Security bit set to one.

Transition P1:P3. Occurs if the FLOGI LS_ACC has both the Security bit and the Virtual Fabrics bit set to zero.

Transition P2:P3. Occurs if the Authentication transaction performed in state P2 completes successfully and the FLOGI LS_ACC had the Virtual Fabrics bit set to zero.

Transition P1:P4. Occurs if the FLOGI LS_ACC has the Security bit set to zero and the Virtual Fabrics bit set to one.

Transition P2:P4. Occurs if the Authentication transaction performed in state P2 completes successfully and the FLOGI LS_ACC had the Virtual Fabrics bit set to one.

State P4: Process EVFP. The PN_Port shall initiate the EVFP processing as described in 8.4.1.

Transition P4:P5^(k). Occurs if 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 P5^(k).

State P5^(k): FLOGI. In this state the FC frames transmitted by the PN_Port are tagged with the VFT_Header carrying VF_ID K. An FLOGI request, tagged with VF_ID K, is transmitted. This FLOGI request shall carry the N_Port_Name of the VN_Port associated with VF_ID K and the operational parameters (e.g., timeout values, Classes of service) of Virtual Fabric K. The Virtual Fabrics bit shall be set to zero in this FLOGI request. The FLOGI LS_ACC assigns an N_Port_ID in Virtual Fabric K to the VN_Port associated with VF_ID K. All parameters of the FLOGI request and the FLOGI LS_ACC transmitted in this state are used as specified in clause 6, except that the following parameters in the Common Service Parameters are ignored and the values transmitted in state P1 are used instead:

- a) Buffer-to-buffer Credit (see 6.6.2.3);
- b) BB_Credit Management (see 6.6.2.4.9);
- c) BB_SC_N (see 6.6.2.4.22); and
- d) Buffer-to-Buffer Receive Data_Field Size (see 6.6.2.4.23).

Transition P5^(k):P6^(k). Occurs if the FLOGI processing in state P5 is completed, if the FLOGI LS_ACC has the Security bit set to one.

State P6^(k): AUTH_ELS. In this state the FC frames transmitted by the PN_Port are tagged with the VFT_Header carrying VF_ID K. The VN_Port associated with VF_ID K performs an Authentication transaction (see FC-SP-2) in Virtual Fabric K, by using its N_Port_Name or any other appropriate

identity. The corresponding VF_Port authenticates by using the Switch_Name of the Virtual Switch associated with VF_ID K or any other appropriate identity.

Transition P5^(k):P7^(k). Occurs if the FLOGI processing in state P5 is completed, if the FLOGI LS_ACC has the Security bit set to zero.

Transition P6^(k):P7^(k). Occurs if the Authentication transaction performed in state P6^(k) completes successfully.

State P7^(k): VN_Port. For VN_Ports in this state the PN_Port operates as VFT tagging PN_Port. FC frames transmitted by the PN_Port are tagged with the VFT_Header carrying VF_ID K. The VN_Port continues its operations. The PN_Port may acquire additional N_Port_IDs in Virtual Fabric K by following the additional N_Port_ID procedure, tagging the frames with a VFT_Header carrying VF_ID K.

Transition P4:P1. Occurs when the EVFP processing determined that VFT tagging is not performed and the PN_Port and PF_Port have the same Port VF_ID. In state P1 the PN_Port shall transmit an FLOGI request with the Virtual Fabrics bit set to zero, in order to acquire an N_Port_ID in the Virtual Fabric identified by the Port VF_ID.

Transition P4:P8. Occurs if the EVFP processing determined that VFT tagging is not performed and the PN_Port and PF_Port do not have the same Port VF_ID.

State P8: Isolated. In this state the PN_Port is offline.

If VFT tagging is enabled on a link, a Link Reset (see FC-FS-6) shall not change the tagging process, while a Link Initialization (see FC-FS-6) shall stop the tagging process and reinitialize the link.

8.3 Configuration Information

A VF capable PN_Port shall maintain the following configuration parameters:

- a) Tagging Administrative Status, used to negotiate the VFT tagging operational mode of the PN_Port (see 4.3.38.2.2);
- b) Port VF_ID (see 4.3.38.2.3); and
- c) Locally-Enabled VF_ID List, used to negotiate the list of Virtual Fabrics operational over the PN_Port (see 4.3.38.2.4).

8.4 Exchange Virtual Fabrics Parameters Processing

8.4.1 Overview

The Exchange Virtual Fabrics Parameters ELS (EVFP) protocol allows a VF capable PN_Port and a VF capable PF_Port 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 link.

An EVFP transaction occurs between an EVFP Initiator and an EVFP Responder. An EVFP transaction (see figure 5) is identified by a unique Transaction Identifier (T_ID), and consists of a synchronizing phase (EVFP_SYNC) followed by a commit phase (EVFP_COMMIT).

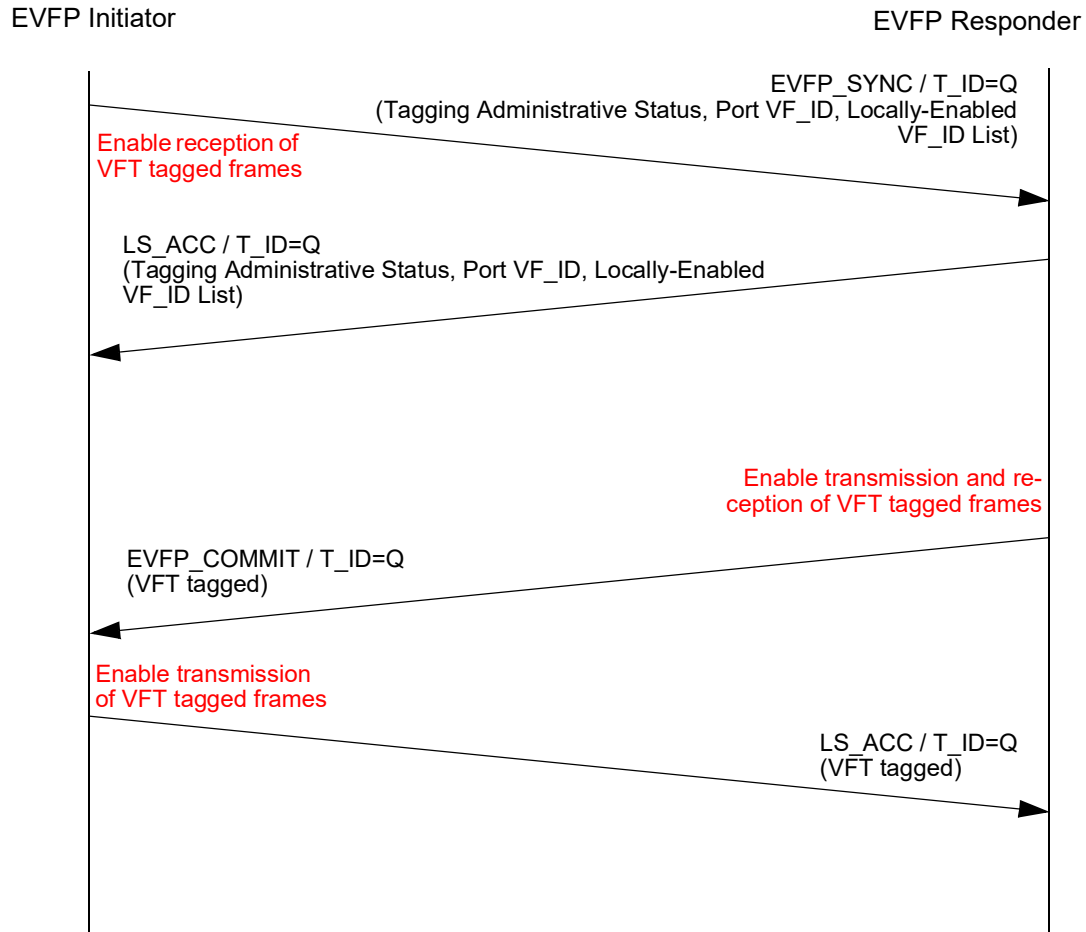


Figure 5 – A Generic EVFP Transaction

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 Initialization of Virtual Fabrics for a PN_Port, proceeds as follows:

