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

INCITS xxx-201x
T11/Project 1745-D/Rev 1.06

FIBRE CHANNEL

INTER-FABRIC ROUTING

(FC-IFR)

REV 1.06

INCITS working draft proposed
American National Standard
for Information Technology

May 12, 2010

Secretariat: Information Technology Industry Council

POINTS OF CONTACT:

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

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

Craig W. Carlson (T11.3 Chair)
QLogic Corporation
4601 Dean Lakes Blvd.
Shakopee, MN 55379
Voice: 952-687-2431
craig.carlson@qlogic.com

Joe Pelissier (FC-IFR Chair)
Cisco Systems, Inc.
18191Sunrise Peaks Lane
Hillsboro, OR 97123
Voice: 503-628-0801
jopeliss@cisco.com

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

BSR INCITS xxx-201x

American National Standard
for Information Technology

Fibre Channel —
Inter-Fabric Routing (FC-IFR)

Secretariat

Information Technology Industry Council

Approved (not yet approved)

American National Standards Institute, Inc.

Abstract

This standard defines the protocols, functions, and mappings for the routing of Fibre Channel frames between physically or logically separated Fabrics.

American National Standard

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

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

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards. The American National Standards Institute does not develop standards and will in no circumstances give an interpretation of any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

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

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard. As of the date of publication of this standard and following calls for the identification of patents that may be required for the implementation of the standard, no such claims have been made. No further patent search is conducted by the developer or publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Published by

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

Printed in the United States of America

Foreword (This Foreword is not part of American National Standard INCITS xxx-201x.)

The Fibre Channel Inter-Fabric Routing (FC-IFR) standard defines the protocols, functions, and mappings for the routing of Fibre Channel frames between physically or logically separated Fabrics.

This standard was developed by Task Group T11.3 of Accredited Standards Committee INCITS during 2004-2010. The standards approval process started in 2009.

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.

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

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

(to be filled in by INCITS)

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

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

Company	Name

(To be filled in prior to submission to INCITS.)

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

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

Company	Name

(To be filled in prior to submission to INCITS.)

Introduction

FC-IFR is one of the Fibre Channel family of standards.

This standard specifies the defines the protocols, functions, and mappings for the routing of Fibre Channel frames between physically or logically separated Fabrics.

Contents **Page**

Figure **Page**

Table **Page**

American National Standard
for Information Technology —

Fibre Channel —
Inter-Fabric Routing (FC-IFR)

1 Scope

The Fibre Channel Inter-Fabric Routing (FC-IFR) standard defines the protocols, functions, and mappings for the routing of Fibre Channel frames between physically or logically separated Fabrics.

This standard is divided into the following clauses:

Clause 1 specifies the scope of this standard.

Clause 2 specifies the normative references that apply to this standard.

Clause 3 specifies the definitions, abbreviations, and conventions used in this standard.

Clause 4 specifies the structure and concepts of an Inter-Fabric Router.

Clause 5 specifies the operational requirements for an Inter-Fabric Router.

Clause 6 specifies the Fabric addressing for an Inter-Fabric Router.

Clause 7 specifies the Internal Link Services for an Inter-Fabric Router.

Clause 8 specifies the zoning protocol for an Inter-Fabric Router.

Clause 9 specifies the operation of redundant Inter-Fabric Routers.

Clause 10 specifies the timers and constants for an Inter-Fabric Router.

2 Normative references

2.1 Overview

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

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

Additional availability contact information is provided below as needed.

2.2 Approved references

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

2.3 References under development

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

T11/Project 1620-D, *Fibre Channel - Link Services - 2 (FC-LS-2)*

T11/Project 1822-D, *Fibre Channel - Switch Fabric - 5 (FC-SW-5)*

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

2.4 Other references

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

3 Definitions and conventions

3.1 Common definitions

3.1.1 address identifier: An address value used to identify the source (S_ID) or destination (D_ID) of a frame (see FC-FS-3).

3.1.2 B_Port: A Bridge Port on a device that implements FC-BB-3_ATM, FC-BB-3_SONET, or FC-BB-3_IP, and connects to an E_Port on an FC Switch.

3.1.3 Bridge: A device that encapsulates/de-encapsulates Fibre Channel frames within another protocol (e.g., Fibre Channel encapsulated within IP).

3.1.4 Destination_Identifier (D_ID): The address identifier used to indicate the targeted destination Nx_Port of the transmitted frame.

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

3.1.6 E_Port_Name: A Name_Identifier (see 3.1.16) that identifies an E_Port (see 3.1.5).

3.1.7 Fabric: The entity that interconnects Nx_Ports attached to it and is capable of routing frames using the D_ID information in the FC-2 frame header (see FC-FS-3).

3.1.8 Fabric_Identifier (F_ID): An identifier assigned to each Fabric (see 6.1).

3.1.9 Fabric_Name: A Name_Identifier (see 3.1.16) associated with a Fabric (see FC-SW-5).

3.1.10 F_Port: A port by which non-loop N_Ports are attached to a Fabric (see FC-SW-5 and FC-FS-3).

3.1.11 F_Port_Name: A Name_Identifier (see 3.1.16) that identifies an F_Port (see 3.1.10).

3.1.12 FC_Port: A port that is capable of transmitting or receiving Fibre Channel frames according to the FC-0, FC-1, and FC-2 levels of the Fibre Channel architecture (see FC-FS-3).

3.1.13 FL_Port: An F_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology (see FC-AL-2).

3.1.14 Inter-Fabric Routing: The process of forwarding frames through a specific Routing Function (see 4.4), including the translation of N_Port_IDs.

3.1.15 Link Control Facility (LCF): A hardware facility that attaches to an end of a link and manages transmission and reception of data. It is contained within each FC_Port. See FC-FS-3.

3.1.16 Name_Identifier: A value with a specified size and format used to identify a Fibre Channel entity (e.g., N_Port, node, F_Port, or Fabric) (see FC-FS-3).

3.1.17 Native Nx_Port: An Nx_Port that is shared via one or more Fabrics as a Proxy N_Port.

3.1.18 Native Fabric: The Fabric where the Native Nx_Port resides.

3.1.19 Node_Name: A Name_Identifier (see 3.1.16) associated with a node (see FC-FS-3).

3.1.20 N_Port: An Nx_Port communicating through an LCF that is not operating a Loop Port State Machine (see FC-AL-2). Services operating at well-known addresses are considered to be N_Ports.

3.1.21 N_Port_ID: A topology unique address identifier of an Nx_Port that may be assigned by the Fabric during the initialization procedure or by other procedures not defined in this standard.

3.1.22 N_Port_Name: A Name_Identifier (see 3.1.16) that identifies an N_Port (see 3.1.20).

3.1.23 NL_Port: An Nx_Port communicating through an LCF that is operating a Loop Port State Machine (see FC-AL-2).

3.1.24 Nx_Port: An end point for Fibre Channel frame communication (see FC-FS-3) where in this standard, the term Nx_Port is used to specify behavior of either N_Ports or Public NL_Ports, and such use neither specifies nor constrains the behavior of Private NL_Ports (see FC-AL-2).

3.1.25 Originator: The logical function associated with an Nx_Port responsible for originating an Exchange.

3.1.26 path selection: The discovery of the best series of Inter-Fabric Routers to utilize to forward a frame from a source to a destination.

3.1.27 Proxy N_Port: A logical N_Port that represents a Native Nx_Port in a Fabric.

3.1.28 Public NL_Port: An NL_Port that attempts a Fabric Login (see FC-AL-2).

3.1.29 Responder: The logical function in an Nx_Port responsible for supporting the Exchange initiated by the Originator in another Nx_Port.

3.1.30 Source_Identifier (S_ID): The address identifier used to indicate the source Nx_Port of the transmitted frame.

3.1.31 Switch: A Fabric element conforming to the Fibre Channel Switch Fabric set of standards (e.g., see FC-SW-5).

3.1.32 Switch_Name: A Name_Identifier (see 3.1.16) that identifies a Switch or a Bridge device. See FC-SW-5.

3.2 Editorial Conventions

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

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

The ISO 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, French, and American conventions is shown in table 1.

Table 1 – ISO, French, and American conventions

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

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

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

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

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

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

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

3.3 List of commonly used acronyms and abbreviations

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

3.3.1 General

DF_ID Destination Fabric Identifier
ESS Exchange Switch Support
F_ID Fabric Identifier
IFR Inter-Fabric Router
IFR_ILS Inter-Fabric Router Internal Link Service
ISL Inter-Switch Link
SF_ID Source Fabric Identifier
SW_ILS Switch Fabric Internal Link Services

3.4 Keywords

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

3.4.2 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.

3.4.3 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.

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

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

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

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

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

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

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

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

4 Structure and concepts

4.1 Inter-Fabric Router overview

This standard describes the architectural model and specification for a Fibre Channel Inter-Fabric Router. An Inter-Fabric Router is a Fibre Channel device that permits attached Fibre Channel Nx_Ports to be shared among multiple independent Fibre Channel Fabrics.

An Inter-Fabric Router, or a series of Inter-Fabric Routers, connect Fabrics and provide the functionality required to present a Native Nx_Port to each Fabric as a Proxy N_Port. The entire interconnection of Fabrics and Inter-Fabric Routers is referred to as a Fibre Channel Inter-Fabric.

Two modes of operation are specified for Inter-Fabric Routers, simple mode and network address translation (NAT) mode. From the perspective of the requirements placed on Inter-Fabric Routers, the NAT mode of operation is a proper subset of the requirements of the simple mode of operation.

The frames that pass between Inter-Fabric Routers that are bound for Native Nx_Ports are referred to as Routable Frames. This term distinguishes these frames from other frames that pass between Inter-Fabric Routers such as the frames related to Inter-Fabric route discovery.

Fabrics operate independently. Consequently, there is no assurance that these Fabrics will select Domain IDs that are unique across all Fabrics within the Inter-Fabric. Therefore, the N_Port_IDs are not necessarily the same between Proxy N_Ports and their corresponding Native Nx_Ports (i.e., different N_Port_IDs are likely to be the case). Inter-Fabric Routers perform the necessary translation of N_Port_IDs within the Fibre Channel frames as they pass through the intervening Inter-Fabric Routers.

No translation of WWNs is performed. Therefore, a Proxy N_Port and its corresponding Native Nx_Port may appear on the Fibre Channel Inter-Fabric with different N_Port_IDs but will always have the same WWN.

Each Fabric within the Inter-Fabric is identified by a 12-bit Fabric Identifier.

4.2 Inter-Fabric Router simple mode operation

Simple mode operation allows for the use of a single Inter-Fabric Router (i.e., no hop Inter-Fabric routing), or two Inter-Fabric Routers (i.e., single hop Inter-Fabric routing) to connect Fabrics.

An example of the supported simple mode Inter-Fabric Router configurations (i.e., no hop and single hop) is shown in figure 1.

