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

INCITS TR-xx-201x
T11/Project 1870-DT

FIBRE CHANNEL

DEVICE ATTACH - 2

(FC-DA-2)

Revision 1.06

INCITS working draft proposed
Technical Report

March 15, 2011

Secretariat: Information Technology Industry Council

NOTE:

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

POINTS OF CONTACT:

Steven Wilson (T11 Chair)
Brocade Communications, Inc.
130 Holger Way
San Jose, CA 95134
Voice: 408-487-8128
Email: swilson@brocade.com

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

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

Bill Martin (FC-DA-2 Facilitator)
Emulex Corporation
7213 Marblethorpe
Roseville, CA 95747
Voice: 916-765-6875
Email: bill.martin@emulex.com

David Peterson (FC-DA-2 Editor)
Brocade Communications
6000 Nathan Lane North
Plymouth, MN 55442
Phone: 612-802-3299
Email: dpeterso@brocade.com

draft proposed INCITS Technical Report

Fibre Channel — Device Attach - 2 (FC-DA-2)

Secretariat

Information Technology Industry Council

Approved ,200x

American National Standards Institute, Inc.

Abstract

This technical report selects and restricts logical options from the Fibre Channel Framing and Signaling, Fibre Channel Protocol for SCSI, Fibre Channel Arbitrated Loop, Fibre Channel Generic Services, and Fibre Channel Single Byte Command Set standards. The intent of this technical report is to facilitate interoperablity between devices whether they are connected in a loop or Fabric topology.

INCITS T11 Technical Report Series

This technical report is one in a series produced by the American National Standards Committee, INCITS, Information Technology. The secretariat for INCITS is held by the Information Technology Industry Council (ITI), 1101 K Street NW Suite 610, Washington DC 20005.

FC-FS-3

As a by-product of the standards development process and the resources of knowledge devoted to it, INCITS from time to time produces Technical Reports. Such Technical Reports are not standards, nor are they intended to be used as such.

INCITS Technical Reports are produced in some cases to disseminate the technical and logical concepts reflected in standards already published or under development. In other cases, they derive from studies in areas where it is found premature to develop a standard due to a still changing technology, or inappropriate to develop a rigorous standard due to the existence of a number of viable options, the choice of which depends on the user's particular requirements. These Technical Reports, thus, provide guidelines, the use of which can result in greater consistency and coherence of information processing systems.

When the draft technical report is completed, the Technical Committee approval process is the same as for a draft standard. Processing by INCITS is also similar to that for a draft standard.

Patent Statement

CAUTION: The developers of this technical report have requested that holder's of patents that may be required for the implementation of this technical report disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this technical report.

As of the date of publication of this technical report and following calls for the identification of patents that may be required for the implementation of this technical report, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any technical report it processes. No representation is made or implied that licenses are not required to

Published by

American National Standards Institute
11 W. 42nd Street, New York, New York 10036

Printed in the United States of America

Foreword (This Foreword is not part of Technical Report INCITS TR-xx-201x.)

The Fibre Channel Device Attach - 2 (FC-DA-2) Technical Report describes a model for end devices (i.e., Nx_Ports), facilitating interoperability in both loop and Fabric environments.

This technical report was developed by Technical Committee T11 of Accredited Standards Committee INCITS during 2007-200x. The final approval process started in 201x.

Requests for interpretation, suggestions for improvements or addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, Information Technology Industry Council, 1101 K Street NW Suite 610, Washington, DC 20005-3922.

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

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

(to be filled in by INCITS)

Technical Committee T11 on Lower Level Interfaces, that reviewed this technical report, had the following members:

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

(to be filled in by INCITS)

Task Group T11.3 on Fibre Channel Protocols, that developed and reviewed this technical report, had the following members:

Craig W. Carlson, Chair
Bill Martin, Vice-Chair
Landon Noll, Secretary

(to be filled in by INCITS)

Table of Contents Page

List of Figures Page

List of Tables

Page

draft proposed INCITS Technical Report
for Information Technology—

Fibre Channel —
Device Attach - 2 (FC-DA-2)

1 Introduction and scope

This technical report defines a model for end devices (i.e., Nx_Ports). This technical report is intended to serve 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 ANSI/INCITS standards.

This technical report is divided into 5 clauses:

Clause 1 is the introduction and scope of this technical report.

Clause 2 enumerates the normative references that apply to this technical report.

Clause 3 describes the definitions, abbreviations, and conventions used in this technical report.

Clause 4 specifies the end device model consisting of arbitrated loop behaviors (Public and Private), initialization, naming, login and login validation, arbitrated loop features, FC framing and signaling features, Link Services, well-known address usage, Name Server queries, N_Port_ID Virtualization, CS_CTL/Priority header field usage, logout, management and discovery, and Fabric event notification.

Clause 5 specifies applicable FC-4 behavior for an end device. Currently FCP, FC-SB-4, and IP behaviors are specified.

This technical report has one annex:

Annex A is an informative annex that provides examples and rationale for the discovery and management processes described in the end device model.

This technical report in combination with technical report FC-MI-3 cover the same scope as, but does not replace, technical report FC-MI. This technical report includes, corrects, and extends to cover more recent standards, the guidelines specified by FC-MI for end devices. Technical report FC-MI-3 serves the same purpose for fabric elements.

Another 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 ANSI/INCITS standards 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 ANSI/INCITS standards require for compliance may be optional in this technical report. Each specification of such an optional feature in this technical report identifies the referenced ANSI/INCITS standards for which the feature is required.

Internal characteristics of conformant implementations are not defined by this technical report. This technical report incorporates features from the standards described in clause 2.

2 Normative references

2.1 Overview

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

For electronic copies of INCITS standards and technical reports listed in 2.2, see the INCITS online store at http://www.incits.org.

2.2 Approved references

ANSI/INCITS 272-1996, *Fibre Channel - Arbitrated Loop (FC-AL)*

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

ANSI/INCITS 355-2001, *Fibre Channel - Switch Fabric - 5 (FC-SW-5)*

ANSI/INCITS TR-30-2002, *Fibre Channel - Methodologies for Interconnects (FC-MI)*

ANSI/INCITS 352-2002, *Fibre Channel - Physical Interface (FC-PI)*

ANSI/INCITS 374-2003, *Fibre Channel - Single-Byte Command Code Sets Mapping Protocol - 3 (FC-SB-3)*

ANSI/INCITS 386-2004, *Fibre Channel - HBA API (FC-HBA)*

2.3 References under development

At the time of publication, the following referenced standards were still under development. For information about obtaining copies of the following documents or for more information on the current status of these documents, contact INCITS at http://www.incits.org or visit the T11 website at http://www.t11.org.

ANSI/INCITS xxx-201x, *Fibre Channel - Generic Services - 7 (FC-GS-7),* T11/Project 2204-D

ANSI/INCITS xxx-201x, *Fibre Channel - Methodologies for Interconnects - 3 (FC-MI-3),* T11/Project 2217-DT

ANSI/INCITS xxx-200x, *Fibre Channel Protocol for SCSI - 4 (FCP-4),* T10/Project 1828-D

ANSI/INCITS xxx-200x, *Fibre Channel - Link Services - 2 (FC-LS-2),* T11/Project 2103-D

ANSI/INCITS xxx-200x, *Fibre Channel - Framing and Signaling - 3 (FC-FS-3),* T11/Project 1861-D

ANSI/INCITS xxx-201x, *Fibre Channel - Security Protocols - 2 (FC-SP-2),* T11/Project 1835-D

3 Definitions and conventions

3.1 Overview

For this technical report, the following definitions, conventions, abbreviations, acronyms, symbols, and keywords apply.

3.2 Definitions

3.2.1 address identifier: An address value used to identify source (S_ID) or destination (D_ID) of a frame.

3.2.2 AL_PA position map: A map that shows device ordering on an Arbitrated Loop (see FC-AL-2).

3.2.3 Arbitrated Loop time out value (AL_Time): Arbitrated Loop time out value (see FC-AL-2).

3.2.4 Application: An entity that makes requests of a Service.

3.2.5 Arbitrated Loop Physical Address (AL_PA): A one-byte address value (see FC-AL-2).

3.2.6 Area: The second level in a three-level addressing hierarchy (see FC-SW-3).

3.2.7 Directory: A repository of information about objects that may be accessed using the Directory Service (see FC-GS-4).

3.2.8 Domain: The highest level in a three-level addressing hierarchy (see FC-SW-3).

3.2.9 Error_Detect_Timeout value (E_D_TOV): An error detection time constant (see FC-FS-3).

3.2.10 F_Port: In this technical report, an F_Port always refers to a port to which non-loop N_Ports are attached to a Fabric, and does not include FL_Ports.

3.2.11 Fabric: An entity that interconnects various Nx_Ports attached to it, and is capable of routing frames using only the D_ID information in an FC-2 frame header.

3.2.12 Fabric Name: An identifier associated with a Fibre Channel Fabric (see FC-FS-3).

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

3.2.14 FL_Port: An L_Port that performs the function of an F_Port when attached to one or more NL_Ports in an Arbitrated Loop topology (see FC-AL-2). The AL_PA of an FL_Port is 00h. In this technical report, an FL_Port always refers to a port to which NL_Ports are attached to a Fabric, and does not include F_Ports.

3.2.15 Fx_Port: A Switch Port capable of operating as an F_Port or FL_Port.

3.2.16 Hub: A device that interconnects L_Ports (see FC-FS-3).

3.2.17 L_Port: An FC_Port that contains Arbitrated Loop functions associated with the Arbitrated Loop topology (see FC-AL-2).

3.2.18 Loop Fabric Address (LFA): An address identifier used to address an FL_Port (see FC-SW-3) for the purpose of loop management (see FC-LS-2).

3.2.19 Name_Identifier: A 64-bit identifier, with a 60-bit value preceded by a 4-bit Network_Address_Authority_Identifier, used to identify entities in Fibre Channel (e.g., Nx_Port, Node, F_Port, or Fabric) (see FC-FS-3).

3.2.20 Name Server: A Fibre Channel Service that allows registration of various attributes (see FC-GS-4).

3.2.21 N_Port: In this technical report, an N_Port always refers to a direct Fabric-attached port, and does not include NL_Ports (see FC-FS-3).

3.2.22 N_Port_Name: A Name_Identifier (see 3.2.19) associated with an Nx_Port (see FC-FS-3).

3.2.23 N_Port Identifier: A Fabric unique address identifier by which an Nx_Port is uniquely known. The identifier is used in the S_ID and D_ID fields of a frame.

3.2.24 NL_Port: An L_Port that is able to perform the function of an N_Port when attached to one or more NL_Ports and zero or more FL_Ports in an Arbitrated Loop topology. In this technical report, an NL_Port always refers to a loop-attached port including both Private NL_Ports and Public NL_Ports, and does not include N_Ports.

3.2.25 Node: A collection of one or more Nx_Ports controlled by a level above FC-2 (see FC-FS-3).

3.2.26 Node_Name: A Name_Identifier (see 3.2.19) associated with a Node (see FC-FS-3).

3.2.27 Non-Participating mode: The operational mode of an L_Port that does not have an AL_PA, but is enabled into the Loop (see FC-AL-2).

3.2.28 Nx_Port: A Port operating as an N_Port or NL_Port.

3.2.29 OLD-PORT: The state of a pair of ports where two devices operate in a point-to-point mode utilizing FC-FS-3 protocols instead of FC-AL-2 protocols (see FC-AL-2 and FC-FS-3).

3.2.30 OPEN_INIT: A state in the LPSM (see FC-AL-2).

3.2.31 Participating mode: The operational mode of an L_Port that has an AL_PA and is enabled into the Loop (see FC-AL-2).

3.2.32 Platform: An association of one or more Nodes for the purpose of discovery and management.

3.2.33 Port: The lowest level in a three-level addressing hierarchy (see FC-SW-3).

3.2.34 Port Identifier: A Name Server object containing the Fibre Channel address identifier assigned to an N_Port or NL_Port (see FC-GS-4).

3.2.35 Private NL_Port: An NL_Port that observes the rules of private loop behavior (see FC-AL-2).

3.2.36 Public NL_Port: An NL_Port that attempts a fabric login and is permitted to open AL_PA=00h. A Public NL_Port observes the rules of either public or private loop behavior (see FC-AL-2).

3.2.37 Receiver_Transmitter_Timeout value (R_T_TOV): A receiver/transmitter time constant (see FC-FS-3).

3.2.38 Resource_Allocation_Timeout value (R_A_TOV): A resource allocation time constant (see FC-FS-3).

3.2.39 Resource_Recovery_Timeout value (RR_TOV): A resource recovery time constant (see FCP-4).

3.2.40 Service: A Service is provided by a Node, accessible through an N_Port that is addressed by a Well-Known Address or an N_Port Identifier (e.g., the Directory Service) (see FC-GS-6).

3.2.41 Speed Negotiation: A process that determines a common operating speed between an FC_Port capable of multiple operating speeds that is connected by a common FC infrastructure to another FC_Port that may also be capable of multiple operating speeds (see FC-FS-3).

3.2.42 Switch: An element that makes up the switching portion of the Fabric (see FC-SW-5).

3.2.43 Switch Port: An E_Port, F_Port, or FL_Port.

3.2.44 Well-Known Address: An address identifier defined to access a Service (see FC-FS-3).

3.3 Editorial conventions

In this technical report 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 priority relationship between the listed items. Numbered lists (e.g., 1-red, 2-blue, 3-green) show a priority ordering between the listed items.

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

Table 1 – ISO and American conventions

ISO/British	ISO/French	American
0.6	0,6	0.6
3.14159265	3,14159265	3.14159265
1 000	1 000	1,000
1 323 462.9	1 323 462,9	1,323,462.9

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

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

If a field or a control bit in a frame is specified as not meaningful, the entity that receives the frame shall not check that field or control bit.

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

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

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

To indicate a bit field, the convention (MSB bit number: LSB bit number) is used (e.g., 7:0).

3.4 Abbreviations and acronyms

Abbreviations and acronyms applicable to this technical report are listed.

