

TECHNICAL REPORT

Information technology –
Part 372: Fibre channel methodologies for interconnects-2 (FC-MI-2)

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-372:2011



THIS PUBLICATION IS COPYRIGHT PROTECTED
Copyright © 2011 ISO/IEC, Geneva, Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either IEC or IEC's member National Committee in the country of the requester.

If you have any questions about ISO/IEC copyright or have an enquiry about obtaining additional rights to this publication, please contact the address below or your local IEC member National Committee for further information.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland
Email: inmail@iec.ch
Web: www.iec.ch

About the IEC

The International Electrotechnical Commission (IEC) is the leading global organization that prepares and publishes International Standards for all electrical, electronic and related technologies.

About IEC publications

The technical content of IEC publications is kept under constant review by the IEC. Please make sure that you have the latest edition, a corrigenda or an amendment might have been published.

- Catalogue of IEC publications: www.iec.ch/searchpub

The IEC on-line Catalogue enables you to search by a variety of criteria (reference number, text, technical committee,...). It also gives information on projects, withdrawn and replaced publications.

- IEC Just Published: www.iec.ch/online_news/justpub

Stay up to date on all new IEC publications. Just Published details twice a month all new publications released. Available on-line and also by email.

- Electropedia: www.electropedia.org

The world's leading online dictionary of electronic and electrical terms containing more than 20 000 terms and definitions in English and French, with equivalent terms in additional languages. Also known as the International Electrotechnical Vocabulary online.

- Customer Service Centre: www.iec.ch/webstore/custserv

If you wish to give us your feedback on this publication or need further assistance, please visit the Customer Service Centre FAQ or contact us:

Email: csc@iec.ch
Tel.: +41 22 919 02 11
Fax: +41 22 919 03 00



TECHNICAL REPORT

**Information technology –
Part 372: Fibre channel methodologies for interconnects-2 (FC-MI-2)**

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

PRICE CODE



CONTENTS

FOREWORD	5
INTRODUCTION	7
1 Scope	8
2 Normative references	8
3 Terms, definitions, abbreviations and other conventions	9
3.1 Overview	9
3.2 Terms and definitions	9
3.3 Editorial conventions	14
3.4 Abbreviations and acronyms	15
3.5 Symbols	16
3.6 Keywords	16
3.7 Applicability and use of this Technical Report	18
3.8 Feature Set table terms, definitions, and abbreviations	18
3.8.1 Overview	18
3.8.2 Feature Set table terms and definitions	19
3.8.3 Feature Set table abbreviations	19
3.9 Feature testing compliance	20
3.10 Timing Constraints	20
4 Structure and concepts	21
5 Loop behaviors	21
5.1 Loop Initialization	21
5.1.1 Power On Behavior	21
5.1.2 Loop Failure	21
5.1.3 Initialization at Power-on	22
5.1.4 FL_Port Time-out during Initialization	22
5.1.5 LIP Generation	22
5.1.6 Response to LIP	22
5.1.7 Origination of LISM Frames	22
5.1.8 Forwarding of LISM frames	22
5.1.9 Address Selection	23
5.1.10 Multi-port Initialization	23
5.1.11 AL_PA Position Map Support	23
5.1.12 Availability after LIP	23
5.2 Post Initialization	24
5.2.1 LIP Generation	24
5.2.2 FL_Port Unfair behavior and Transfer behavior	24
5.2.3 Responses to OPN	24
5.2.4 No response to OPN	25
6 Fx_Port Behaviors	26
6.1 Fx_Port Login responses	26
6.1.1 FLOGI ACC features and parameters for Fx_Ports	26
6.1.2 Nx_Port PLOGI	26
6.1.3 Fx_Port Common Service Parameters (FLOGI ACC)	27
6.1.4 Fx_Port Class 2 Service Parameters (FLOGI ACC)	29

6.1.5 Fx_Port Class 3 Service Parameters (FLOGI ACC)	30
6.2 Link Services	32
6.2.1 Basic Link Services	32
6.2.2 ELS requirements for well-known addresses	32
6.2.3 Extended Link Service Replies	32
6.3 FC-AL-2 features for FL_Ports	33
6.4 Loop Fabric Address	35
7 Fabric Behaviors	36
7.1 Overview	36
7.2 Switch-to-Switch Requirements	37
7.2.1 Overview	37
7.2.2 Switch Port Types	40
7.2.3 Exchange Link Parameters (ELP)	40
7.2.4 Principal Switch Selection	42
7.2.5 Fabric Shortest Path First (FSPF)	42
7.2.6 Distributed Services	43
7.2.7 Zoning	48
7.2.8 Distributed Event Notification	50
7.2.9 Additional Switch Requirements	50
7.3 Fabric Service Requirements	51
7.3.1 Overview	51
7.3.2 Name Server	51
7.3.3 Fabric Configuration Server	55
7.3.4 Fabric Zone Server	56
7.3.5 Registered State Change Notification	58
7.4 Domain Controller and Well-Known Addresses	60
7.4.1 Domain controller and well-known address support requirements for fabrics	60
7.4.2 N_Port domain controller and well-known address usage	61
7.4.3 Domain controller and well-known address (WKA) ELS login and address assignment requirements	63
8 Discovery and Management	67
8.1 Overview	67
8.2 Management of Interconnect Components	67
8.2.1 Overview	67
8.2.2 Switch	67
8.2.3 Managed Hub	68
8.2.4 Gateway/Bridge	68
9 Conformance Environments	69
Bibliography	70

Tables

Table 1 ISO/IEC and American Conventions	15
Table 2 Feature Set table terms and definitions	19
Table 3 Feature Set table key abbreviations	19
Table 4 Feature testing compliance relationship to definitions	20
Table 5 FLOGI ACC features and parameters for Fx_Ports	26
Table 6 Fx_Port Common Service Parameters (FLOGI ACC)	27
Table 7 Class 2 Service Parameters (FLOGI ACC)	29
Table 8 Class 3 Service Parameters (FLOGI ACC)	30
Table 9 Extended Link Service Replies	32
Table 10 FC-AL features for FL_Ports	33
Table 11 FC-SW-3 Fabric Support Summary	37
Table 12 Switch Port Type Usage	40
Table 13 ELP SW_ILS Parameters	40
Table 14 Flow Control Parameters	42
Table 15 FC-SW-3 Defined Name Server Requests	43
Table 16 FC-GS-4 Defined Name Server Requests	43
Table 17 FC-GS-4 Common Requests	45
Table 18 Name Server Object Usage Summary	46
Table 19 FC-SW-3 Defined Management Server Requests	46
Table 20 FC-GS-4 Defined Management Server Requests	47
Table 21 Zoning Support	48
Table 22 FC-GS-4 Fabric Support	51
Table 23 Name Server Request Support	51
Table 24 Common Request Support	53
Table 25 Fabric Configuration Server Request Support	55
Table 26 Fabric Zone Server Basic Zoning Request Support	56
Table 27 Fabric Zone Server Enhanced Zoning Request Support	57
Table 28 Domain Controller and Well-Known Address Support Requirements for Fabrics	60
Table 29 ELS requirements for domain controller and well-known addresses	61
Table 30 Broadcast WKA (FFFFFFh) ELS login requirements	63
Table 31 Fabric F_Port WKA (FFFFFFEh) ELS address assignment requirements	64
Table 32 Fabric Controller WKA (FFFFFFDh) ELS login requirements	65
Table 33 Directory Service WKA (FFFFFFCh) ELS login requirements	65
Table 34 Management Service WKA (FFFFFFAh) ELS login requirements	66
Table 35 Domain Controller (FFFCxxh) ELS login requirements	66
Table 36 Loop Fabric Address ELS login requirements	67
Table 37 Switch Support Summary	67
Table 38 Conformance Environments	69

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 372: Fibre channel methodologies for interconnects-2 (FC-MI-2)

FOREWORD

- 1) ISO (International Organization for Standardization) and IEC (International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards. Their preparation is entrusted to technical committees; any ISO and IEC member body interested in the subject dealt with may participate in this preparatory work. International governmental and non-governmental organizations liaising with ISO and IEC also participate in this preparation.
- 2) In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.
- 3) The formal decisions or agreements of IEC and ISO on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC and ISO member bodies.
- 4) IEC, ISO and ISO/IEC publications have the form of recommendations for international use and are accepted by IEC and ISO member bodies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC, ISO and ISO/IEC publications is accurate, IEC or ISO cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 5) In order to promote international uniformity, IEC and ISO member bodies undertake to apply IEC, ISO and ISO/IEC publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any ISO/IEC publication and the corresponding national or regional publication should be clearly indicated in the latter.
- 6) ISO and IEC provide no marking procedure to indicate their approval and cannot be rendered responsible for any equipment declared to be in conformity with an ISO/IEC publication.
- 7) All users should ensure that they have the latest edition of this publication.
- 8) No liability shall attach to IEC or ISO or its directors, employees, servants or agents including individual experts and members of their technical committees and IEC or ISO member bodies for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication of, use of, or reliance upon, this ISO/IEC publication or any other IEC, ISO or ISO/IEC publications.
- 9) Attention is drawn to the normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 10) Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 14165-372, which is a technical report, has been prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

The list of all currently available parts of the ISO/IEC 14165 series, under the general title *Information technology – Fibre Channel*, can be found on the IEC web site.

This Technical Report has been approved by vote of the member bodies and the voting results may be obtained from the address given on the second title page.

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-372:2011

INTRODUCTION

This technical report specifies common methodologies for both Arbitrated Loop and Switched environments. The goal of this technical report is to facilitate interoperability between devices whether they are connected in a loop or Fabric topology.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-372:2011

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 372: Fibre channel methodologies for interconnects-2 (FC-MI-2)

1 Scope

This part of ISO/IEC 14165 is intended to document interoperability behavior for Fabric elements (i.e., E_Port, F_Port, FL_Port). This Technical Report includes a wide range of issues such as link initialization, error detection, error recovery, Fabric operation, management capabilities, and zoning.

This Technical Report serves as an implementation guide, whose primary objective is to maximize the likelihood of interoperability between conforming implementations. This Technical Report prohibits or requires some features that are in the referenced documents.

A second objective of this Technical Report is to simplify implementations and their associated documentation, testing, and support requirements. As a result there may be some optional features of the referenced documents that are not mutually exclusive, but are prohibited or required for the purpose of this simplification. Features that some but not all of the referenced documents require for compliance may be optional in this report. Each specification of such an optional feature in this report identifies the referenced document for which the feature is required.

Internal characteristics of conformant implementations are not defined by this Technical Report, but it incorporates features from the documents cited in clause 2.

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the of the referenced document (including any amendments) applies.

The provision of the referenced specification other than ISO/IEC, IEC, ISO and ITU documents, as identified in this clause, are valid within the context of this Technical Report. The reference of such a specification within this Technical Report does not give it any further status within ISO/IEC. In particular, it does not give the referenced specification the status of an International Standard.