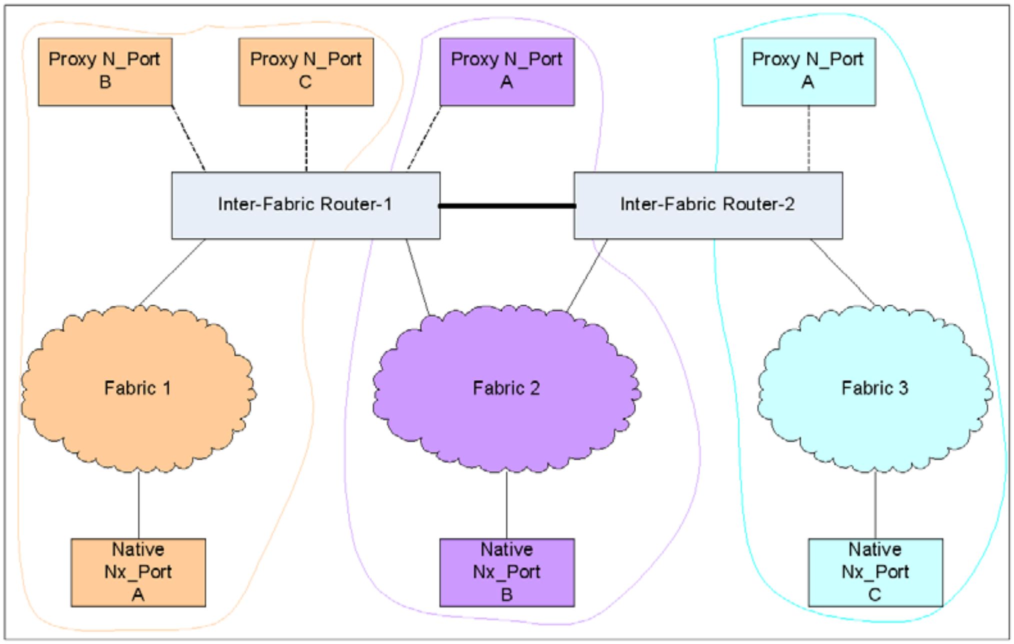

Figure 1 – Example simple mode Inter-Fabric Router configuration

In figure 1, the individual Fabrics are represented as Fabrics 1, 2, and 3. Three Native Nx_Ports are shown, A, B, and C, and their Native Fabrics are Fabric 1, Fabric 2, and Fabric 3, respectively. In this example, Native Nx_Port A, which is attached to Fabric 1, is communicating with Native Nx_Ports B and C. Inter-Fabric Router-1 presents Proxy N_Port B and Proxy N_Port C to Fabric 1. Any Fibre Channel frames originated by Native Nx_Port A and addressed to Proxy N_Port B transit through Fabric 1, Inter-Fabric Router-1, and Fabric 2 to Native Nx_Port B (i.e., no hop communication). Any Fibre Channel frames originated by Native Nx_Port A and addressed to Proxy N_Port C transit through Fabric 1, Inter-Fabric Router-1, Fabric 2, Inter-Fabric Router-2, and Fabric 3 to Native Nx_Port B (i.e., single hop communication). Since Fibre Channel operations are generally bi-directional, Native Nx_Ports B and C in this example must also be able to transmit frames to Native Nx_Port A. To accomplish this, Inter-Fabric Routers 1 and 2 each present a Proxy N_Port A to Fabrics 2 and 3 respectively. These Proxy N_Ports represent Native Nx_Port A in Fabrics 2 and 3. Consequently, Native Nx_Ports B and C direct their frames to Proxy N_Port A. These frames are then forwarded to Native Nx_Port A on behalf of Native Nx_Ports B and C by the Inter-Fabric Routers.

Multiple paths are supported through the Inter-Fabric using redundant Inter-Fabric Routers. A routable frame may transit through a single intermediate Fabric. Multiple parallel Fabrics are supported and the frame may be directed through any of these intermediate Fabrics with the restriction that only a single intermediate Fabric may be transited by any given routable frame.

The first Inter-Fabric Router encountered, referred to as the ingress Inter-Fabric Router, is responsible for the translation of the N_Port_ID of the Proxy N_Port (i.e., the D_ID in the Frame_Header) to that of its corresponding Native Nx_Port in the Native Fabric. In addition, the ingress Inter-Fabric Router translates N_Port_IDs that appear in the frame payload as needed (see 5.10.2). The last Inter-Fabric Router encountered, referred to as the egress Inter-Fabric Router, translates the S_ID in the Frame_Header to that of the Proxy N_Port in the Native Fabric that represents the Nx_Port that originated the frame.

Depending on the topology of the Fibre Channel Inter-Fabric and the source and destination of a given frame, a given Inter-Fabric Router is required to perform the ingress Inter-Fabric Router role, the egress Inter-Fabric Router role, or both the ingress and egress Inter-Fabric Router role.

Between Inter-Fabric Routers, an IFR_Header (see FC-FS-3) is prepended to permit intervening Inter-Fabric Routers along the path to properly forward the frame. The ingress Inter-Fabric Router prepends the IFR_Header and the egress Inter-Fabric Router removes the IFR_Header.

Additionally, an Enc_Header (see FC-FS-3) is prepended to the IFR_Header between Inter-Fabric Routers. The Enc_Header enables intervening legacy Fibre Channel Switches to properly forward the frame to the next hop Inter-Fabric Router.

Neither the IFR_Header nor the Enc_Header appears on any frame prior to the ingress Inter-Fabric Router or after the egress Inter-Fabric Router.

4.3 Inter-Fabric Router NAT mode operation

NAT mode operation allows for the use of one or more Inter-Fabric Routers to connect Fabrics. No IFR_Header or Enc_Header is required for NAT mode operation.

An example of a supported NAT mode Inter-Fabric Router configuration (i.e., no hop and single hop) is shown in figure 2.

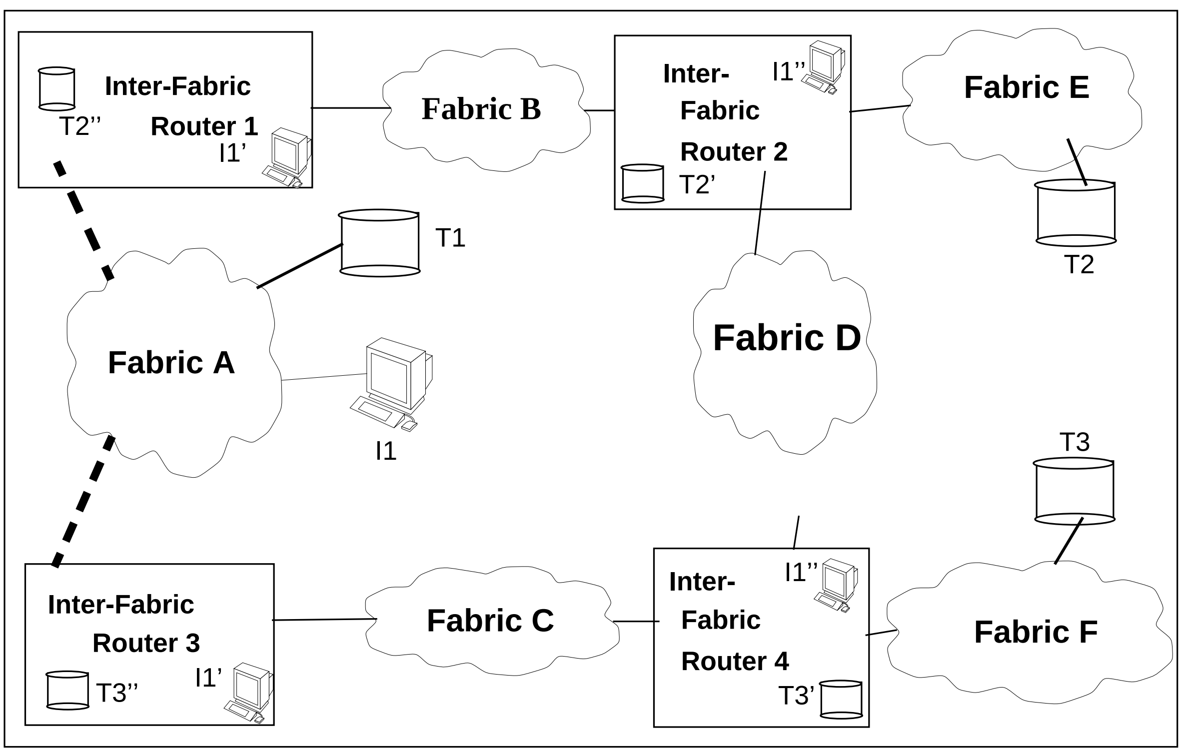

Figure 2 – Example NAT mode Inter-Fabric Router configuration

In figure 2, all intervening Fabrics contain Proxy N_Ports. Inter-Fabric Router 2 presents Proxy N_Port T2', representing Native Nx_Port T2, to Fabric B. Likewise, Inter-Fabric Router 4 presents Proxy N_Port T3', representing Native N_Port T3, to Fabric C. Inter-Fabric Router 1 presents Proxy

N_Port T2”, representing proxy N_Port T2’, to Fabric A. Any Fibre Channel frames originated by I1 and destined for Proxy N_Port T2” transit through Fabric A, Inter-Fabric Router 1, Fabric B, Inter-Fabric Router 2, and Fabric E to Native Nx_Port T2. Likewise, Inter-Fabric Router 3 presents Proxy N_Port T3”, representing Proxy N_Port T3’, to Fabric A. Any frames destined for Proxy N_Port T3” transit through Fabric A, Inter-Fabric Router 3, Fabric C, Inter-Fabric Router 4, and Fabric F, to Native Nx_Port T3. Device I1 does not originate frames directly to Native Nx_Ports T2 and T3. These frames are instead directed to Proxy N_Ports T2” and T3”. The Inter-Fabric Routers forward the frames on to the respective Proxy Nx_Ports, if any, and ultimately the Native Nx_Ports on I1’s behalf. Since Fiber Channel operations are generally bi-directional, Native Nx_Ports T2 and T3 in this example must also be able to transmit frames to I1. To accomplish this, Inter-Fabric Routers 1 and 3 present a Proxy N_Port I1’ to Fabrics B and C, and Inter-Fabric Routers 2 and 4 each present a Proxy N_Port I1” to Fabrics E and F respectively. These Proxy N_Ports represent the Native Nx_Port I1 in Fabrics E and F. Consequently, Native Nx_Ports T2 and T3 direct their frames to Proxy N_Port I1”. These frames are then forwarded to Native Nx_Port I1 on Native Nx_Ports T2 and T3 behalf by the Inter-Fabric Routers.

Multiple intermediate fabrics are supported by presenting a new Proxy N_Port at each Fabric. Thus in figure 2, with the establishment of the appropriate Proxy N_Ports, a frame may flow from T3, through Inter-Fabric Router 4, Fabric D, Inter-Fabric Router 2, Fabric B, Inter-Fabric Router 1, and on to I1. This would enable an alternate path to respond to a failure in Fabric C or Inter-Fabric Router 3, for example.

Inter-Fabric Routers perform the necessary translation of N_Port_IDs within the Fibre Channel frames as they pass through the intervening Inter-Fabric Routers. Each Inter-Fabric Router performs both the ingress and egress Inter-Fabric Router role translations. The N_Port_ID of the Proxy N_Port is translated to the next hop Proxy N_Port_ID or the native N_Port_ID if this is the last hop. The S_ID is translated to the Proxy N_Port_ID representing the previous N_Port, Proxy or Native as appropriate. All of the N_Port_ID’s within the frame are also translated accordingly. This is exactly the same operation a simple mode Inter-Fabric Router would perform if it were the only Inter-Fabric Router on the path.

4.4 Inter-Fabric Router logical components

4.4.1 Overview

Figure 3 illustrates the logical internal components of an Inter-Fabric Router that apply to both the simple mode and NAT mode of operation. An Inter-Fabric Router consists of a Routing Function and a number of Front Domain Switches and Translate Domain Switches.

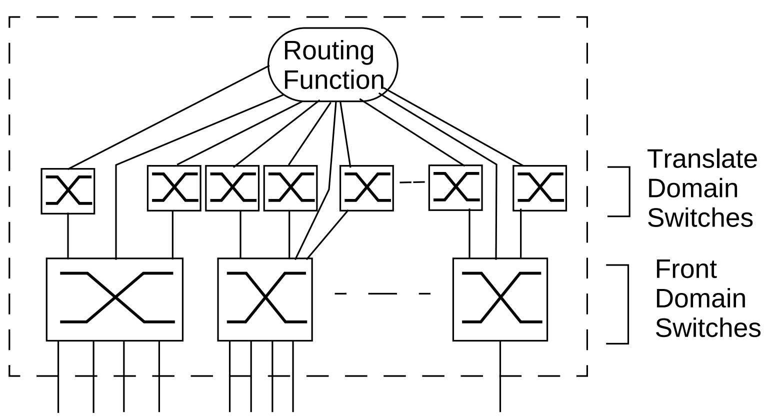

Figure 3 – Inter-Fabric Router logical components

4.4.2 Front Domain Switches

The Front Domain Switches within an Inter-Fabric Router provide connectivity to the Fibre Channel Fabrics that are interconnected by the Inter-Fabric Router. An implication of this is that Inter-Fabric Routers appear as Fibre Channel Switches to the Fabric to which they attach. Each Front Domain Switch exposes one or more ports to each interconnected Fabric. The ports may be statically assigned to the Front Domain Switches or may be configurable using Virtual Fabrics.