- 1) The EVFP Initiator shall start the EVFP transaction by sending the EVFP_SYNC message (see 4.3.38.2) to the EVFP Responder. In the EVFP_SYNC message, the EVFP Initiator shall specify the Transaction Identifier, and shall send its Core N_Port_Name if PN_Port or its Core Switch_Name if PF_Port, together with its Tagging Administrative Status (see 4.3.38.2.2), Port VF_ID (see 4.3.38.2.3) and Locally-Enabled VF_ID List (see 4.3.38.2.4). On sending the EVFP_SYNC message the EVFP Initiator enables the reception of VFT tagged frames;
- 2) The EVFP Responder shall reply with an LS_ACC carrying its Core Switch_Name if PF_Port or its Core N_Port_Name if PN_Port, together with its Tagging Administrative Status, Port VF_ID and Locally-Enabled VF_ID List. Then the EVFP Responder shall determine if VFT Tagging has to be enabled on the link, according to table 127. If VFT Tagging has to be enabled, the EVFP Responder shall go to step 3. If VFT Tagging has 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 the EVFP protocol terminates and the EVFP Responder goes in Isolated state (transition P4:P8, see 8.2); or
- B) if the peer's Port VF_ID is equal to the local Port VF_ID, on completion of the Exchange the EVFP protocol terminates and the EVFP Responder goes in state P1 (transition P4:P1, see 8.2).

On receiving the EVFP_SYNC LS_ACC, the EVFP Initiator shall determine if VFT Tagging has to be enabled on the link, according to table 127. If VFT Tagging has to be enabled, on completion of the Exchange the EVFP Initiator shall go to step 4. If VFT Tagging has 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 the EVFP protocol terminates and the EVFP Initiator goes in Isolated state (transition P4:P8, see 8.2); or
 - B) if the peer's Port VF_ID is equal to the local Port VF_ID, on completion of the Exchange the EVFP protocol terminates and the EVFP Initiator goes in state P1 (transition P4:P1, see 8.2);
- 3) On completion of the EVFP_SYNC Exchange, 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 4.2.56.2.4. 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 4.2.56.3), 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 4.2.56.2.4. 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 LS_ACC message tagged with its Port VF_ID.

If 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. If the EVFP transaction is completed the processing continues independently for each Virtual Fabric operational on the link, as shown by transitions P4:P5(k) (see 8.2). If the computed set of VF_IDs operational on the link is NULL, the involved FC_Ports remain in state P4 (see 8.2) until a new EVFP transaction is performed in the Control VF_ID.

If an PN_Port and an PF_Port start an EVFP transaction at the same time, or if an FC_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 PN_Port shall remain the EVFP Initiator, while the PF_Port shall become the EVFP Responder. The FC_Port that remains the EVFP Initiator shall reply to the received EVFP_SYNC message with an 'EVFP collision' LS_RJT (see 4.3.38.1.2). The FC_Port that becomes the EVFP Responder shall reply to the received EVFP_SYNC message and abort its own transaction upon receipt of the LS_RJT.

The EVFP protocol is used also if some PN_Port or PF_Port configuration information (see 8.3) is changed by a management action. The EVFP messages may be carried in FC frames tagged with the Port VF_ID if the EVFP protocol begins while the link is not performing VFT tagging (see 8.4.1). The EVFP messages are carried in FC frames tagged with the Control VF_ID if the EVFP protocol begins while the link is performing VFT tagging (see 8.4.2 and 8.4.3).

8.4.2 Changing the VFT Tagging Mode

If a management action changes the Administrative Tagging Mode of a VF capable PN_Port or VF capable PF_Port that determined during initialization the peer supports the EVFP protocol, the FC_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 PN_Port or PF_Port is currently performing tagging, all EVFP protocol messages shall be tagged with the Control VF_ID. If the PN_Port or PF_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 4.3.38.2.2), Port VF_ID, and the Locally-Enabled VF_ID List; and
- 2) The EVFP Responder shall reply with an LS_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 127. The EVFP Responder:
 - A) if VFT Tagging has not to be changed, on completion of the Exchange terminates the EVFP protocol; or
 - B) if VFT Tagging has to be changed, on completion of the Exchange shall perform a link initialization.

On receiving the EVFP_SYNC LS_ACC, the EVFP Initiator shall determine if VFT Tagging has to be changed on the link, according to table 127. The EVFP Initiator:

- A) if VFT Tagging has not to be changed, on completion of the Exchange terminates the EVFP protocol; or
- B) if VFT Tagging has to be changed, shall participate in the link initialization initiated by the EVFP Responder.

8.4.3 Adding or Removing Virtual Fabrics

If a management action changes the Locally-Enabled VF_ID List of a tagging PN_Port or tagging PF_Port, the FC_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 4.3.38.2.4);
- 2) The EVFP Responder shall reply with an LS_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 4.3.38.2.4):

- A) if the operational VF_ID List did not change, on completion of the Exchange in the Control VF_ID terminates the EVFP protocol; or
- B) if the operational VF_ID List did change, on completion of the Exchange in the Control VF_ID goes to step 3.

On receiving the EVFP_SYNC LS_ACC in the Control VF_ID, the EVFP Initiator, depending on the resulting operational VF_ID List:

- A) if the operational VF_ID List did not change, on completion of the Exchange in the Control VF_ID terminates the EVFP protocol; or
 - B) if the operational VF_ID List did change, on completion of the Exchange in the Control VF_ID goes to step 4.
- 3) On completion of the EVFP_SYNC Exchange 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. The removal of a Virtual Fabric should be treated as an implicit logout. 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. The removal of a Virtual Fabric should be treated as an implicit logout. Then the EVFP Initiator shall send an EVFP_COMMIT LS_ACC message.

If 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. If the EVFP transaction is completed, the updated operational VF_ID List is operative.

8.4.4 Changing the Port VF_ID

If a management action changes the Port VF_ID of a tagging PN_Port or tagging PF_Port, no changes are applied to the link.

If a management action changes the Port VF_ID of a non tagging PN_Port or tagging PF_Port, the FC_Port shall perform a link initialization.

If a management action changes the Port VF_ID of a PN_Port or PF_Port in Isolated state (i.e., state P8), the FC_Port shall go in state P1 (see 8.2).

9 Priority Tagging

9.1 Overview

The Priority Tagging mechanism uses values of the Priority field in the Fibre Channel header as VE identifiers local to an N_Port_ID of a Virtual Entities Manager (VEM). The VEM tags the traffic of a specific VE with an 8-bit local VE ID that is in the Priority field of the FC header. This enables the Fabric and the devices to identify, monitor, and handle FC traffic based on VE tags.