ELS	Extended Link Service
FC	Fibre Channel
FC-CT	Common Transport for Fibre Channel Service
FCP	Fibre Channel Protocol for SCSI
IP	Internet Protocol
IU	Information Unit
LFA	Loop Fabric Address
LIP	Loop Initialization Primitive Sequence
LIPr	Loop Initialization - reset L_Port
LSB	Least Significant Bit
MSB	Most Significant Bit
NFS	Network File Server
NPIV	N_Port_ID Virtualization
ULP	Upper Level Protocol
2xR_A_TOV	two times R_A_TOV

3.5 Symbols

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

||	concatenation
=	assignment
!=	not equal
>	greater than
<=	less than or equal to

3.6 Keywords

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

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

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

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

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

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

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

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

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

3.6.10 restricted: A keyword referring to bits, bytes, words, and fields that are set aside for use in other Fibre Channel standards referenced by this technical report. 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.

3.6.11 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 technical report.

3.6.12 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 sup-

port 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.

In the event of conflicts between this technical report and other technical reports, the resolution of those conflicts is beyond the scope of this technical report.

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 of Public or Private NL_Ports.

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 that 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 or feature is not meaningful.

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 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 End device model

4.1 Loop behaviors

4.1.1 Loop initialization

4.1.1.1 Power-on behavior

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

NOTE 1 – The delay of 200 microseconds and the enforcement of limits on disruption recognizes that: a) many NL_Ports by design, disrupt during power-on; and b) most hubs isolate an attached NL_Port until it has presented at least 200 microseconds of FC-PI-compliant signal. The expectation is that any disruption caused by an NL_Port at power-on is harmless until after it has presented 200 microseconds of valid signal, because until then it is isolated from the loop by its hub.

FC-FS-3 specifies the disruption of a loop that may result from attaching to it an NL_Port configured to execute Speed Negotiation. An NL_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-3 for definitions and derivations of t_disrupt1 and t_disrupt2.

4.1.1.2 Loop failure

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

4.1.1.3 Initialization at power-on

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

An NL_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 NL_Port originating or passing through LISM frames, such an NL_Port is subject to the limits of this subclause.

4.1.1.4 NL_Port time-out during initialization

It may be necessary for an NL_Port to request loop initialization multiple times. An NL_Port operating in compliance with FC-AL may repeatedly request initialization using the following rules:

1. if loop initialization has not completed within 2 seconds;
2. the second time the NL_Port shall increase the time-out to greater than 24 seconds; and
3. the third and subsequent times the NL_Port shall increase the time-out to greater than 128 seconds.

An NL_Port operating in compliance with FC-AL-2 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 in this subclause for an NL_Port operating in compliance with FC-AL.

4.1.1.5 Loop initialization LIP generation

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

4.1.1.6 Response to LIP

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

An NL_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 NL_Port originating or passing through LISM frames, such an NL_Port is subject to the limits of this subclause.

4.1.1.7 Origination of LISM frames

An NL_Port that supports LIM and has not received a higher priority LISM shall originate LISM frames with a maximum of 5 milliseconds between LISM frames. An NL_Port operating in compliance with FC-AL may source IDLEs for 15 milliseconds before it begins sourcing LISMs.

4.1.1.8 Forwarding of LISM frames

An NL_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) 5 milliseconds; or

b) 2 times the delay between received LISM frames.

4.1.1.9 Address selection

After any of the following an NL_Port shall take an address during the LIHA or LISA phase, if an address is available:

a) a power cycle or equivalent reset of the device;

b) LIPr directed to the NL_Port;

c) another NL_Port takes the address in the LIFA or LIPA phase before that NL_Port may take its address (e.g., two loop segments were joined with overlapping addresses); or

d) when the NL_Port transitions from non-participating to participating.

After an NL_Port has taken an address, if the NL_Port has succeeded in FLOGI utilizing that address, the NL_Port shall take its address in the LIFA phase if its address is available. If the NL_Port has not yet performed FLOGI but has originated a PLOGI request or reply utilizing that address, the NL_Port shall take its address in the LIPA phase if its address is available.

If an NL_Port is unable to successfully perform FLOGI after either of the following, an NL_Port shall take an address in the LIPA phase of the next loop initialization, if its address is available:

a) a LISA frame with the Login Required bit set to one is received; or

b) a FAN is not received within E_D_TOV.

If an NL_Port is using the Fibre Channel Port Control (19h) SCSI mode page with the Disable Soft Address (DSA) bit set to one, then it shall exit initialization as non-participating if its hard address is unavailable.

4.1.1.10 Multi-port initialization

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

a) a LIP, other than LIPr, on one NL_Port shall not cause a LIP on any of the other NL_Ports; and

NOTE 2 – Because a LIPr is the equivalent of a power on reset, it may impact both ports on a multi-port device.

b) initialization on one NL_Port shall not cause any other NL_Port port to lose any frames, state, or otherwise impact traffic on its loop.

4.1.1.11 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 NL_Ports operating in compliance with FC-AL-2; and

b) NL_Ports that do not support the AL_PA position map may be present on the loop, therefore NL_Ports shall not rely on AL_PA position map support for functionality.

If all NL_Ports on a loop support AL_PA position map, then the AL_PA position map shall be used to determine what L_Ports are available, and attempting to open L_Ports that were not in the last AL_PA position map is prohibited.

4.1.1.12 Availability after LIP

In order to minimize disruptions on a loop after LIP, NL_Ports that have acquired an address in a previous initialization cycle, shall not become non-participating after a LIP is received if their address or another address is still available and it is not the intention of the NL_Port to go non-participating.

If an NL_Port goes non-participating when a LIP is received and does not participate in loop initialization and discovery after loop initialization, the NL_port shall implicitly logout with all other NL_Ports.

4.1.2 Post initialization

4.1.2.1 LIP generation

4.1.2.1.1 Proper LIP generation

Once an NL_Port has initialized, an NL_Port should minimize the generation of LIPs. Situations where an NL_Port may generate a LIP are as follows:

a) unable to win arbitration for greater than LP_TOV;

b) loss of word sync for greater than R_T_TOV;

c) a loss of signal is detected; or

d) a CLS in response to CLS has not been received within LP_TOV.

4.1.2.1.2 Improper LIP generation

An NL_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; or

e) no R_RDY or CLS received in response to an OPN (see 4.1.4).

4.1.2.2 Multi-port behavior

In the absence of the receipt of commands specified to cause changes in the operation on other Nx_Ports of the device, all Nx_Ports of a multi-port device shall continue to operate as specified in this technical report, independent of what traffic is present on other Nx_Ports.

A multi-port NL_Port device, once initialized and in the absence of any command requiring the device to perform some form of initialization, shall adhere to the following behavior:

a) the establishment of a loop circuit and frame transmission or reception on any NL_Port of the device shall not cause the failure of Exchanges on any of the device's other NL_Ports;

b) the establishment of a loop circuit and frame transmission or reception on any NL_Port of the device shall not interfere with the repeating of Transmission Words or proper generation of Fill Words on the device's other NL_Ports;

c) the establishment of a loop circuit and frame transmission or reception on any NL_Port of the device shall not result in any of the device's other NL_Ports going non-participating;

d) the establishment of a loop circuit and frame transmission or reception on any NL_Port of the device shall not result in failure to process Primitive Signals or Primitive Sequences on any of the device's other NL_Ports (e.g., processing of OPN, LIP, or CLS); and

e) the establishment of a loop circuit and frame transmission or reception on any NL_Port of the device shall not result in any of the device's other NL_Ports originating a LIP.

4.1.3 Receipt of CLS in response to OPN

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. The proper response to this condition is for the originator of the OPN to forward the CLS and when the arbitration window is again available arbitrate and attempt to OPN this device. Because this condition is cleared by this procedure, the Originator of the OPN shall not originate a LIP to clear this condition.

4.1.4 No response to OPN

The lack of any response to an OPN indicates a failure occurred (e.g., the OPN was corrupted on the loop, or the R_RDY or CLS from the device was corrupted). These conditions may be cleared by

transmitting a CLS and then attempting to re-establish the circuit. 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 may be cleared by this procedure, the originator of the OPN shall not originate a LIP to clear this condition. If the CLS is lost, then behavior as described in 4.1.2.1.1 item d) is applicable.

4.1.5 Broadcast and multicast

An FL_Port shall not receive frames delivered to the Local Loop by an NL_Port using OPNfr or OPNyr (see FC-AL-2). An FL_Port may receive broadcast and multicast frames from an attached NL_Port during a Loop Tenancy initiated by OPNyx or OPNyy, and shall deliver these frames to the addressed destination group, including the Local Loop if any ports on the Local Loop are part of the addressed destination group.

4.1.6 NL_Port operation

An NL_Port attached to a Public Loop is expected to be either:

a) a Private NL_Port, which communicates only with other NL_Ports on the Local Loop, and does not communicate with the FL_Port; or

b) a Public NL_Port, which may communicate with other NL_Ports on the Local Loop, and may communicate with the FL_Port for purposes of communicating with Remote Ports.

NL_Ports shall be able to receive frames in a multiplexed manner (i.e., frames from one Sequence may be mixed with frames from other Sequences, and possibly from other sources, even within the same Loop Tenancy). An NL_Port may, however, expect frames to be delivered by the Fabric in-order from a single source (i.e., they are delivered to the destination in the order they were received from the source).

4.2 Public Loop and Private Loop behavior

4.2.1 Overview

A port that only exhibits Private Loop behavior is called a Private NL_Port. A port that exhibits Public behavior is called a Public NL_Port, even though it may communicate with Private NL_Ports. For example, a Public NL_Port in figure 1 may be an NFS server which communicates with NFS Clients residing directly on the fabric using IP, and with SCSI Targets that are Local Ports on the same loop using FCP.

Devices with only Private NL_Ports are called private loop devices. A private loop device is Prohibited from providing or requesting fabric services. Devices with at least one Public NL_Port are called public loop devices.

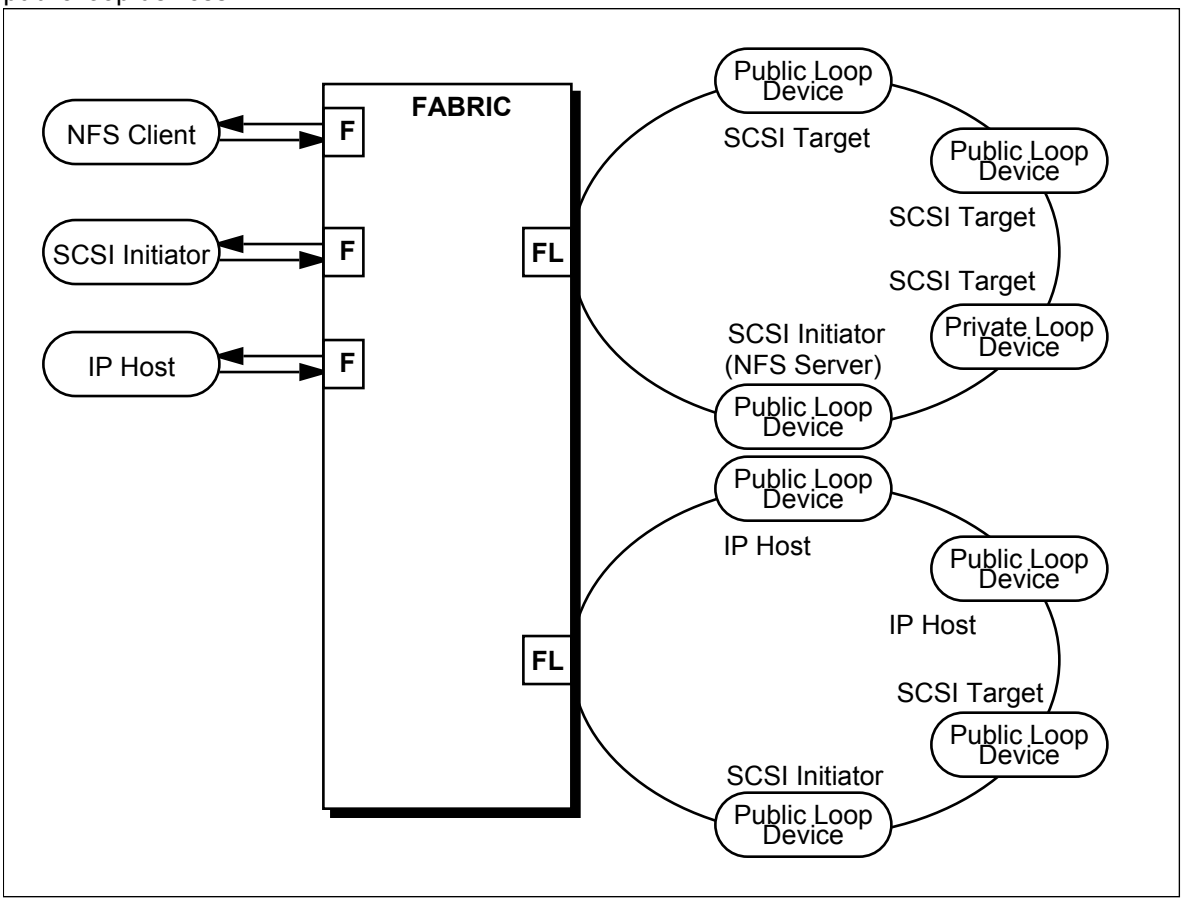

Figure 1 — Public Loop and Private Loop device coexistence

4.2.2 Public NL_Port and Private NL_Port behavior

F/NL_Port operation is prohibited. Public NL_Port and Private NL_Port behavior is described in table 5.

Table 5 – Public and Private NL_Port behavior

Behavior	Public NL_Port	Private NL_Port
Domain + Area of device's NL_Port ID=0000h	Prohibited	Required
NL_Port may Open AL_PA=00h	Required	Prohibited
NL_Port may be Opened by FL_Port	Required	Prohibited
Public NL_Port may Open any Local NL_Port	Allowed	Allowed
Private NL_Port may Open any Local NL_Port	Allowed	Allowed

4.2.3 Public and Private device addressing

An N_Port identifier is divided into three 8 bit fields. In order of bit significance these are: Domain, Area, and Port. The value of the Port field in the N_Port Identifier for an NL_Port shall be the AL_PA of the NL_Port. No two ports on the same Local Loop shall share the same AL_PA. All Public NL_Ports on the same Local Loop shall have the same Domain || Area.

The Domain || Area of a Private NL_Port is zero and a Private NL_Port shall accept frames from NL_Ports that have any value for their Domain || Area in the S_ID of the frame header. This does not imply that Private NL_Ports should expect to receive frames across a Fabric.