ISO/IEC 14165-115, *Information technology – Fibre channel – Part 115: Physical Interfaces (FC-PI)* ²

ISO/IEC 14165-122, *Information technology – Fibre Channel – Part 122: Arbitrated loop-2 (FC-AL-2)* ³

ISO/IEC 14165-133, *Information technology – Fibre channel – Part 133: Fibre Channel Switch Fabric-3 (FC-SW-3)* ⁴

T11/Project 1620D, *Fibre Channel - Link Services (FC-LS)* ⁵

INCITS TR-36-2003, *Fibre Channel - Device Attach (FC-DA)* ⁶

2. ANSI INCITS 352-2002, *Fibre Channel - Physical Interfaces (FC-PI)*

3. ANSI INCITS 332-1999, *Fibre Channel - Arbitrated Loop (FC-AL-2)*

4. ANSI INCITS 384-2004, *Fibre Channel - Switch Fabric - 3 (FC-SW-3)*

5. ISO/IEC 14165-261, *Information technology – Fibre Channel – Part 261: Link services (FC-LS)* (under consideration)

6. ISO/IEC 14165-341, *Information technology – Fibre Channel – Part 341: Device attach (FC-DA)* (under consideration)

ISO/IEC 14165-414, Information technology – Fibre Channel – Part 414: Generic services-4 (FC-GS-4) ⁷

ISO/IEC 14165-431, *Information technology – Fibre Channel – Part 431: Security Protocols (FC-SP)* (to be published) ⁸

Internet Engineering Task Force RFC 791, Internet Protocol, September 1981

Internet Engineering Task Force RFC 2373, IP Version 6 Addressing Architecture, July 1998

Internet Engineering Task Force RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, December 1998

Internet Engineering Task Force RFC 3410, Introduction and Applicability Statements for Internet Standard Management Framework, December 2002

Internet Engineering Task Force RFC 4044, Fibre Channel Management MIB, May 2005

3 Terms, definitions, abbreviations and other conventions

3.1 Overview

For the purpose of this document the following definitions, conventions, abbreviations, acronyms, and symbols apply.

3.2 Terms and definitions

3.2.1

address identifier

address value used to identify source (S_ID) or destination (D_ID) of a frame

3.2.2

AL_PA bit map

bit map that shows which ports are present on an Arbitrated Loop

NOTE See FC-AL-2.

3.2.3

AL_PA position map

map that shows port ordering on an Arbitrated Loop

NOTE See FC-AL-2.

3.2.4

Arbitrated Loop time out value

AL_Time

time interval defined by FC-AL-2

3.2.5

Application

entity that makes requests of a Service

3.2.6

Arbitrated Loop Physical Address

AL_PA

one-byte address value

NOTE See FC-AL-2.

7. ANSI INCITS 387-2004, *Fibre Channel - Generic Services - 4 (FC-GS-4)*

8. T11/Project 1570D, *Fibre Channel - Security Protocols (FC-SP)*

3.2.7

AR0

special AR (see 3.2.9) containing only the switch backbone network, that in general may consist of point-to-point links, broadcast links, and switching devices

NOTE See FC-AL-2.

3.2.8

Area Identifier

second or middle level of the three-level addressing hierarchy, found in bits 15 through 8 of an address identifier

NOTE (see 3.2.1) and FC-SW-3.

3.2.9

Autonomous Region

AR

one or more Fibre Channel address domains consisting of switches that all run a common routing protocol

NOTE See FC-SW-3.

3.2.10

B_Port

Interconnect_Port used to connect bridge devices with E_Ports on a switch

NOTE See FC-SW-3.

3.2.11

Domain Identifier

highest level of the three-level addressing hierarchy, found in bits 23 through 16 of an address identifier (see 3.2.1)

NOTE See FC-SW-3.

3.2.12

E_Port

Fabric Expansion port that attaches to another Interconnect_Port to create an Inter-Switch Link

NOTE See FC-SW-3.

3.2.13

Entry Switch

role that a switch assumes with respect to a distributed service request

NOTE The switch that is attached to an Nx_Port making a service request assumes the role of an entry switch with respect to that request. See FC-SW-3.

3.2.14

Error_Detect_Timeout value

E_D_TOV

time interval defined in FC-FS-2

3.2.15

F_Port

port through which non-loop N_Ports are attached to a Fabric, and it does not include FL_Ports

3.2.16

Fabric

Fibre Channel frame transport infrastructure that includes switches and interconnects various Nx_Ports attached to it

3.2.17**Fabric F_Port**

entity at the well-known address FFFFFFFh

NOTE See FC-SW-3.

3.2.18**Fabric_Name**

identifier associated with a Fibre Channel Fabric

NOTE See FC-FS-2.

3.2.19**Fabric Shortest Path First****FSPF**

link state protocol used for Path Selection

NOTE See FC-SW-3.

3.2.20**FC_Port**

port transmitting or receiving Fibre Channel frames

NOTE FC_Port includes any port type defined in the Fibre Channel family of standards (e.g., N_Port, F_Port, FL_Port, E_Port).

3.2.21**FL_Port**

L_Port that is able to perform the function of an F_Port, attached via a link to one or more NL_Ports in an Arbitrated Loop topology (see FC-AL -2)

NOTE The AL_PA of an FL_Port is 00h. In this Technical Report, an FL_Port is assumed to always refer to a port to which NL_Ports are attached to a Fabric; it does not include F_Ports.

3.2.22**Fx_Port**

Switch Port (see 3.2.52) capable of operating as an F_Port or FL_Port

3.2.23**Gateway/Bridge**

any device that interfaces FC to some other interface (see FC-FS-2)

NOTE The definition of such a device is outside the scope of this report.

3.2.24**Hub**

interconnect element used with Arbitrated Loops

NOTE A Hub may be a stand-alone enclosure, or part of an integrated system (e.g., such as the connection point within a JBOD). Typically a Hub is transparent to any connected Fibre Channel devices.

3.2.25**Interconnect_Port**

E_Port or a B_Port

3.2.26**Loop Fabric Address****LFA**

address identifier used to address an FL_Port (see FC-SW-3) for the purpose of loop management

NOTE See 6.4 and FC-LS.

3.2.27**L_Port**

port that contains Arbitrated Loop functions associated with the Arbitrated Loop topology

NOTE See FC-AL-2.

3.2.28

Managed Hub

Hub (see 3.2.24) that provides either in-band or out-of-band management functions

3.2.29

Name Server

Server (see 3.2.48) that allows registration and reporting of various objects

NOTE See FC-GS-4.

3.2.30

N_Port

direct Fabric-attached port that does not include NL_Ports

NOTE See FC-FS-2.

3.2.31

N_Port Identifier

1 Fabric unique address identifier by which an N_Port is known

NOTE The identifier is used in the S_ID and D_ID fields of a frame (see FC-FS-2).

2 Name Server object (see FC-GS-4) containing the Fibre Channel address identifier assigned to an N_Port or NL_Port

3.2.32

NL_Port

L_Port that is able to perform the function of an N_Port, attached to one or more NL_Ports and zero or more FL_Ports in an Arbitrated Loop topology

NOTE In this technical report, an NL_Port is assumed to always refer to a loop-attached port including both Private NL_Ports and Public NL_Ports, and does not include N_Ports.

3.2.33

Node_Name

identifier associated with a Fibre Channel node

NOTE See FC-FS-2.

3.2.34

Non-Participating mode

operational mode of an L_Port that does not have an AL_PA, but is enabled into the Loop

NOTE See FC-AL-2.

3.2.35

Nx_Port

port operating as an N_Port or NL_Port

3.2.36

OLD-PORT

state on a set of ports where two devices operate in a point-to-point mode utilizing FC-FS-2 protocols instead of FC-AL-2 protocols

NOTE See FC-AL-2 and FC-FS-2.

3.2.37

OPEN_INIT

state in the LPSM as defined in FC-AL-2

3.2.38**N_Port_Name**

identifier associated with an FC_Port

NOTE See FC-FS-2.

3.2.39**Partial Response**

response from a Distributed Service that may not have a complete set of data

NOTE See FC-SW-3.

3.2.40**Participating mode**

operational mode of an L_Port that has an AL_PA and is enabled into the Loop

NOTE See FC-AL-2.

3.2.41**Platform**

association of one or more nodes for the purpose of discovery and management

3.2.42**Port**

N_Port, NL_Port, F_Port, FL_Port, B_Port, or E_Port, used in a context where the distinction between specific port types is clarified by other text or not significant

3.2.43**Port Identifier**

lowest level of the three-level addressing hierarchy, found in bits 7 through 0 of an address identifier (see 3.2.1)

NOTE See FC-SW-3.

3.2.44**Private NL_Port**

NL_Port that is observing the rules of private loop behavior

NOTE See FC-AL-2 and FC-DA.

3.2.45**Public NL_Port**

NL_Port that attempts a fabric login and may transfer frames through the FL_Port

NOTE A Public NL_Port observes the rules of public loop behavior and/or private loop behavior. See FC-AL-2 and FC-DA.

3.2.46**Receiver_Transmitter_Timeout value****R_T_TOV**

time interval defined in FC-FS-2

3.2.47**Resource_Allocation_Timeout value****R_A_TOV**

time interval defined in FC-FS-2

3.2.48**Server**

entity that accepts CT requests and provides CT responses

NOTE 1 A Server is accessed via a Service (e.g., the Name Server is accessed using the Directory Service).

NOTE 2 See FC-GS-4.

3.2.49

Service

service provided by a node, accessible via an N_Port that is addressed by a Well-Known Address (e.g., the Directory Service and the Alias Service)

NOTE A Service provides access to one or more Servers, see FC-GS-4.

3.2.50

Speed Negotiation

process that allows an FC_Port capable of multiple operating speeds and connected to another FC_Port that may or may not have the same capability to arrive at the optimum speed of operation

NOTE See FC-FS-2.

3.2.51

Switch

element that makes up the switching portion of the Fabric

NOTE See FC-SW-3.

3.2.52

Switch Port

E_Port, B_Port, F_Port, or FL_Port

3.2.53

Symbolic Name

user-defined name for an object, up to 255 characters in length

NOTE The Directory Service (see FC-GS-4) does not guarantee uniqueness of its value.

3.2.54

Well-Known Address

WKA

address identifier defined to access a Service

NOTE See FC-FS-2.

3.2.55

Zone

group of Zone Members

NOTE Members of a Zone are made aware of each other, but not made aware of Zone Members outside the Zone. See FC-GS-4 and FC-SW-3.

3.2.56

Zone Member

specification of a device to be included in a Zone

NOTE See FC-GS-4 and FC-SW-3.

3.2.57

Zone Set

set of Zones that are used in combination

NOTE See FC-GS-4 and FC-SW-3.

3.3 Editorial conventions

In FC-MI-2, a number of conditions, mechanisms, sequences, parameters, events, states, or similar terms are printed with the first letter of each word in uppercase and the rest lowercase (e.g., Exchange, Class). Any lowercase uses of these words have the normal technical English meanings.

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

The ISO/IEC convention of numbering is used (i.e., the thousands and higher multiples are separated by a space and a comma is used as the decimal point.) A comparison of the American and ISO conventions are shown in table 1.

Table 1 – ISO/IEC and American Conventions

ISO/IEC	American
0,6	0.6
1 000	1,000
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 clauses.