Figure 6 shows the relationship among local VE IDs (VEs), N_Port_IDs (VN_Ports) and physical ports (PN_Ports) in a system supporting server virtualization. Each PN_Port is able to instantiate multiple VN_Ports through the NPIV mechanism. Each VN_Port has its own N_Port_ID and is able to identify multiple VEs through local VE ID values (i.e., values local to the VN_Port's N_Port_ID). The pair {VN_Port's N_Port_ID, Local VE ID} is the Fabric VE ID of a specific VE.

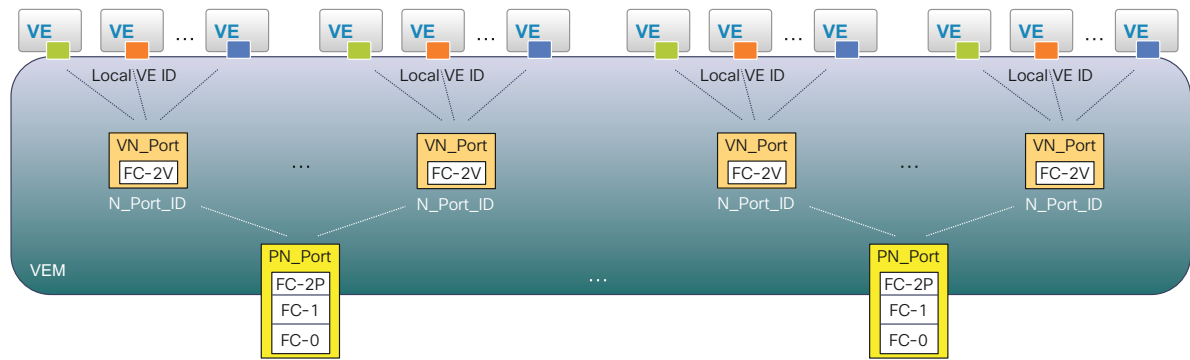


Figure 6 – Example of Relationship among Local VE IDs, N_Port_IDs, and PN_Ports

Devices and the Fabric are able to identify the traffic of a specific VE through the Fabric VE ID. The Fabric maintains the mapping between global VE IDs and Fabric VE IDs through the VE Identification Server (see FC-GS-9). Table 232 shows the mappings maintained by the VE Identification Server.

Table 232 – VE Identification Server Mappings

Global VE ID	Fabric VE ID	
Global VE ID #1	N_Port_ID #1	Local VE ID #1
Global VE ID #2	N_Port_ID #2	Local VE ID #2
...		
Global VE ID #n	N_Port_ID #n	Local VE ID #n

Figure 7 shows an example of VE traffic identification using the Priority Tagging mechanism.

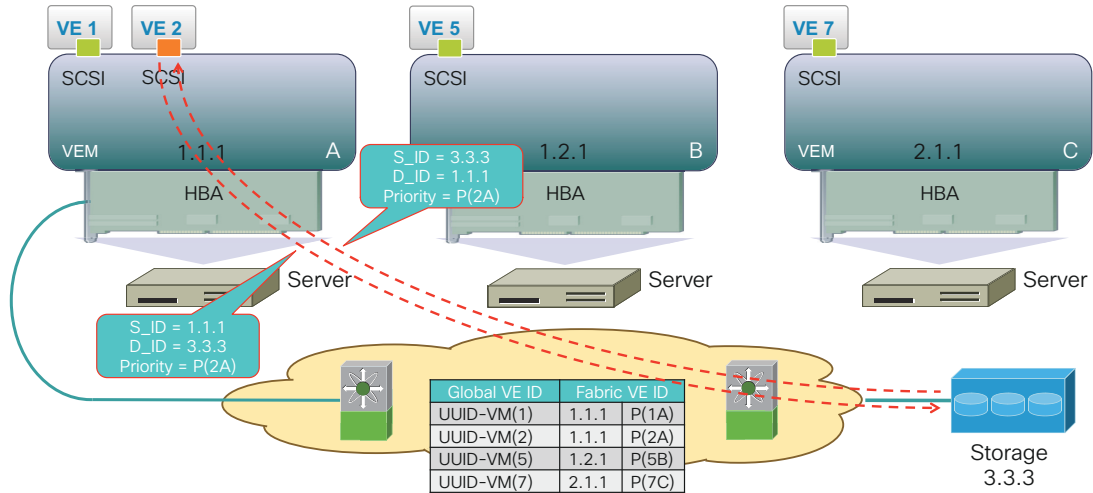


Figure 7 – Example of Priority Tagging of the Traffic of a VE

As shown in figure 7, the VE Identification Server maps the global VE IDs of VEs 1, 2, 5, and 7 into their respective Fabric VE IDs. For example, the FC frames that VE 2 originates to the storage are identified by their S_ID and Priority field value and the FC frames that it receives from the storage are identified by their D_ID and Priority field value (i.e., the Fabric VE ID of VE 2). A query to the VE Identification Server enables the resolution of the Fabric VE ID of VE 2 into its global VE ID.

The VE Identification Server mappings are updated by the VEMs. Each time a VE is instantiated, deinstantiated, or moved, the involved VEM(s) shall update the VE Identification Server accordingly, to keep the mapping information current.

9.2 Operations

Support for the Priority Tagging mechanism shall be indicated via the Priority Tagging Supported bit in the Login Common Service Parameters (see 6.6.2.4.15).

An N_Port supporting Priority Tagging (i.e., setting to one the Priority Tagging Supported bit in the FLOGI request) and the instantiation of VEs, upon completing FLOGI with a Switch also supporting Priority Tagging (i.e., setting to one the Priority Tagging Supported bit in the FLOGI LS_ACC) shall inquire the Fabric about the Priority field values useable as local VE IDs through the QFPA ELS (see 4.3.50). The QFPA LS_ACC provides to the N_Port the ranges of Priority field values useable as local VE IDs and the QoS priority processing the Fabric provides for each range. The N_Port may use the available Priority field values in the manner that it deems appropriate. Priority field value zero shall not be used as a local VE ID for VEs that require traffic identification.

Upon instantiating a VE, an N_Port supporting Priority Tagging shall select one of the available Priority field values as the local VE ID for that VE and update the VE Identification Server by registering with the Fabric the correspondent global VE ID to Fabric VE ID mapping using a UVEM ELS (see 4.3.51). If the N_Port is logged in to a device that supports Priority Tagging (i.e., a device that set to one the Priority Tagging Supported bit in the PLOGI LS_ACC), then the N_Port should notify the device about the instantiated VE by registering with the device the correspondent global VE ID to Fabric VE ID mapping using a UVEM ELS.

Upon deinstantiating a VE, an N_Port supporting Priority Tagging shall deselect the local VE ID for that VE and update the VE Identification Server by deregistering with the Fabric the correspondent global VE ID to Fabric VE ID mapping using a UVEM ELS. If the N_Port is logged in to a device that supports Priority Tagging (i.e., a device that set to one the Priority Tagging Supported bit in the PLOGI LS_ACC), then the N_Port should notify the device about the deinstantiated VE by deregistering with the device the correspondent global VE ID to Fabric VE ID mapping using a UVEM ELS.

A VE move from one N_Port to another is handled as a VE deinstantiation on the origin N_Port and a VE instantiation on the destination N_Port.

9.3 Scalability

Priority field values enable a VEM to identify the traffic of up to 255 VEs per VN_Port. A PN_Port may instantiate multiple VN_Ports (i.e., acquire additional N_Port_IDs) through NPIV if the traffic of more than 255 VEs need to be identified. Large systems supporting many VEs are usually equipped with multiple N_Ports, each providing the ability to identify up to 255 VEs (see figure 6).