Table 6 describes the required addressing behavior for Public and Private devices (i.e., how S_ID and D_ID and OPNs shall be created when sending frames). The S_ID and D_ID are created by the sender of the frame, and the sender performs an OPN on the source loop. The FL_Port performs the

OPN on the destination loop. If both source and destination are on the same Local Loop, the OPN shall be performed by the frame sender only.

Table 6 – Public and Private Device addressing

Frame sent from[a,d]	Frame received by[a,d]	OPN AL_PD		S_ID[a]		D_ID[a]	
		source loop	dest. loop	Bits 23:8	Bits 7:0	Bits 23:8	Bits 7:0
Local Public NL_Port	Fabric-Attached N_Port	00h	-	Local D&A	Source AL_PA	Address Identifier	
Fabric-Attached N_Port	Remote Public NL_Port	-	D_ID bits 7:0	Address Identifier		Remote D&A	Dest. AL_PA
Local Public NL_Port	Remote Public NL_Port	00h	D_ID bits 7:0	Local D&A	Source AL_PA	Remote D&A	Dest. AL_PA
Remote Public NL_Port	Local Public NL_Port	00h	D_ID bits 7:0	Remote D&A	Source AL_PA	Local D&A	Dest. AL_PA
Local Public NL_Port	Local Public NL_Port	D_ID bits 7:0[b]		Local D&A	Source AL_PA	Local D&A or 0000h[c]	Dest. AL_PA
Local Public NL_Port	Local Private NL_Port	D_ID bits 7:0		Local D&A	Source AL_PA	0000h	Dest. AL_PA
Local Private NL_Port	Local Public NL_Port	D_ID bits 7:0		0000h	Source AL_PA	Local D&A or 0000h[c]	Dest. AL_PA
Local Private NL_Port	Local Private NL_Port	D_ID bits 7:0		0000h	Source AL_PA	0000h	Dest. AL_PA

a D&A refers to the Domain and Area. Local means the Domain and Area on the Local Loop. Remote means the Domain and Area of the Remote Loop.

b The behavior of an FL_Port when it receives a unicast frame from a Local NL_Port destined for another Local NL_Port is not defined by this technical report.

c A Public NL_Port shall process PLOGI, LOGO, and ABTS frames with D_ID equal to 0000h || AL_PA regardless of whether the NL_Port has a valid Domain and Area from a Fabric Login. This allows Private NL_Ports to discover Public NL_Ports on the same loop and Public NL_Ports to use a common process for logging into Private and Public NL_Ports on the same loop. Public NL_Ports may discard any other frames addressed to 0000h || AL_PA.

d A Public NL_Port is an NL_Port that has successfully completed Fabric Login. A NL_Port that was not able to successfully complete Fabric Login follows the rules for a Private NL_Port.

4.3 Nx_Port initialization

Nx_Ports shall initialize as described in this subclause. Figure 2 shows a flow diagram of the process. N_Ports supporting NPIV shall initialize as described in 4.13.2.

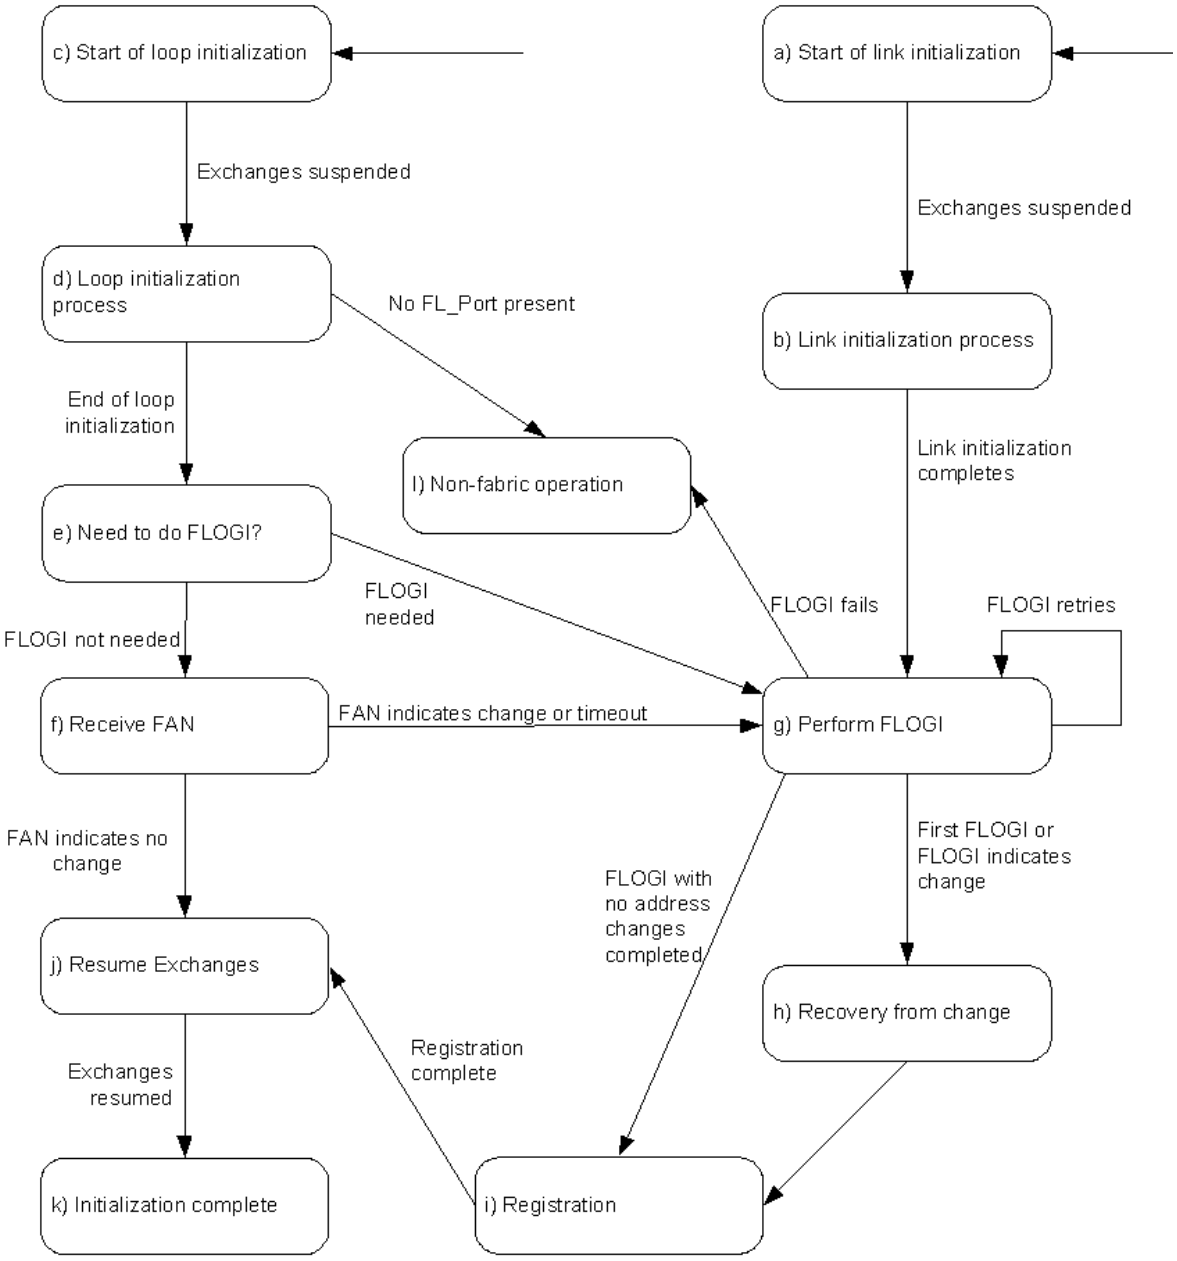

Figure 2 — Nx_Port initialization flow

a) **Start of link initialization**. Whenever an N_Port receives an OLS, NOS, loss of synchronization for greater than R_T_TOV, or loss of signal, the N_Port shall begin the Link Initialization Process (see FC-FS-3). The N_Port shall suspend all open Exchanges with all other Nx_Ports, implicitly logout with the Fabric, and shall proceed to step (b).

b) **Link initialization process.** Upon completion of the Link Initialization process the N_Port shall proceed to step (g).

c) **Start of loop initialization**. Whenever an NL_Port receives a LIP, loss of synchronization for greater than R_T_TOV, or loss of signal, it shall begin the Loop Initialization Process (see FC-AL-2). The NL_Port shall suspend all open Exchanges with all other Nx_Ports, and the NL_Port shall proceed to step (d).

d) **Loop initialization process**. If the NL_Port had an AL_PA prior to the LIP, and the NL_Port had completed FLOGI, then the NL_Port shall attempt to acquire its Fabric-assigned AL_PA during the LIFA Loop Initialization Sequence. If the NL_Port had an AL_PA prior to the LIP, and the NL_Port had not completed FLOGI, then the NL_Port shall attempt to acquire its previously-assigned AL_PA during the LIPA Loop Initialization Sequence. If the NL_Port did not have an AL_PA prior to the LIP, or if the NL_Port was unable to acquire an AL_PA during LIFA or LIPA, the NL_Port shall attempt to acquire an AL_PA during the LIHA or LISA Loop Initialization Sequences, as appropriate. If it is determined that there is an FL_Port, the NL_Port shall proceed to step (e). Otherwise the NL_Port shall proceed to step (l).

e) **Need to do FLOGI?** Upon completion of Loop Initialization, the NL_Port shall implicitly log out with the Fabric and perform FLOGI as described in step (g), if one or more of the following is true:

   A) the L_bit was set to one in the LISA Sequence during Loop Initialization;

   B) the NL_Port did not acquire an AL_PA in the LIFA phase of loop initialization; or

   C) the N_Port_Name or Node_Name of the NL_Port changed to a different value than it had prior to the Loop Initialization.

   Otherwise, the NL_Port shall proceed to step (f).

f) **Receive FAN**. The NL_Port shall wait a minimum of E_D_TOV to receive the FAN (Fabric Address Notification). The NL_Port shall discard all frames received except ADISC/PDISC and PLOGI requests until the FAN is received. The NL_Port shall either receive any ADISC/PDISC and/or PLOGI requests and delay the reply until the FAN is received, or the NL_Port shall respond to ADISC/PDISC and PLOGI request Sequences with an LS_RJT Reply Sequence, with a Reason code of “Unable to perform Command Request at this time”. After receiving the FAN, if the NL_Port determines that the FL_Port has the same Loop Fabric Address, F_Port_Name and Fabric_Name that the FL_Port had before initialization, the NL_Port shall proceed to step (j). Otherwise, the NL_Port shall implicitly log out with the Fabric and perform FLOGI as described in step (g). If the NL_Port does not receive a FAN within E_D_TOV of the completion of Loop Initialization Protocol, it shall also implicitly log out with the Fabric and perform FLOGI as described in step (g).

NOTE 3 – The behavior described in item (f) in which frames, regardless of Class of Service, are discarded prior to FAN being received does not normally happen if the FL_Port is present and sending FAN. In that case, the FAN is always the first frame received following loop initialization.

g) **Perform FLOGI**. On completion of step (b) or if FLOGI is required as described in step (e) or step (f), the Nx_Port shall perform this step.

   The Nx_Port shall attempt to send FLOGI to its local Fx_Port. The Nx_Port Originator shall set the D_ID to the well-known F_Port_ID (i.e., FFFFFEh) and the S_ID to:

A) 000000h; or

B) 0000YYh where YY is set to any value chosen by the N_Port Originator or is set to the AL_PA of the NL_Port Originator.

The payload of the FLOGI Request Sequence shall contain an N_Port_Name and Node_Name associated with the Nx_Port initiating the login. The Common Service Parameters, Class 2 Service Parameters, and Class 3 Service Parameters shall follow the rules as defined in 4.6.

Until the FLOGI completes, if the NL_Port receives a FAN ELS, the NL_Port shall ignore the FAN ELS.

Until the FLOGI completes, if the Nx_Port receives a PLOGI or ADISC, the Nx_Port shall either delay the reply until the FLOGI is complete or respond with an LS_RJT Reply Sequence, with a reason code of “Unable to perform command request at this time”.

If the FLOGI Request attempt failed (e.g., by receiving the OPN sent or receiving an LS_RJT), or if the FLOGI ACC Sequence is not received within 2xR_A_TOV of the FLOGI Request, the Nx_Port shall either:

A) retry step (g);

B) proceed to step (l); or

C) wait for a Loop Initialization in which the L_bit is set to one in the LISA Sequence or a link initialization.

The Nx_Port shall proceed to step (i) if:

A) the FLOGI is completed successfully;

B) the Nx_Port has previously completed a FLOGI;

C) the Nx_Port has the same N_Port_ID, N_Port_Name, and Node_Name that the Nx_Port had before loop initialization or link initialization; and

D) the FL_Port has the same Loop Fabric Address, F_Port_Name and Fabric_Name that the FL_Port had before loop initiaization or the F_Port has the same F_Port_Name and Fabric_Name that the F_Port had before link initialization.

The Nx_Port shall proceed to step (h) if the FLOGI is completed successfully and:

A) the Nx_Port has not previously completed a FLOGI;

B) the Nx_Port does not have the same N_Port_ID, N_Port_Name, and Node_Name that the Nx_Port had before loop initialization or link initialization; or

C) the FL_Port does not have the same Loop Fabric Address, F_Port_Name and Fabric_Name that the FL_Port had before loop initialization or the F_Port does not have the same F_Port_Name and Fabric_Name that the F_Port had before link initialization.

NOTE 4 – It is possible for an Nx_Port to receive ELS frames before any PLOGI frames are received.

h) **Recovery from address change**. An Nx_Port performing this step has determined that its own addressing information and/or that of the Fx_Port have changed, or that it did not have complete addressing information, prior to this initialization. The Nx_Port shall discard all suspended and queued Exchanges in a manner consistent with the FC-4, implicitly logout with all other Nx_Ports, and notify the FC-4. If the Clean Address bit in the FLOGI Accept (LS_ACC) is set to one, Nx_Port shall proceed to step (i); otherwise, the Nx_Port shall wait for R_A_TOV before originating any new Exchanges. During this time, the Nx_Port shall discard all frames received except for the PLOGI and ADISC request Sequences, and LOGO LS_ACC reply Sequences. The Nx_Port shall respond to all PLOGI request Sequences with an LS_RJT Reply Sequence, with a Reason code of “Unable to perform Command Request at this time”. The Nx_Port shall respond to all ADISC request Sequences with a LOGO request Sequence to the Nx_Port that initiated the ADISC request Sequence. After R_A_TOV, the Nx_Port shall proceed to step (i).

i) **Registration**. The Nx_Port shall perform PLOGI explicitly with the Directory Server (i.e., FFFFFCh) and shall perform registration with the Name Server (i.e., GS_Subtype code 02h). The Name Server Registration requests may be performed in any order. The Nx_Port shall proceed to step (j).

j) **Resume Exchanges**. The Nx_Port shall resume all suspended Exchanges, if any, with other Nx_Ports that are known to be public. If an NL_Port has any suspended Exchanges with Private Loop devices, the NL_Port shall perform the Private Loop Exchange authentication process described in 5.1.14, as either the originator or recipient, as appropriate. The Nx_Port shall proceed to step (k).