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

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

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

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

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

3.4 Abbreviations and acronyms

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

ABTS	Abort Sequence
ACC	Accept
ADISC	Discover Address
AL_PA	Arbitrated Loop Physical Address, see FC-AL-2
AL_PD	Arbitrated Loop Destination Address, see FC-AL-2
AL_Time	Arbitrated Loop Time out value, see FC-AL-2
AR	Autonomous Region, see FC-SW-3
CLS	Close Primitive Signal, see FC-AL-2
CT	Common Transport, see FC-GS-4
CT_IU	Common Transport Information Unit, see FC-GS-4
D_ID	Destination address identifier
E_D_TOV	Error_Detect_Timeout value
EFP	Exchange Fabric Parameters, see FC-SW-3
ELP	Exchange Link Parameters, see FC-SW-3
ELS	Extended Link Service, see FC-SW-3
FAN	Fabric Address Notification Extended Link Service
FLOGI	F_Port Login
FSPF	Fabric Shortest Path First, see FC-SW-3
GS	Fibre Channel Generic Service

HBA	Host Bus Adapter
IDLE	Idle Word
IP	Internet Protocol, see RFC 791, RFC 2373, and RFC 2460
ISL	Inter-Switch Link, see FC-SW-3
IU	Information Unit
LFA	Loop Fabric Address
LIFA	Loop Initialization Fabric Assigned, see FC-AL-2
LINIT	Loop Initialization Request, see FC-LS
LIP	Loop Initialization Primitive Sequence, see FC-AL-2
LIRR	Link Incident Report Registration, see FC-LS
LISM	Loop Initialization Select Master, see FC-AL-2
LS_ACC	Link Service Accept
LSTS	Loop Status, see FC-LS
LOGO	Logout
LPSM	Loop Port State Machine, see FC-AL-2
LP_TOV	Loop Time out Value, see FC-AL-2
LS_RJT	Link Service Reject
NPIV	N_Port_ID Virtualization, see FC-LS
OPN	Open Primitive Signal, see FC-AL-2
PLOGI	N_Port Login
RLIR	Registered Link-Incident Record
RLS	Read Link Error Status Block
RNID	Request Node Identification Data Extended Link Service
RPS	Read Port Status block
RSCN	Registered State Change Notification
R_T_TOV	Receiver_Transmitter_Timeout value
S_ID	Source address identifier
SAN	Storage Area Network
SNMP	Simple Network Management Protocol, see RFC 3410
SW_ILS	Switch Internal Link Service, see FC-SW-3
UDP	User Datagram Protocol, see RFC 768
ULP	Upper Level Protocol
WAN	Wide Area Network
WKA	Well-Known Address

3.5 Symbols

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

|| concatenation

3.6 Keywords

expected

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

ignored

A keyword used to describe an unused bit, byte, word, field or code value. The content- 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.

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.

mandatory

A keyword indicating an item that is required to be implemented as defined in this report.

may

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

may not

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

obsolete

A keyword indicating that an item was defined in prior Fibre Channel standards but has been removed from a subsequent Fibre Channel standard.

optional

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

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. Recipients are not required to check reserved bits, bytes, words or fields for zero values. In defined fields, receipt of reserved code values shall be reported as an error.

restricted

A keyword referring to bits, bytes, words, and fields that are set aside for use in other 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 Technical Report.

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 report.

should

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

3.7 Applicability and use of this Technical Report

This Technical Report specifies which features shall be used (i.e., required) and which features shall not be used (i.e., prohibited) by interoperating compliant Fibre Channel implementations. Use of some features is optional (i.e., allowed). The use of such functions is either negotiated in a fixed and standard manner or the availability of the functions for use shall be determined in a standard manner.

The relationship between use as specified in this Technical Report and support as implemented by a product is subtle. If this Technical Report specifies that a feature shall be used, then a compliant implementation shall support it. In some cases, this Technical Report is asymmetric: to ensure interoperability when an optional feature is used, this Technical Report requires support for the infrastructure required to use the feature without specifying that the feature be used to conform to this Technical Report.

The requirements of this Technical Report are a proper subset of the various relevant standards. This Technical Report prohibits use of many features and options in these standards. Use of prohibited features may prevent interoperability with Fibre Channel devices complying to this Technical Report. This Technical Report does not prohibit implementation of features, only their use. Functions that are mandatory in the appropriate base standard are assumed to be implemented. Implementations may support features whose use is prohibited by this Technical Report and such prohibited features may be required for compliance with the relevant standards or other Technical Reports.

3.8 Feature Set table terms, definitions, and abbreviations

3.8.1 Overview

Features in this Technical Report are summarized in the form of Feature Set tables. These tables indicate whether the feature is Required, Prohibited, Invocable, or Allowed for compliance with this Technical Report; or whether a parameter is Required to be a particular value for compliance with this Technical Report. Features or parameters that are not listed do not affect interoperability.

In several tables within this Technical Report, there are references to notes associated with the table. These notes are normative and are mandatory requirements of this Technical Report.

3.8.2 Feature Set table terms and definitions

Terms and definitions are used to define usage of reference features or options provided by the applicable standards are described in table 2.

Table 2 – Feature Set table terms and definitions

Term	Definition
Prohibited	A feature that shall not be used between Fibre Channel devices compliant with this Technical Report. An implementation may use the feature to communicate with non-compliant implementations. This Technical Report does not prohibit the implementation of features, only their use between Fibre Channel devices compliant with this Technical Report. Use of a prohibited feature may prevent interoperability with Fibre Channel devices complying to this Technical Report.
Required	A feature or parameter value that shall be used between all Fibre Channel devices compliant with this Technical Report. Fibre Channel devices compliant with this Technical Report are required to implement the feature. An implementation may use the feature to communicate with non-compliant implementations. If a Fibre Channel device does not implement a required feature that device may not be interoperable with Fibre Channel devices complying to this Technical Report.
Allowed	A feature or parameter value that may be used between Fibre Channel devices compliant with this Technical Report. Fibre Channel devices compliant with this Technical Report are not required to implement the feature, but if they do, the feature shall be implemented as described in the applicable standard. The potential user of a feature may determine if the recipient supports that feature via a Required discovery process or a minimal response by the recipient.
Invocable	A feature or parameter that is required to be implemented by a device to which a request may be sent, but it is not required to be used by a requesting device.

3.8.3 Feature Set table abbreviations

Table 3 contains the key of table abbreviations used within this Technical Report.

Table 3 – Feature Set table key abbreviations

Key	Definition
P	Prohibited
R	Required
A	Allowed
I	Invocable
X	This parameter has no required value; any value is allowed.
-	This parameter is ignored or this feature is not applicable.

3.9 Feature testing compliance

While not a compliance vehicle in itself, this Technical Report should be used as the basis for compliance testing. As such, the relationship for the definitions of prohibited, required, allowed, and invocable (see 3.8) to compliance testing is described in table 4.

Table 4 – Feature testing compliance relationship to definitions

Term	Definition
Prohibited	A feature that is not allowed to be used in a compliant implementation. A compliance test shall only verify that the prohibited feature is not requested by a device. A compliance test shall not generate a request for the prohibited feature, since if it does so, the response is outside the scope of this Technical Report.
Required	A feature that is required to be implemented by a compliant device. A compliance test is allowed to verify that the feature is correctly requested, and responded to as specified in this Technical Report or other referenced standards and technical reports.
Allowed	If a device claims to implement the feature, the feature may be tested. If tested, the feature shall conform to the discovery mechanism and function as specified in this Technical Report and other referenced standards and technical reports. If a device claims to not implement a feature, the feature shall not be tested.
Invocable	A feature that is required to be implemented by a device to which a request may be directed, but it is not required to be used by a requesting device. A compliance test is allowed to verify that the request is correctly responded to as specified in this Technical Report or other referenced standards and technical reports.

3.10 Timing constraints

All timings defined in this Technical Report are meant to limit the amount of time a device takes to accomplish a task. These timings shall be measured on an unloaded system. A heavily loaded system may exhibit timings in excess of those specified. The term unloaded means that system activity is induced only in direct invocation of the feature under test.

4 Structure and concepts

This Technical Report defines an interoperability environment in which a compliant device adheres to certain sets of behavior (FC_Port Behavior) that allow for interoperability across a wide variety of topologies and management methods. This environment encompasses three areas.

- a) Management Behavior: The set of behaviors required to create an interoperable management environment.
- b) Loop Behavior: The set of behaviors required to create an interoperable Arbitrated Loop environment.
- c) Fabric Behavior: The set of behaviors required to create an interoperable Fabric environment.

FC_Port Behavior is defined as Loop Behavior, Fabric Behavior, and Management Behavior taken together.

All behaviors are interrelated (e.g., a Fabric may be required not only to meet Fabric Behavior but Loop and Management Behavior as well, or an adapter may be required to meet both Loop and Management Behavior). The level of behavior support in any particular device depends on the supported Management, Loop, FC_Port and Fabric functions.

A device may be non-compliant with this Technical Report, and still be compliant with the various referenced Fibre Channel Standards. However, devices not compliant with this Technical Report may not be interoperable.

5 Loop behaviors

5.1 Loop Initialization

5.1.1 Power On Behavior

During power on, after the transmitter has been enabled and 200 μ s of valid Fibre Channel signal at or above FC-PI minimum amplitude levels has been driven by the transmitter, an FL_Port that is not executing Speed Negotiation shall not disrupt the operation of the loop for more than 95 ms and shall begin forwarding transmission words or begin initialization within 95 ms.

NOTE 1 - The delay of 200 μ s of enforcement of limits on disruption recognizes that many HBAs and devices by design disrupt during power-on, and that most hubs isolate an attached device or HBA until it has presented at least 200 μ s of an FC-PI-compliant signal. The expectation is that any disruption caused by a device or HBA at power-on is harmless until after it has presented 200 μ s of valid signal, because until then it is isolated from the loop by its hub.

FC-FS-2 specifies the disruption of a loop that may result from attaching to it an FL_Port configured to execute Speed Negotiation. An FL_Port that is executing Speed Negotiation disrupts up to three times over a period not exceeding t_{disrupt2} . No period of disruption exceeds t_{disrupt1} . See FC-FS-2 for definitions and derivations of t_{disrupt1} and t_{disrupt2} .

5.1.2 Loop failure

An FL_Port shall not issue a LIP for loop failure based on loss of synchronization before the loss of synchronization exceeds R_T_TOV .

5.1.3 Initialization at power on

FL_Ports shall request only one loop initialization at power on, unless the FL_Port attempts to enter OLD-PORT state, after which only one additional loop initialization may be requested.

An FL_Port that is executing Speed Negotiation may disrupt the loop causing the loop to initialize more than once, however, having completed Speed Negotiation as evidenced by the FL_Port originating or passing through LISM frames, such an FL_Port is subject to the limits of this subclause.

5.1.4 FL_Port time-out during initialization

It may be necessary for an FL_Port to request loop initialization multiple times. An FC-AL FL_Port may request initialization again using the following rules:

- a) if loop initialization has not completed within 2 s;
- b) the second time the device shall increase the initialization completion timeout to greater than 24 s;
- c) the third and subsequent times the device shall increase the initialization completion timeout to greater than 128 s.

FC-AL-2 FL_Ports shall follow the FC-AL-2 initialization timeout requirements. If there is an upper level timer running on the initialization process the upper level shall follow the rules above for FC-AL FL_Ports.

5.1.5 LIP generation

An FL_Port requesting loop initialization shall generate a minimum of 12 LIPs.

5.1.6 Response to LIP

An FL_Port shall recognize and forward 12 of the LIPs received within 5 ms, unless it is already in OPEN_INIT and ignoring LIPs for AL_Time.

An FL_Port that is executing Speed Negotiation may issue continuous LIPs of type unrelated to the type it receives, however, having completed Speed Negotiation as evidenced by the FL_Port originating or passing through LISM frames, such an FL_Port is subject to the limits of this subclause.

5.1.7 Origination of LISM Frames

An FL_Port that has not received a higher priority LISM shall originate LISM frames with a maximum of 5 ms between LISM frames. FC-AL FL_Ports may source IDLEs for 15 ms before they begin sourcing LISMs.

5.1.8 Forwarding of LISM frames

An FL_Port that has received a higher priority LISM than what it is currently transmitting shall forward the last LISM recognized, provided only higher priority LISMs are received, with a delay of no more than the maximum of either

- a) 25 ms, or
- b) 2 times the delay between received LISM frames.

Subsequent LISM frames shall be forwarded, provided only higher priority LISMs are received, with a delay of no more than the maximum of either

- a) 5 ms, or
- b) 2 times the delay between received LISM frames.

5.1.9 Address selection

After any of the following, the FL_Port shall attempt to take address 00h during the LIFA phase, as follows:

- a) a power cycle or equivalent reset of the FL_Port;
- b) receipt of any LIP, or
- c) when the FL_Port transitions from Non-Participating to Participating.