Two Front Domain Switches that belong to the same Inter-Fabric Router may belong to the same or different Fabrics. Two Front Domain Switches belonging to different Fabrics may have the same Domain_ID.

4.4.3 Translate Domain Switches

A Translate Domain Switch is created for each set of Proxy N_Ports whose corresponding Native Nx_Ports exist within a single physical Fabric, or a set of physical Fabrics. The link cost of the ISLs between the Front Domain Switches and the Translate Domain Switches may be adjusted to create a preferential path between an Nx_Port on the local Fabric and a Proxy N_Port that is presented by two Inter-Fabric Routers to the same Fabric (see clause 9).

4.4.4 Routing Function

The Routing Function provides for the forwarding of frames between independent Fibre Channel Fabrics including the discovery of Inter-Fabric Routers and Fabrics, and the translation of N_Port_ID’s to present the Proxy N_Ports to the local Fabrics.

5 Inter-Fabric Router operational requirements

5.1 Overview

An Inter-Fabric Router shall only support Class 2 service and Class 3 service.

5.2 Inter-Fabric Router architectural model

Figure 4 illustrates the architectural model of an Inter-Fabric Router. This model does not specify a required internal implementation. The externally observable behavior of an Inter-Fabric Router shall be consistent with this architectural model. The following sections specify the detailed requirements of the architectural components of the Inter-Fabric Router.

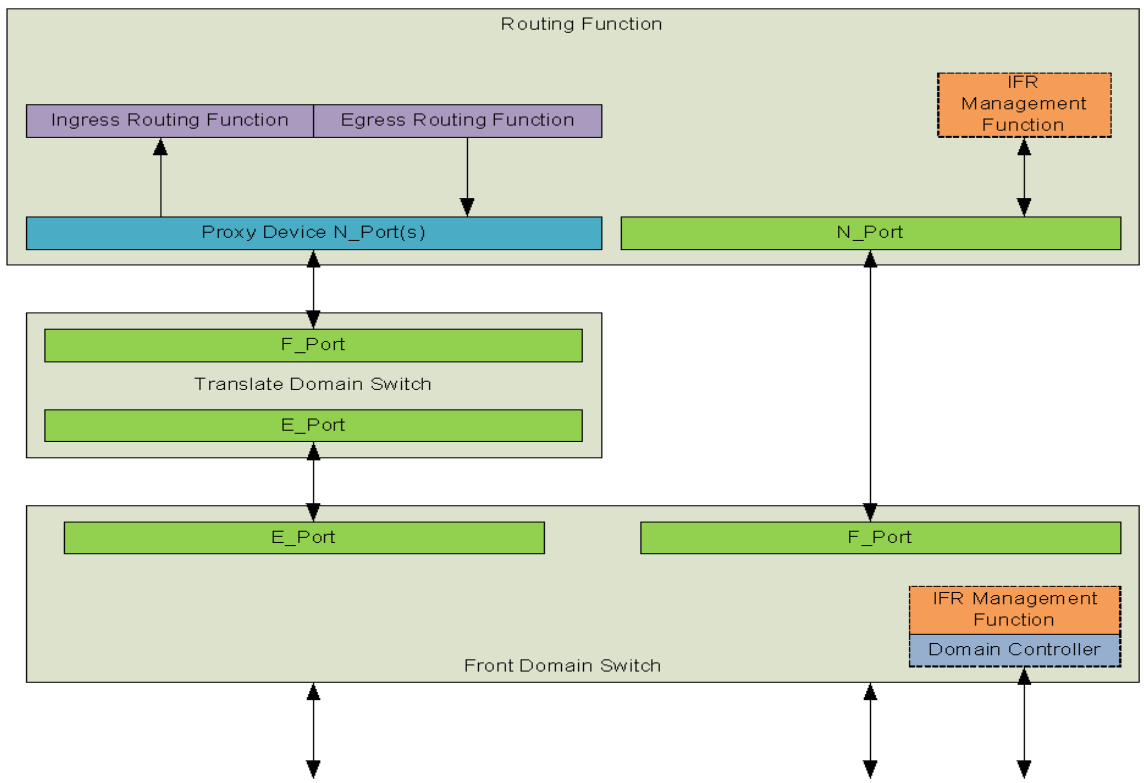

Figure 4 – Inter-Fabric Router architectural model

5.3 Operation of Inter-Fabric Router internal Switches

The Switches internal to the Inter-Fabric Router shall, to the extent externally observable, operate in accordance with the requirements of FC-SW-5.

The number of Front Domain Switches supported by an Inter-Fabric Router shall be greater than one if the Front Domain Switch does not support Virtual Fabric Tagging or shall be at least one if the Front Domain Switch does support Virtual Fabric Tagging.

The number of Translate Domain Switches supported by an Inter-Fabric Router shall be at least one.

5.4 Interconnection of Inter-Fabric Router components

Each Front Domain Switch shall expose at least one port external to the Inter-Fabric Router.

Each Front Domain Switch shall provide one link to the Routing Function. The Front Domain Switch shall operate the port connected to this link as an F_Port. The Routing Function shall operate the port connected to this link as an N_Port. This link shall be used by the Routing Function to communicate with the associated Fabric as follows:

a) all frames that originate from the Routing Function that are not destined for Proxy N_Ports or Native Nx_Ports; and
b) all frames that are received and addressed to the IFR Management Function (see 5.5), as opposed to a Proxy N_Port.

A link shall be created between each Translate Domain Switch and a single Front Domain Switch connected to the Fabric. The ports on both ends of this link shall operate as E_Ports.

Each Translate Domain Switch shall provide one link to the Routing Function. The Translate Domain Switch shall operate the port connected to this link as an F_Port that is capable of Multiple N_Port_ID Assignment (see FC-LS-2). The Routing Function shall operate the port connected to this link as an N_Port that is capable of Multiple N_Port_ID Support (see FC-LS-2). This link shall be used by the Routing Function to communicate with the associated Fabric as follows:

a) all frames that originate from the Routing Function as a frame sourced from a Proxy N_Port (i.e., routed frames from the egress Routing Function); and
b) all frames that are received and addressed to a Proxy N_Port (i.e., frames to be sent to the ingress Routing Function).

5.5 IFR Management Function

The IFR Management Function shall reside in either the Routing Function or the Front Domain Switch. The IFR Management Function is responsible for the processing of IFR_ILS requests and responses.

5.6 Inter-Fabric Router discovery

5.6.1 IFR discovery protocol overview

The IFR discovery protocol is based on Name Server registration (see 5.6.2) and subsequent queries (see 5.6.3), and the IF_ILS_HLO IFR_ILS (see 7.2.3). The Name Server is used to discover neighbor Inter-Fabric Routers (see 5.6.4), and the IF_ILS_HLO IFR_ILS is used to establish adjacency with the neighbor Inter-Fabric Router and monitor the health of the link between adjacent Inter-Fabric Routers. The IFR discovery protocol allows use of either an N_Port associated with the Front Domain or the Domain Controller of the Front Domain for the IFR Management Function even though the Domain Controller N_Port_ID is not registered with the Name Server.

5.6.2 Name Server registration

As part of initialization, an IFR Management Function shall register its N_Port_ID, Node_Name, N_Port_Name, and FC-4 TYPEs it supports with the Name Server. An IFR Management Function shall register FC-4 TYPE 25h (i.e., Fibre Channel Inter-Fabric Router Services) with the Name Server. The registered N_Port_ID shall be an address identifier of a Routing Function N_Port associated with the Front Domain.

5.6.3 Name Server requirements

If a query for Fibre Channel Inter-Fabric Router Services (i.e., FC-4 TYPE 25h) is received by the Name Server, the query shall not be subject to Zoning and an RSCN or Event Notification shall be sent if an IFR Management Function is added to the Name Server.

5.6.4 IFR neighbor discovery

IFR neighbor discovery is performed by querying the Name Server using the FC-4 TYPE associated with the IFR Management Function (i.e., 25h) to determine, at minimum, the N_Port_ID and Node_Name of all neighbor Inter-Fabric Routers.

At this point, an Inter-Fabric Router is aware of the N_Port_IDs and the Node_Names of the neighbor IFR Management Functions, if any.

For each discovered neighbor IFR Management Function, an IFR Management Function shall:

a) send an IF_ILS_HLO to the neighbor IFR Management Function with the D_ID field set to its discovered N_Port_ID and the S_ID field set to the actual N_Port_ID used by the originating IFR Management Function; and
b) send an IF_ILS_HLO to the neighbor IFR Management Function with the D_ID field set to the Domain Controller N_Port_ID associated with the discovered N_Port_ID and the S_ID field set to the actual N_Port_ID used by the originating IFR Management Function.

The receiving IFR Management Function shall:

a) respond to the IF_ILS_HLO sent to the address identifier used by the receiving IFR Management Function;
b) discard the other IF_ILS_HLO; and
c) continue with IFR adjacency establishment as specified in (see 5.6.5).

5.6.5 IFR adjacency establishment

IFR adjacency establishment is performed by sending IF_ILS_HLO messages to a neighbor IFR Management Function according to the following rules:

a) initially send IF_ILS_HLO messages (see 5.6.4) at a frequency of at least every IF_ILS_HLO_INT for a period of one minute. This is the accelerated discovery phase;
b) following the accelerated discovery phase, IF_ILS_HLO messages are sent at a frequency of at least every IF_ILS_HLO_INT*10; and
c) the accelerated phase is reinitiated when:
   A) no IF_ILS_HLO message is received from a known Inter-Fabric Router for IFR_TOV;
   B) the first IF_ILS_HLO message is received;
   C) a new IFR Management Function has been discovered via the Name Server; or
   D) an IFR Management Function has been removed from the Name Server.

IFR adjacency is established when the neighbor IFR Management Functions have synchronized their IFR topology databases (see 5.6.6).

If an IF_ILS_HLO message is not received from a known Inter-Fabric Router for IFR_TOV, the adjacency is de-established.

5.6.6 IFR discovery state machine

5.6.6.1 Overview

For each detected remote IFR Management Function, there is an instance of the IFR discovery state machine specified in figure 5.

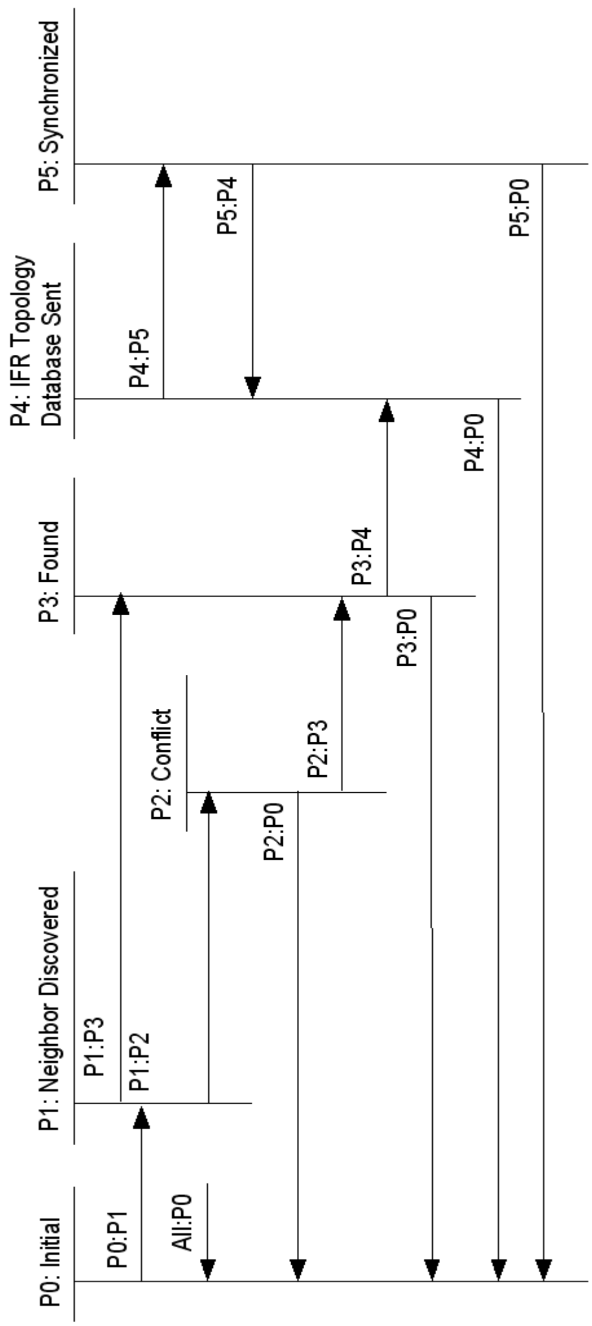

Figure 5 – IFR discovery state machine

5.6.6.2 IFR discovery state machine states and transitions

State P0: Initial. A remote IFR Management Function is available. Perform Name Server query.