NOTE 5 – It is possible for an Nx_Port to receive ELS frames before any PLOGI frames are received.

k) **Completion of Nx_Port initialization**. This completes initialization for the Nx_Port. The Nx_Port may proceed to originate and respond to Exchanges, and login with other Nx_Ports as needed.

l) **Fall-back to non-Fabric operation**. The Nx_Port is not connected to a Fabric.

4.4 Nx_Port node and port naming

Nx_Ports complying with this technical report shall have a Node_Name. Node_Names shall be unique among all nodes, however, multiple Nx_Ports may have the same Node_Name if they are associated with the same node. To ensure uniqueness of these names, all names shall use one of the following registered name formats (see FC-FS-3):

a) IEEE 48-bit format;

b) IEEE extended format;

c) IEEE registered format; or

d) EUI-64 mapped.

4.5 NL_Port login validation

All NL_Ports shall acquire a valid AL_PA before performing FLOGI with the Fabric or PLOGI with other NL_Ports.

NL_Ports shall retain the F_Port_Name and Fabric_Name of the Fabric from the FLOGI and associate that information with the Loop Fabric Address of that Fabric. This information shall be retained for as long as the NL_Port is logged in with the Fabric.

NL_Ports shall validate the current Fabric login following every Loop Initialization by comparing the Loop Fabric Address, F_Port_Name, and Fabric_Name received during FLOGI with those reported by the received FAN. All identifiers reported by the FAN shall match the values reported during FLOGI or a configuration change has occurred (i.e., an implicit logout of all Nx_Ports is required and all open Exchanges shall be terminated).

NL_Ports shall retain the Node_Name and N_Port_Name of the other port from each PLOGI and associate that information with the Address Identifier of that NL_Port. This information shall be retained for as long as the NL_Port is logged in with the other port.

Private NL_Ports shall validate NL_Port logins and Public NL_Ports shall validate logins with Private NL_Ports by comparing the N_Port_Name, Node_Name, and Address Identifier received during the PLOGI with those received ADISC/PDISC/LS_ACC that follows the Loop Initialization. All identifiers reported in ADISC/PDISC/LS_ACC shall match the values reported during PLOGI or a configuration change has occurred (i.e., an implicit logout of the affected NL_Port is required and all open Exchanges with that NL_Port shall be terminated).

4.6 Nx_Port login

4.6.1 Class of service support

Table 7 specifies Nx_Port Class of service support with usage defined by this technical report.

Table 7 – Nx_Port Class of service support

Class of service	Nx_Port Originator	Nx_Port Responder
Class 2	A	A
Class 3	I	R

4.6.2 Class of service support for FLOGI and PLOGI

Table 8 specifies Nx_Port Class of service support for FLOGI and PLOGI with usage defined by this technical report.

Table 8 – Nx_Port Class of service support for FLOGI and PLOGI

Class of service	Nx_Port Originator	Nx_Port Responder
FLOGI		
Class 2	A	-[a]
Class 3	I	-[a]
PLOGI		
Class 2	A	A
Class 3	I	R

a See FC-MI-3 for Fx_Port Responder Class of service support.

4.6.3 FLOGI parameters

Table 9 lists Nx_Port Common Service Parameters for FLOGI with usage defined by this technical report. The parameters are valid for both Class 2 and Class 3 delivery service. See FC-MI-3 for the parameters returned by an Fx_Port.

Table 9 – Nx_Port Common Service Parameters (FLOGI)

Common Service Parameter	Nx_Port Originator
FC-PH	
Highest Version	20h[a]
Lowest Version	20h[a]
Buffer-to-Buffer Credit	X
Common Features	
Multiple N_Port_ID Support	X
Virtual Fabrics bit	X
Valid Vendor Version Level	0
N_Port/F_Port	0
BB_Credit Management	X[b]
Name Server Session Begin	X
Query Data Buffer conditions	I
Security bit (see FC-SP-2)	X
Clock Synchronization Primitive Capable	0
R_T_TOV Value	0
Dynamic Half Duplex Supported	X
Payload bit	0
BB_SC_N	X[d]
Buffer-to-Buffer Receive Data Field Size (minimum)	256[c]

a Even though the version numbers are obsolete, this field shall be set as specified.
b For NL_Ports operating in loop mode this value shall be set to one, for N_Ports this value shall be set to zero.
c FC-FS-3 specifies the default Buffer-to-Buffer Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.
d 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.

4.6.4 PLOGI parameters

Table 10 lists Nx_Port Common Service Parameters for PLOGI with usage defined by this technical report. The parameters are valid for both Class 2 and Class 3 delivery service.

Table 10 – Nx_Port Common Service Parameters (PLOGI) (Sheet 1 of 2)

Common Service Parameter	Nx_Port Originator	Nx_Port Responder
FC-PH Version		
Highest Version	20h[a]	20h[a]
Lowest Version	20h[a]	20h[a]
Buffer-to-Buffer Credit	X[d]	X[d]
Common Features		
Continuously Increasing Relative Offset	X[b]	X[b]
Random Relative Offset	X[b]	X[b]
Valid Vendor Version Level	0	0
N_Port/F_Port	0	0
BB_Credit Management	X[c]	X[c]
E_D_TOV Resolution	X[d]	X[d]
Query Data Buffer conditions	I	I
Security bit (see FC-SP-2)	X	X
Clock Synchronization Primitive Capable	0	0
R_T_TOV Value	0	0
Dynamic Half Duplex Supported	X	X
SEQ_CNT	X[e]	X[e]

a Even though the version numbers are obsolete, this field shall be set as specified.
b Use of this field is FC-4 specific.
c For NL_Ports operating in loop mode this value shall be set to one, for N_Ports this value shall be set to zero. For NL_Ports operating as point-to-point devices, this value shall be set to zero.
d Only meaningful when performing PLOGI with other local NL_Ports or a directly attached N_Port.
e The use of a continuously increasing Sequence count is required for streamed Sequences in Class 2 or Class 3. Otherwise the use of continuously increasing Sequence count is optional. Nx_Ports that optionally choose to use the additional information provided by knowledge that Sequence counts are continuously increasing may detect certain types of link failures more immediately than those that do not use this information.
f As specified by the Fabric if present. If no Fabric is present, the default value shall be 2000 milliseconds.
g FC-FS-3 specifies the default Buffer-to-Buffer Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.
h 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.

Table 10 – Nx_Port Common Service Parameters (PLOGI) (Sheet 2 of 2)

Common Service Parameter	Nx_Port Originator	Nx_Port Responder
Payload bit	0	0
BB_SC_N	X[h]	X[h]
Buffer-to-Buffer Receive Data Field Size (minimum)	256[g]	256[g]
Nx_Port Total Concurrent Sequences (minimum)	1	1
Relative Offset by Information Category		
Information Category 1 (Solicited Data)	X[b]	X[b]
All other Information Categories	X[b]	X[b]
E_D_TOV Value	X[d,f]	X[d,f]

a Even though the version numbers are obsolete, this field shall be set as specified.
b Use of this field is FC-4 specific.
c For NL_Ports operating in loop mode this value shall be set to one, for N_Ports this value shall be set to zero. For NL_Ports operating as point-to-point devices, this value shall be set to zero.
d Only meaningful when performing PLOGI with other local NL_Ports or a directly attached N_Port.
e The use of a continuously increasing Sequence count is required for streamed Sequences in Class 2 or Class 3. Otherwise the use of continuously increasing Sequence count is optional. Nx_Ports that optionally choose to use the additional information provided by knowledge that Sequence counts are continuously increasing may detect certain types of link failures more immediately than those that do not use this information.
f As specified by the Fabric if present. If no Fabric is present, the default value shall be 2000 milliseconds.
g FC-FS-3 specifies the default Buffer-to-Buffer Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.
h 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.

4.6.5 Nx_Port Class 2 Service Parameters (PLOGI)

If Class 2 is supported, table 11 lists Class 2 Service Parameters for PLOGI with usage defined by this technical report.

Table 11 – Class 2 Service Parameters (PLOGI) (Sheet 1 of 2)

Class 2 Service Parameter	Value
Class validity	1
Service Options	
Sequential Delivery	-

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
b FC-FS-3 specifies the default Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.

Table 11 – Class 2 Service Parameters (PLOGI) (Sheet 2 of 2)

Class 2 Service Parameter	Value
Priority/Preemption	X[a]
Preference	X
DiffServ QoS	0
Initiator Control	
ACK_0 capable	A
ACK generation assistance	A
Clock synchronization ELS capable	P
Recipient Control	
ACK_0 capable	A
X_ID interlock	X[a]
Clock synchronization ELS capable	A
Reserved	0
Receive Data Field Size (minimum)	256[b]
Concurrent Sequences (minimum)	1
N_Port End-to-end Credit (minimum)	1
Open Sequences per Exchange (minimum)	1

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
b FC-FS-3 specifies the default Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.

4.6.6 Nx_Port Class 2 Service Parameters (FLOGI)

If Class 2 is supported, table 12 lists Class 2 Service Parameters for FLOGI with usage defined by this technical report.

Table 12 – Class 2 Service Parameters (FLOGI) (Sheet 1 of 2)

Class 2 Service Parameter	Value
Class validity	1
Service Options	
Sequential Delivery	1
Priority/Preemption	X[a]
Preference	X
DiffServ QoS	0
Initiator Control	

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.

Table 12 – Class 2 Service Parameters (FLOGI) (Sheet 2 of 2)

Class 2 Service Parameter	Value
ACK_0 capable	-
ACK generation assistance	-
Clock synchronization ELS capable	P
Recipient Control	
ACK_0 capable	-
X_ID interlock	-
Error policy support	-
Categories per Sequence	-
Clock synchronization ELS capable	A
Reserved	0
Receive Data Field Size (minimum)	-
Concurrent Sequences (minimum)	-
N_Port End-to-end Credit (minimum)	-
Open Sequences per Exchange (minimum)	-

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.

4.6.7 Nx_Port Class 3 Service Parameters (PLOGI)

Table 13 lists Class 3 Service Parameters for PLOGI with usage defined by this technical report.

Table 13 – Class 3 Service Parameters (PLOGI) (Sheet 1 of 2)

Class 3 Service Parameter	Value
Class validity	1
Service Options	
Sequential Delivery	-
Priority/Preemption	X[a]
Preference	X
DiffServ QoS	0
Initiator Control	
ACK_0 capable	-
ACK generation assistance	-

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
b FC-FS-3 specifies the default Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.

Table 13 – Class 3 Service Parameters (PLOGI) (Sheet 2 of 2)

Class 3 Service Parameter	Value
Clock synchronization ELS capable	P
Recipient Control	
ACK_0 capable	-
X_ID interlock	-
Error policy support	00b
Categories per Sequence	00b
Clock synchronization ELS capable	A
Reserved	0
Receive Data Field Size (minimum)	256[b]
Concurrent Sequences (minimum)	1
N_Port End-to-end Credit	-
Open Sequences per Exchange (minimum)	1

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
b FC-FS-3 specifies the default Receive Data Field Size to be set to 128 bytes. This technical report specifies a larger value for login.

4.6.8 Nx_Port Class 3 Service Parameters (FLOGI)

Table 14 lists Class 3 Service Parameters for FLOGI with usage defined by this technical report.

Table 14 – Class 3 Service Parameters (FLOGI) (Sheet 1 of 2)

Class 3 Service Parameter	Value
Class validity	1
Service Options	
Sequential Delivery	1
Priority/Preemption	X[a]
Preference	X
DiffServ QoS	0
Initiator Control	
ACK_0 capable	-
ACK generation assistance	-
Clock synchronization ELS capable	P
Recipient Control	

a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.

Table 14 – Class 3 Service Parameters (FLOGI) (Sheet 2 of 2)

Class 3 Service Parameter	Value
ACK_0 capable	-
X_ID interlock	-
Error policy support	-
Categories per Sequence	-
Clock synchronization ELS capable	A
Reserved	0
Receive Data Field Size (minimum)	-
Concurrent Sequences (minimum)	-
N_Port End-to-end Credit	-
Open Sequences per Exchange (minimum)	-
a FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.	

4.7 FC-AL-2 features for NL_Ports

Table 15 lists Fibre Channel Arbitrated Loop features for NL_Ports with usage defined by this technical report. In table 15, “NL_Port Originator” refers to an NL_Port sending an FC-AL-2 feature to an FL_Port or another Local NL_Port, and “NL_Port Responder” refers to an NL_Port receiving an FC-AL-2 feature from an FL_Port or another Local NL_Port. See FC-MI-3 for FC-AL and FC-AL-2 requirements for FL_Ports.