5.1.10 Multi-port initialization

In order to minimize disruptions, a multi-port FL_Port device has the following initialization requirements:

- a) a LIP on one FL_Port shall not cause a LIP on any of the other FL_Ports, and
- b) initialization on one FL_Port shall not cause any other FL_Port to lose any frames, state, or otherwise impact traffic on its loop, except data being transferred from the initializing FL_Port to another FL_Port.

5.1.11 AL_PA position map support

AL_PA position map support has the following rules:

- a) AL_PA position map support shall be provided by all FC-AL-2 FL_Ports;
- b) devices that do not support the AL_PA position map may be present on the loop, therefore FL_Ports shall not rely on AL_PA position map support for functionality; and
- c) if all devices on a loop support AL_PA position map, then AL_PA position map shall be used to determine what devices are available. Attempting to open devices that were not in the last AL_PA position map is prohibited.

5.1.12 Availability after LIP

For FL_Ports, the following shall apply after a LIP:

- a) an FL_Port that has taken an address during a loop initialization shall be capable of processing all supported Fabric services upon completion of loop initialization; and
- b) after an FL_Port has taken an address, unless it goes Non-Participating for any reason, it shall take an address in all subsequent loop initializations.

5.2 Post initialization

5.2.1 LIP generation

5.2.1.1 Proper LIP generation

Once an FL_Port has initialized, an FL_Port should minimize the generation of LIPs. Situations where it is permitted, but not required for an FL_Port to generate a LIP are the following:

- a) unable to win arbitration for greater than LP_TOV;
- b) loss of word sync for greater than R_T_TOV;
- c) loss of signal is detected;
- d) a CLS in response to CLS has not been received within LP_TOV; or
- e) an FL_Port receives an OPN that it sent to an NL_Port.

5.2.1.2 Improper LIP generation

An FL_Port shall not generate a LIP in response to the following:

- a) link error counter overrun;
- b) CRC error detection;
- c) reception of an unexpected frame;
- d) reception of a CLS in response to an OPN; and
- e) no R_RDY or CLS received in response to an OPN.

5.2.2 FL_Port Unfair behavior and Transfer behavior

The FL_Port delivers frames by arbitrating for the loop, opening the destination NL_Port, transmitting the frames, and closing the loop. As allowed by FC-AL-2, the FL_Port may arbitrate unfairly (i.e., the FL_Port may enter the ARBITRATING state at any time without regard to access fairness).

Once the FL_Port has acquired control of the loop, it may retain that control to deliver frames to more than one destination NL_Port by using the TRANSFER state. When the FL_Port transmits and receives a CLS with a destination NL_Port, it may maintain control of the loop and open another destination NL_Port. The FL_Port should weigh this control against starvation of other ports on the loop. In the case of the FAN ELS, the FL_Port is required to deliver frames to all attached Public NL_Ports using the TRANSFER state with access set true.

NOTE 2 – One possible arbitration strategy is to retain control of the loop until an NL_Port arbitrates, then release the loop after completing the current loop tenancy.

5.2.3 Responses to OPN

If an FL_Port sends an OPN, and receives that same OPN, the NL_Port addressed by the OPN is no longer participating on the loop. The FL_Port shall perform an implicit logout of the NL_Port. The FL_Port may perform a Loop Initialization.

The reception of a CLS in response to an OPN is an indication that an NL_Port is busy. This is not an error condition, and shall not cause error recovery to occur. If the FL_Port is operating with an NL_Port with a Login BB_Credit of zero, the NL_Port may be busy, and the FL_Port shall forward the CLS, and retry the OPN. Because the busy condition is handled by the procedure in this subclause, the FL_Port shall not originate a LIP to clear this condition. The normal frame timeout limits shall be observed if the NL_Port is busy for sufficient time to cause a frame timeout to occur. If the busy condition persists for greater than LP_TOV, the FL_Port may originate a LIP to initialize the loop.

5.2.4 No response to OPN

The lack of any response to an OPN can indicate a number of failures (e.g., the OPN was corrupted on the loop, or the R_RDY or CLS from the device was corrupted). The proper response to this condition is for the originator of the OPN to originate a CLS and after the CLS is received by the originator, when the arbitration window is again available arbitrate and attempt to OPN this device. Because this error condition is cleared by this procedure, the originator of the OPN shall not originate a LIP to clear this condition.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-372:2011

6 Fx_Port behaviors

6.1 Fx_Port Login responses

6.1.1 FLOGI ACC features and parameters for Fx_Ports

Table 5 lists features and parameters for FLOGI ACC with usage defined by this Technical Report.

Table 5 – FLOGI ACC features and parameters for Fx_Ports

Feature/Parameter	Fx_Port Responder
FLOGI ACC S_ID	
N_Port	FFFFFFEh
NL_Port	FFFFFFEh
FLOGI ACC D_ID	
N_Port	ddaaXX ^a
NL_Port	ddaa AL_PA ^b
Class of service for FLOGI^c	
Class 2	A
Class 3	R
Supported Classes of Service	
Class 1	P
Class 2	A
Class 3	R
Class 4 and up	P
<p>^a The fabric shall assign the N_Port_ID according to the rules in FC-LS. Multiple N_Ports may be assigned addresses with the same aa.</p> <p>^b The fabric shall build the assigned ID by using the requestor's Port_ID (AL_PA) for the low order byte (i.e., the AL_PA was received in the FLOGI request as the low order byte of the S_ID). The fabric shall assign the Domain_ID (dd) and Area_ID (aa) to a valid value (see FC-SW-3). The value of 'ddaa' shall be the same for all Public NL_Ports on a given loop. The fabric shall return the value of the low order byte of the S_ID of the FLOGI request in the low order byte of the D_ID.</p> <p>^c FLOGI ACC shall be returned in the same class of service in which the FLOGI was received, if that class of service is supported by the fabric. If the class of service of the request is not supported, an LS_RJT shall be returned in the class of service of the request.</p>	

6.1.2 Nx_Port PLOGI

Fabric services in a switch shall respond to PLOGIs as specified in FC-DA.

6.1.3 Fx_Port Common Service Parameters (FLOGI ACC)

Table 6 lists Fx_Port Common Service Parameters for FLOGI ACC with usage defined by this Technical Report. The parameters are valid for both Class 2 and Class 3 delivery service.

Table 6 – Fx_Port Common Service Parameters (FLOGI ACC) (part 1 of 2)

Common Service Parameter	Fx_Port responder
FC-LS	
FC-PH Version	2020h ^a
Buffer-to-Buffer Credit (minimum for FL_Port)	0 ^b
Buffer-to-Buffer Credit (minimum for F_Port)	1
Common features	
Valid Vendor Version Level	0
N_Port/F_Port	1
BB_Credit Management	X ^c
Buffer-to-Buffer Receive Data Field Size (minimum)	256 bytes
Clean Address bit	X ^d
Multiple N_Port_ID Assignment	X ^g
E_D_TOV Resolution	0 ^e
Multicast supported by Fabric	X
Broadcast supported by Fabric	X ^h
Hunt Group routing supported by Fabric	X
Query Data Buffer conditions	0
<p>^a Although FC-LS has obsoleted the version numbers, this field shall be set as specified.</p> <p>^b An FL_Port is required to support a value of zero or greater for this parameter.</p> <p>^c For FL_Ports operating in loop mode this value shall be one, for F_Ports this value shall be zero.</p> <p>^d The Clean Address bit functionality shall be supported.</p> <p>^e These shall be used as default values. A compliant implementation may change these values, but it is up to the implementor to ensure interoperability.</p> <p>^f These values may need to be modified when used in a WAN environment.</p> <p>^g The Multiple N_Port_ID Assignment bit shall be set as defined in FC-LS.</p> <p>^h For Fabrics that support RFC 2625 and/or RFC 3831, this bit shall be set to one.</p> <p>ⁱ If the payload bit is set to one in a FLOGI request an FC-MI-2 compliant device may respond with the payload bit set to one or reject the request. FC-DA compliant devices do not set the payload bit to one.</p> <p>^j If BB_Credit Recovery is supported, then BB_SC_N shall be set to 8. If BB_Credit Recovery is not supported, then BB_SC_N shall be set to zero.</p>	

Table 6 – Fx_Port Common Service Parameters (FLOGI ACC) (part 2 of 2)

Common Service Parameter	Fx_Port responder
Security bit (See FC-SP)	X
Virtual Fabrics bit	X
Clock Synchronization Primitive Capable	X
R_T_TOV Value	0
Dynamic Half Duplex Supported	X
Payload bit	X ⁱ
BB_SC_N	X ^j
R_A_TOV	10 000 ms ^{e, f}
E_D_TOV	2 000 ms ^{e, f}
<p>^a Although FC-LS has obsoleted the version numbers, this field shall be set as specified.</p> <p>^b An FL_Port is required to support a value of zero or greater for this parameter.</p> <p>^c For FL_Ports operating in loop mode this value shall be one, for F_Ports this value shall be zero.</p> <p>^d The Clean Address bit functionality shall be supported.</p> <p>^e These shall be used as default values. A compliant implementation may change these values, but it is up to the implementor to ensure interoperability.</p> <p>^f These values may need to be modified when used in a WAN environment.</p> <p>^g The Multiple N_Port_ID Assignment bit shall be set as defined in FC-LS.</p> <p>^h For Fabrics that support RFC 2625 and/or RFC 3831, this bit shall be set to one.</p> <p>ⁱ If the payload bit is set to one in a FLOGI request an FC-MI-2 compliant device may respond with the payload bit set to one or reject the request. FC-DA compliant devices do not set the payload bit to one.</p> <p>^j If BB_Credit Recovery is supported, then BB_SC_N shall be set to 8. If BB_Credit Recovery is not supported, then BB_SC_N shall be set to zero.</p>	

6.1.4 Fx_Port Class 2 Service Parameters (FLOGI ACC)

If Class 2 is supported, Table 7 lists Class 2 Service Parameters for FLOGI ACC response with usage defined by this Technical Report.

Table 7 – Class 2 Service Parameters (FLOGI ACC) (part 1 of 2)

Class 2 Service Parameter	Value
Class validity	1
Service options	
Intermix Mode	-
Stacked Connect Requests	-
Sequential Delivery	1
Dedicated Simplex	-
Camp-On ^a	-
Buffered Class 1	-
Priority	X ^b
Preference	X ^b
DiffServ QoS	0
Initiator control	
X_ID reassignment	-
Initial Responder Process_Associator	-
ACK_0 capable	-
ACK_N capable ^a	-
ACK generation assistance	-
Data compression capable ^a	-
Data compression History buffer size ^a	-
Data encryption capable ^a	-
Clock synchronization capable	0
Recipient Control	
ACK_0 capable	-
ACK_N capable ^a	-
X_ID interlock	-
^a This feature has been obsoleted in FC-LS, and may be reclaimed for use in a future standard. ^b This bit shall only be set to one if the entire fabric supports the feature.	

Table 7 – Class 2 Service Parameters (FLOGI ACC) (part 2 of 2)

Class 2 Service Parameter	Value
Error Policy Supported	-
Categories per Sequence	-
Data compression capable ^a	-
Data compression History buffer size ^a	-
Data decryption capable ^a	-
Clock synchronization capable	A
Reserved - Fabric specific	0
Receive data field size (minimum)	-
Concurrent Sequences (minimum)	-
N_Port End-to-end Credit (minimum)	-
Open Sequences per Exchange (minimum)	-
^a This feature has been obsoleted in FC-LS, and may be reclaimed for use in a future standard. ^b This bit shall only be set to one if the entire fabric supports the feature.	

6.1.5 Fx_Port Class 3 Service Parameters (FLOGI ACC)

Table 8 lists Class 3 Service Parameters for FLOGI ACC response with usage defined by this Technical Report.