If a VE does not require traffic identification, then the value zero shall be used as its local VE ID.

Multiple VEs may use the same local VE ID. If the same local VE ID is used by multiple VEs then those VEs are identified as a group and their individual traffic is not able to be identified.

Annex A: Fabric notification information and examples (informative)

A.1 Overview

The Fabric notifications functions (e.g., EDC ELS, FPIN ELS) provide additional information to participating devices to assist in determining the nature of impacts to frame processing operations. This information is provided by the Fabric and attached devices using implementation specific tools. Devices are required to register with the Fabric to receive notifications and signals and the Fabric only delivers notifications or signals to registered devices.

Notifications are a mechanism for providing feedback about events occurring throughout the Fabric; thus, the scope of notifications is Fabric wide. Signals are a mechanism for providing feedback regarding transmission resources between attached devices; thus, the scope of signals is link by link. Notifications are provided in the form of ELSs (see clause 4) and signals are provided in the form of primitives (see FC-FS-6).

A.2 Causes of congestion

Notifications indicating congestion is present include an indication of the cause of the congestion such as lost buffer credits, credit-stalled device behavior, or oversubscription.

Lost credits are caused by links experiencing physical errors that affect the transmission of frames or buffer credits (i.e., a frame is corrupted or an R_RDY signal is corrupted). Lost credits may prevent a transmitter from sending frames at full rate and results in degraded throughput. The effects of lost credits increase in severity as more credits are depleted. If all credits are lost and frames are not transmitted for E_D_TOV, then the Link Reset Protocol is performed (see FC-FS-6). Identification of lost credit is implementation dependent and may include detection of credit latency at a port or queue latency at ports upstream from the port. Mitigation for lost credits is achieved by using the BB_Credit Recovery protocol (see FC-FS-6).

A credit-stalled condition occurs if a receiving port stops returning credits to the transmitter causing credit depletion. This results in the transmitter to stop transmitting due to reaching zero transmit credits. The reasons for credit-stalled conditions are varied and the delays range widely from small delays to large delays (e.g., triggers E_D_TOV processing). Credit-stalled conditions may affect flows unrelated to the credit-stalled port due to backpressure or congestion spreading in the Fabric. Application performance is impacted if the credit-stall is long enough to cause large queue latencies (e.g., greater than 200 ms). If the credit-stalling condition persists for E_D_TOV, then the Link Reset Protocol is performed (see FC-FS-6). Identification of credit-stalled conditions is implementation dependent and may be identified by credit latency or frame loss at an upstream port.

Oversubscription occurs if a port or a link is asked to handle more frames than it is able to efficiently process, causing frames to back up into the Fabric. Oversubscription conditions may affect flows unrelated to the oversubscribed port due to backpressure or congestion spreading in the Fabric. Oversubscription also occurs on frame flows from multiple devices to a common link (e.g., exceeds the capacity of an ISL). Oversubscription is commonly caused by a bandwidth mismatch between the source and destination ports (e.g., a 32GFC port sending to an 8GFC port) and may affect flows that share the same path through the Fabric. Some reasons for oversubscription include a device asking for more data than its interface is able to process effectively (e.g., an initiator exhibiting read oversubscription behavior) or a device zoned with more peers than its configuration is able to efficiently manage (e.g., over-utilization of a target due to too many initiators). Identification of oversubscription is implementation dependent and may be identified by queue latency on upstream port(s) and high bandwidth utilization at a downstream port.

A.3 General use and device behaviors

A.3.1 Overview

The FPIN ELS descriptors that identify the location and effect of the detected event in the Fabric contain the Detecting Port Name field, the Attached Port Name field, and the Port Name List field. The Detecting Port Name field is set to the Port_Name of the FC_Port that detected the event. The Attached Port Name field is set to the Port_Name of the FC_Port attached to the FC_Port that detected the event. The Port Name List field is set to the N_Port_Names of the N_Ports affected by the event and zoned with the recipient of the FPIN ELS.

A.3.2 Port name lists

If the Fabric detects an event on a link between the Fabric and a device, then the Fabric sets the Port Name List field to the N_Port_Names affected by the event and zoned with the recipient of the FPIN ELS.

If the device detects an event on a link between the device and the Fabric, then the device sets the Port Name List field to the N_Port_Names affected by the event. When the Fabric distributes the FPIN ELS from the device, then the Fabric sets the Port Name List field to the N_Port_Names affected by the event and zoned with the recipient of the FPIN ELS.

If the Fabric detects an event on a link between Fabric ports (i.e., E_Ports), then the Fabric sets the Port Name List field to the N_Port_Names affected by the event and zoned with the recipient of the FPIN ELS.

In essence, the Port Name List (i.e., PNL) is the list of Port_Names affected by the event and zoned with the device receiving the FPIN ELS. The device receiving the FPIN ELS uses the PNL to determine what actions to take (i.e., if any). For example, the PNL of an FPIN ELS sent to a target device for an event on a host link may include the list of NPIV devices behind the host port that are zoned with the target port. The target device uses the PNL to determine the impacted devices behind the host port and to decide what action to take for those devices.

A.3.3 FPIN Generation and Transmission

A.3.3.1 Overview

If the Fabric detects an event on a link between the Fabric and a device, then the Fabric sends the FPIN ELS to those devices zoned with the Attached Port_Name and to those devices zoned with the N_Port_Names in the Port Name List. The Fabric only delivers the FPIN ELS to devices registered to receive the FPIN ELS.

If the device detects an event on a link between the device and the Fabric, then the device sets the Port Name List field to the N_Port_Names affected by the event and sends the FPIN ELS to the Fabric. The Fabric sends an FPIN ELS to those devices zoned with the Detecting Port_Name and to those devices zoned with the N_Port_Names in the Port Name List. The Fabric only delivers the FPIN ELS to devices registered to receive the FPIN ELS.

If the Fabric detects an event on a link between Fabric ports (i.e., E_Ports), then the Fabric sends the FPIN ELS to those devices zoned with the N_Port_Names in the Port Name List. The Fabric only delivers the FPIN ELS to devices registered to receive the FPIN ELS.

A.3.3.2 Port name list examples

A.3.3.2.1 Basic port name list example

In this example (see figure 8), N_Port A and N_Port X are registered to receive FPIN notifications and are zoned with each other. The example illustrates the Fabric behavior when sending an FPIN ELS in response to an event that occurs on the link between F_Port 0 and N_Port A and detected by the Fabric:

- a) The Fabric sets the Detecting Port Name field to F_Port 0 and the Attached Port Name field to N_Port A;
- b) The Fabric sets the Port Name List field to a list containing N_Port X for the FPIN ELS sent to N_Port A; and,
- c) The Fabric sets the Port Name List field to a list containing N_Port A for the FPIN ELS sent to N_Port X.