Table 15 – FC-AL-2 features for NL_Ports (Sheet 1 of 3)

Feature	NL_Port Originator	NL_Port Responder
Open Full Duplex (OPNyx)		
Open NL_Port on Local Loop	I	R
Open FL_Port	I	-
Open Half Duplex (OPNyy)		
Open NL_Port on Local Loop	I	R
Open FL_Port	I	-
Send frames to FL_Port for more than one D_ID in a single Loop Tenancy	I	-
Receive frames from FL_Port for more than one S_ID in a single Loop Tenancy	-	R
Unfairness[a]	I	R
Transfer mode[a] (use of TRANSFER loop state)	I	R

a NL_Ports shall tolerate the use of Unfairness and Transfer Mode by other L_Ports.
b NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. An FL_Port shall not receive frames associated with OPNfr or OPNyr (see FC-AL-2). Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop. Use of broadcast and multicast is specific to the FC-4. The method by which an NL_Port may discover the ability of other NL_Ports to receive OPNfr or OPNyr frames is not defined by this technical report.
c Alternate BB_Credit management is mandatory in FC-AL.
d NL_Ports shall interoperate with NL_Ports and FL_Ports which advertise any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit>0, but the originator of frames is not Required to take advantage of it.
e This LIP may be issued by an NL_Port to request an AL_PA if it has none, or if it is unable to win arbitration within LP_TOV. NL_Port response to a LIP is described in 4.1.1.6.
f This LIP may be issued by an NL_Port that detects a Link Failure (see FC-FS-3).
g An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in an FC-4 specific manner. In the absence of an FC-4 defined reset, a Public NL_Port shall respond to this form of LIP in the same manner as LIP(F7,AL_PS).
h An NL_Port that receives MRK shall attempt to forward the MRK. An NL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).
i 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). Note that this could cause the frames to be delivered to the Local Loop if any NL_Ports on the Local Loop are part of the broadcast or multicast group.
j The only method Allowed by this technical report of delivering broadcast or multicast 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 or multicast frames via OPNfr. Use of broadcast and multicast, and the number of frames that may be sent following OPNfr, are both specific to the FC-4.

Table 15 – FC-AL-2 features for NL_Ports (Sheet 2 of 3)

Feature	NL_Port Originator	NL_Port Responder
Broadcast and Multicast via Broadcast Replicate (OPNfr)		
To NL_Ports (only) on the Local Loop[b]	A	-
From NL_Ports (only) on the Local Loop[b]	-	A
To Remote Ports[i]	P	-
From Remote Ports[j]	-	A
Broadcast and Multicast via Selective Replicate (OPNyr)		
To NL_Ports (only) on the Local Loop[b]	A	-
From NL_Ports (only) on the Local Loop[b]	-	A
To Remote Ports[i]	P	-
From Remote Ports[j]	-	P
Alternate BB_Credit model[c]	R	R
Login_BB_Credit=0[d]	R	R
Login_BB_Credit>0[d]	A	A

a NL_Ports shall tolerate the use of Unfairness and Transfer Mode by other L_Ports.
b NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. An FL_Port shall not receive frames associated with OPNfr or OPNyr (see FC-AL-2). Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop. Use of broadcast and multicast is specific to the FC-4. The method by which an NL_Port may discover the ability of other NL_Ports to receive OPNfr or OPNyr frames is not defined by this technical report.
c Alternate BB_Credit management is mandatory in FC-AL.
d NL_Ports shall interoperate with NL_Ports and FL_Ports which advertise any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit>0, but the originator of frames is not Required to take advantage of it.
e This LIP may be issued by an NL_Port to request an AL_PA if it has none, or if it is unable to win arbitration within LP_TOV. NL_Port response to a LIP is described in 4.1.1.6.
f This LIP may be issued by an NL_Port that detects a Link Failure (see FC-FS-3).
g An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in an FC-4 specific manner. In the absence of an FC-4 defined reset, a Public NL_Port shall respond to this form of LIP in the same manner as LIP(F7,AL_PS).
h An NL_Port that receives MRK shall attempt to forward the MRK. An NL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).
i 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). Note that this could cause the frames to be delivered to the Local Loop if any NL_Ports on the Local Loop are part of the broadcast or multicast group.
j The only method Allowed by this technical report of delivering broadcast or multicast 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 or multicast frames via OPNfr. Use of broadcast and multicast, and the number of frames that may be sent following OPNfr, are both specific to the FC-4.

Table 15 – FC-AL-2 features for NL_Ports (Sheet 3 of 3)

Feature	NL_Port Originator	NL_Port Responder
LPEyx/LPByx	A	A
LIP		
(F7, F7) and (F7, AL_PS) [initializing][e]	I	R
(F8, F7) and (F8, AL_PS) [loop failure][f]	I	R
(AL_PD, AL_PS) [selective hard reset][g]	I	R
MRK[h]	P	R

a NL_Ports shall tolerate the use of Unfairness and Transfer Mode by other L_Ports.
b NL_Ports may use Broadcast Replicate (OPNfr) or Selective Replicate (OPNyr) to communicate on a Local Loop. An FL_Port shall not receive frames associated with OPNfr or OPNyr (see FC-AL-2). Therefore, frames originated on a Local Loop by an NL_Port using OPNfr or OPNyr shall not propagate beyond the Local Loop. Use of broadcast and multicast is specific to the FC-4. The method by which an NL_Port may discover the ability of other NL_Ports to receive OPNfr or OPNyr frames is not defined by this technical report.
c Alternate BB_Credit management is mandatory in FC-AL.
d NL_Ports shall interoperate with NL_Ports and FL_Ports which advertise any Login_BB_Credit. A recipient of frames may login with Login_BB_Credit>0, but the originator of frames is not Required to take advantage of it.
e This LIP may be issued by an NL_Port to request an AL_PA if it has none, or if it is unable to win arbitration within LP_TOV. NL_Port response to a LIP is described in 4.1.1.6.
f This LIP may be issued by an NL_Port that detects a Link Failure (see FC-FS-3).
g An NL_Port or FL_Port may invoke this form of LIP to reset an NL_Port in an FC-4 specific manner. In the absence of an FC-4 defined reset, a Public NL_Port shall respond to this form of LIP in the same manner as LIP(F7,AL_PS).
h An NL_Port that receives MRK shall attempt to forward the MRK. An NL_Port may remove a MRK if necessary for clock skew management (see FC-AL-2).
i 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). Note that this could cause the frames to be delivered to the Local Loop if any NL_Ports on the Local Loop are part of the broadcast or multicast group.
j The only method Allowed by this technical report of delivering broadcast or multicast 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 or multicast frames via OPNfr. Use of broadcast and multicast, and the number of frames that may be sent following OPNfr, are both specific to the FC-4.

4.8 Other Nx_Port FC-FS-3 and FC-LS-2 features

Table 16 lists other FC-FS-3 and FC-LS-2 features not covered in previous clauses for Nx_Ports with usage defined by this technical report. In table 16, “Nx_Port Originator” refers to an Nx_Port sending

a Frame or Sequence to another Nx_Port, and “Nx_Port Responder” refers to an Nx_Port receiving a Frame or Sequence from another Nx_Port.

Table 16 – Other FC-FS-3 and FC-LS-2 features for Nx_Ports (Sheet 1 of 2)

Feature	Nx_Port Originator	Nx_Port Responder
Routing Control (R_CTL)		
FC-4 Device_Data frame	R	R
Extended Link_Data frame	R	R
Basic Link_ Data frame	R	R
Link_Control frame[a]	R	R
TYPE field set to any value	A	A
Frame Control (F_CTL)		
Abort Sequence Condition bits by Sequence Initiator	X[e]	-
Abort Sequence Condition bits by Sequence Recipient	-	X[e]
ACK_Form in Class 2[b]		
ACK_Form = no assist	A	A
ACK_Form = ACK_1 expected	A	A
ACK_Form = ACK_0 expected	A	A
Sequence Retransmission	P	P
Nonzero Continue Sequence Condition values	P	-
Ignore nonzero Continue Sequence Condition values	-	R
Link Control frames[a]		
ACK_0	A	A
ACK_1	I	R
F_BSY	P	R
F_RJT	P	R
P_BSY	I	R
P_RJT	I	R
LCR	R	R
Optional Headers (all)[c]	A	A

a Link Control frames are required only for Class 2. A Class 3-only Nx_Port shall be able to return a P_RJT in Class 2.
b Values of 01b and 11b are allowed based on the result of login.
c The use of optional headers is specific to the FC-4 or security policy requirements.
d The size of a frame that is not the last frame of an Information Category shall be a multiple of four (see FC-FS-3). A frame length of zero is allowed.
e FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.

Table 16 – Other FC-FS-3 and FC-LS-2 features for Nx_Ports (Sheet 2 of 2)

Feature	Nx_Port Originator	Nx_Port Responder
Payload size when frame is not the last frame of a sequence[d]	(size MOD 4) == 0	(size MOD 4) == 0

a Link Control frames are required only for Class 2. A Class 3-only Nx_Port shall be able to return a P_RJT in Class 2.
b Values of 01b and 11b are allowed based on the result of login.
c The use of optional headers is specific to the FC-4 or security policy requirements.
d The size of a frame that is not the last frame of an Information Category shall be a multiple of four (see FC-FS-3). A frame length of zero is allowed.
e FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.

4.9 Nx_Port Link Services

4.9.1 Basic Link Services

Table 17 specifies the support required by this technical report for Basic Link Services. Any Basic Link Service not referenced in table 17 shall be prohibited.

Table 17 – Basic Link Services

Name	Abbr.	Nx_Port Originator Support	Nx_Port Responder Support
No Operation	NOP	P	-
Abort Sequence	ABTS	I	R
Basic Accept	BA_ACC	I	R
Basic Reject	BA_RJT	I	R

4.9.2 Extended Link Service requests

Table 18 specifies the support required by this technical report for Extended Link Services. Any Extended Link Service not referenced in table 18 shall be prohibited.

Table 18 – Extended Link Service requests (Sheet 1 of 3)

Name	Abbr.	OpCode	Nx_Port Originator Support	Nx_Port Responder Support
N_Port Login	PLOGI	03h	I	R
F_Port Login	FLOGI	04h	R[a]	-
Logout	LOGO	05h	I	R
Request Sequence Initiative	RSI	0Ah	P	-
Establish Streaming	ESTS	0Bh	P	-
Estimate Credit	ESTC	0Ch	P	-
Advise Credit	ADVC	0Dh	P	-
Read Timeout Value	RTV	0Eh	P	-
Read Link Error Status Block	RLS	0Fh	I	R
Echo	ECHO	10h	A[i]	A
Test	TEST	11h	P	-
Reinstate Recovery Qualifier	RRQ	12h	I	R
Read Exchange Concise	REC	13h	-[b]	-[b]
Process Login	PRLI	20h	-[b]	-[b]
Process Logout	PRLO	21h	-[b]	-[b]
Test Process Login State	TPLS	23h	P	-
Third Party Process Logout	TPRLO	24h	-[b]	-[b]
Discover N_Port Service Parameters	PDISC	50h	I[c]	R[c]
Discover F_Port Service Parameters	FDISC	51h	I	-
Discover Address	ADISC	52h	I	R
Read Port Status Block	RPS	56h	A	A

a This is required for Public NL_Ports if a fabric is present and for N_Ports.
b FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
c Use of ADISC is recommended over PDISC.
d There is no reply sequence for this ELS.
e An NL_Port shall not respond with an LS_ACC to a received FAN (see 4.3).
f Only required if the Nx_Port has registered for this service.
g An Nx_Port may LS_RJT this ELS unless the FC-4 layer requires its support.
h An Nx_Port may LS_RJT this ELS if it does not support LIRR.
i Payload size should be limited to less than or equal to 104 bytes.
j A port is required to accept an RPBC if it has set the Query Data Buffer conditions bit to one in its login with the Nx_Port that sent the RPBC; otherwise, it may reject the RPBC.

Table 18 – Extended Link Service requests (Sheet 2 of 3)

Name	Abbr.	OpCode	Nx_Port Originator Support	Nx_Port Responder Support
Read Port List	RPL	57h	A	A
Report Port Buffer Conditions	RPBC	58h	I	R[j]
Fabric Address Notification	FAN	60h	-[e]	
Registered State Change Notification	RSCN	61h	I	R[f]
State Change Registration	SCR	62h	I	R[g]
Report Node FC-4 Types	RNFT	63h	I	R
Clock Synchronization Request	CSR	68h	P	-
Clock Synchronization Update	CSU	69h	P	-
Loop Initialize	LINIT	70h	P	-
Loop Status	LSTS	72h	I	-
Request Node Identification Data	RNID	78h	I	R
Registered Link Incident Report	RLIR	79h	I	R[f]
Link Incident Record Registration	LIRR	7Ah	I	R[h]
Scan Remote Loop	SRL	7Bh	P	-
Set Bit-error Reporting Parameters	SBRP	7Ch	A	-
Report Port Speed Capabilities	RPSC	7Dh	A	A
Query Security Attributes	QSA	7Eh	-[b]	-[b]
Exchange Virtual Fabric Parameters	EVFP	7Fh	A	A
Link Keep Alive	LKA	80h	A	A
Link Cable Beacon	LCB	81h	A	A
Authentication ELS	AUTH_ELS	90h	A	A
Registered Fabric Change Notification	RFCN	97h	-[b]	-[b]
Define FFI Domain Topology Map	FFI_DTM	A0h	P	P
Request FFI Domain Topology Map	FFI_RTM	A1h	P	P
FFI AE Principal Switch Selector	FFI_PSS	A2h	P	P

a This is required for Public NL_Ports if a fabric is present and for N_Ports.
b FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
c Use of ADISC is recommended over PDISC.
d There is no reply sequence for this ELS.
e An NL_Port shall not respond with an LS_ACC to a received FAN (see 4.3).
f Only required if the Nx_Port has registered for this service.
g An Nx_Port may LS_RJT this ELS unless the FC-4 layer requires its support.
h An Nx_Port may LS_RJT this ELS if it does not support LIRR.
i Payload size should be limited to less than or equal to 104 bytes.
j A port is required to accept an RPBC if it has set the Query Data Buffer conditions bit to one in its login with the Nx_Port that sent the RPBC; otherwise, it may reject the RPBC.

Table 18 – Extended Link Service requests (Sheet 3 of 3)

Name	Abbr.	OpCode	Nx_Port Originator Support	Nx_Port Responder Support
FFI Map Update Registration	FFI_MUR	A3h	P	P
FFI Registered Map Update Notification	FFI_RMUN	A4h	P	P
FFI Suspend Map Updates	FFI_SMU	A5h	P	P
FFI Resume Map Updates	FFI_RMU	A6h	P	P

a This is required for Public NL_Ports if a fabric is present and for N_Ports.
b FC-4 implementation specific. Refer to the appropriate FC-4 for support requirements.
c Use of ADISC is recommended over PDISC.
d There is no reply sequence for this ELS.
e An NL_Port shall not respond with an LS_ACC to a received FAN (see 4.3).
f Only required if the Nx_Port has registered for this service.
g An Nx_Port may LS_RJT this ELS unless the FC-4 layer requires its support.
h An Nx_Port may LS_RJT this ELS if it does not support LIRR.
i Payload size should be limited to less than or equal to 104 bytes.
j A port is required to accept an RPBC if it has set the Query Data Buffer conditions bit to one in its login with the Nx_Port that sent the RPBC; otherwise, it may reject the RPBC.