Transition ALL:P0. A remote IFR Management Function is detected, an initialization event occurs, or the remote IFR Management Function entry has been removed from the Name Server.

Transition P0:P1. This transition occurs due to the detection of the remote IFR Management Function entry in the Name Server.

State P1: Neighbor Discovered. The remote IFR Management Function has been discovered using the Name Server.

Transition P1:P2. An IF_ILS_HLO with an F_ID that is not the same as the local Inter-Fabric Router is received from the remote IFR Management Function.

Transition P1:P3. This transition is made when an IF_ILS_HLO with the same F_ID as the local Inter-Fabric Router is received from the remote IFR Management Function.

State P2: Conflict. In this state, the F_ID of the remote Inter-Fabric Router does not match the F_ID of the local Inter-Fabric Router. The IFR topology database and the IF_Zone database are not shared with the neighbor Inter-Fabric Router. Database entries from the remote Inter-Fabric Router shall be removed.

Transition P2:P0. No IF_ILS_HLO is received from the remote IFR Management Function within IFR_TOV.

Transition P2:P3. This transition is made when an IF_ILS_HLO with the same F_ID as the local Inter-Fabric Router is received from the remote IFR Management Function.

State P3: Found. The local Inter-Fabric Router and the remote Inter-Fabric Router have the same F_ID. The local Inter-Fabric Router shall transmit the IFR topology database to the remote Inter-Fabric Router.

Transition P3:P0. No IF_ILS_HLO is received from the remote IFR Management Function within IFR_TOV.

Transition P3:P4. This transition occurs when the local Inter-Fabric Router has completed transmission of the IFR topology database to the remote Inter-Fabric Router.

State P4: IFR Topology Database Sent. The IFR topology database associated with the local Inter-Fabric Router has been sent to the remote Inter-Fabric Router. Inter-Fabric Router is waiting for an acknowledgement.

Transition P4:P0. No IF_ILS_HLO has been received from the remote IFR Management Function within IFR_TOV.

Transition P4:P5. This transition occurs when an acknowledgement from the remote IFR Management Function is received.

State P5: Synchronized. The IFR topology database associated with the remote Inter-Fabric Router is fully synchronized with the IFR topology database of the local Inter-Fabric Router.

Transition P5:P4. The IFR topology database has been sent to the remote IFR Management Function.

Transition P5:P0. No IF_ILS_HLO has been received from the remote IFR Management Function within IFR_TOV.

5.6.7 Inter-Fabric Router switch support discovery

After completion of IFR discovery, the Domain Controller associated with the IFR Management Function shall transmit an Exchange Switch Support (ESS) SW_ILS containing a Switch Support Capability object (see FC-SW-5), at minimum, to each Domain Controller in the Inter-Fabric to determine if all switches in the inter-fabric contain the proper functionality for the Inter-Fabric Router to operate. If an Inter-Fabric Router determines that the inter-fabric does not contain the proper functionality for the Inter-Fabric Router to operate, the Inter-Fabric Router may transition to the isolated state (see FC-SW-5).

5.7 Class 2/F support

5.7.1 Overview

An issue exists in that if an Enc_Header is added to a frame and a Class 2/F F_RJT or F_BSY is generated by a legacy switch within the Inter-Fabric, the Inter-Fabric Router that receives the Class 2/F F_RJT or F_BSY does not have the proper addressing information needed to return the frame to the device that originated the Class 2/F frame that was rejected or busied. An example is shown in figure 6.

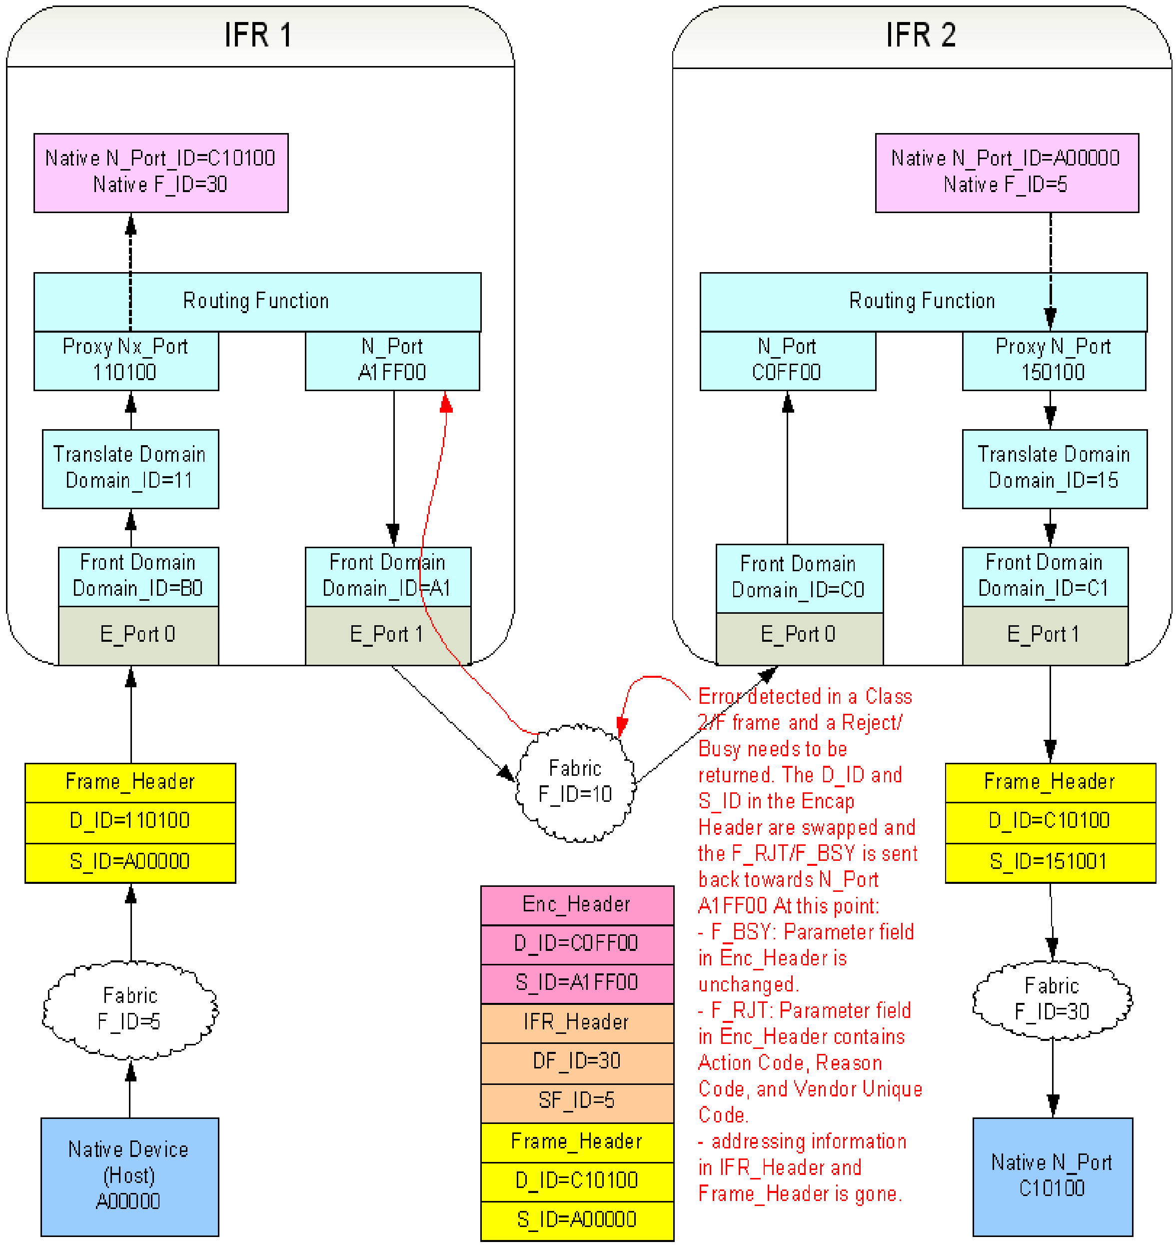

Figure 6 – Inter-Fabric Router Class 2/F F_RJT/F_BSY example

5.7.2 F_RJT and F_BSY processing for Class 2/F

If an Inter-Fabric Router receives a Class 2/F F_RJT or Class 2/F F_BSY, the Inter-Fabric Router shall:

a) discard the Class 2/F F_RJT or Class 2/F F_BSY if it does not have the proper address identifiers needed to forward the frame to the entity that originated the Class 2/F frame; or
b) process the frame as specified in FC-SW-5.

5.8 Enc_Header and IFR_Header processing

5.8.1 Overview

The format of the Enc_Header and IFR_Header is specified in FC-FS-3.

5.8.2 Enc_Header field processing

The Enc_Header fields are set as follows:

a) the D_ID field shall be set to the N_Port_ID of the next hop Routing Function;
b) the S_ID field shall be set to the N_Port_ID of the Routing Function that is sending the frame;
c) the R_CTL field shall be set to 52h; and
d) all other fields shall be set to the values contained in the associated Frame_Header.

5.8.3 IFR_Header field processing

5.8.3.1 DF_ID field

The DF_ID field shall be set by the ingress Inter-Fabric Router to the identifier of the destination fabric.

5.8.3.2 SF_ID field

The SF_ID field shall be set by the ingress Inter-Fabric Router to the identifier of the source fabric.

5.8.3.3 Exp_Time field

The Expiration Time (Exp_Time) field shall be set by the ingress Inter-Fabric Router to the time at which Inter-Fabric Routers receiving the IFR_Header shall discard the received frame. The Exp_Time field shall be set to a valid value regardless of the setting of the Expiration Timestamp Valid (ETV) bit (see FC-FS-3).

The use of the Exp_Time value as an expiration timer requires that all Inter-Fabric Routers that process the IFR_Header have the same time value +/- one count.

The Exp_Time field value shall be compared to the equivalent of bits 37 to 30 (i.e., 0.25 sec increments) in the Network Time Protocol 64-bit timestamp field (see RFC 2030).

The ingress Inter-Fabric Router shall not set the Exp_Time field to a value:

a) that results in the the difference between the Exp_Time value and the exp_timestamp (see FC-FS-3) value being greater than the R_A_TOV value used by the source fabric; or
b) more than 126 x 0.25 (i.e., 31.5 seconds) greater than the exp_timestamp value of the ingress Inter-Fabric Router.

If the ETV bit is set to one, an Inter-Fabric Router shall discard a received frame if the result of (exp_timestamp - Exp_Time) is greater than 2 and less than 127 using modulo 256 arithmetic.

5.8.3.4 Hop_Cnt field

The ingress Inter-Fabric Router shall set the initial Hop_Cnt value.

If the Hop Count Valid (HCV) bit is set to one, an Inter-Fabric Router shall process the frame as follows:

a) if an Inter-Fabric Router receives a frame with a Hop_Cnt field value of 01h, the frame shall be discarded;
b) if an Inter-Fabric Router receives a frame with a Hop_Cnt field value of 00h, the frame shall not be discarded; or
c) each Inter-Fabric Router that receives and forwards the frame shall decrement the Hop_Cnt field value by one.