Table 8 – Class 3 Service Parameters (FLOGI ACC) (part 1 of 2)

Class 3 Service Parameter	Value
Class validity	1
Service Options	
Intermix Mode	-
Stacked Connect Requests	-
Sequential Delivery	1
Dedicated Simplex	-
ACK_N capable ^a	-
Buffered Class 1	-
Priority	X ^b
^a This feature has been obsoleted in FC-LS, and may be reclaimed for use in a future standard. ^b This bit shall only be set to one if the entire fabric supports the feature.	

Table 8 – Class 3 Service Parameters (FLOGI ACC) (part 2 of 2)

Class 3 Service Parameter	Value
Preference	X ^b
DiffServ QoS	0
Initiator Control	
X_ID reassignment	-
Initial Responder Process_Associator	-
ACK_0 capable	-
ACK_N capable ^a	-
ACK generation assistance	-
Data compression capable ^a	-
Data compression History buffer size ^a	-
Data decryption capable ^a	-
Clock synchronization capable	0
Recipient control	
ACK_0 capable	-
ACK_N capable ^a	-
X_ID interlock	-
Error Policy Supported	-
Categories per Sequence	-
Data compression capable ^a	-
Data compression History buffer size ^a	-
Data decryption capable ^a	-
Clock synchronization capable	A
Reserved - Fabric specific	0
Receive data field size (minimum)	-
Concurrent Sequences (minimum)	-
N_Port End-to-end Credit	-
Open Sequences per Exchange (minimum)	-
^a This feature has been obsoleted in FC-LS, and may be reclaimed for use in a future standard. ^b This bit shall only be set to one if the entire fabric supports the feature.	

6.2 Link services

6.2.1 Basic Link Services

Fabric services in a switch shall support Basic Link Services as specified in FC-DA.

6.2.2 ELS requirements for well-known addresses

See FC-DA for N_Port well-known address responder requirements.

If Zoning is active in the fabric, an ELS response from a well-known address shall only include data relating to Nx_Ports that are in the same Zone(s) as the requesting Nx_Port.

6.2.3 Extended Link Service Replies

Table 9 specifies the support required by this Technical Report for Extended Link Service Replies. This table only applies to those supported Extended Link Service requests that have valid replies as indicated in FC-DA.

Table 9 – Extended Link Service Replies

Name	Abbr	OpCode	Originator Support	Responder Support
Accept	LS_ACC	-	-	R
Link Service Reject	LS_RJT	-	-	R

6.3 FC-AL-2 features for FL_Ports

Table 10 lists Fibre Channel Arbitrated Loop features for FL_Ports with usage defined by this Technical Report. In table 10, the FL_Port Originator column refers to an FL_Port sending an FC-AL feature to an NL_Port, and the FL_Port Responder column refers to an FL_Port receiving an FC-AL feature from an NL_Port.

Table 10 – FC-AL features for FL_Ports (part 1 of 3)

Feature	FL_Port Originator	FL_Port Responder
Open Full Duplex	R	R
Open Half Duplex	P	R
Send frames from multiple S_IDs in a single Loop Tenancy	I	-
Receive frames to multiple D_IDs in a single Loop Tenancy		R
Unfairness	I	-
Transfer mode (use of TRANSFER loop state)	R ^a	-
<p>^a Required uses of Transfer mode are described in 5.2.2.</p> <p>^b Alternate BB_Credit management is mandatory in FC-AL-2.</p> <p>^c FL_Ports shall interoperate with NL_Ports that support any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit greater than zero, but the originator of frames is not required to take advantage of the non- zero Login_BB_Credit.</p> <p>^d NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. As required by FC-AL-2, the FL_Port shall not receive frames associated with OPNfr or OPNyr. Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop.</p> <p>^e The only method allowed by this Technical Report for delivering broadcast from the FL_Port to the attached NL_Ports is via an OPNfr sent by the FL_Port. The FL_Port is not required to perform a discovery process to determine whether the attached NL_Ports are able to receive broadcast frames via OPNfr. Use of broadcast, and the number of frames that may be sent following OPNfr, are both specific to the ULP.</p> <p>^f An NL_Port with broadcast or multicast frames intended for Remote Ports shall open the FL_Port via OPNyx or OPNyy or be opened by the FL_Port via OPNyx, and send the frames to the FL_Port for forwarding by the fabric (see FC-AL-2). This causes broadcast frames to be delivered to the Local Loop and multicast frames if any NL_Ports on the Local Loop are part of the multicast group.</p> <p>^g LPEfx allows resetting the bypass circuits of NL_Ports that have been bypassed and have lost their AL_PAs due to a LIP and therefore are not enabled using an addressed LPEyx.</p> <p>^h This LIP may be issued by an FL_Port to request an AL_PA if it is non-participating, or for any of the reasons listed in 5.2.1.1. The FL_Port response to a LIP is described in 5.1.6.</p> <p>ⁱ This LIP may be issued by an FL_Port which detects a Link Failure (see FC-FS-2).</p> <p>^j An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in a ULP-specific manner. An FL_Port shall respond to this form of LIP in the same manner as LIP(F7, AL_PS).</p> <p>^k An FL_Port that receives MRK shall attempt to forward the MRK. An FL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).</p>		

Table 10 – FC-AL features for FL_Ports (part 2 of 3)

Feature	FL_Port Originator	FL_Port Responder
Alternate BB_Credit model ^b	R	R
Login_BB_Credit equal to 0 ^c	I	R
Login_BB_Credit greater than 0 ^c	A	A
Broadcast and Multicast via Broadcast Replicate (OPNfr) Transmit to NL_Ports on the Local Loop ^{e, f} Receive from NL_Ports on the Local Loop ^{d, f}	A -	- P
Broadcast and Multicast via Selective Replicate (OPNyr) Transmit to NL_Ports on the Local Loop ^e Receive from NL_Ports on the Local Loop ^d	P -	- P
<p>^a Required uses of Transfer mode are described in 5.2.2.</p> <p>^b Alternate BB_Credit management is mandatory in FC-AL-2.</p> <p>^c FL_Ports shall interoperate with NL_Ports that support any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit greater than zero, but the originator of frames is not required to take advantage of the non- zero Login_BB_Credit.</p> <p>^d NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. As required by FC-AL-2, the FL_Port shall not receive frames associated with OPNfr or OPNyr. Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop.</p> <p>^e The only method allowed by this Technical Report for delivering broadcast from the FL_Port to the attached NL_Ports is via an OPNfr sent by the FL_Port. The FL_Port is not required to perform a discovery process to determine whether the attached NL_Ports are able to receive broadcast frames via OPNfr. Use of broadcast, and the number of frames that may be sent following OPNfr, are both specific to the ULP.</p> <p>^f An NL_Port with broadcast or multicast frames intended for Remote Ports shall open the FL_Port via OPNyx or OPNyy or be opened by the FL_Port via OPNyx, and send the frames to the FL_Port for forwarding by the fabric (see FC-AL-2). This causes broadcast frames to be delivered to the Local Loop and multicast frames if any NL_Ports on the Local Loop are part of the multicast group.</p> <p>^g LPEfx allows resetting the bypass circuits of NL_Ports that have been bypassed and have lost their AL_PAs due to a LIP and therefore are not enabled using an addressed LPEyx.</p> <p>^h This LIP may be issued by an FL_Port to request an AL_PA if it is non-participating, or for any of the reasons listed in 5.2.1.1. The FL_Port response to a LIP is described in 5.1.6.</p> <p>ⁱ This LIP may be issued by an FL_Port which detects a Link Failure (see FC-FS-2).</p> <p>^j An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in a ULP-specific manner. An FL_Port shall respond to this form of LIP in the same manner as LIP(F7, AL_PS).</p> <p>^k An FL_Port that receives MRK shall attempt to forward the MRK. An FL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).</p>		

Table 10 – FC-AL features for FL_Ports (part 3 of 3)

Feature	FL_Port Originator	FL_Port Responder
LPEyx, LPByx, or LPEfx ^g	I	P
LIP		
(F7, F7) and (F7, AL_PS) [initializing] ^h	I	R
(F8, F7) and (F8, AL_PS) [loop failure] ⁱ	I	R
(AL_PD, AL_PS) [selective hard reset] ^j	I	R
MRK ^k	P	R
<p>^a Required uses of Transfer mode are described in 5.2.2.</p> <p>^b Alternate BB_Credit management is mandatory in FC-AL-2.</p> <p>^c FL_Ports shall interoperate with NL_Ports that support any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit greater than zero, but the originator of frames is not required to take advantage of the non- zero Login_BB_Credit.</p> <p>^d NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. As required by FC-AL-2, the FL_Port shall not receive frames associated with OPNfr or OPNyr. Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop.</p> <p>^e The only method allowed by this Technical Report for delivering broadcast from the FL_Port to the attached NL_Ports is via an OPNfr sent by the FL_Port. The FL_Port is not required to perform a discovery process to determine whether the attached NL_Ports are able to receive broadcast frames via OPNfr. Use of broadcast, and the number of frames that may be sent following OPNfr, are both specific to the ULP.</p> <p>^f An NL_Port with broadcast or multicast frames intended for Remote Ports shall open the FL_Port via OPNyx or OPNy or be opened by the FL_Port via OPNyx, and send the frames to the FL_Port for forwarding by the fabric (see FC-AL-2). This causes broadcast frames to be delivered to the Local Loop and multicast frames if any NL_Ports on the Local Loop are part of the multicast group.</p> <p>^g LPEfx allows resetting the bypass circuits of NL_Ports that have been bypassed and have lost their AL_PAs due to a LIP and therefore are not enabled using an addressed LPEyx.</p> <p>^h This LIP may be issued by an FL_Port to request an AL_PA if it is non-participating, or for any of the reasons listed in 5.2.1.1. The FL_Port response to a LIP is described in 5.1.6.</p> <p>ⁱ This LIP may be issued by an FL_Port which detects a Link Failure (see FC-FS-2).</p> <p>^j An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in a ULP-specific manner. An FL_Port shall respond to this form of LIP in the same manner as LIP(F7, AL_PS).</p> <p>^k An FL_Port that receives MRK shall attempt to forward the MRK. An FL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).</p>		

6.4 Loop Fabric Address

The Loop Fabric Address 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 Loop Fabric Address. Devices compliant with FC-DA are prohibited from sending LINIT to the Loop Fabric Address. As such, devices compliant with this Technical Report are only required to support LSTS sent to the Loop Fabric Address.

7 Fabric behaviors

7.1 Overview

This clause defines the interoperability requirements for a Fabric. A Fabric consists of a Fabric containing a single Switch, or containing two or more Switches connected via E_Ports. The purpose of this clause is to define the minimum functionality that needs to be supported in order to deploy a Fabric.

This clause defines the following two aspects of Fabrics:

- a) required Switch-to-Switch (E_Port) behaviors. These are based primarily on functions defined in FC-SW-3; and
- b) required Fabric services provided to an attached Nx_Port. These are based primarily on functions defined in FC-GS-4.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 14165-372:2011

7.2 Switch-to-Switch Requirements

7.2.1 Overview

Table 11 summarizes support requirements for features defined in FC-SW-3. A compliant device shall implement all features as specified in table 11. FC-SW-3 clauses that are descriptive are not included in table 11.