4.9.3 Extended Link Service replies

Table 19 specifies the support required by this technical report for Extended Link Service replies. Table 21 only applies to those supported ELS requests that have valid replies as indicated in table 18.

Table 19 – Extended Link Service replies

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

4.9.4 Link Error Status Block (LESB) support requirements

Table 20 specifies the support required by this technical report for the reporting of LESB values.

Table 20 – LESB support requirements

Word	Error Count	Support
0	Link Failure Count	R
1	Loss-of-Synchronization Count	R
2	Loss-of-Signal Count	A
3	Primitive Sequence Protocol Error	A
4	Invalid transmission Word	R
5	Invalid CRC Count	R

4.10 Well-known address usage by Nx_Ports

Table 21 lists the well-known address ELS requirements for Nx_Ports with usage defined by this technical report. In table 21, “Nx_Port Originator” refers to an Nx_Port sending a request Sequence to an N_Port at a well-known address. Extended Link Services not listed are Prohibited or are FC-4 specific.

Table 21 – ELS requirements for well-known addresses (Sheet 1 of 2)

Feature	Nx_Port Originator
Well-known address FFFFFFh (Broadcast)	
FARP-REQ	A
Well-known address FFFFFEh (Fabric F_Port)	
ECHO	I
FDISC	I
FLOGI	R[a]
LOGO	I
RLS	I
RRQ	I
Well-known address FFFFFDh (Fabric Controller)	
ECHO	I
LOGO	I
PDISC	A
PLOGI	I
RNID	I

a Explicit FLOGI is required only if a fabric is present.
b If the Management Service is used, explicit PLOGI is required.

Table 21 – ELS requirements for well-known addresses (Sheet 2 of 2)

Feature	Nx_Port Originator
RRQ	I
RSCN	I
SCR	I
Well-known address FFFFFCh (Directory Service)	
ECHO	I
LOGO	I
PDISC	A
PLOGI	R
RRQ	I
Well-known address FFFFFAh (Management Service)	
ECHO	I
LIRR	I
LOGO	I
PDISC	A
PLOGI	I[b]
RLIR	I
RRQ	I
Well-known address FFFCXXh (Domain Controller)	
ECHO	I
LOGO	I
PDISC	A
PLOGI	I
RLS	I
RPSC	A
RRQ	I
SBRP	A

a Explicit FLOGI is required only if a fabric is present.
b If the Management Service is used, explicit PLOGI is required.

4.11 FC-CT Common Request support

FC-CT Common Request support is summarized in table 22.

Table 22 – FC-CT Common Request support

OpCode	Feature	Support
7FF8h	Get More Information (GMI)	I
7FF9h	Server Session Begin (SSB)	I
7FFAh	Server Session End (SSE)	I
7FFBh	Asynchronous Notification (ASYNC)	A

4.12 Nx_Port Name Server Request support

Nx_Port Name Server Request support in summarized in table 23.

Table 23 – Nx_Port Name Server Request support (Sheet 1 of 2)

OpCode	Feature	Support
0300h	De-register All (DA_ID)	I
0100h	Get All Next (GA_NXT)	I
0114h	Get Class of Service (GCS_ID)	I
011Fh	Get FC-4 Features (GFF_ID)	I
0117h	Get FC-4 Types (GFT_ID)	I
011Dh	Get Hard Address (GHA_ID)	P
0101h	Get Identifiers - Scope (GID_A)	P
0113h	Get Node Name (GNN_ID)	I
0173h	Get Node Names (GNN_FT)	I
0180h	Get Node Names (GNN_FF)	I
01D1h	Get Permanent Port Name - Port Identifier (GPPN_ID)	I
0121h	Get Port Identifier (GID_PN)	I
01F2h	Get Port Identifier (GID_DP)	P
01C1h	Get Port Identifiers - Fabric Port Name (GID_FPN)	I
01F1h	Get Port Identifiers (GID_FF)	I
0171h	Get Port Identifiers (GID_FT)	I
0131h	Get Port Identifiers (GID_NN)	I
01A1h	Get Port Identifiers (GID_PT)	I
0112h	Get Port Name (GPN_ID)	I

a An Nx_Port is required to register its FC-4 Type(s).
b An Nx_Port shall register FC-4 Features for each FC-4 mapping it supports that defines a use for the FC-4 Features object.

Table 23 – Nx_Port Name Server Request support (Sheet 2 of 2)

OpCode	Feature	Support
0181h	Get Port Names (GPN_FF)	I
0172h	Get Port Names (GPN_FT)	I
0132h	Get Port Names (GPN_NN)	I
0182h	Get Port Names (GPN_SDFCP)	A
011Ah	Get Port Type (GPT_ID)	I
0139h	Get Symbolic Node Name (GSNN_NN)	I
0118h	Get Symbolic Port Name (GSPN_ID)	I
0214h	Register Class of Service (RCS_ID)	I
021Fh	Register FC-4 Features (RFF_ID)	I[b]
0217h	Register FC-4 Types (RFT_ID)	R[a]
021Dh	Register Hard Address (RHA_ID)	P
0213h	Register Node Name (RNN_ID)	P
0239h	Register Symbolic Node Name (RSNN_NN)	I
0218h	Register Symbolic Port Name (RSPN_ID)	I

a An Nx_Port is required to register its FC-4 Type(s).
b An Nx_Port shall register FC-4 Features for each FC-4 mapping it supports that defines a use for the FC-4 Features object.

4.13 N_Port_ID Virtualization

4.13.1 Overview

This subclause specifies the implementation guidelines for FC end devices to support N_Port_ID Virtualization (NPIV). NPIV provides a Fibre Channel facility for assigning an N_Port_ID to each VN_Port supported by a PN_Port.

4.13.2 N_Port_ID Virtualization acquisition procedure

The NPIV acquisition procedure is the method by which a PN_Port attached to a fabric acquires additional N_Port_ID’s from the Fabric. NPIV is not supported in an arbitrated loop topology. Figure 3 shows a flow diagram of the procedure.

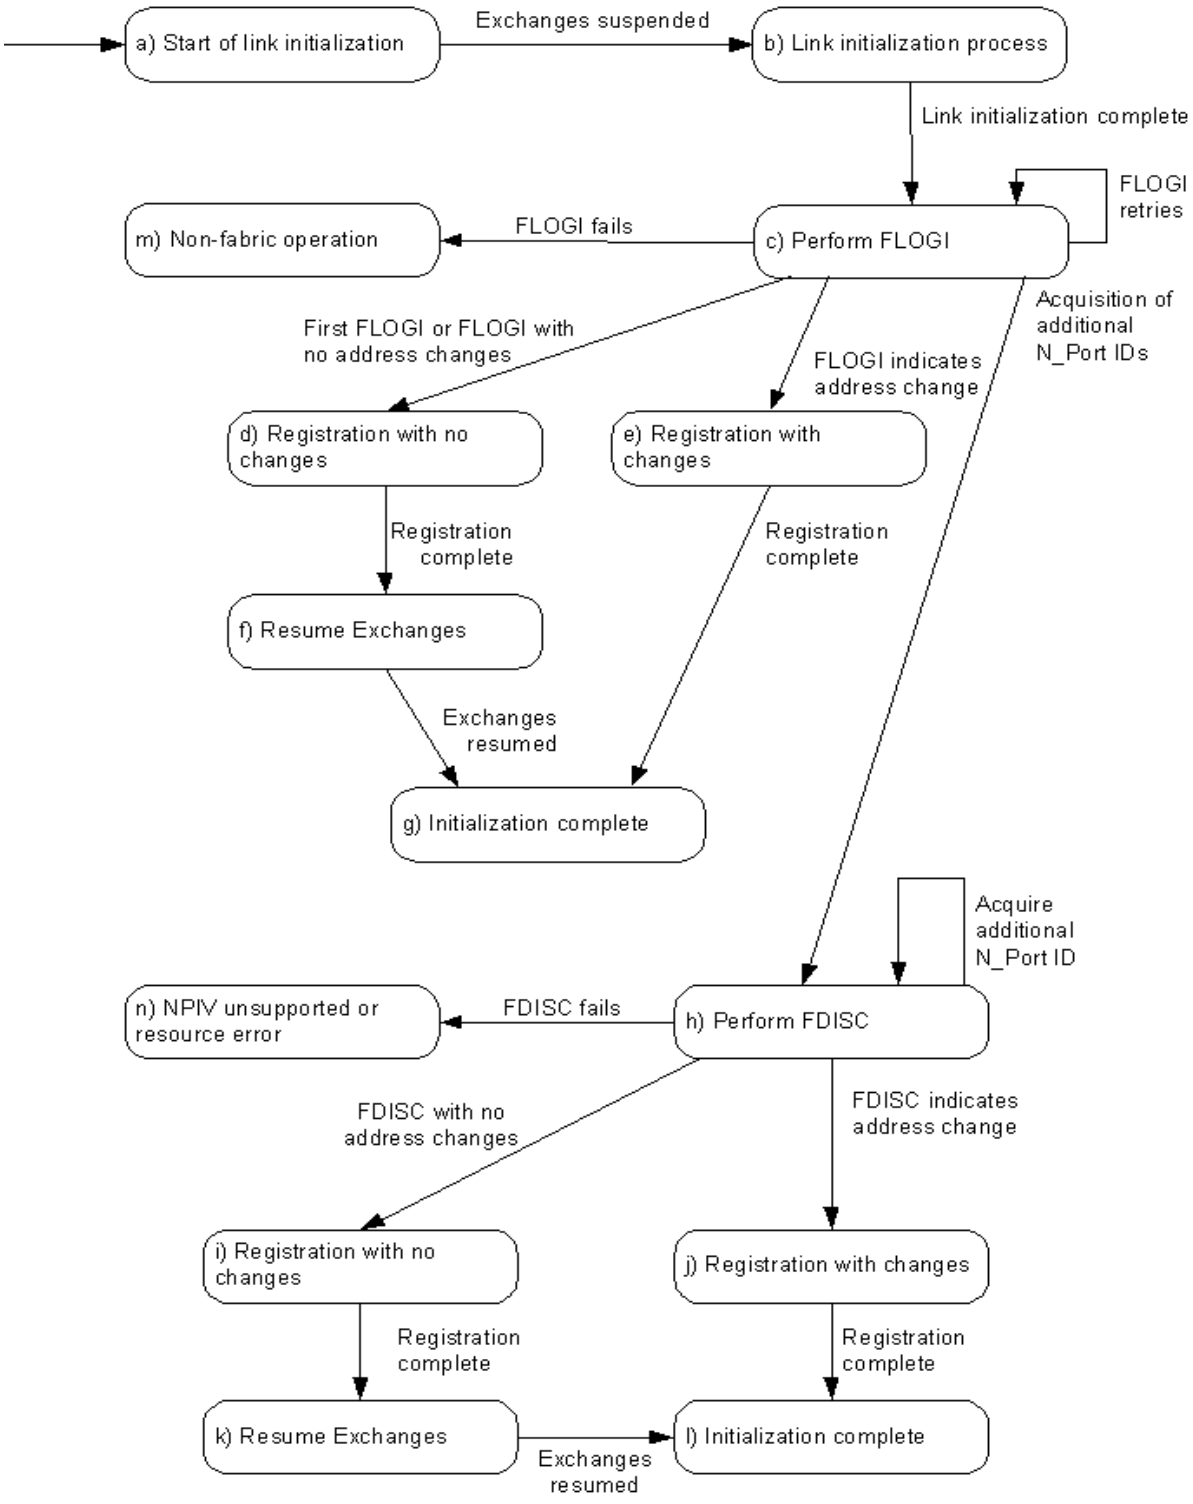

Figure 3 — N_Port_ID acquisition flow

a) **Start of link initialization**. Whenever a PN_Port receives an OLS, NOS, loss of synchronization for greater than R_T_TOV, or loss of signal, the PN_Port shall begin the Link Initialization Process (see FC-FS-3). Each VN_Port associated with that PN_Port shall suspend all open Exchanges and implicitly log out from the Fabric. The PN_Port shall proceed to step (b).

b) **Link initialization process.** Upon completion of the Link Initialization process the PN_Port shall proceed to step (c).

c) **Acquisition of FLOGI assigned N_Port_ID - Perform FLOGI**. The N_Port Controller shall attempt to send FLOGI to its local F_Port Controller. The Originator shall set the D_ID to the F_Port Controller(i.e., FFFFFEh) and the S_ID to:

   A) 000000h; or

   B) 0000YYh where YY is set to any value chosen by the Originator.

   The payload of the FLOGI Request Sequence shall contain an N_Port_Name and Node_Name associated with the N_Port requesting the login. The Common Service Parameters, Class 2 and Class 3 Service Parameters shall follow the rules as defined in 4.6.

   Until the FLOGI completes, the reply to any PLOGI received by the PN_Port shall either be:

   A) delayed until the FLOGI is completed; or

   B) terminated by an LS_RJT Reply Sequence, with a reason code of “Unable to perform command request”.

   If the FLOGI Request attempt failed (e.g., an LS_RJT is received), or if the FLOGI LS_ACC Sequence is not received within 2xR_A_TOV of the FLOGI Request, the N_Port Controller shall either:

   A) retry step (c); or

   B) proceed to step (m).

   The N_Port Controller shall instantiate a VN_Port that shall proceed to step (d) if:

   A) the FLOGI is completed successfully;

   B) the VN_Port has the same N_Port_ID, N_Port_Name, and Node_Name that the VN_Port had before link initialization; and

   C) the VF_Port has the same F_Port_Name and Fabric_Name that the VF_Port had before link initialization.

   The N_Port Controller shall instantiate a VN_Port that shall proceed to step (e) if the FLOGI is completed successfully and:

   A) the VN_Port did not have a N_Port_ID assigned prior to link initialization;

   B) the VN_Port does not have the same N_Port_ID, N_Port_Name, and Node_Name that the VN_Port had before link initialization; or

C) the VF_Port does not have the same F_Port_Name and Fabric_Name that the VF_Port had before link initialization.

NOTE 6 – It is possible for a VN_Port to receive ELS frames before any PLOGI frames are received.