5.9 Operation of the Routing Function

5.9.1 Overview

The Routing Function appears as an N_Port to the Translate Domain Switches and as an N_Port to the Front Domain Switches. This function presents the Proxy N_Ports to the Fabric. In addition, the Routing Function performs N_Port_ID translation and forwards frames toward their ultimate physical destination. As described previously, the operation of the Routing Function is broken up into two roles:

a) ingress Routing Function role: covers the Routing Function operations required of the first Routing Function in the path of a given frame; or
b) egress Routing Function role: covers the Routing Function operations required of the last Routing Function in the path of a given frame.

5.9.2 Ingress Routing Function role

A Routing Function that receives a frame destined to a Proxy N_Port performs the ingress Routing Function.

Frames received by the ingress Routing Function are addressed to a Proxy N_Port (i.e., the D_ID field contains the N_Port_ID of the Proxy N_Port). The S_ID field is set to the N_Port_ID of the device that generated the original frame.

The ingress Routing Function shall perform the following functions:

a) translate the Frame_Header D_ID to the N_Port_ID of the Native Nx_Port;
b) translate N_Port_IDs embedded in the frame payload as specified in 5.10;
c) store Exchange context required for translating embedded N_Port_IDs as specified in 5.10.5;
d) add an IFR_Header and Enc_Header as specified in 5.8 if the Routing Function is not performing the combined ingress and egress Routing Function role; and
e) forward the frame to the next hop Routing Function or to the egress Routing Function if this is the only Routing Function on the path.

Since the frame contains the address of a Proxy N_Port, it enters the Inter-Fabric Router through a Front Domain Switch attached to the Fabric. The frame is forwarded by the Front Domain Switch to the Translate Domain Switch that supports the Fabric in which the Native Nx_Port resides. The Translate Domain Switch then forwards the frame to the Routing Function.

The Routing Function shall maintain associations between the N_Port_IDs of each Proxy N_Port it presents, the N_Port_ID of its corresponding Native Nx_Port, and the Fabric Identifier of the Fabric on which the Native Nx_Port physically resides. These associations are established using the IFR discovery state machine (see 5.6.6). Using this information, the Routing Function shall replace the D_ID in the frame header with that of the corresponding Native Nx_Port.

If this Inter-Fabric Router also hosts the Proxy N_Port representing the source of this frame in the Native Fabric, then this Routing Function shall perform the egress Routing Function Role, as specified in 5.9.3. Otherwise, the frame shall be forwarded by the Routing Function to its N_Port connected to one of the Local Domain Switches that is connected to the next hop Fabric. The next hop Fabric and the next hop Inter-Fabric Router shall be determined as those that provide the least cost to the Native Fabric. At this point, normal switching operations deliver the frame to the next hop Inter-Fabric Router and its Routing Function.

5.9.3 Egress Routing Function role

A Routing Function that receives a Routable frame destined to a Fabric for which the Routing Function provides the Proxy N_Port for the source of the frame shall perform the egress Routing Function.

The first header in these frames will be either a Frame_Header or an Enc_Header. If it is a Frame_Header, then the frame is forwarded by the Front Domain Switch to the Routing Function as required by 5.3. If the first header of the frame is an Enc_Header, then the D_ID field of the Enc_Header will be that of the N_Port of the attached Routing Function and, consequently, the Front Domain Switch will forward the frame to the Routing Function using normal switching procedures.

The egress Routing Function shall perform the following functions:

a) validate the FC headers and discard the frame if required (see FC-FS-3);
b) translate the Frame_Header S_ID to the N_Port_ID of the Proxy N_Port that, in the Native Fabric, represents the port that originated the frame in the Inter-Fabric;
c) remove the IFR_Header and the Enc_Header, if present; and
d) forward the frame to the Native Fabric where normal switching operation will deliver the frame to the Native Nx_Port.

The egress Routing Function shall perform the validation of the Frame_Header and possibly discard the frame as required by FC-FS-3.

If the frame is not discarded, the Routing Function shall maintain associations between the N_Port_ID's of each Proxy N_Port it presents, the N_Port_ID of its corresponding Native Nx_Port, and the Fabric Identifier of the Fabric on which the Native Nx_Port physically resides. These associations are established using the Inter-Fabric Router Zone Protocol (see clause 8). Using this information, the Routing Function shall replace the source N_Port_ID in the Frame_Header with that of the Proxy N_Port that represents the source of the frame in the Native Fabric.

If valid, the frame shall be forwarded by the Routing Function to the Translate Domain Switch associated with the Proxy N_Port for the original source of the frame. From this point, normal switching operations forward the frame to the Native Nx_Port.

5.9.4 Combined ingress and egress Routing Function role

In certain topologies a Routing Function performs both the role of the ingress and egress Routing Function. This role is the effective combination of each role individually rendering certain operations moot (e.g., adding and removing the Enc_Header and IFR_Header). To be compliant with this standard, as specified in 5.2, an Inter-Fabric Router shall present externally observable behavior consistent with the requirements of this clause. How this is accomplished is implementation dependent. Therefore, the operation of these moot requirements, such as whether or not the Frame_Header is actually added and removed in this combined case, is beyond the scope of this standard.

5.10 Link Service processing

5.10.1 Overview

Certain Link Service Requests and Accepts contain N_Port_IDs embedded in the frame payload. As such, N_Port_ID translation of the embedded N_Port_IDs shall be performed by the ingress Routing Function as specified in 5.10.2 and 5.10.3.

5.10.2 ELS with embedded N_Port_ID processing requirements

Table 2 specifies the processing requirements for ELSs that contain embedded N_Port_IDs.

Table 2 – ELS with embedded N_Port_IDs

ELS	Description	Process embedded N_Port_ID in IFR
LOGO	Payload - word 1	Y
RCS	Payload - word 1	N[a]
RCS LS_ACC	Payload - word 1	N
RLS	Payload - word 1: If the D_ID is a Domain Controller well known address (FFFCxxh), the N_Port_ID field shall be set to an N_Port_ID within the associated domain. The LESB requested is for the F_Port that the N_Port_ID is logged in with.	Y
RRQ	Payload - word 1: Used for Class 2.	Y
RSI	Payload - word 1 (obsolete)	N[a]
FAN	Payload - word 1: Used to report the Loop Fabric Address for the FL_Port.	N[a]
RSCN	Affected Port ID page(s)	N[a]
RLIR	Common Link Incident Record - Incident Port_ID field (obsolete)	N
GAID (obsolete)	NP_List Entry Format	N[a]
FACT (obsolete)	Nx_Port List	N[a]
FDACT (obsolete)	Nx_Port List	N[a]

a) The Inter-Fabric Router shall reply with an LS_RJT with the reason code set to “Command not supported” and the reason code explanation set to “Request not supported”.

Table 2 – ELS with embedded N_Port_IDs

ELS	Description	Process embedded N_Port_ID in IFR
ADISC	Payload - word 1: Hard Address of Originator field shall be set to all zeroes. word 6: N_Port_ID of Originator	Y
ADISC LS_ACC	Payload - word 1: Hard Address of Responder field shall be set to all zeroes. word 6: N_Port_ID of Responder	Y
TPRLO	Logout Parameter Page - word 3: (obsolete)	N
TPRLO LS_ACC	Logout Parameter Page - word 3: (obsolete)	N
RPBC	Originator S_ID (obsolete)	N
SRL	Flag parameter: If set to 01h, specifies the address identifier of the FL_Port. (obsolete)	N[a]
REC	Payload - word 1: Exchange Originator S_ID	Y
REC LS_ACC	Payload - word 2: Originator Address Identifier word 3: Responder Address Identifier	Y
a) The Inter-Fabric Router shall reply with an LS_RJT with the reason code set to “Command not supported” and the reason code explanation set to “Request not supported”.		

5.10.3 FC-4 Link Service with embedded N_Port_ID processing requirements

Table 3 specifies the processing requirements for FC-4 Link Services that contain embedded N_Port_IDs.

Table 3 – FC-4 Link Service with embedded N_Port_ID processing requirements

FC-4 Link Service	Description	Process embedded N_Port_ID in IFR
REC	Payload - word 1: Excahnge Originator S_ID	Y
REC LS_ACC	Payload - word 2: Originator Address Identifier word 3: Responder Address Identifier	Y

5.10.4 ELS Receive Data_Field Size processing

An Inter-Fabric Router shall monitor the following ELS frames:

a) Port Login (PLOGI);
b) Port Discovery (PDISC); and
c) Report Buffer Conditions (RPBC).

If the Buffer-to-Buffer Receive Data_Field Size field value, Receive Data_Field Size field value, or ELS Receive Data Field Size field value in the received frame is greater than 2048, then the Inter-Fabric Router shall set the field value to 2048 prior to forwarding the frame to the destination.

NOTE 1 – The Buffer-to-Buffer Receive Data_Field Size field value, Receive Data_Field Size field value, or ELS Receive Data Field Size field value is set to 2048 to ensure interoperability with legacy devices that are not receive larger frame sizes due to the addition of Extended Headers (e.g., the Enc_Header and IFR_Header).

5.10.5 Link Service Exchange Context storage requirements

To monitor and potentially modify specific Link Service LS_ACC payloads (e.g., embedded N_Port_IDs and the Receive Data Field Size value), the ingress Routing Function shall store and distribute the Exchange context (see clause 9) to all redundant Inter-Fabric Routers associated with the Proxy N_Port, the following ELS and FC-4 Link Service Requests:

a) Port Login (PLOGI);
b) Port Discovery (PDISC);
c) Report Buffer Conditions (RPBC);
d) Discover Address (ADISC); and
e) Read Exchange Concise (REC).

6 Fabric addressing

6.1 Fabric_Identifiers

Each Fabric is identified by a twelve bit Fabric_Identifier (F_ID). The F_ID is an administravely configured value and the values 0 and FF0h through FFFh are reserved.

6.2 Worldwide_Name and Domain_ID usage

Each Translate Domain Switch, Front Domain Switch, and the Routing Function within the Inter-Fabric Router operate as independent entities. Consequently, each shall identify themselves with a unique Worldwide_Name. Likewise, each Translate Domain Switch and Front Domain Switch shall obtain a Domain_ID that is unique within the scope of the Fabrics that they belong to using the normal Principal Switch Domain_ID assignment procedures (see FC-SW-5).

6.3 Routing Function N_Port_IDs

Each port of the Routing Function presents itself as an N_Port attached to the Fabric. As such, each of these ports shall obtain an N_Port_ID with the Domain_ID portion matching that of its associated Front Domain Switch or Translate Domain Switch. Since the interface between the Front Domain Switches, Translate Domain Switches, and the Routing Function is internal to the Inter-Fabric Router, the mechanism by which this N_Port_ID is obtained is implementation specific and not within the scope of this standard.

7 Inter-Fabric Router Internal Link Services (IFR_ILS)

7.1 IFR_ILS Overview

This subclause describes Link Services that operate internal to the Fabric between the Routing Functions of Inter-Fabric Routers. All IFR_ILS frames shall be transmitted using the Class F service. The Frame_Header field values for all IFR_ILS frames is specified in table 4. All other Frame_Header fields shall be set as appropriate according to the rules defined in FC-FS-3.

Table 4 – Frame_Header field values for IFR_ILS frames

Frame_Header field	Value
R_CTL	Set to 02h for all request frames, and to 03h for all reply frames.
CS_CTL	Set to 00h.
D_ID	Set as indicated for the specific IFR_ILS.
S_ID	Set as indicated for the specific IFR_ILS.
TYPE	Set to 25h, indicating Fibre Channel Inter-Fabric Router Services.

Unless otherwise specified, the rules for Inter-Fabric Router Internal Link Services are the same as for Extended Link Services as specified in FC-FS-3 (e.g., Sequence and Exchange management, error detection and recovery). Time-out values for specific IFR_ILS's and the actions following a time-out expiration are specified in clause 10.

7.2 Inter-Fabric IFR_ILS (IF_ILS) requests

7.2.1 Inter-Fabric IFR_ILS request overview

The IF_ILS is used to communicate information between Inter-Fabric Routers. IF_ILS requests are used to:

a) discover all Inter-Fabric Routers;
b) exchange IFR topology databases; and
c) exchange zoning information.