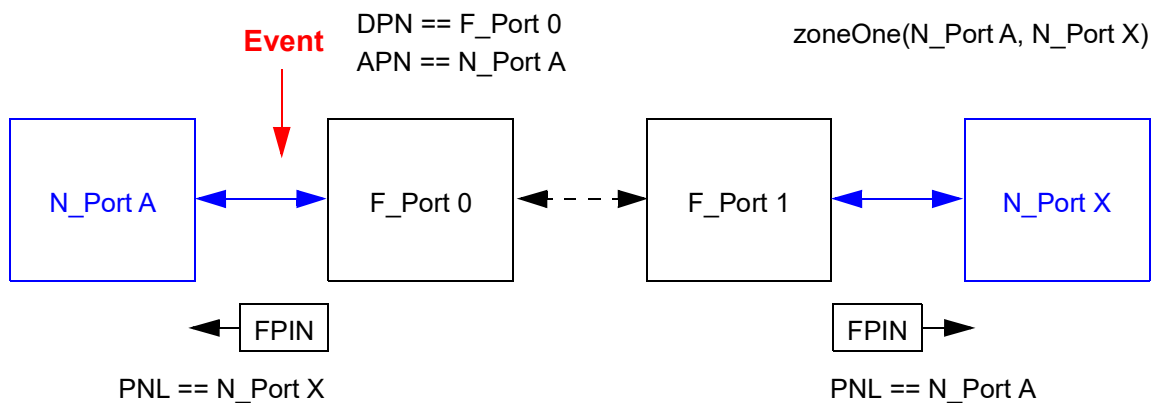


Figure 8 – Basic port name list example

A.3.3.2.2 Typical port name list example

In this example (see figure 9), N_Port A, N_Port B, N_Port X, and N_Port Y are registered to receive FPIN notifications. In addition, N_Port A is zoned with N_Port X and N_Port Y, and N_Port B is zoned with N_Port X and N_Port Y. The example illustrates the Fabric behavior when sending an FPIN ELS in response to an event that occurs on the link between F_Port 0 and N_Port A:

- a) The Fabric sets the Detecting Port Name field to F_Port 0 and the Attached Port Name field to N_Port A;
- b) The Fabric sets the Port Name List field to a list containing N_Port X, and N_Port Y for the FPIN ELS sent to N_Port A;
- c) The Fabric sets the Port Name List field to a list containing N_Port A for the FPIN ELS sent to N_Port X; and
- d) The Fabric sets the Port Name List field to a list containing N_Port A for the FPIN ELS sent to N_Port Y.

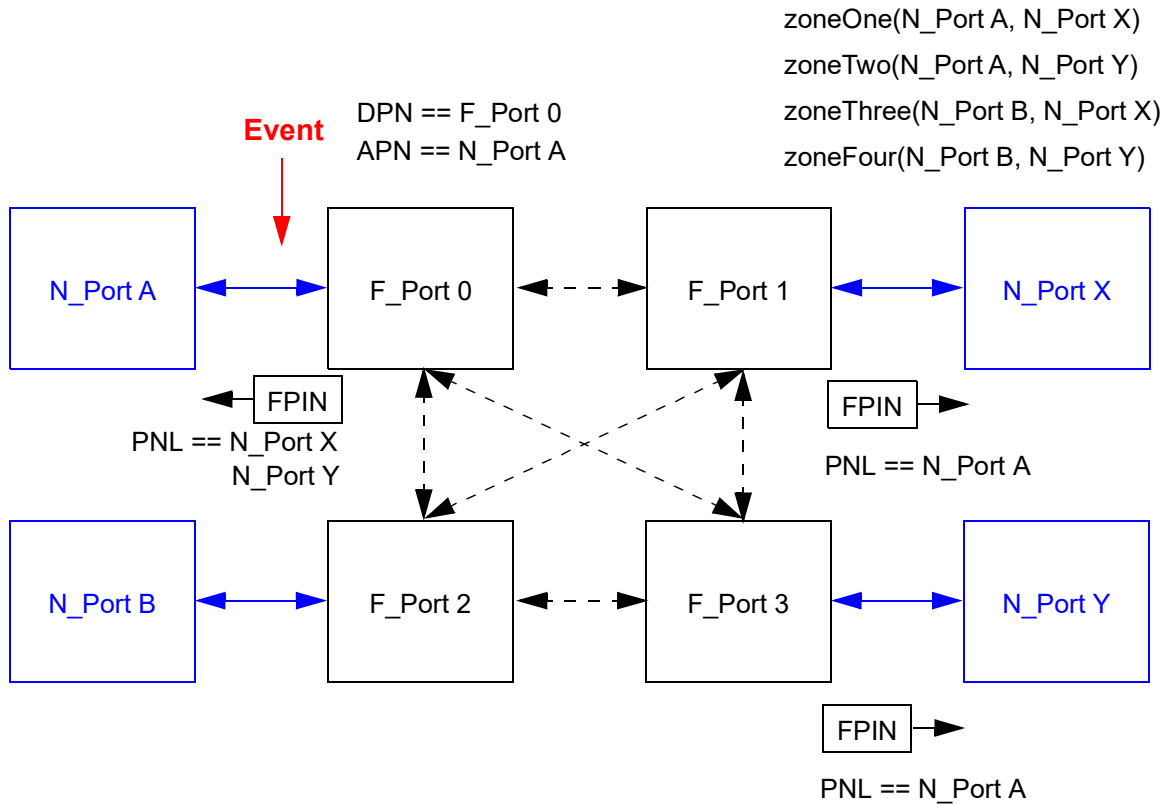


Figure 9 – Typical port name list example

A.3.3.2.3 NPIV port name list example

In this example (see figure 10), N_Port A, NPIV a, NPIV b, and N_Port X are registered to receive FPIN notifications. In addition, N_Port X is zoned with NPIV a and NPIV b. The example illustrates the Fabric behavior when sending an FPIN ELS in response to an event that occurs on the link between F_Port 0 and N_Port A:

- a) The Fabric sets the Detecting Port Name field of an FPIN ELS to F_Port 0 and the Attached Port Name field to N_Port A;
- b) The Fabric sets the Port Name List field to the empty list for the FPIN ELS sent to N_Port A;
- c) The Fabric sets the Port Name List field to a list containing N_Port X for the FPIN ELS sent to NPIV a;
- d) The Fabric sets the Port Name List field to a list containing N_Port X for the FPIN ELS sent to NPIV b; and,
- e) The Fabric sets the Port Name List field to a list containing NPIV a and NPIV b for the FPIN ELS sent to N_Port X.