Table 11 – FC-SW-3 Fabric support summary (part 1 of 3)

Description	Support
Structure and concepts	
Fabric Addressing	R
Switch Ports	
For specific Port Type support, see table 12.	
Class F Service	R
Switch Internal Link Services	
SW_ACC	R
Switch Reject (SW_RJT)	R
Exchange Link Parameters (ELP) (see 7.2.3)	R ^a
Exchange Fabric Parameters (EFP)	R ^{a, b}
Domain_ID Assigned (DIA)	R ^a
Request Domain_ID (RDI)	R ^a
Hello (HLO)	R ^c
<p>^a The ELP, EFP, RDI, DIA, and BF SW_ILS frames shall not be used except as part of the initialization and Principal Switch selection state machines as specified in FC-SW-3.</p> <p>^b The Multicast_ID_List fields of the EFP shall not be used.</p> <p>^c FSPF SW_ILS frames shall not be used except as part of the FSPF state machines as specified in FC-SW-3.</p> <p>^d Only the Switch Link Record (i.e., LSR Type 01h) is required in LSU and LSA SW_ILS frames.</p> <p>^e A Fabric is prohibited from autonomously generating an RCF, but an outside administrative function may request a switch to generate an RCF. Such an administrative function is outside the scope of this Technical Report.</p> <p>^f The ESC SW_ILS is allowed here for vendor-specific reasons. The ESC SW_ILS shall not be used to indicate FSPF-Backbone.</p> <p>^g All switches in the Fabric shall use AR0.</p> <p>^h The MR SW_ILS is required for Basic Zoning and allowed for Enhanced Zoning.</p> <p>ⁱ Switches adhering to FC-SW-2 respond unpredictably if an Enhanced Zoning MR is received.</p> <p>^j The DRLIR SW_ILS is required for FC-SB-3 support and allowed in other environments.</p> <p>^k See 7.2.4 for details of Principal Switch Selection requirements.</p>	

Table 11 – FC-SW-3 Fabric support summary (part 2 of 3)

Description	Support
Link State Update (LSU)	R ^{d, c}
Link State Acknowledgement (LSA)	R ^{d, c}
Build Fabric (BF)	R ^a
Reconfigure Fabric (RCF)	A ^e
Disconnect Class 1 (DSCN)	P
Exchange Switch Capabilities (ESC)	A ^f
Exchange Switch Support (ESS)	I
Merge Request (MR) ⁱ	R ^h
Merge Request Resource Allocation (MRRA)	I
Acquire Change Authorization Request (ACA)	R
Release Change Authorization Request (RCA)	R
Stage Fabric Configuration Request (SFC)	R
Update Fabric Configuration Request (UFC)	R
Distributed Registered Link Incident Records (DRLIR)	R ^j
Inter-Switch Registered State Change Notifications (SW_RSCN)	R
<p>^a The ELP, EFP, RDI, DIA, and BF SW_ILS frames shall not be used except as part of the initialization and Principal Switch selection state machines as specified in FC-SW-3.</p> <p>^b The Multicast_ID_List fields of the EFP shall not be used.</p> <p>^c FSPF SW_ILS frames shall not be used except as part of the FSPF state machines as specified in FC-SW-3.</p> <p>^d Only the Switch Link Record (i.e. LSR Type 01h) is required in LSU and LSA SW_ILS frames.</p> <p>^e A Fabric is prohibited from autonomously generating an RCF, but an outside administrative function may request a switch to generate an RCF. Such an administrative function is outside the scope of this Technical Report.</p> <p>^f The ESC SW_ILS is allowed here for vendor-specific reasons. The ESC SW_ILS shall not be used to indicate FSPF-Backbone.</p> <p>^g All switches in the Fabric shall use AR0.</p> <p>^h The MR SW_ILS is required for Basic Zoning and allowed for Enhanced Zoning.</p> <p>ⁱ Switches adhering to FC-SW-2 respond unpredictably if an Enhanced Zoning MR is received.</p> <p>^j The DRLIR SW_ILS is required for FC-SB-3 support and allowed in other environments.</p> <p>^k See 7.2.4 for details of Principal Switch Selection requirements.</p>	

Table 11 – FC-SW-3 Fabric support summary (part 3 of 3)

Description	Support
Fabric configuration	
Switch Port Initialization	R
Principal Switch Selection	R ^k
Address Distribution	R
E_port and Fabric Isolation	R
B_Port Operation	A ^a
Routing Protocols	
Fabric Shortest Path First (FSPF)	R
FSPF-Backbone Routing Protocol	p ^g
Distributed Services (see 7.2.6)	
Switch Zone Exchange and Merge (see 7.2.7)	
Distributed Event Notification (see 7.2.8)	
<p>^a The ELP, EFP, RDI, DIA, and BF SW_ILS frames shall not be used except as part of the initialization and Principal Switch selection state machines as specified in FC-SW-3.</p> <p>^b The Multicast_ID_List fields of the EFP shall not be used.</p> <p>^c FSPF SW_ILS frames shall not be used except as part of the FSPF state machines as specified in FC-SW-3.</p> <p>^d Only the Switch Link Record (i.e., LSR Type 01h) is required in LSU and LSA SW_ILS frames.</p> <p>^e A Fabric is prohibited from autonomously generating an RCF, but an outside administrative function may request a switch to generate an RCF. Such an administrative function is outside the scope of this Technical Report.</p> <p>^f The ESC SW_ILS is allowed here for vendor-specific reasons. The ESC SW_ILS shall not be used to indicate FSPF-Backbone.</p> <p>^g All switches in the Fabric shall use AR0.</p> <p>^h The MR SW_ILS is required for Basic Zoning and allowed for Enhanced Zoning.</p> <p>ⁱ Switches adhering to FC-SW-2 respond unpredictably if an Enhanced Zoning MR is received.</p> <p>^j The DRLIR SW_ILS is required for FC-SB-3 support and allowed in other environments.</p> <p>^k See 7.2.4 for details of Principal Switch Selection requirements.</p>	

7.2.2 Switch Port Types

Switch Port Type usage is shown in table 12.

Table 12 – Switch Port Type Usage

Switch Port Type	Support
F_Port	A
FL_Port	A
E_Port	R
B_Port	A

Switch Ports have the following requirements.

- A Switch may have a mix of different Switch Port types. Not all Switch Ports on a Switch are required to support the same Switch Port types.
- A single Switch Port may support more than one Switch Port type. If a single Switch Port supports multiple Switch Port types, the Switch Port type shall be discovered as specified in FC-SW-3.
- A Switch shall support E_Port functionality on at least one Switch Port.
- A Switch shall support at least one of F_Port, or FL_Port operation. A switch may support both F_Port and FL_Port.

7.2.3 Exchange Link Parameters (ELP)

The ELP SW_ILS as defined in FC-SW-3 shall be used with the parameters specified in table 13.

Table 13 – ELP SW_ILS Parameters (part 1 of 2)

Item	Value
Revision	3
Flags	0 ^a
BB_SC_N	0 or 8
<ol style="list-style-type: none"> B_Port functionality is allowed, but its use is outside the scope of this Technical Report. Support of Class F for an E_Port is required as specified by FC-SW-3. Use of Class 1 is Prohibited by this Technical Report. This value shall be used in FLOGI ACC. These shall be used as default values. A compliant implementation may change these values, but it is up to the implementor to ensure interoperability. These values may need to be modified when used in a WAN environment. 	

Table 13 – ELP SW_ILS Parameters (part 2 of 2)

Item		Value
R_A_TOV		10 000 ms ^{d, e, f}
E_D_TOV		2 000 ms ^{d, e, f}
Class F Params	Class Valid	1 ^b
	X_ID Interlock	0
	Receive Data Field Size	2 112 bytes
	Concurrent Sequences	1
	End-to-End Credit	1
	Open Sequences per Exchange	1
Class 1 Params		0 ^c
Class 2 Params	Class Valid	X
	Sequential Delivery	1
	Receive Data Field Size	2 112 bytes
Class 3 Params	Class Valid	1
	Sequential Delivery	1
	Receive Data Field Size	2 112 bytes
ISL Flow Control Mode		2
Flow Control Parameter Length		20 bytes
Flow Control Parameters		(see table 14)
<p>^a B_Port functionality is allowed, but its use is outside the scope of this Technical Report.</p> <p>^b Support of Class F for an E_Port is required as specified by FC-SW-3.</p> <p>^c Use of Class 1 is Prohibited by this Technical Report.</p> <p>^d This value shall be used in FLOGI ACC.</p> <p>^e These shall be used as default values. A compliant implementation may change these values, but it is up to the implementor to ensure interoperability.</p> <p>^f These values may need to be modified when used in a WAN environment.</p>		

To ensure consistent support of Classes of Service over the entire Fabric, the ELP SW_ILS shall be accepted by the responding Switch only if the two Switches support the same Classes of Service, otherwise it shall be rejected with a SW_RJT Reason Code “Logical Error” and SW_RJT Reason Code Explanation “Class N Service Parameter error”, and the involved E_Ports shall isolate.

To ensure consistent timeout values over the entire Fabric, the ELP SW_ILS shall be accepted by the responding Switch only if the two Switches have the same value for R_A_TOV and the same value for E_D_TOV, otherwise it shall be rejected with a SW_RJT Reason Code “Logical Error” and SW_RJT Reason Code Explanation “R_A_TOV or E_D_TOV mismatch”, and the involved E_Ports shall isolate.

Required values for the Flow Control Parameters field of the ELP are shown in table 14. These values only apply to the ISL Flow Control Mode value as specified in table 13. The parameter values and format for any other ISL Flow Control Modes are outside the scope of this Technical Report.

Table 14 – Flow Control Parameters

Item	Size (Bytes)	Value
BB_Credit	4	x^a
Compatibility Parameter 1	4	2 112 ^b
Compatibility Parameter 2	4	10 000 ^b
Compatibility Parameter 3	4	2 000 ^b
Compatibility Parameter 4	4	0 ^b
<p>^a BB_Credit may be set to any value supported by the Originating E_Port. This BB_Credit value shall be used only after the ELP and LR have been completed. Before that, the BB_Credit shall be one as specified by FC-SW-3.</p> <p>^b These values are required for backward compatibility and interoperability reasons.</p>		

7.2.4 Principal Switch selection

Each Switch in the Fabric shall support Principal Switch selection as specified in FC-SW-3 except as modified by the following requirements:

- at a minimum, each Switch shall support origination of an EFP with a Principal Switch Priority of FFh. This implies that each Switch is required to participate in the Principal Switch selection process but is not required to be capable of becoming the Principal Switch; and
- a Switch may support origination of an EFP with a Principal Switch Priority of 01h to FEh. Supporting a Principal Switch Priority of FEh means that the Switch is capable of becoming a Principal Switch.

7.2.5 Fabric Shortest Path First (FSPF)

Each Switch in the Fabric shall support FSPF as specified in FC-SW-3. The following requirements apply to FSPF implementation:

- each Switch in a Fabric shall use only AR0 (i.e., the AR Number field of the FSPF Header as specified in FC-SW-3 shall be set to zero);

- b) the LSR Type field in the Link State Header shall be set to 01h indicating Switch Link Record; and
- c) the Link Type field in the Link Descriptor shall be set to 01h indicating Point-to-Point Link.

7.2.6 Distributed Services

7.2.6.1 Overview

Distributed services as defined in FC-SW-3 allow Switches in a Fabric to share information required for support of the CT based well-known services. Switch-to-Switch support requirements for these Services are defined in 7.2.6.

7.2.6.2 Distributed Name Server

7.2.6.2.1 Switch-to-Switch Name Server requests

Switch-to-Switch Name Server requests shall be supported as shown in table 15, table 16, and table 17.