For IF_ILS requests the S_ID shall be set to the N_Port_ID of the originating IFR Management Function and the D_ID shall be set to the N_Port_ID of the destination IFR Management Function.

7.2.2 IF_ILS_REQ Header format

The format of the IF_ILS_REQ Header is specified in table 5.

Table 5 – IF_ILS_REQ Header format

Item	Size (bytes)
Command	4
Subcommand	4
Payload Version	2
Capability Version	2
Command Specific Payload	n

Command: The Command field shall be set to 44000000h.

Subcommand: The Subcommand field specifies the subcommand for the IF_ILS_REQ. Subcommand values are specified in table 6.

Table 6 – IF_ILS_REQ Subcommand values

Value	Description	Abbr.	Reference
00030001h	IF_ILS Hello	IF_ILS_HLO	7.2.3
00030002h	IF_ILS Topology Database Send	IF_ILS_TDS	7.2.4
00030003h	IF_ILS Topology Database ACK	IF_ILS_TD_ACK	7.2.5
00040001h	IF_ILS Zone Database Send	IF_ILS_ZDS	7.2.6
00040002h	IF_ILS Zone Database ACK	IF_ILS_ZD_ACK	7.2.6

Payload Version: The Payload Version field specifies the version of the payload that follows the IF_ILS_REQ Header. The Payload Version field shall be set to 0001h.

Capability Version: The Capability Version field specifies the highest version number of the command specific payload that the originator is able to process.

Command Specific Payload: The Command Specific Payload field contains the payload associated the the IF_ILS_REQ subcommand.

7.2.3 IF_ILS Hello (IF_ILS_HLO) subcommand

The IF_ILS_HLO subcommand is used to determine when two-way communication is established with a neighbor Inter-Fabric Router. The exchange of IFR Management Function N_Port_IDs is also used to determine the health of the path.

The IF_ILS_HLO subcommand shall be sent as a unidirectional Exchange.

The format of the IF_ILS_HLO payload is specified in table 7.

Table 7 – IF_ILS_HLO payload format

Item	Size (bytes)
Domain_ID	1
Fabric Identifier	2
Reserved	1

Domain_ID: The Domain_ID field shall specify the Domain Identifier associated with the originating IFR Management Function port.

Fabric Identifier: The Fabric Identifier field shall specify the Fabric Identifier of the Fabric associated with the originating IFR Management Function port.

7.2.4 IF_ILS Topology Database Send (IF_ILS_TDS) subcommand

7.2.4.1 IF_ILS_TDS overview

The IF_ILS_TDS subcommand is used to exchange IFR topology databases between Inter-Fabric Routers.

The format of the IF_ILS_TDS payload is specified in table 8.

Table 8 – IF_ILS_TDS payload format

Item	Size (bytes)
Fabric_ID	2
Reserved	1
Domain_ID	1
Generation Count	4
Number of IFR Nx_Port Descriptors (n)	4
IFR Nx_Port Descriptor 1	see 7.2.4.2
IFR Nx_Port Descriptor n	see 7.2.4.2

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier for the Inter-Fabric Router.

Domain_ID: The Domain_ID field specifies the Domain_ID of the Inter-Fabric Router.

Generation Count: The Generation Count field specifies a number representing when the IFR topology database was last updated. The generation count shall be a monotonically increasing value that is assigned by the entity modifying or updating the IFR topology database. The generation count shall be incremented by one whenever the IFR topology database is changed.

Number of IFR Nx_Port Descriptors: The Number of IFR Nx_Port Descriptors field specifies the number of IFR Nx_Port Descriptors contained in the IF_ILS_RDS payload.

7.2.4.2 IFR Nx_Port Descriptor

An IFR Nx_Port Descriptor specifies information required to enable inter-fabric routing between a pair of Native Nx_Ports. One IFR Nx_Port Descriptor shall be contained in the IFR topology database for each Proxy Nx_Port presented by the Inter-Fabric Router in the Inter-Fabric.

The IFR Nx_Port Descriptor format is specified in table 9.

Table 9 – IFR Nx_Port Descriptor format

Item	Size (bytes)
Proxy Fabric_ID	2
Reserved	2
Native Fabric - Principal Switch Name	8
Proxy N_Port_ID	4
Cost	4

Proxy Nx_Port Fabric_ID: The Proxy Nx_Port Fabric_ID field specifies the Fabric Identifier associated with the Proxy Nx_Port.

Native Fabric - Principal Switch Name: The Native Fabric - Principal Switch Name specifies the name of the Principal Switch for the Native Fabric associated with the Proxy Nx_Port.

Proxy N_Port_ID: The Proxy N_Port_ID field specifies the address identifier of the Proxy Nx_Port.

Cost: The Cost field specifies the cost of the route through the Proxy Nx_Port.

7.2.5 IF_ILS Topology Database ACK (IF_ILS_TD_ACK) subcommand

The format of the IF_ILS_TD_ACK subcommand payload is specified in table 10.

Table 10 – IF_ILS_TD_ACK payload format

Item	Size (bytes)
Fabric_ID	2
Reserved	1
Domain_ID	1
Generation Count	4

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier for the Inter-Fabric Router.

Domain_ID: The Domain_ID field specifies the Domain_ID of the Inter-Fabric Router.

Generation Count: The Generation Count field specifies a number representing when the IFR topology database was last updated.

7.2.6 IF_ILS Zone Database Send (IF_ILS_ZDS) subcommand

7.2.6.1 IF_ILS_ZDS overview

The IF_ILS_ZDS subcommand allows Inter-Fabric zoning information to be synchronized among the Inter-Fabric Routers in the Inter-Fabric.

The format of the IF_ILS_ZDS payload is specified in table 11.

Table 11 – IF_ILS_ZDS payload format

Item	Size (bytes)
Number of IF Zone Descriptors	4
IF_Zone Descriptor 1	see 7.2.6.2
IF_Zone Descriptor n	see 7.2.6.2

Number of IF_Zone Descriptors: The Number of IF_Zone Descriptors field specifies the number of IF_Zone Descriptors contained in the IF_ILS_ZDS payload.

7.2.6.2 IF_Zone Descriptor

The format of the IF_Zone Descriptor is specified in table 12.

Table 12 – IF_Zone Descriptor format

Item	Size (bytes)
IF_Zone Name	60
Fabric_ID	2
Reserved	2
Generation Count	4
Number of IF Device Descriptors (n)	4
IF Device Descriptor 1	see 7.2.6.3
IF Device Descriptor n	see 7.2.6.3

IF_Zone Name: The IF_Zone Name field specifies the name of the IF_Zone.

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier of the IF_Zone's Native Fabric (i.e., the Fabric where the IF_Zone is defined).

Generation Count: The Generation Count field specifies a number representing when the IF_Zone was last updated.

Number of IF Device Descriptors: The Number of IF Device Descriptors field specifies the number of IF Device Descriptors contained in the IF_Zone Descriptor payload.

7.2.6.3 IF Device Descriptor

The format of the IF Device Descriptor is specified in table 13.

Table 13 – IF Device Descriptor format

Item	Size (bytes)
Port_Name	8
Fabric_ID	2
Reserved	2
Proxy Nx_Port Address Identifier	4
Native Nx_Port Address Identifier	4
Port State	4
Node_Name	8
Class of Service	4
Port Type	4
FC-4 Support	8
Fabric Port_Name	8
Device Fabric_ID	2
Reserved	2
Generation Count	4
RSCN State	240

Port_Name: The Port_Name field specifies the Port_Name of the device.

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier of the Fabric where the device is defined in the IF_Zone.

Proxy Nx_Port Address Identifier: When a device is located on a Fabric that is different from the IF_Zone's Native Fabric, the Proxy Nx_Port Address Identifier field specifies the Proxy Nx_Port address identifier for that device from the perspective of the IF_Zone's Native Fabric.

Native Nx_Port Address Identifier: When a device is located on a Fabric that is the same Fabric as the IF_Zone's Native Fabric, the Native Nx_Port Address Identifier field specifies the Nx_Port address identifier for that device. The Native Nx_Port Address Identifier field is only used when the device is physically attached to the IF_Zone's Native Fabric.

Port State: The Port State field specifies the state of the device. Valid states are specified in table 14.

Table 14 – Port State field values

Value	State
1	CONFIGURED
2	EXIST
3	INITIALIZING
4	IMPORTED

Node_Name: The Node_Name associated with the device port.

Class of Service: The Class of Service field specifies the class of service associated with the device port.

Port Type: The Port Type field specifies the port type for the device port.

FC-4 Support: The FC-4 Support field specifies the FC-4s that are supported at the device port.

Fabric Port_Name: The Fabric Port_Name field specifies the Port_Name of the Switch port that the device port is connected to.

Device Fabric_ID: The Device Fabric_ID field specifies the Fabric Identifier of the device port's Native Fabric.

Generation Count: The Generation Count field specifies a number representing when the device descriptor information was last updated.

RSCN State: The RSCN State field specifies an indexed table that specifies whether RSCNs have been sent to a particular Domain.

7.2.7 IF_ILS Zone Database ACK (IF_ILS_ZD_ACK) subcommand

7.2.7.1 IF_ILS_ZD_ACK overview

The IF_ILS_ZD_ACK subcommand is used to synchronize Inter-Fabric zoning information among Inter-Fabric Routers in the Inter-Fabric.

The format of the IF_ILS_ZD_ACK payload is specified in table 15.

Table 15 – IF_ILS_ZD_ACK payload format

Item	Size (bytes)
Number of IF_Zone ACK Descriptors (n)	4
IF_Zone ACK Descriptor 1	see 7.2.7.2
IF_Zone ACK Descriptor n	see 7.2.7.2

Number of IF_Zone ACK Descriptors: The Number of IF_Zone ACK Descriptors field specifies the number of IF_Zone ACK Descriptors contained in the IF_ILS_ZD_ACK payload.

7.2.7.2 IF_Zone ACK Descriptor

The format of the IF_Zone ACK Descriptor is specified in table 16.

Table 16 – IF_ILS_ZD_ACK Descriptor format

Item	Size (bytes)
IF_Zone Name	60
Fabric_ID	2
Reserved	2
Generation Count	4

IF_Zone Name: The IF_Zone Name field specifies the name of the IF_Zone.

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier for the IF_Zone's Native Fabric (i.e., the Fabric where the IF_Zone is defined).

Generation Count: The Generation Count field specifies a number representing when the IF_Zone was last updated.

7.3 Front Domain IFR_ILS (FD_ILS) requests

7.3.1 Front Domain ILS request overview

The FD_ILS is used to communicate information between Front Domains that attach to the same Fabric and between Inter-Fabric Routers attached to the same Fabric. The FD_ILS is used to:

a) discover all Front Domains attached to a Fabric and have access to the same remote Fabric;
b) assign a Front Domain to be the owner of the remote Fabric;
c) exchange zoning information between Front Domains; and
d) exchange context for embedded N_Port processing.

For FD_ILS requests the S_ID shall be set to the Domain Controller N_Port_ID of the originating Front Domain. The D_ID shall be set to the Domain Controller ID of the destination Front Domain.

7.3.2 FD_ILS_REQ Header format

The format of the FD_ILS_REQ Header is specified in table 17.

Table 17 – FD_ILS_REQ Header format

Item	Size (bytes)
Command	4
Subcommand	4
Payload Version	2
Capability Version	2
Command Specific Payload	n

Command: The Command field shall be set to 45000000h.

Subcommand: The Subcommand field specifies the subcommand for the FD_ILS_REQ. Subcommand values are specified in table 18.

Table 18 – FD_ILS_REQ Subcommand values

Value	Description	Abbr.	Reference
00020001h	Front Domain Hello	FD_ILS_HLO	7.3.3
00020004h	Front Domain Zone Database Send	FD_ILS_ZDS	7.3.4
00020005h	Front Domain Zone Database ACK	FD_ILS_ZD_ACK	7.3.5
00020006h	Front Domain Staging Context	FD_ILS_SC	7.3.6
00020007h	Front Domain Staging Context ACK	FD_ILS_SC_ACK	7.3.7
00020008h	Front Domain Staging Context Delete	FD_ILS_SCD	7.3.8

Payload Version: The Payload Version field specifies the version of the payload that follows the FD_ILS_REQ Header. The Payload Version field shall be set to 0001h.