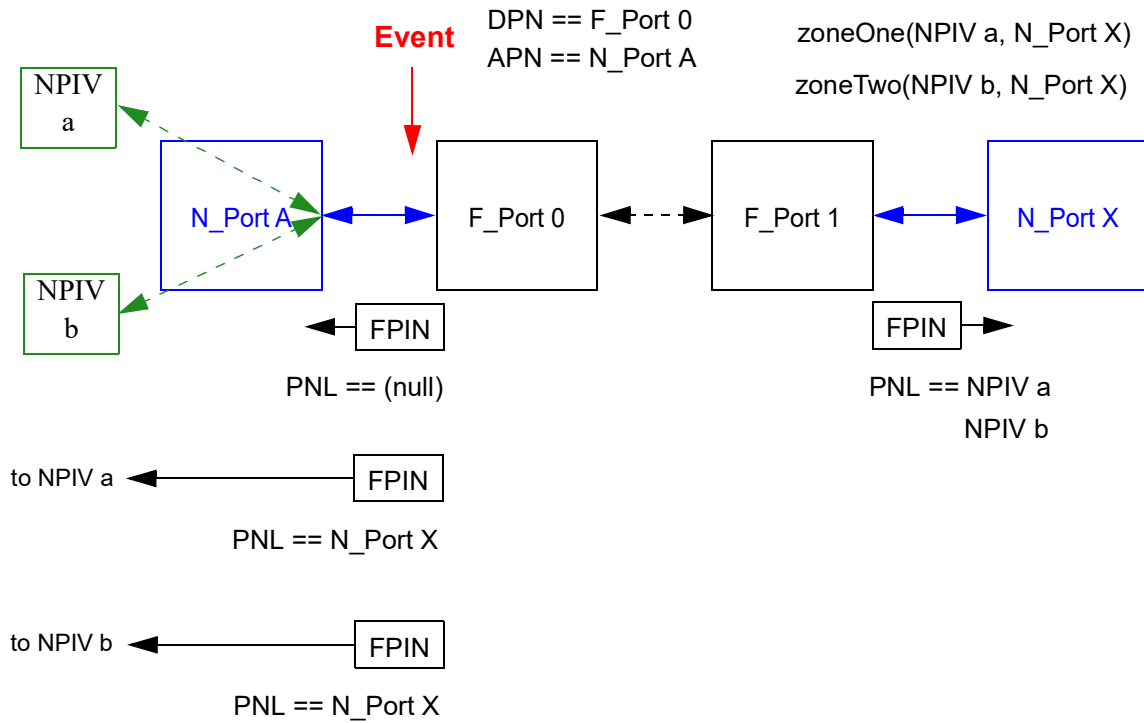


Figure 10 – NPIV port name list example

A.3.3.2.4 ISL port name list example

In this example (see figure 11), N_Port A, N_Port B, N_Port X, and N_Port Y are registered to receive FPIN notifications. In addition, N_Port A is zoned with N_Port X and N_Port Y, and N_Port B is zoned with N_Port X. The example illustrates the Fabric behavior when sending an FPIN ELS in response to an event that occurs on the link between E_Port W and E_Port Z:

- a) The Fabric sets the Detecting Port Name field of an FPIN ELS to E_Port Z and the Attached Port Name field to E_Port W;
- b) The Fabric sets the Port Name List field to a list containing N_Port X and N_Port Y for the FPIN ELS sent to N_Port A;
- c) The Fabric sets the Port Name List field to a list containing N_Port X for the FPIN ELS sent to N_Port B;
- d) The Fabric sets the Port Name List field to a list containing N_Port A and N_Port B for the FPIN ELS sent to N_Port X; and,
- e) The Fabric sets the Port Name List field to a list containing N_Port A for the FPIN ELS sent to N_Port Y.

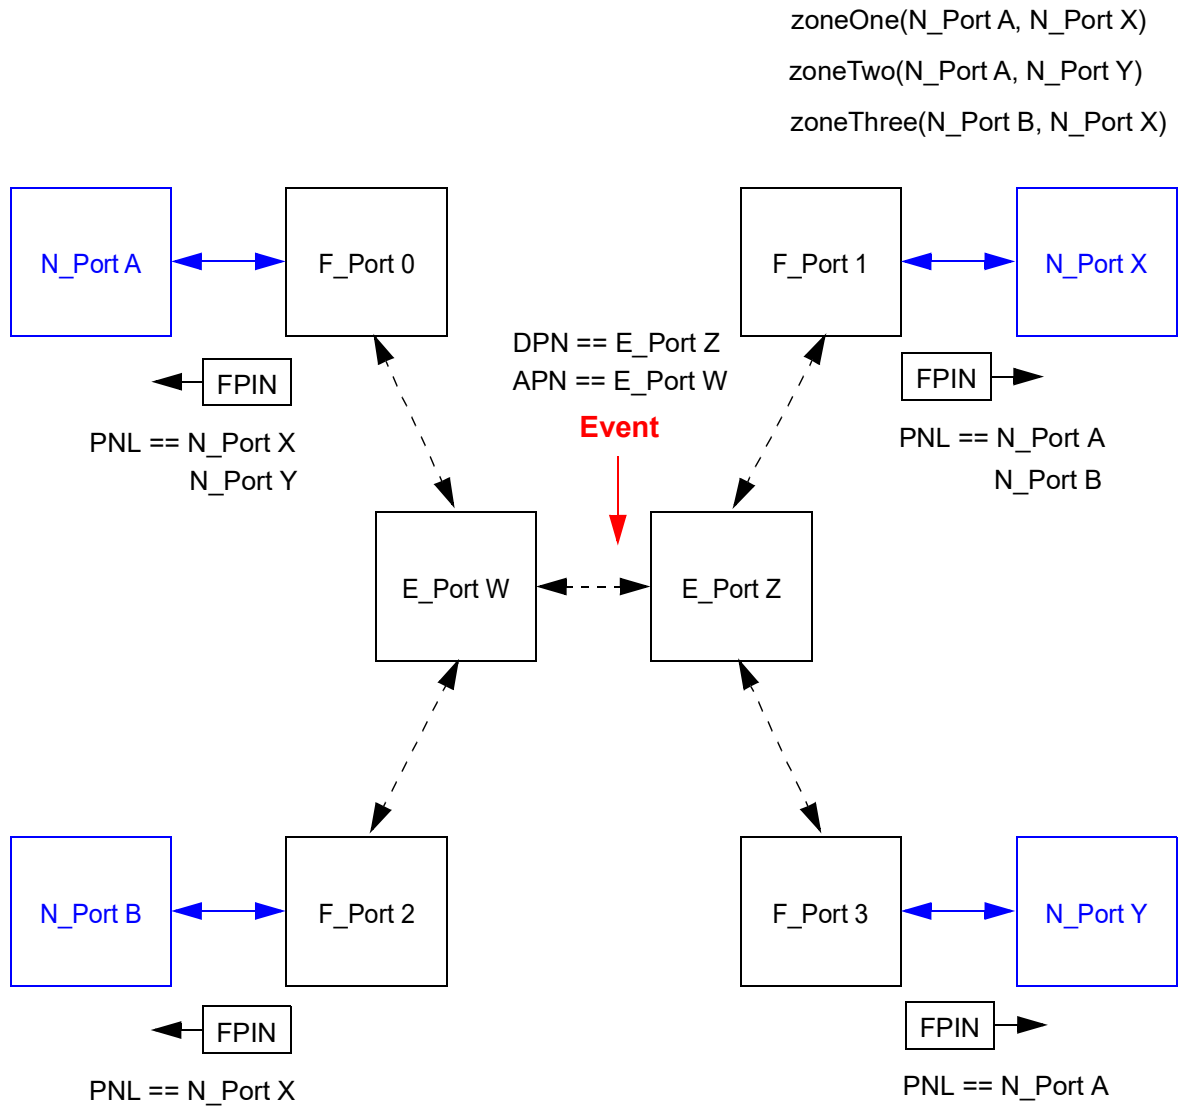


Figure 11 – ISL port name list example

A.3.4 Detecting ports and attached ports

Generally, the FPIN ELS descriptors identify the location of the detected event in the Fabric. The Detecting Port Name field is set to the Port_Name of the FC_Port that detected the event and the Attached Port Name field is set to the Port_Name of the FC_Port attached to the FC_Port that detected the event. If these fields are not in the descriptor (e.g., Congestion Notification descriptor), the event occurred at the device port (i.e., the local Nx_Port to Fx_Port connection).

If the Fabric detects an event on a link between the Fabric and a device, then the attached Port_Name is known by the devices receiving the FPIN ELS (i.e., they are zoned with the attached Port_Name). If the device detects an event on a link between the Fabric and the device, then the detecting Port_Name is known by the devices receiving the FPIN ELS (i.e., they are zoned with the

detecting Port_Name). If the Fabric detects an event on a link between Fabric ports, then the detecting and attached Port_Names are not known by the devices receiving the FPIN ELS (i.e., the event occurred on an ISL port, which is not a member of a zone).

NOTE 24 – If the Fabric or the device detects an event on a link with NPIV connections, then the devices receiving the FPIN ELS may not be zoned with the Attached Port_Name for events detected by the Fabric or with the Detecting Port_Name for events detected by the device.