Table 15 – FC-SW-3 defined Name Server requests

Feature	Support
Get Entry based on Port Identifier (GE_ID)	I
Get Entry based on Port Name (GE_PN)	I
Get Entry based on Node Name (GE_NN)	I
Get Entries based on IP Address - Node (GE_IP)	P
Get Entries based on FC-4 TYPEs (GE_FT)	I
Get Entries based on Port Type (GE_PT)	I
Get Entries based on Zone Member (GE_ZM)	P
Get Entries based on Zone Name (GE_ZN)	P
Get Entries based on IP Address - Port (GE_IPP)	I
Get Entries based on FC-4 Feature (GE_FF)	I
Get Entries based on Fabric Port_Name (GE_FPN)	P
Remove All (RA)	P

Table 16 – FC-GS-4 defined Name Server requests (part 1 of 3)

Feature	Support
Get All Next (GA_NXT)	I
Get Identifiers - Scope (GID_A)	P
Get Port Name (GPN_ID)	I
Get Node Name (GNN_ID)	I
Get Class of Service (GCS_ID)	I
Get FC-4 Types (GFT_ID)	I

Table 16 – FC-GS-4 defined Name Server requests (part 2 of 3)

Feature	Support
Get Symbolic Port Name (GSPN_ID)	I
Get Port Type (GPT_ID)	I
Get IP address (GIPP_ID)	P
Get Fabric Port Name (GFPN_ID)	I
Get Hard Address (GHA_ID)	P
Get FC-4 Descriptors (GFD_ID)	P
Get FC-4 Features (GFF_ID)	I
Get Port Identifier (GID_PN)	I
Get IP Address (GIPP_PN)	P
Get Port Identifiers (GID_NN)	I
Get Port Names (GPN_NN)	I
Get IP Address - Node Name (GIP_NN)	P
Get Initial Process Associator (GIPA_NN)	P
Get Symbolic Node Name (GSNN_NN)	I
Get Node Name (GNN_IP)	P
Get Initial Process Associator (GIPA_IP)	P
Get Port Identifiers (GID_FT)	I
Get Port Names (GPN_FT)	I
Get Node Names (GNN_FT)	I
Get Port Identifiers (GID_PT)	I
Get Port Identifier (GID_IPP)	P
Get Port Name (GPN_IPP)	P
Get Port Identifiers (GID_FF)	I
Get Port Identifiers - Fabric Port Name (GID_FPN)	I
Get Permanent Port Name - Port Identifier (GPPN_ID)	I
Register Port Name (RPN_ID)	P
Register Node Name (RNN_ID)	P
Register Class of Service (RCS_ID)	P
Register FC-4 Types (RFT_ID)	P
Register Symbolic Port Name (RSPN_ID)	P

Table 16 – FC-GS-4 defined Name Server requests (part 3 of 3)

Feature	Support
Register Port Type (RPT_ID)	P
Register IP Address - Port (RIPP_ID)	P
Register Hard Address (RHA_ID)	P
Register FC-4 Descriptor (RFD_ID)	P
Register FC-4 Features (RFF_ID)	P
Register IP Address - Node (RIP_NN)	P
Register Initial Process Associator (RIPA_NN)	P
Register Symbolic Node Name (RSNN_NN)	I
Deregister All (DA_ID)	P

Table 17 – FC-GS-4 Common Requests

Feature	Support
Get More Information (GMI)	A

The requests in table 15, table 16, and table 17 have the following requirements.

- The requests specified in table 15 are defined in FC-SW-3 for the special purpose of Switch-to-Switch exchange of Name Server data. The response payload of these requests contains the Name Server Object defined in FC-SW-3. The Name Server Object shall be used as specified in 7.2.6.2.2.
- The requests specified in table 16 and table 17 are defined in FC-GS-4 and allowed within FC-SW-3 for Switch-to-Switch Name Server requests. These requests and their responses shall be sent according to the rules defined in FC-SW-3.
- Switch-to-Switch Name Server requests shall be processed independent of Zoning. All filtering of Name Server data to support zoning shall be done at the Entry Switch. GE_ZM and GE_ZN are exceptions to this requirement since they may be required to perform zoning operations on a responding Switch.

7.2.6.2.2 Name Server Object Usage

For all GE_XXX (i.e., 4yyh) Name Server to Name Server requests, as defined in FC-SW-3, the Name Server Object used in the response shall be as follows.

- If the responder has either a non-null Port Symbolic Name or a non-null Node Symbolic Name, and a null FC-4 Descriptor, and a null FC-4 Features, then the Name Server Entry Object with an Entry Object Format Indicator of 00h shall be used by the responder.
- If the responder has a null Port Symbolic Name, a null Node Symbolic Name, a null FC-4 Descriptor, and a null FC-4 Features, then the Name Server Entry Object with an Entry Object Format Indicator of 01h shall be used by the responder.

- c) If the responder has either a non-null Port Symbolic Name or a non-null Node Symbolic Name, and either a non-null FC-4 Descriptor or non-null FC-4 Features, then the Name Server Entry Object with an Entry Object Format Indicator of 02h shall be used by the responder.
- d) If the responder has a null Port Symbolic Name and a null Node Symbolic Name and either a non-null FC-4 Descriptor or non-null FC-4 Features, then the Name Server Entry Object with an Entry Object Format Indicator of 03h shall be used by the responder.

The rules for Name Server Object usage are summarized in table 18.

Table 18 – Name Server Object Usage summary

	Null FC-4 Features and Null FC-4 Descriptor	non-Null FC-4 Features and/or non-Null FC-4 Descriptor
non-Null Symbolic Node_Name and/or non-Null Symbolic N_Port_Name	00h	02h
Null Symbolic Node_Name and Null Symbolic N_Port_Name	01h	03h

7.2.6.2.3 Distributed Name Server response

The following rules shall apply to responses to Distributed Name Server requests:

- a) a reject of a 1-to-1 switch request, as specified by FC-SW-3, shall result in a reject of the original Nx_Port request;
- b) a reject of a 1-to-many or 1-to-all request, as specified by FC-SW-3, shall result in either
 - A) a reject of the original Nx_Port request, or
 - B) a Partial Response to the original requesting Nx_Port. If a Partial Response is sent, the Partial Response bit in the CT Header shall be set to one;
- c) the original request shall be rejected with a Reject CT_IU if there is no response to a 1-to-all request; and
- d) a response shall be generated for an original request if there is an answer, even if it is a partial answer, to a 1-to-all request.

7.2.6.3 Distributed Management Server

Switch-to-Switch Management Server requests shall be supported as shown in table 19 and table 20.

Table 19 – FC-SW-3 Defined Management Server requests

Feature	Support
Get Management Server Capabilities (GCAP)	R

Table 20 – FC-GS-4 Defined Management Server requests (part 1 of 2)

Feature	Support
Get Topology Information (GTIN)	P
Get Interconnect Element List (GIEL) ^a	P
Get Interconnect Element Type (GIET) ^a	P
Get Domain Identifier (GDID) ^a	P
Get Management Identifier (GMID)	I
Get Fabric Name (GFN)	P
Get Interconnect Element Logical Name (GIELN)	I
Get Interconnect Element Management Address List (GMAL)	I
Get Interconnect Element Information List (GIELI)	I
Get Port List (GPL)	I
Get Port Type (GPT)	I
Get Physical Port Number (GPPN)	I
Get Attached Port Name List (GAPNL)	I
Get Port State (GPS)	I
Get Attached Topology Information (GATIN)	P
Get Platform Node Name List (GPLNL)	I
Get Platform Type (GPLT)	I
Get Platform Management Address List (GPLML)	I
Get Platform Name - Node Name (GNPL)	I
Get Platform List (GPL)	I
Get Node Identification Data - Node Name (GNID)	P
Get Port Speed Capabilities (GPSC)	I
Get Switch Enforcement Status (GSES)	I
Get Platform Attribute Block (GPAB)	I
Get Platform FCP Type (GPFCP)	I
Get Platform OS LUN Mappings (GPLI)	I
Register Interconnect Element Logical Name (RIELN)	R ^b
Register Platform (RPL)	R ^b
^a This only applies to Interconnect Elements of type Switch. ^b This is required to allow local data copies to be maintained on remote switches as specified in FC-SW-3.	

Table 20 – FC-GS-4 Defined Management Server requests (part 2 of 2)

Feature	Support
Register Platform Node Name (RPLN)	R ^b
Register Platform Type (RPLT)	R ^b
Register Platform Management Address (RPLM)	R ^b
Register Platform Attribute Block (RPAB)	R ^b
Register Platform FCP Type (RPFCP)	R ^b
Register Platform OS LUN Mappings (RPLI)	R ^b
Deregister Platform (DPL)	R ^b
Deregister Platform Node Name (DPLN)	R ^b
Deregister Platform Management Address List (DPLML)	R ^b
Deregister Platform Management Address (DPLM)	R ^b
Deregister Platform OS LUN Mappings (DPLI)	R ^b
Deregister Platform Attribute Block (DPAB)	R ^b
^a This only applies to Interconnect Elements of type Switch. ^b This is required to allow local data copies to be maintained on remote switches as specified in FC-SW-3.	

7.2.7 Zoning

7.2.7.1 Minimum Zoning configuration rules

Zoning, as defined in FC-SW-3, shall be supported as shown in table 21. Administration of Zones may be vendor-specific and is therefore outside the scope of this Technical Report.

Table 21 – Zoning support

Item	Support
Zoning Management Support	
Basic Zoning Management	R
Enhanced zoning Management	A
Zoning Enforcement	
Soft Zoning Enforcement	R
Hard Zoning Enforcement	A
Broadcast Zoning Enforcement	A
Zoning Data Structures Support	
^a Alias Name zone member type support is Prohibited in Basic Zoning.	

Table 21 – Zoning support

Item	Support
Active Zone Set	R
Zone Set Database	A
Zone Object support	
Zone Set Object	A
Zone Object	R
Zone Alias Object	A
Zone Reference Object	A
Zone Attribute Object	A
Vendor-Specific Object	P
Zone Member Types support	
N_Port_Name	R
Domain_ID & Physical Port	P
N_Port_ID	A
Alias Name	A ^a
Node_Name	A
F_Port_Name	A
Vendor-Specific	P
Zone Attribute Types	
Protocol	A
Broadcast Zone	A
Hard Zone	A
Vendor Specific	P
^a Alias Name zone member type support is Prohibited in Basic Zoning.	

7.2.7.2 Zone Objects

For Zone Objects and Zone Members as defined in FC-SW-3, the following applies:

- Protocol field of the Zone Object shall be set to zero; and
- Flag field of the Zone Member shall be set to zero.

7.2.8 Distributed Event Notification

7.2.8.1 Switch-to-Switch Registered State Change Notifications (SW_RSCN)

A Fabric detected SW_RSCN shall be generated under the conditions specified in 7.3.5.4 with the exception that an SW_RSCN shall not be generated for Fabric reconfiguration that is detectable by each switch in the Fabric (i.e., Zone and FSPF database changes).

7.2.8.2 Distribute Registered Link Incident Records (DRLIR)

Each Switch in the Fabric shall support the DRLIR request as defined in FC-SW-3.

7.2.9 Additional Switch requirements

A Switch is required to support sequential delivery. Sequential delivery within a Switch has the following implications.

- a) Frames between a single entry Switch Port and exit Switch Port pair of a Fabric shall be delivered at the exit Switch Port in the same order they were received on the entry Switch Port.
- b) A Switch is not responsible for maintaining sequence or frame order relative to other sequences or frames if an Nx_Port retransmits a sequence or frame due to error recovery policy.
- c) The method by which a Switch implements sequential delivery between its own ports is outside the scope of this Technical Report.
- d) If there is more than one Switch in a Fabric, then sequential delivery shall be maintained by using the same path for frames between any pair of ports on the Fabric, on a per class basis.
- e) Frames are only required to be delivered sequentially within a single class of service.
- f) Class F frames sent on inter-switch links are not required to be delivered sequentially.
- g) If a path disappears or is introduced into a Fabric, sequential delivery may be temporarily suspended.

NOTE 1 – An example of why sequential delivery may be suspended as described in g) is a frame that has partially traversed a path to a destination may be rerouted due to a change in the routing table causing it to arrive later than a frame transmitted at a later time to the same destination. It is for this reason that R_A_TOV is applied to frames in a Fabric. After R_A_TOV, any frames still traveling non-existent paths will have been delivered, or discarded.