The N_Port Controller may attempt to acquire additional N_Port_IDs by proceeding to step (h).

d) **Registration with no address change.** The VN_Port may optionally login (PLOGI) with the Directory Server (i.e., FFFFFCh) and perform registration, de-registration, and queries with the Name Server (i.e., GS_Subtype code 02h). The Name Server requests may be performed in any order. The VN_Port shall proceed to step (f).

e) **Registration with address change**. In this state:

1. all VN_Ports associated with the PN_Port shall discard all suspended and queued Exchanges, if any, in a manner consistent with the FC-4, implicitly logout with all other Nx_Ports, and notify the FC-4s;

2. the VN_Port shall perform PLOGI explicitly with the Directory Server (i.e., FFFFFCh) and shall perform registration with the Name Server. The Name Server Registration request may be performed in any order; and

3. if the Clean Address bit in the FLOGI Accept (LS_ACC) is set to one, the VN_Port shall proceed to step (g), otherwise, the VN_Port shall wait for R_A_TOV before originating any new Exchanges. During this time, the VN_Port shall discard all frames received except for the PLOGI Request Sequences and LOGO LS_ACC reply Sequences. The VN_Port shall respond to all PLOGI Request Sequences with an LS_RJT Reply Sequence, with a reason code of “Unable to perform command request”. The VN_Port shall proceed to step (g).

f) **Resume Exchanges**. If there are any suspended Exchanges with other Nx_Ports for the N_Port_ID acquired by FLOGI, the VN_Port shall resume them. The VN_Port shall proceed to step (g).

g) **Completion of VN_Port initialization**. This completes initialization for the VN_Port. The VN_Port may proceed to originate and respond to Exchanges and login with other Nx_Ports as needed for the N_Port_ID acquired by FLOGI.

h) **Acquisition of FDISC assigned N_Port_ID - Perform FDISC**. The N_Port Controller shall perform this step to acquire an additional N_Port_ID.

The N_Port Controller shall send FDISC to its local F_Port Controller. The D_ID shall be set to the F_Port Controller (i.e., FFFFFEh) and the S_ID to 000000h. The payload of the FDISC Request Sequence shall contain an N_Port_Name and Node_Name associated with the VN_Port for which an N_Port_ID is being requested. The Service Parameters should be identical to the parameters defined in the last FLOGI payload originated by the N_Port Controller.

If an LS_RJT Reply Sequence is received with a reason code of “Unable to perform command request” and a reason code explanation of “Login required”, indicating that there are no N_Port_IDs assigned to the N_Port and a FLOGI is required before an FDISC is permitted, then the N_Port Controller may proceed to step (c).

If the FDISC Request attempt failed for other reasons, including:

A) receiving an LS_RJT Reply Sequence with a reason code of “Command not supported” and a reason code explanation of “Request not supported”, indicating that the Fabric does not support NPIV;

B) receiving an LS_RJT Reply Sequence with a reason code of “Unable to perform command request” and a reason code explanation of “Insufficient resources to support Login”, indicating that the Fabric is unable to support additional FDISC assigned N_Port_IDs; or

C) not receiving an LS_ACC Reply Sequence within 2 x R_A_TOV from the FDISC Request,

then the N_Port Controller shall proceed to step (n) or retry step (h).

If an LS_ACC Reply Sequence is received, then the D_ID field contains the additional N_Port_ID.

The N_Port Controller shall instantiate a VN_Port that shall proceed to step (i) if:

A) the FDISC is completed successfully;

B) the N_Port_ID assigned by FDISC is the same N_Port_ID as was previously assigned to the VN_Port;

C) the VN_Port has the same N_Port_Name and Node_Name that it had before the FDISC; and

D) the VF_Port has the same F_Port_Name and Fabric_Name that it had before the FDISC.

The N_Port Controller shall instantiate a VN_Port that shall proceed to step (j) if the FDISC is completed successfully and:

A) the N_Port_ID assigned by FDISC is not the same N_Port_ID as was previously assigned to the VN_Port;

B) the VN_Port does not have the same N_Port_Name or Node_Name that it had before the FDISC; or

C) the VF_Port does not have the same F_Port_Name or Fabric_Name that it had before the FDISC.

i) **Registration with no address change.** The VN_Port, associated with the FDISC assigned N_Port_ID, may optionally login (i.e., PLOGI) with the Directory Server (i.e., FFFFFCh) and perform registration, de-registration and queries with the Name Server (i.e., GS_Subtype code 02h). The Name Server requests may be performed in any order. The VN_Port shall proceed to step (k).

j) **Registration with address change**. A VN_Port performing this step has determined that the identifying information for an N_Port_ID assigned to the VN_Port by this FDISC and/or the identifying information of the VF_Port have changed since the last assignment of the N_Port_ID to the VN_Port. The VN_Port shall perform the following:

1. the VN_Port shall discard all suspended and queued Exchanges in a manner consistent with the FC-4, implicitly logout with all Nx_Ports, and notify the FC-4;

2. the VN_Port shall perform PLOGI explicitly with the Directory Server (i.e., FFFFFCh) and

shall perform registration with the Name Server (i.e., GS_Subtype code 02h). The Name Server Registration requests may be performed in any order; and

3. if the Clean Address bit in the FDISC Accept (LS_ACC) is set to one, the VN_Port shall proceed to step (l); otherwise, the VN_Port shall wait for R_A_TOV before originating any new Exchanges. During this time, all frames received with a destination address of the VN_Port shall be discarded except for the PLOGI Request Sequences and LOGO LS_ACC reply Sequences. The VN_Port shall respond to all PLOGI Request Sequences with a reason code of “Unable to perform command request”. The VN_Port shall proceed to step (l).

k) **Resume Exchanges**. If there are any suspended Exchanges with other Nx_Ports for the VN_Port, the VN_Port shall resume them. The VN_Port shall proceed to step (l).

l) **Initialization complete**. The VN_Port may proceed to login with other Nx_Ports and originate and respond to Exchanges as needed.

m) **Fall-back to non-Fabric operation**. The PN_Port is not connected to a Fabric, therefore all VN_Ports with previously assigned N_Port_IDs shall discard all suspended and queued exchanges in a manner consistent with the FC-4, implicitly logout with all other Nx_Ports, and notify the FC-4s.

n) **NPIV unsupported or resource error.** Each VN_Port without a successfully assigned N_Port_ID shall discard all suspended and queued Exchanges and implicitly log out with all other Nx_Ports. Each VN_Port with a successfully assigned N_Port_ID may proceed to originate and respond to Exchanges and login with other Nx_Ports as needed.

4.13.3 Buffer-to-buffer flow control management

NPIV shall support buffer-to-buffer flow control (see FC-FS-3). The total number of receive buffers is inclusive of all the Sequence Initiators and Recipients associated with all N_Port_IDs. The total number of transmitted frames, for which R_RDY responses are outstanding, is inclusive of all the Sequence Initiators and Recipients associated with all N_Port_ID’s.

4.14 CS_CTL/Priority header field usage

4.14.1 Overview

The meaning of the CS_CTL header field is controlled by the CS_CTL/Priority Enable bit (F_CTL, bit 17) (see FC-FS-3). When the CS_CTL/Priority Enable bit is set to zero, word 1, bits 31-24 of the FC header shall be interpreted as CS_CTL information. When the CS_CTL/Priority Enable bit is set to one, word 1, bits 31-24 of the FC header shall be interpreted as Priority information.

4.14.2 Priority

When supported, Priority allows end devices and/or Switches in a Fabric to resolve resource contention or to determine the order in which to deliver frames. The Priority field in the Frame Header indicates the Priority assigned to the frame. Devices compliant with this technical report shall implement Priority, if supported, (see FC-FS-3) with the following restrictions:

a) the Sequence Initiator shall set the Priority, for the duration of the Sequence Initiative, to the same value for all frames; and

b) if multiple sequences occur, it is the responsibility of the Sequence Initiator to set the Priority of each frame in each Sequence to the same value.

4.14.3 CS_CTL

The CS_CTL field contains management information for the used class of service. Table 24 specifies the CS_CTL field for Class 2 and 3.

Table 24 – CS_CTL field

Bit	Abbr.	Meaning
31	PREF	0 = Frame is delivered with no Preference 1 = Frame may be delivered with Preference
30		Reserved for additional Preference function
29-24	DSCP	Differentiated Services Code Point

Devices compliant with this technical report shall implement PREF or DSCP, if supported, (see FC-FS-3) with the following restrictions:

a) PREF and DSCP shall be mutually exclusive. If PREF is set to one, then DSCP shall be set to zero; if DSCP is set to a nonzero value, then PREF shall be set to zero. The handling by the Fabric of frames with both PREF and DSCP set to a nonzero value is unspecified;

b) the Sequence Initiator shall set the PREF or DSCP fields, for the duration of the Sequence Initiative, to the same value for all frames of a Sequence; and

c) if multiple sequences occur, it is the responsibility of the Sequence Initiator to set the PREF or DSCP fields of each frame in each Sequence to the same value.

4.15 Logout procedure

The destination Logout procedure (see FC-LS-2) provides a method for removing service between an N_Port and Nx_Ports or for removing an N_Port_ID that was previously assigned by the Fabric. Logout releases resources associated with maintaining service with a Fabric and/or destination Nx_Port.

4.16 Nx_Port management and discovery

4.16.1 Overview

This subclause specifies interoperability and implementation guidelines for Nx_Ports (i.e., end devices) to allow Fibre Channel end devices to be discovered and managed by a wide range of SAN applications.

4.16.2 End devices

4.16.2.1 Overview

Fibre channel end devices are divided into subclasses for the establishment of discovery and management implementation guidelines. An end device may be an:

a) end device node;

b) platform; or

c) HBA device driver.

4.16.2.2 End device node

End device node (e.g., host bus adapters, and storage devices) discovery and management support is summarized in table 25.

Table 25 – End device node support summary

Item	Support
Responds to RNID	R[a]
Name Server Registration	R[b]
Accepts LIRR with Common Format	R[c]
Originates RLIR as required by accepted LIRRs	R[c]

a A data format of DFh shall be supported at a minimum.
b If a Fabric is attached then the Name Server registration shall be supported as defined in 4.12.
c Applicable to Switched Fabric only.

4.16.2.3 Platform

Platform support for discovery and management is summarized in table 26. In addition, any end device node that is contained within a platform shall meet the requirements in 4.16.2.2.

Table 26 – Platform support summary

Item[a]	Support
Platform Name	R
Platform Type	R[b]
Platform Management Address(es)	R
Platform Node Name(s)	R

a This platform information shall be registered with the Fabric Configuration Server.
b Only platform types gateway, host, storage subsystem, storage access device, or NAS server are required to be supported (see FC-GS-6).

4.16.2.4 HBA device driver

An HBA device driver shall meet the compliance requirements specified in FC-HBA.

4.17 Fabric event notification

An Nx_Port shall not send PLOGIs to an Nx_Port based solely on the reception of an RSCN.

Following the receipt of a domain format RSCN, an N_Port_ID shall validate any existing N_Port Logins with N_Port_IDs in the affected Domain. The mechanism by which this validation is done is FC-4 specific.

See FC-MI-3 for a description of RSCNs generated by Fabric elements.

5 FC-4 specific behavior

5.1 FCP Nx_Port behavior

5.1.1 Device binding

Nx_Port address identifiers may change (e.g., due to Fabric re-initialization events). An application shall manage its Fibre Channel resources using persistent unique identifiers (e.g., Name_Identifiers).

5.1.2 Registered state change notification

Following the receipt of a domain format RSCN, an FCP initiator Nx_Port that is registered to receive RSCN shall validate any existing N_Port Login using ADISC or Name Server queries.

Following the receipt of an RSCN, an FCP initiator Nx_Port that is registered to receive RSCN shall use ADISC or Name Server queries to validate existing N_Port logins of Nx_Ports identified in an affected Port_ID page in the RSCN.

5.1.3 FCP Nx_Port Common Service Parameters for PLOGI

Table 27 lists FC-4 specific Nx_Port PLOGI Common Service Parameters for FCP usage as defined by this technical report.

Table 27 – Nx_Port PLOGI Common Service Parameters

Parameter	FCP initiator	FCP target
Continuously Increasing Relative Offset	1	1
Random Relative Offset	0	0
Relative Offset by Information Category		
Information Category 1 (Solicited Data)	R	R
All other Information Categories[a]	A	A

a) Support of relative offset is not required for any Information Category other than Information Category 1. The FCP initiator is allowed to set the bit to one, but the FCP target may ignore it. The relative offset for the first frame of an IU shall be zero for all other Information Categories.

5.1.4 FCP Class 2 and Class 3 Service parameters for FLOGI

Table 28 lists FC-4 specific FLOGI Class 2 and Class 3 Service Parameters for FCP usage for as defined by this technical report.

Table 28 – FCP FLOGI Class 2 and Class 3 Service Parameters

Parameter	FCP initiator	FCP target
Priority/Preemption	0	0

5.1.5 FCP Class 2 Service Parameters for PLOGI - sequential-access devices

Table 29 lists FC-4 specific PLOGI Class 2 Service Parameters for sequential-access devices using FCP as defined by this technical report.

Table 29 – FCP PLOGI Class 2 Service Parameters for sequential-access devices

Parameter	FCP initiator	FCP target
Priority/Preemption	0	0
X_ID interlock	1	X

5.1.6 FCP Class 3 Service Parameters for PLOGI

Table 30 lists FC-4 specific PLOGI Class 3 Service Parameters for FCP usage as defined by this technical report.

Table 30 – FCP PLOGI Class 3 Service Parameters

Parameter	FCP initiator	FCP target
Priority/Preemption	0	0

5.1.7 FCP usage of F_CTL Abort Sequence Condition bits

5.1.7.1 FCP usage of F_CTL Abort Sequence Condition bits by Sequence Initiator

5.1.7.1.1 FCP Class 2 F_CTL Abort Sequence Condition bits by Sequence Initiator - sequential-access devices

Table 31 lists FC-4 specific Class 2 F_CTL Abort Sequence Condition bits usage by Sequence Initiators for sequential-access devices using FCP as defined by this technical report.

Table 31 – FCP Class 2 F_CTL Abort Sequence Condition bits by Sequence Initiator - sequential-access device

Abort Sequence Condition	FCP initiator
Abort, discard multiple Sequences (00b)	A
Abort, discard a single Sequence (01b)	R
Process policy with infinite buffers (10b)	P

5.1.7.1.2 FCP Class 3 F_CTL Abort Sequence Condition bits by Sequence Initiator

Table 32 lists FC-4 specific Class 3 F_CTL Abort Sequence Condition bits usage by Sequence Initiators using FCP as defined by this technical report. Where a difference in usage not specified, the requirement for direct-access and sequential-access devices is the same.