A.3.5 Congestion events

A.3.5.1 Overview

An implementation may use both Congestion notifications and Congestion Signals to detect congestion events and to distinguish between the types and duration of congestion conditions. Using both mechanisms provides a method of optimizing the mitigation responses to congestion in order to maximize throughput while minimizing Fabric impact.

A.3.5.2 Congestion Signal detection and persistence

Congestion Signals indicate the presence of congestion between two directly attached FC_Ports (see FC-FS-6). The detecting FC_Port monitors for resource consumption due to congestion and transmits a Congestion Signal if the resource consumption is above a threshold. Thus, Congestion Signals provide a near instantaneous indication of the presence of congestion. Congestion Signals are particularly useful in identifying congestion due to credit stall when transmission of ELSs is impacted.

An implementation may use Congestion Signals as a means of detecting instantaneous congestion. Implementations should monitor the reception of Congestion Signals over time to determine if the congestion condition represents sustained or intermittent congestion.

A.3.5.3 Congestion notification detection and persistence

Congestion notifications indicate the presence of congestion to both the congested FC_Port (i.e., Congestion notification, see 4.3.54.7.5) and the FC_Ports zoned with the congested FC_Port (i.e., Peer congestion notification, see 4.3.54.7.4). The detecting FC_Port monitors for congestion conditions and generates the congestion notifications if the congestion condition exceeds the Event Threshold. Thus, congestion notifications are an indication that congestion existed for the period indicated by the Event Threshold.

Upon detection of a congestion event, the detecting port initiates the generation of FPIN ELSs to the attached port and to the peers of the attached port (i.e., the FC_Ports in the zone membership list(s) associated with the PN_Port attached to the detecting FC_Port). The FPIN ELS sent to the attached port contains a Congestion notification descriptor and the FPIN ELS sent to the peers of the attached port contains a Peer Congestion notification descriptor. While the congestion condition persists, the detecting port generates the FPIN ELSs at a rate specified in the Event Threshold field of the respective descriptors. The congestion condition clears when the detecting port generates an FPIN ELS with an Event Type of "Congestion cleared" or does not generate a subsequent FPIN ELS for two times the Event Threshold.

An implementation may use Congestion notifications as a means to detecting sustained congestion. Implementations should monitor the reception of Congestion notifications over time to determine when the congestion condition subsides.

A.3.5.4 Congestion notification Event Type evaluation

The Event Type reported in the FPIN ELS provides supplemental information and an implementation may update the Event Type in subsequent FPIN ELSs. For example, an implementation may initially report the congestion condition as “Oversubscription” then later change the condition to “Lost credit” for the same congestion event. Thus, an FPIN ELS with an Event Type of “Congestion cleared” received by the device indicates the attached port is no longer congested regardless of the previous values in the Event Type field.

A.4 Device participation

A.4.1 Overview

Devices respond to notifications and signals based on the implementation characteristics and limitations of the device; therefore, participation in the Fabric notifications operations may vary widely and still provide significant improvements in Fabric operation. Participation may range from basic participation (e.g., registration only), logging, or behavior modifications. The participating Fabric may provide a method for recording the notifications and signals sent to registered devices (e.g., logging). This subclause (i.e., A.4) provides a set of guidelines based on various levels of participation.

A.4.2 Registration only

Registration only participation in Fabric notifications consists of registering for each of the notification descriptors of the FPIN ELS and registering to receive signals as described in the EDC ELS. The device provides no other functions than to receive the notifications and signals.

This level of participation in Fabric notifications improves problem determination capabilities for the Fabric by including the added information of the distributed notifications and signals.

A.4.3 Logging

In addition to the registration only behavior described in A.4.2, the device records reception of notifications and records a count of signals received using device dependent methods (e.g., recording if one or more signals are received over some time interval).

This level of participation in Fabric notifications improves problem determination as described in A.4.2 and it provides a means to correlate the Fabric logs with the device logs to further isolate events.

A.4.4 Device behavior modifications

A.4.4.1 Overview

The device participating at this level modifies its behavior in response to the received notifications or signals (see examples in A.4.4.2 and A.5). Changes in device behavior are executed cautiously as small adjustments to the normal operation of the device and are limited to minimize the impact on the device applications. The objective is to reduce the impact on the Fabric and neighboring devices.

This level of participation in Fabric notifications improves problem determination as described in clause A.4.3 and automatically adjusts to changing characteristics in the Fabric.

NOTE 25 – The examples provide guidelines for participating devices with the understanding that non-linear, real-world solutions require specific analysis and evaluation unique to the device implementation.

A.4.4.2 Device behavior modification examples

A.4.4.2.1 Congestion actions

A.4.4.2.1.1 Congestion signal actions

Congestion signals indicate a port's resources are consumed above defined thresholds. This condition is detected by the transmitter and signifies a congestion condition (see FC-FS-6). Congestion signals are delivered at the rate indicated by the Congestions Signal Capability descriptor in the EDC ELS (see 4.3.52) and persist while the resource consumption condition persists above defined thresholds.

A device may modify frame processing behaviors if it is receiving congestion signals. The device should monitor the signals to determine if the congestion condition persists long enough to be affected by the device given the device's reaction time. The device should reduce the rate of requests for data or reduce the size of the data requests commensurate with the persistence and severity of the congestion signals. The reduction in the rate of requests for data should be in small percentages. After each reduction, the device should monitor the subsequent congestion signals before making further adjustments. The overall reduction should be limited to a level necessary to support the device operations. If the congestion signals subside, the device may increase the frame processing request rate and monitor for increases in subsequent congestion signals.

A.4.4.2.1.2 Congestion notification actions

Congestion notifications indicate if a port is processing frames slower than the rate the frames are arriving at that port. Peer Congestion notifications indicate to a peer port (i.e., a port zoned with the congested port) that the congested port is processing frames slower than they are arriving. The Congestion and Peer Congestion notifications persist while the congestion condition persists and are delivered at the rate indicated by the value in the Event Threshold field in the Congestion and Peer Congestion descriptors. The congestion condition is cleared if the Congestion and Peer Congestion notifications are not received after two times the value in the Event Threshold field after the last notification, or if a Congestion or Peer Congestion notification is received with an Event Type field set to Clear.

A device may modify frame processing behaviors if it is receiving Congestion notifications. The device should monitor the notifications to determine if the congestion condition persists long enough to be affected by the device given the device's reaction time. The device should reduce the rate of requests for data or reduce the size of the data requests (e.g., transfer ready or read commands) commensurate with the severity of the congestion indicated in the notification. The reduction in the rate of requests for data should be in small percentages. After each reduction, the device should monitor the subsequent Congestion notifications before making further adjustments. The overall reduction should be limited to a level necessary to support the device operations. If the Congestion notifications subside, the device may increase the rate of requests for data or the size of the data requests and monitor for increases in subsequent Congestion notifications.

A.4.4.2.2 Link Integrity notification actions

Link Integrity notifications indicate a link to a device, a link to a zoned device or a network path within the Fabric (e.g., an ISL) has encountered a specified type of error. The device response to these notifications depends on the type of link integrity event detected, the availability of an alternate port on the same device, the role of the device (i.e., initiator or target) or the phase of the device connection (e.g., during active frame processing). The device may choose to record the event with no additional changes in behavior.

If a device acting as an initiator receives a link integrity event on a port, then the initiator port may record the event as a component of the path status. If an alternative path is available, then the initiator port may choose to favor the alternative path until a transition of the impacted path is detected (i.e., an RSCN indicating the path state has transitioned). If an alternative path is not available, then the initiator port may choose to recover a Sequence or Exchange sooner, or to increase the number of retries before declaring the path unavailable.