7.3 Fabric Service requirements

7.3.1 Overview

Table 22 specifies support requirements for features defined in FC-GS-4. A compliant device shall implement all features as specified in table 22. FC-GS-4 features that are not included in table 22 are Prohibited.

Table 22 – FC-GS-4 Fabric support

Description	Support
Common Transport for Generic Services (CT)	R ^a
Directory Service	
Name Server	R ^b
Management Service	
Fabric Configuration Server	R
Unzoned Name Server	R
Fabric Zone Server	R
^a Synchronous mode only. ^b For supported Name Server Objects, see specific requests supported in table 23.	

7.3.2 Name Server

7.3.2.1 Name Server Request support

Nx_Port originated Name Server requests shall be supported by each switch element in a Fabric as shown in table 23 and table 24. All other requests are Prohibited.

Table 23 – Name Server request support (part 1 of 3)

Feature	Support
Get All Next (GA_NXT)	R
Get Identifiers - Scope (GID_A)	P
Get Port Name (GPN_ID)	R
Get Node Name (GNN_ID)	R
Get Class of Service (GCS_ID)	R
Get FC-4 Types (GFT_ID)	R
Get Symbolic Port Name (GSPN_ID)	R
^a A Fabric shall implicitly register the N_Port_Name and Node_Name of an Nx_Port based on FLOGI payload. ^b FC_Port type information shall be implicitly provided by the Fabric.	

Table 23 – Name Server request support (part 2 of 3)

Feature	Support
Get Port Type (GPT_ID)	R
Get IP address (GIPP_ID)	P
Get Fabric Port Name (GFPP_ID)	R
Get Hard Address (GHA_ID)	P
Get FC-4 Descriptors (GFD_ID)	A
Get FC-4 Features (GFF_ID)	R
Get Port Identifier (GID_PN)	R
Get IP Address (GIPP_PN)	P
Get Port Identifiers (GID_NN)	R
Get Port Names (GPN_NN)	R
Get IP Address - Node Name (GIP_NN)	P
Get Initial Process Associator (GIPA_NN)	P
Get Symbolic Node Name (GSNN_NN)	R
Get Node Name (GNN_IP)	P
Get Initial Process Associator (GIPA_IP)	P
Get Port Identifiers (GID_FT)	R
Get Port Names (GPN_FT)	R
Get Node Names (GNN_FT)	R
Get Port Identifiers (GID_PT)	R
Get Port Identifier (GID_IPP)	P
Get Port Name (GPN_IPP)	P
Get Port Identifiers (GID_FF)	R
Get Port Identifiers - Fabric Port Name (GID_FPN)	R
Get Permanent Port Name - Port Identifier (GPPN_ID)	R
Get Port Identifier (Domain/Port) (GID_DP)	A
Register Port Name (RPN_ID)	p ^a
Register Node Name (RNN_ID)	p ^a
<p>^a A Fabric shall implicitly register the N_Port_Name and Node_Name of an Nx_Port based on FLOGI payload.</p> <p>^b FC_Port type information shall be implicitly provided by the Fabric.</p>	

Table 23 – Name Server request support (part 3 of 3)

Feature	Support
Register Class of Service (RCS_ID)	R
Register FC-4 Types (RFT_ID)	R
Register Symbolic Port Name (RSPN_ID)	R
Register Port Type (RPT_ID)	p ^b
Register IP Address - Port (RIPP_ID)	R
Register Hard Address (RHA_ID)	P
Register FC-4 Descriptor (RFD_ID)	A
Register FC-4 Features (RFF_ID)	R
Register IP Address - Node (RIP_NN)	P
Register Initial Process Associator (RIPA_NN)	P
Register Symbolic Node Name (RSNN_NN)	R
De-register All (DA_ID)	R
^a A Fabric shall implicitly register the N_Port_Name and Node_Name of an Nx_Port based on FLOGI payload. ^b FC_Port type information shall be implicitly provided by the Fabric.	

Table 24 – Common request support

Feature	Support
Get More Information (GMI)	R
Server Session Begin (SSB)	R
Server Session End (SSE)	R

7.3.2.2 Name Server Object registration

7.3.2.2.1 Required Name Server Object registration

The following Name Server Objects shall be implicitly registered by the Switch.

- Node_Name and N_Port_Name objects of an Nx_Port. This registration is performed based on the FLOGI or FDISC payload.
- Class of service supported by an Nx_Port. This registration is performed based on the FLOGI or FDISC payload.
- FC_Port type of an Nx_Port based on results of Switch Port initialization. FC_Port type registration shall be performed as follows:

- A) if the Switch Port is operating as an FL_Port, then any attached Nx_Port shall be registered as an NL_Port in the Name Server; or
- B) if the Switch Port is operating as an F_Port, then the attached Nx_Port shall be registered as an N_Port in the Name Server.

7.3.2.2.2 Prohibited Name Server Object registration

A Switch shall not perform implicit Name Server Object registration for the following:

- a) Switch Ports (see 3.2.52); or
- b) FC-4 Type based on registration of FC-4 Features or FC-4 Descriptors.

7.3.2.3 Name Server Object removal

7.3.2.3.1 Full removal

If any of the following is true, the Name Server shall remove the affected Port Identifier and all objects associated with the Port Identifier from the Name Server Database and send an SW_RSCN SW_ILS to all other switches in the fabric:

- a) an N_Port attached to an F_Port goes to the OFFLINE state or LINK FAILURE state;
- b) an OPN sent from an FL_Port to an NL_Port returns to the FL_Port. This includes the OPN for sending a FAN after initialization;
- c) an NL_Port is no longer present at the end of a loop initialization. This may be detected by using the AL_PA bit map, item b) in this subclause, or the AL_PA position map;
- d) a DA_ID is received for the Nx_Port;
- e) a LOGO ELS is sent to FFFFFFFEh and the fabric supports NPIV; or

NOTE 2 – When a Fabric supporting NPIV receives a LOGO ELS to D_ID FFFFFFFEh from a source Port Identifier, it logs out the source Port Identifier from the Fabric. See FC-LS.

- f) a LOGO ELS sent to D_ID FFFFFFFEh is accepted by the Fabric F_Port.

NOTE 3 – A switch that does not support NPIV may not accept a LOGO ELS to D_ID FFFFFFFEh and thus does not remove the affected port identifier from the Name Server Database.

7.3.2.3.2 Partial removal

When a FLOGI is received in which either the Port Name or Node Name is different from what is currently registered for the Port Identifier, all Name Server objects are removed and implicit registration as defined in 7.3.2.2.1 is performed. All objects associated with all other Port Identifiers assigned to the N_Port by the Fabric F_Port are also removed.

When a FLOGI is received in which the Port Name and Node Name are the same as what is currently registered for the Port Identifier, the Port Identifier and its objects shall not be removed. All objects associated with all other Port Identifiers assigned to the N_Port by the Fabric F_Port are removed.

7.3.2.4 Name Server resource utilization

The name server shall generate an explicit LOGO if it logs a Nx_Port out due to resource limitations.

7.3.3 Fabric Configuration Server

As specified in table 22, support of the Fabric Configuration Server is required. Nx_Port originated Fabric Configuration Server requests shall be supported as shown in table 25.

Table 25 – Fabric Configuration Server Request Support (part 1 of 2)

Item	Support
Get Topology Information (GTIN)	P
Get Interconnect Element List (GIEL)	R
Get Interconnect Element Type (GIET)	R
Get Interconnect Element Domain Identifier (GDID)	R
Get Interconnect Element Management Identifier (GMID)	R
Get Interconnect Element Fabric Name (GFN)	R
Get Interconnect Element Logical Name (GIELN)	R
Get Interconnect Element Management Address List (GMAL)	R
Get Interconnect Element Information List (GIEIL)	R
Get Port List (GPL)	R
Get Port Type (GPT)	R
Get Physical Port Number (GPPN)	R
Get Attached Port Name List (GAPNL)	R
Get Port State (GPS)	R
Get Attached Topology Information (GATIN)	P
Get Platform Node Name List (GPLNL)	R
Get Platform Type (GPLT)	R
Get Platform Management Address List (GPLML)	R
Get Platform Name - Node Name (GNPL)	R
Get Platform Name List (GPNL)	R
Get Node Identification Data - Node Name (GNID)	P
Get Port Speed Capabilities (GPSC)	R
Get Switch Enforcement Status (GSES)	R
Get Platform Attribute Block (GPAB)	R
Get Platform FCP Type (GPFCP)	R
Get Platform OS LUN Mappings (GPLI)	R
Register Interconnect Element Logical Name (RIELN)	R

Table 25 – Fabric Configuration Server Request Support (part 2 of 2)

Item	Support
Register Platform (RPL)	R
Register Platform Node Name (RPLN)	R
Register Platform Type (RPLT)	R
Register Platform Management Address (RPLM)	R
Register Platform Attribute Block (RPAB)	R
Register Platform FCP Type (RPFCP)	R
Register Platform OS LUN Mappings (RPLI)	R
Deregister Platform (DPL)	R
Deregister Platform Node Name (DPLN)	R
Deregister Platform Management Address List (DPLML)	R
Deregister Platform Management Address (DPLML)	R
Deregister Platform OS LUN Mappings (DPLI)	R
Deregister Platform Attribute Block (DPAB)	R

7.3.4 Fabric Zone Server

As specified in table 22 support of the Fabric Zone Server is required. Nx_Port originated Fabric Zone Server Basic Zoning requests shall be supported as shown in table 26.

Table 26 – Fabric Zone Server Basic Zoning request support (part 1 of 2)

Item	Support
Get Capabilities (GZC)	R
Get Enforcement State (GEST)	R
Get Zone Set List (GZSN)	A
Get Zone List (GZD)	A
Get Zone Member List (GZM)	A
Get Active Zone Set (GAZS)	R
Get Zone Set (GZS)	A
Add Zone Set (ADZS)	A ^a
Activate Zone Set Direct (AZSD)	R ^a
Activate Zone Set (AZS)	A
^a Use of the Node_Name Zone Member Identifier type (04h) is prohibited.	

Table 26 – Fabric Zone Server Basic Zoning request support (part 2 of 2)

Item	Support
Deactivate Zone Set (DZS)	R
Add Zone Members (AZM)	A ^a
Add Zone (AZD)	A
Remove Zone Members (RZM)	A ^a
Remove Zone (RZD)	A
Remove Zone Set (RZS)	A
^a Use of the Node_Name Zone Member Identifier type (04h) is prohibited.	

If the Fabric Zone Server supports Enhanced Zoning, Nx_Port originated Fabric Zone Server Enhanced Zoning requests shall be supported as shown in table 27.

Table 27 – Fabric Zone Server Enhanced Zoning request support (part 1 of 2)

Item	Support
Control Zoning Management requests	
Get Fabric Enhanced Zoning Support (GFEZ)	R
Set Fabric Enhanced Zoning Support (SFEZ)	R
Zoning Session requests	
Server Session Begin (SSB)	R
Server Session End (SSE)	R
Commit Zone Changes (CMIT)	R
Enhanced Zoning requests	
Get Activation Results (GAR)	R
Get Zone Attribute Object Name (GZA)	R
Get Zone Attribute Block (GZAB)	R
Get Zone Set List - Enhanced (GZSE)	R
Get Zone List - Enhanced (GZDE)	R
Get Zone Member List - Enhanced (GZME)	R
Get Zone Attribute Object List (GZAL)	R
Get Alias List (GAL)	R
Get Alias Member List (GAM)	R
Set Zone Attribute Object Name (SZA)	R
Set Zone Attribute Block (SZAB)	R