Table 32 – FCP Class 3 F_CTL Abort Sequence Condition bits by Sequence Initiator

Abort Sequence Condition	FCP initiator
Abort, discard multiple Sequences (00b) direct-access device	R
Abort, discard multiple Sequences (00b) sequential-access device	A
Abort, discard a single Sequence (01b) direct-access device	P
Abort, discard a single Sequence (01b) sequential-access device	R
Process policy with infinite buffers (10b)	P

5.1.7.2 FCP usage of Class 2 F_CTL Abort Sequence Condition bits by Sequence Recipient - sequential-access devices

Table 33 lists FC-4 specific Class 2 F_CTL Abort Sequence Condition bits usage by Sequence Recipients for sequential-access devices using FCP as defined by this technical report.

Table 33 – FCP Class 2 F_CTL Abort Sequence Condition bits by Sequence Recipient

Abort Sequence Condition	FCP target
Continue Sequence (00b)	P
Abort Sequence, send ABTS (01b)	R
Stop Sequence (10b)	P

5.1.8 FCP usage of Process Login parameters

Table 34 lists FCP usage of Process Login parameters defined by this technical report. It does not define the values that are required to be reported during Process Login. Where a difference in usage is not specified, the requirement for direct-access and sequential-access devices is the same.

Table 34 – FCP PRLI parameters

Parameter	Word	Bit(s)	Use by FCP initiator	Use by FCP target
TYPE Code	0	31-24	R	R
Common Service Parameters	0	31-24	P	P
TYPE Code Extension	0	23-16	P	P
Establish Image Pair (bit 13) = 1	0	13	R	R
Establish Image Pair (bit 13) = 0	0	13	I	R
Enhanced Discovery = 1	3	11	A	A
REC Support = 1	3	10	A	A
Task Retry Identification Requested = 1 random access device	3	9	P	P
Task Retry Identification Requested = 1 sequential access device	3	9	I	R
Retry = 1 random access device	3	8	P	P
Retry = 1 sequential access device	3	8	I	R
Confirmed Completion Allowed = 1 random access device	3	7	P	P
Confirmed Completion Allowed = 1 sequential access device	3	7	I	R
Data Overlay Allowed = 1 random access device	3	6	P	P
Data Overlay Allowed = 1 sequential access device	3	6	I	R
Initiator Function =1	3	5	R	A
Target Function = 1	3	4	A	R
Read XFER_RDY Disabled = 1	3	1	R	R
Write XFER_RDY Disabled = 1	3	0	P	P

5.1.9 FCP Extended Link Services

Table 35 lists FC-4 specific Extended Link Services for FCP usage as defined by this technical report. Where a difference in usage is not specified, the requirement for direct-access and sequential-access devices is the same.

Table 35 – FCP Extended Link Service support

Name	Originated by FCP initiator	Response from FCP target	Originated by FCP target	Response from FCP initiator
Process Login (PRLI)	I	R	P	-
Process Logout (PRLO)	I	R	I	R
Third Party Process Logout (TPRLO)	I	R	P	-
Read Exchange Concise (REC) random access device	A	-	P	-
Read Exchange Concise (REC) sequential access device	I	R	I	R[a]
a Required only if FCP_CONF is supported.				

5.1.10 FC-4 Link Service Requests and Responses for FCP

Table 36 lists specific FC-4 Link Services for FCP as defined by this technical report.

Table 36 – FC-4 Link Service Requests and Responses for FCP

Name	Originated by FCP initiator	Response from FCP target	Originated by FCP target	Response from FCP initiator
Sequence Retransmission Request SRR (Request)				
random access device	P	-	P	-
sequential access device	I	R	P	-

5.1.11 FCP_RSP payload fields

Table 37 lists the FCP_RSP payload fields with usage as defined by this technical report.

Table 37 – FCP_RSP payload fields

Feature	Processed by FCP initiator	Support by direct-access FCP target	Support by sequential-access FCP target
FCP_CONF_REQ	A	P	R
a If present, this field shall contain the FCP_RSP code in an 8 byte field in the format specified by FCP-4.			

Table 37 – FCP_RSP payload fields (Continued)

Feature	Processed by FCP initiator	Support by direct-access FCP target	Support by sequential-access FCP target
FCP_RESID_UNDER	R	I	I
FCP_RESID_OVER	R	I	I
FCP_SNS_LEN_VALID	R	I	I
FCP_RSP_LEN_VALID	R	I	I
FCP_RESID	R	I	I
FCP_SNS_LEN (length of FCP_SNS_INFO field)	R	<=96	<=128
FCP_RSP_LEN (length of FCP_RSP_INFO field)	R	0 or 8	0 or 8
FCP_RSP_INFO[a]	R	I	I
FCP_SNS_INFO	R	I	I
Maximum value in the SCSI sense data field “Additional sense length”	R	<=78	<=120

a If present, this field shall contain the FCP_RSP code in an 8 byte field in the format specified by FCP-4.

5.1.12 FC-AL-2 requirements for FCP

Table 38 lists FC-AL-2 requirements for NL_Ports using FCP as defined by this technical report.

Table 38 – FC-AL features for FCP NL_Ports

Feature	NL_Port Originator	NL_Port Responder
Broadcast and Multicast via Broadcast Replicate (OPNfr)	A	A
Broadcast and Multicast via Selective Replicate (OPNyr)	A	A

5.1.13 FCP device discovery

To facilitate discovery, an FCP initiator device shall process (i.e., not reject) a received PRLI request with the ESTABLISH IMAGE PAIR bit set to zero (see FCP-4).

5.1.14 FCP Private NL_Port Exchange authentication

5.1.14.1 FCP Private NL_Port initiator Exchange authentication

Following transmission of the CLS signalling completion of loop initialization, FCP initiators are required to authenticate each FCP target with which they have completed PLOGI using the following process:

a) an FCP initiator shall suspend execution of all open tasks with an FCP target following loop initialization until Exchange authentication is complete. Following transmission or forwarding of a LIP, only frames related to FCP target Exchange authentication (i.e., ADISC, PDISC, or PLOGI) shall be transmitted to that FCP target before frames for other Exchanges to that FCP target are transmitted;

b) for each FCP target that it has successfully completed PLOGI, an FCP initiator shall originate an ADISC or PDISC to that FCP target such that the ADISC or PDISC request Sequence arrives at each FCP target within RR_TOV of completing loop initialization;

c) for each FCP target that returns an ACC, if the [N_Port Identifier:N_Port_Name:Node_Name] triplet for that FCP target:

   A) does not match the triplet of a logged-in FCP target, the FCP initiator shall transmit LOGO to that FCP target; or

   B) does match the triplet of a logged-in FCP target, the FCP initiator may resume all tasks with that FCP target.

d) FCP initiator Exchange authentication is complete when the ACC to an ADISC or PDISC has been received. If the FCP initiator fails to receive an ACC within R_A_TOV, it shall implicitly logout the FCP target and terminate all tasks for that FCP target;

e) following loop initialization, any frames received, other than ACC to an ADISC, PDISC, or PLOGI, or a LOGO request, from an FCP target that has not been authenticated shall be discarded (i.e., those frames shall not be considered part of any valid Exchange, and only R_RDYs to maintain BB_Credit shall be sent in response to these frames); and

f) if an FCP initiator receives a LOGO from an Nx_Port during this procedure, it shall terminate all open Exchanges with that Nx_Port and send an ACC.

5.1.14.2 FCP Private NL_Port target Exchange authentication

Following completion of loop initialization, FCP targets are required to wait for Exchange authentication from the FCP initiators with which they have completed PLOGI using the following process:

a) for each FCP initiator that has open tasks, an FCP target shall suspend tasks associated with that FCP initiator until an ADISC or PDISC is received from that FCP initiator. Following transmission or forwarding of a LIP, only frames related to FCP initiator Exchange authentication (i.e., ACC to an ADISC or PDISC) shall be transmitted to that FCP initiator before frames for other Exchanges to that FCP initiator are transmitted;

b) for each FCP initiator that sends an ADISC or PDISC, if the [N_Port Identifier:N_Port_Name:Node_Name] triplet for that FCP initiator:

   A) does not match the triplet of a logged-in FCP initiator, the FCP target shall transmit LOGO to that FCP initiator and terminate all Exchanges associated with that FCP initiator; or

   B) does match the triplet of a logged-in FCP initiator, the FCP target shall send the ACC to the Extended Link Service. The FCP target may then resume all tasks with that FCP initiator.

c) if an FCP target does not receive a PDISC or ADISC from each logged-in FCP initiator within RR_TOV of completing loop initialization, then it may implicitly logout that FCP initiator and terminate all Exchanges associated with that FCP initiator;

d) FCP target Exchange authentication is complete when the ACC to an ADISC or PDISC has been transmitted; and

e) following loop initialization, any frames received, other than ADISC, PDISC, or PLOGI, from an FCP initiator that has not been authenticated shall be discarded. Only R_RDYs to maintain BB_Credit shall be sent in response to those frames.

5.2 FC-SB-4 N_Port behavior

5.2.1 FC-SB-4 Extended Link Services

Table 39 lists FC-4 specific Extended Link Services for FC-SB-4 usage as defined by this technical report.

Table 39 – FC-SB-4 Extended Link Service support

Name	Originated by FC-SB-4 Channel	Response by FC-SB-4 Control Unit	Originated by FC-SB-4 Control Unit	Response by FC-SB-4 Channel
Link Incident Record Registration (LIRR)[b]	R	A[a]	P	-
Query Security Attributes (QSA)	R[c]	-	A	-
Read Link Error Status Block (RLS)	-	-	R[d]	-
Registered Fabric Change Notification (RFCN)	P	-	P	R
Registered Link Incident Report (RLIR)[b]	P	-	R	R
Registered State Change (RSCN)	R	R	R	R
Request Node Identification Data (RNID)[b]	R	R	R	R
State Change Registration (SCR)	R	R	R	R
Process Login (PRLI)	I[e]	R[e]	P	-
Process Logout (PLOGI)	R[f]	R[f]	R[f]	R[f]

a Support is recommended for FC-SB-4 Control Units.
b FC-SB-4 (type 18h) specific node-identification data and link-incident record format shall be supported (see FC-SB-4).
c Support is required for FC-SB-4 channels when one or more of it's Control Units are located in a different switch domain from that of the channel.
d Required if a FC-SB-4 Control Unit that supports error code transfer receives a Purge Path IU (see FC-SB-4) with a nonzero error code in a fabric topology.
e Invocable if support for PRLI is indicated in the RNID response from the Control Unit.
f Required if PRLI has been originated by the channel.

5.2.2 FC-SB-4 Class of Service support

Classes 2 and 3 shall be supported.

5.2.3 FC-SB-4 Class Service Parameters support

Priority/Preemption is allowed.

5.3 IP Nx_Port behavior

When IPv4 or IPv6 is enabled for an Nx_Port, the Nx_Port shall register the FC-4 type '05h' with the Name Server.

Annex A: Discovery and Management: Examples & Rationale

(Informative)

A.1 Examples & Rationale

A.1.1 Overview

This annex presents examples for the processes described in 4.16.

A.1.2 RNID responses

A.1.2.1 Overview

RNID allows a node to provide basic identity information about itself. The node may also report a management address if it has one. Some example responses:

A.1.2.2 Disk Drive

Response (hex)	Description
02 00 00 00	preamble for RNID accept
DF	Node Identification Data Format
10	Common Identification Data Length
00	reserved
34	Specific Identification Data Length
50 05 07 60 06 41 83 F3	N_Port Name
50 05 07 60 06 01 83 F3	Node_Name
xx xx xx xx xx xx xx xx	Vendor Specific
xx xx xx xx xx xx xx xx	Vendor Specific
00 00 00 09	Associated Type
00 00 00 00	Physical Port Number
00 00 00 01	Number of Attached Nodes
00	Node Management
00	IP Version
00 00	UDP/TCP Port Number
00 00 00 00 00 00 00 00	IP Address
00 00 00 00 00 00 00 00	IP Address continued
00 00	reserved
xx xx	Vendor Specific

xx - response value not defined

A.1.2.3 Storage Subsystem

(with out-band management capabilities)

Response (hex)	Description
02 00 00 00	preamble for RNID accept
DF	Node Identification Data Format
10	Common Identification Data Length
00	reserved
34	Specific Identification Data Length
50 05 07 60 06 41 81 A0	N_Port Name

```
50 05 07 60 06 01 80 C3       Node_Name
xx xx xx xx xx xx xx xx       Vendor Specific
xx xx xx xx xx xx xx xx       Vendor Specific
00 00 00 0B                   Associated Type
00 00 00 01                   Physical Port Number
00 00 00 01                   Number of Attached Nodes
00                            Node Management
01                            IP Version
00 A1                         UDP/TCP Port Number
00 00 00 00 00 00 00 00       IP Address
00 00 00 00 09 05 7B D5       IP Address continued
00 00                         reserved
xx xx                         Vendor Specific
```

xx - response value not defined

A.1.3 Platform registration

Platform registration is important for systems such as hosts and storage subsystems that may:

a) contain more than one FC node;

b) have multiple management interfaces; or

c) participate in more than one SAN.

This registration allows a management application to fully discover and identify systems attached to a SAN beyond the basic node identification provided by RNID. Registration also enables a management application to automatically surface platform-specific management interfaces.

Platform registration by end devices that contain only one FC node and have no out-band management interfaces is optional.

See table A.1 for two examples:

Table A.1 – Platform Registration Example

Item	Size (Bytes)	Example 1: Storage Subsystem	Example 2: Host
CT_IU preamble	16		
Platform name	256	stor1.rchland.ibm.com	rserver1.austin.ibm.com
Platform type	4	0000000Bh	0000000Ah
Number of management address entries	4	00000002h	00000001h
Management address 1	256	snmp://9.5.126.97[:161]	snmp://9.5.156.39[:161]
Management address 2	256	http://9.5.126.97[:1008]/storman	-
Number of platform node name entries	4	00000002h	00000004h
Platform node name 1	8	500507600601803Ch	200000E08B018F3Eh
Platform node name 2	8	5005076006018041h	50050760170020E0h
Platform node name 3	8	-	5015342D0702054Ah
Platform node name 4	8	-	200000E01235CA17h