If a device acting as a target receives a link integrity event on a port, then the target port may record the event as a component of the port status. If an alternative port providing access to the same target port resources is available, then the target port may choose to disable the port until a transition of the port is detected (i.e., an offline/online transition for a local port or an RSCN indicating the remote port state has transitioned). If an alternative port is not available, then the target port may choose to recover a Sequence or Exchange sooner, or to increase the number of retries before determining the port is unavailable.

If a device detects a condition affecting the reliability of the link, then the device may send a link integrity event to the attached Fabric Controller for distribution to the device's peers.

A.4.4.2.3 Delivery notification actions

Delivery notifications may indicate a command frame was be delivered by the Fabric or processed by a device. If a port receives a Delivery notification, then the port may circumvent the command retry timeout process and instigate actions that lead to faster Sequence or Exchange recovery. Similar actions may be taken for response, status, or last frame in Sequence frames depending on the timeout characteristics and recovery behaviors of the ULP.

A.5 Examples

A.5.1 Overview

This subclause provides examples of participation in Fabric Notifications.

A.5.2 Improving multi-path solutions

A.5.2.1 Overview

Multi-path solutions are effective at recovering from transient errors through retry or logical reset operations that are relatively non-disruptive to the system. They are also effective at mitigating persistent physical errors by declaring the path associated with the physical error as inactive or marginal and selecting alternate paths that are still healthy. Multi-path solutions have difficulty mitigating the effects of physical errors that are persistently intermittent because they are unable to tell if the impact was due to a transient error or a physical error.

A.5.2.2 Fabric notifications actions

If persistently intermittent error events occur in the Fabric, Fabric notifications provide the ability for multi-path solutions to receive explicit notification of these events by registering for the Link Integrity notification. If the Fabric or end device detects the occurrence of the intermittent physical issue, then link integrity notifications are sent to the registered devices affected by the event. The link integrity event provides the multi-path solution with the location and cause of the physical error allowing the multi-path solution to "route around" the affected path by selecting good, alternate paths.

A.5.3 Eliminating Target Credit Stall

A.5.3.1 Overview

Target Credit Stall is a special case of the credit-stall condition and may be the result of internal architectural limitations, internal resource constraints, or processing conditions. A Target Credit Stall condition may occur if unsolicited command frames (e.g., ULP commands, ELSs, etc) arrive at the target faster than the ability to process them. Unsolicited frames are typically queued for processing, which consumes internal resources. The consumption of internal resources may cause the process to impact the adapter resources, which in turn, causes the process to consume Fabric resources and produce the Target Credit Stall phenomena.

Mitigating the problem of Target Credit Stall is implementation dependent and may include design modifications to minimize the occurrences of the condition. Modifications may include increasing the available resources for processing unsolicited frames, elevating the priority of processing unsolicited frames to handle them more quickly, or providing offload resources to process frames later. In addition, an implementation may use Fabric notifications to supplement traditional pacing techniques as described below.

A.5.3.2 Queue full and device busy actions

Some ULPs define pacing methods that indicate if target resources are consumed and provide a technique for causing initiators to hold off additional requests (e.g., SCSI TASK SET FULL or BUSY, Control Unit Busy, etc). With this method, initiators reduce the demand on the target, and then gradually restore it to previous levels if no further pacing occurs.

A.5.3.3 Fabric notifications actions

Fabric notifications provides the ability for a target port to indicate if it encounters internal resource constraints leading to credit-stall behavior. If the target port detects a condition indicating the internal resources are becoming consumed (i.e., the queuing of unsolicited frames has reached a high water mark), then that target port may send a Peer Congestion notification with a Resource Contention Event Modifier, which may result in the reduction or suspension of the rate of incoming unsolicited frames (see 4.3.54.7.4). The target port continues to send the Peer Congestion notification at a frequency indicated in the Event Threshold field of the Peer Congestion descriptor while consumption of internal resources persists.

If the target port determines the consumption of internal resources has subsided, then the target port discontinues sending the Peer Congestion notification or sends the Peer Congestion notification with an Event Type "Congestion cleared" to indicate the target port's internal resources are available.

A.5.3.4 Frame discard TOV actions

If the target port continues to experience internal resource consumption leading to the Target Credit Stall behavior and the target port supports the Frame Discard Timeout (F_D_TOV), For example, , then that target port may discard frames that have been longer than the F_D_TOV.

If the target port continues to experience internal resource consumption leading to the Target Credit Stall behavior and the target port supports the Frame Discard Timeout (F_D_TOV), then the target port may discard frames that have been queued longer than the F_D_TOV (see FC-FS-6). This action frees the internal resources allowing the target port to return buffer credit to the Fabric, which results in reducing backpressure on the Fabric. For example, if a target port is unable to queue more unsolicited frames, then that target port may discard frames that have been queued longer than the F_D_TOV.

A.5.4 Speed matching to minimize oversubscription

A.5.4.1 Overview

Speed mismatches between sending and receiving FC_Ports may be a cause of oversubscription. If oversubscription is detected, then the sending port may evaluate the receiving port's port speed (see GPSC in FC-GS-9, RDP (see 4.3.49), and RPSC (see 4.3.36)) and employ pacing to send data at the receiving port's port speed.

A.5.4.2 Fabric notifications actions

If the sending port receives the Peer Congestion notification, then that port may evaluate the receiving port's port speed to determine if it is lower than the sending port's port speed. If a speed mismatch exists (i.e., the sending port's port speed is greater than the receiving port's port speed), then the sending port may adjust the rate data is sent to match the receiving port's port speed.

If the sending port continues to receive Peer Congestion notifications, then it is likely the receiving port is congested for other reasons (e.g., over utilized, see A.3.4). In this case, the sending port may take additional actions to alleviate the congestion condition (see A.4.4.2.1.2).

A.5.5 Handling first burst congestion spikes

A.5.5.1 Overview

The first burst function provides an optimized method for transmitting immediate data from an initiator to a target for protocols that support the function (see FC-SB-6 and FC-NVMe-2). Devices that support the first burst function may encounter oversubscription conditions due to resource contention for the available first burst data buffers if multiple initiators are sending first burst operations or if a single initiator is sending multiple first burst operations. Once the target first burst resources are consumed, the process backs up into the adapter resources and once the adapter resources are consumed, the process backs up into the Fabric producing either an oversubscription condition or a credit-stall condition similar to Target Credit Stall (see A.5.3).

Solutions that provide mitigating capabilities for first burst oversubscription provide mechanisms to discard the first burst data and indicate to the initiator to use transfer ready mode (see FC-SB-6). Other solutions are able to utilize congestion signals and Fabric notifications to mitigate the first burst oversubscription condition.

A.5.5.2 Device actions

If the target port detects congestion signals or Congestion notifications and determines the consumption of internal first burst resources is above a threshold, then that target port may send a Peer Congestion notification with a Resource Contention Event Modifier which may result in the reduction or suspension of the rate of first burst operations (see 4.3.54.7.4). The target port continues to send the Peer Congestion notification while the internal first burst resources are consumed at a frequency indicated in the Event Threshold field of the Peer Congestion descriptor. If the target port determines the consumption of internal first burst resources is below a threshold, then the target port discontinues sending the Peer Congestion notification or sends the Peer Congestion notification with an Event Type "Congestion cleared" to indicate the target port's internal first burst resources are available.