Capability Version: The Capability Version field specifies the highest version number of the command specific payload that the originator is able to process.

Command Specific Payload: The Command Specific Payload field contains the payload associated the the FD_ILS_REQ subcommand.

7.3.3 Front Domain Hello (FD_ILS_HLO) subcommand

7.3.3.1 FD_ILS_HLO subcommand overview

The FD_ILS_HLO subcommand is used to inform Front Domains in a given Fabric of other Front Domains that are part of the same Fabric. The FD_ILS_HLO subcommand also indicates which remote Fabrics are reachable using a given Front Domain.

The format of the FD_ILS_HLO payload is specified in table 19.

Table 19 – FD_ILS_HLO payload format

Item	Size (bytes)
Front Domain_ID	1
Reserved	3
Principal Switch_Name	8
IFR_Name	8
Attached E_Port Name	8
Number of Fabric Descriptors (n)	4
Fabric Descriptor 1	see 7.3.3.2
Fabric Descriptor n	see 7.3.3.2

Front Domain_ID: The Front Domain_ID field specifies the Domain_ID of the Front Domain associated with a given E_Port presented by the Inter-Fabric Router that is attached to the Fabric. A Front Domain is established within the Inter-Fabric Router for attachment to a Fabric.

Principal Switch_Name: The Principal Switch_Name field name specifies a Name_Identifier that identifies the Principal Switch in the associated Fabric.

IFR_Name: A Name_Identifier that identifies the Inter-Fabric Router containing the Front Domain.

Attached E_Port_Name: The Attached E_Port_Name field specifies a Name_Identifier that identifies the E_Port on the Front Domain that attaches to the Fabric.

Number of Fabric Descriptors: The Number of Fabric Descriptors field specifies the number of Fabric Descriptors contained in the payload.

7.3.3.2 Fabric Desriptor format

The Fabric Descriptor format is specified in table 20.

Table 20 – Fabric Descriptor format

Item	Size (bytes)
Fabric_ID	2
Translate Domain_ID	1
Old Translate Domain_ID	1
Fabric Reachable Flag	1
Low Domain_ID	1
Owner Domain_ID	1
Reserved	1
Translate Domain_Name	8
Number of Translate Port Descriptors (n)	4
Translate Port Descriptor 1	12
Translate Port Descriptor n	12

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier for the remote Fabric.

Translate Domain_ID: The Translate Domain_ID field specifies the Domain_ID of the Translate Domain that represents the remote Fabric for the specified Front Domain.

Old Translate Domain:ID: The Old Translate Domain_ID field specifies the cached Domain_ID of the old Translate Domain if the current Translate Domain_ID is invalid.

Fabric Reachable Flag: The Fabric Reachable Flag field specifies whether the remote Fabric is reachable from this Front Domain. The values for the Fabric Reachable Flag field are specified in table 21.

Table 21 – Fabric Reachable Flag field values

Value	Description
00h	Remote Fabric is not reachable.
01h	Remote Fabric is reachable.
All others	Reserved

Low Domain_ID: The Low Domain_ID field specifies the Front Domain with the lowest Domain_ID that is able to reach the remote Fabric from this Fabric.

Owner Domain_ID: The Owner Domain_ID field specifies the Domain ID of the Front Domain that owns the remote Fabric with respect to this Fabric.

Translate Domain_Name: The Translate Domain_Name field specifies the Switch_Name of the Translate Domain that represents the remote Fabric for the specified Front Domain.

Number of Translate Port Descriptors: The number of Translate Port Descriptors field specifies the number of Translate Port Descriptors contained in the Fabric Descriptor.

7.3.3.3 Translate Port Descriptors

The format of the Translate Port Descriptor is specified in table 22.

Table 22 – Translate Port Descriptor format

Item	Size (bytes)
Port Index	4
Port_Name	8

Port Index: The Port Index field specifies the index assigned to the Translate Port.

Port_Name: The Port_Name field specifies the Name_Identifier of the Translate Port.

7.3.4 Front Domain Zone Database Send (FD_ILS_ZDS) subcommand

The FD_ILS_ZDS subcommand allows Inter-Fabric zoning information to be synchronized among Front Domains in the Fabric.

The format of the FD_ILS_ZDS payload is specified in table 23.

Table 23 – FD_ILS_ZDS payload format

Item	Size (bytes)
Number of IF_Zone Descriptors (n)	4
IF_Zone Descriptor 1	see 7.2.6.2
IF_Zone Descriptor n	see 7.2.6.2

Number of IF_Zone Descriptors: The Number of IF_Zone Descriptors field specifies the number of IF_Zone Descriptors contained in the FD_ILS_ZDS subcommand payload.

7.3.5 Front Domain Zone Database ACK (FD_ILS_ZD_ACK) subcommand

7.3.5.1 FD_ILS_ZD_ACK subcommand overview

The FD_ILS_ZD_ACK subcommand provides an acknowledgement to the FD_ILS_ZDS subcommand.

The FD_ILS_ZD_ACK payload format is specified in table 24.

Table 24 – FD_ILS_ZD_ACK subcommand payload format

Item	Size (bytes)
Number of IF_Zone ACK Descriptors (n)	4
IF_Zone ACK Descriptor 1	see 7.3.5.2
IF_Zone ACK Descriptor n	see 7.3.5.2

Number of IF_Zone ACK Descriptors: The Number of IF_Zone ACK Descriptors field specifies the number of IF_Zone ACK Descriptors contained in the FD_ILS_ZD_ACK subcommand payload.

7.3.5.2 IF_Zone ACK Descriptor

The IF_Zone ACK Descriptor format is specified in table 25.

Table 25 – IF_Zone ACK Descriptor format

Item	Size (bytes)
IF_Zone Name	60
Fabric_ID	2
Reserved	2
Number of Device Descriptors	4
Device Descriptor 1	see 7.3.5.3
Device Descriptor n	see 7.3.5.3

IF_Zone Name: The IF_Zone Name field specifies the name of the IF_Zone.

Fabric_ID: The Fabric_ID field specifies the Fabric Identifier for the IF_Zone"s Native Fabric (i.e., the Fabric where the IF_Zone is defined).

Number of Device Descriptors: The Number of Device Descriptors field specifies the number of Device Descriptors contained in the IF_Zone ACK Descriptor payload.

7.3.5.3 Device Descriptor

The Device Descriptor format is specified in table 26.

Table 26 – Device Descriptor format

Item	Size (bytes)
Port_Name	8
Generation Count	4

Port Index: The Port Index field specifies the name of the device port.

Generation Count: The Generation Count field specifies a number representing when the device descriptor was last updated.

7.3.6 Front Domain Staging Context (FD_ILS_SC) subcommand

The FD_ILS_SC subcommand allows context information required for embedded N_Port_ID processing to be exchanged between Front Domains attached to the same Fabric.

The FD_ILS_SC payload format is specified in table 27.

Table 27 – FD_ILS_SC payload format

Item	Size (bytes)
Generation Count	8
Contect Token	2
ELS Length (n)	2
Reserved	2
Result	2
Reserved	12
ELS	n

Generation Count: The Generation Count field specifies a number to fully distinguish ELSs of the same type (e.g., multiple PLOGIs).

Context Token: The Context Token field contains the context token. The context token is meaningful only to the originator of the FD_ILS_SC subcommand.

ELS Length: The ELS Length field contains the length of the ELS in bytes.

Result: The Result field shall be set to 0001h.

ELS: The ELS field contains the original ELS that was received by the Front Domain.

7.3.7 Front Domain Staging Context ACK (FD_ILS_ASC_ACK) subcommand

The FD_ILS_SC_ACK subcommand provides an acknowledgement to the FD_ILS_SC subcommand.

The FD_ILS_SC_ACK payload format is specified in table 28.

Table 28 – FD_ILS_SC_ACK payload format

Item	Size (bytes)
Generation Count	8
Contect Token	2
ELS Length (n)	2
Reserved	2
Result	2
Reserved	12
ELS	n

Generation Count: The Generation Count field specifies a number to fully distinguish ELSs of the same type (e.g., multiple PLOGIs).

Context Token: The Context Token field contains the context token. The context token is meaningful only to the originator of the FD_ILS_SC subcommand.

ELS Length: The ELS Length field contains the length of the ELS in bytes.

Result: The Result field shall be set to 0001h.

ELS: The ELS field contains the original ELS that was received by the Front Domain.

7.3.8 Front Domain Staging Context Delete (FD_ILS_SCD) subcommand

The FD_ILS_SCD subcommand provides a method to release the context for a particular ELS.

The FD_ILS_SCD payload format is specified in table 29.

Table 29 – FD_ILS_SCD payload format

Item	Size (bytes)
Generation Count	8
Contect Token	2
ELS Length (n)	2
Reserved	2
Result	2
Reserved	12
ELS	n

Generation Count: The Generation Count field specifies a number to fully distinguish ELSs of the same type (e.g., multiple PLOGIs).

Context Token: The Context Token field contains the context token. The context token is meaningful only to the originator of the FD_ILS_SC subcommand.

ELS Length: The ELS Length field contains the length of the ELS in bytes.

Result: The Result field shall be set to 0001h.

ELS: The ELS field contains the original ELS that was received by the Front Domain.

8 Inter-Fabric Router Zone Protocol (IFR_ZP)

8.1 Overview

In order for Nx_Ports to communicate with one another in a network consisting of multiple Fibre Channel Fabrics, they shall first be designated to do so using the Inter-Fabric Router Zone Protocol (IFR_ZP).

For example, assume Fabric A contains Nx_Port DA and Fabric B contains Nx_Port DB. Fabric A contains an IF_Zone entry that allows DA and DB to communicate and Fabric B contains an IF_Zone entry that allows the same. After at least two Inter-Fabric Routers have verified that both Fabrics allow the Nx_Ports DA and DB to communicate and that the Nx_Ports are online in their respective Fabrics, then the Inter-Fabric Routers proxy Nx_Port DB to Fabric A and Nx_Port DA to Fabric B.

Furthermore, IF_Zones shall be synchronized among all Inter-Fabric Routers. The pair-match algorithm (see 8.1.2.3), which specifies the import criteria for an Nx_Port, is triggered based on Nx_Port states (see 8.1.2.2) contained in the IF_Zone.

Figure 2 specifies how these operations are seen externally. Fabric 1 is connected to Fabric 2 using Inter-Fabric Routers 1 and 2. Inter-Fabric Router 1 is directly attached to Fabric 1 and Inter-Fabric Router 2 is directly attached to Fabric 2.

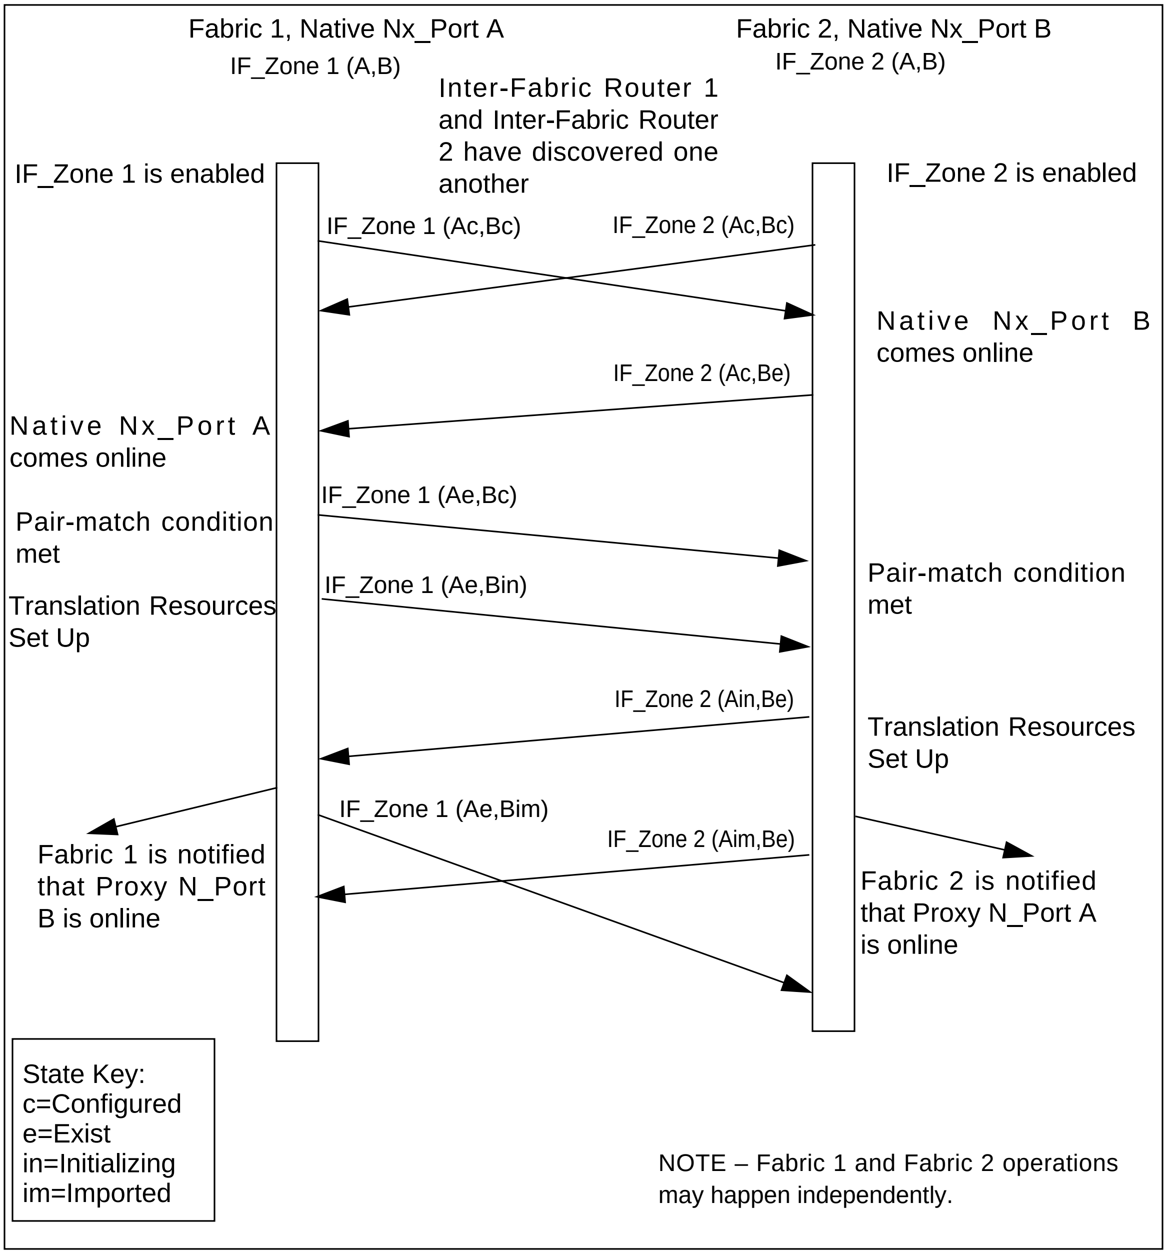

Figure 7 – Example IFR_ZP operation

8.1.1 IF_Zone synchronization

IFR_ZP shall be used to synchronize IF_Zone information across all Inter-Fabric Routers for both simple mode and NAT mode. IF_Zone synchronization occurs between adjacent Inter-Fabric Routers across a Fabric, and between Front Domains that attach to a Fabric when the Front Domains are located in different Inter-Fabric Routers.

IF_Zones are synchronized between Inter-Fabric Routers so that all Inter-Fabric Routers have the most current IF_Zone configurations for all Fabrics in the configuration. All Inter-Fabric Routers have access to IF_Zone configurations even though they may not be connected to the Fabrics where the IF_Zones are defined.

IF_Zone distribution occurs between Inter-Fabric Routers after they have become adjacent. When an Inter-Fabric Router receives IF_Zone information, the generation count is checked. If the latest acknowledged generation count for the IF_Zone associated with the remote Inter-Fabric Router is different than the current generation count for the IF_Zone associated with the local Inter-Fabric Router, then a complete IF_Zone update is sent to the remote Inter-Fabric Router. After the IF_Zone update is received, the remote Inter-Fabric Router acknowledges receipt by sending back the generation count associated with the newly obtained IF_Zone information. The local Inter-Fabric Router maintains this generation count for the remote Inter-Fabric Router. If the remote Inter-Fabric Router becomes not available, or its F_ID is different from the local Inter-Fabric Router, then the local Inter-Fabric Router clears the cached generation count. If not acknowledged, the request is retried at the interval of IF_ZONE_TOV.

When multiple Front Domains attach to the same Fabric, one Front Domain assumes the role of the Owning Front Domain (see 9). IF_Zone information is exchanged between Front Domains to assure a consistent view of Nx_Port information is maintained by each Front Domain. When a Front Domain assumes the role of owner for a Fabric, it ensures consistency of the Translate Domain and Node_Name consistent with the attached Fabric. The Owning Front Domain also maintains Nx_Port information such as proxy N_Port_IDs and Nx_Port state within the attached Fabric.

During the IF_Zone synchronization process, the Owning Front Domain shall send a partial IF_Zone update containing only imported Nx_Ports to a non-owning Front Domain. Again the generation count is checked and if different, the IF_Zone information is updated appropriately. The non-owning Front Domain acknowledges by returning the generation count to the Owning Front Domain. The Owning Front Domain maintains the received generation count on behalf of the other non-owning Front Domain.

The Owning Front Domain maintains the generation count until the Front Domain is removed from the attached Fabric, or if a new owner is selected. If not acknowledged, the request is retried at the interval of IF_ZONE_TOV.

During IF_Zone synchronization the following IFR_ILS requests are available:

a) Front Domain Zone Database Send (see 7.3.4); and
b) Front Domain Zone Database ACK (see 7.3.5).

8.1.1.1 IF_Zone synchronization exceptions

The Symbolic Port_Name and Symbolic Node_Name are not specified as part of IF_Zone synchronization due to their large size and because they are typically associated with management.

When an Inter-Fabric Router receives a Name Server request for either the Symbolic Port_Name or the Symbolic Node_Name, the request shall be routed to the destination Fabric for processing and the subsequent response shall be routed back to the originator.

8.1.2 IFR Nx_Ports and Inter-Fabric Router Zones

8.1.2.1 IFR Nx_Port roles

In the Inter-Fabric Router Zone environment, there are two roles that an IFR Nx_Port assumes. An IFR Nx_Port's Native Fabric is the Fabric where the IFR Nx_Port is physically attached. In this case the IFR Nx_Port assumes the role of a Native Nx_Port. Fabrics other than the Native Fabric where the IFR Nx_Port is not physically attached are known as remote Fabrics. From the perspective of a remote Fabric, the IFR Nx_Port assumes the role of a Proxy Nx_Port. For each Native Nx_Port, there may be multiple associated Proxy Nx_Ports in multiple remote Fabrics.

8.1.2.2 IFR Nx_Port states

The valid states for Native Nx_Ports and Proxy Nx_Ports in the Inter-Fabric Router Zone environment are as follows:

a) Configured - An IFR Nx_Port in this state has been defined in an IF_Zone database associated with the IFR Nx_Port's Native Fabric, or in an IF_Zone database associated with a remote Fabric. Both Native Nx_Ports and Proxy Nx_Ports may be in the configured state;
b) Exist - A Native Nx_Port in this state is online;
c) Initializing - A Proxy Nx_Port is in this state when a pair-match condition involving the IFR Nx_Port is met. A Proxy Nx_Port in this state is initializing the resources for the Proxy Nx_Port; and
d) Imported - A Proxy Nx_Port in this state is online and the pair-match condition is still met. The Fabric has been notified of IFR Nx_Port availability (e.g., SW_RSCNs have been sent).

8.1.2.3 Pair-match condition

Before IFR Nx_Ports attached to different Fabrics may communicate, they both shall be imported to the appropriate remote Fabrics. A pair-match condition between the two IFR Nx_Ports shall be met before IFR Nx_Ports are imported and communication is possible. The pair-match condition from the perspective of an Inter-Fabric Router in Fabric x is met if:

a) IFR Nx_Port xd is online in Fabric x and and IFR Nx_Port yd is online in Fabric y;
b) IF_Zone x has IFR Nx_Ports xd and yd defined (i.e., the Native Nx_Port pair) in Fabric x and IF_Zone y has IFR Nx_Ports xd and yd defined (i.e., the opposite pair) in Fabric y; and;
c) IF_Zone x has an IFR Nx_Port state of Exist for IFR Nx_Port x and IF_Zone y has a IFR Nx_Port state of Exist for IFR Nx_Port y.

8.1.2.4 Inter-Fabric Router Zone IFR Nx_Port state machine

The state machine specifying the IFR Nx_Port state and pair-match condition as it pertains to IF_Zones is specified in figure 8. States P0 and P1 are applicable to the Native Nx_Port. States P0, P2, and P3 are applicable to the Proxy N_Port.

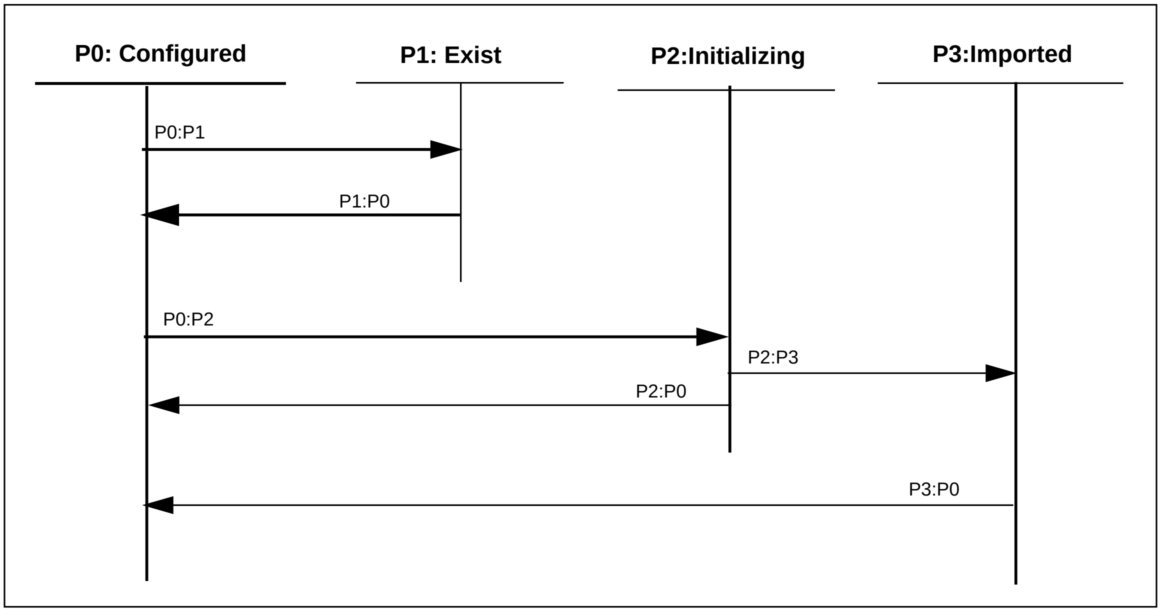

Figure 8 – Inter-Fabric Router Zone IFR Nx_Port state machine

8.1.2.5 Inter-Fabric Router Zone IFR Nx_Port state machine states and transitions

State P0: Configured. The IFR Nx_Port is defined in the Inter-Fabric Router Zone. The Configured state serves as the initial state for an IFR Nx_Port in the Inter-Fabric Router Zone.

Transition P0:P1. This transition is made when a Native Nx_Port is online or an online event has been received. An IFR Nx_Port online event may be either the receipt of an SW_RSCN or the response to a Name Service query.

Transition P0:P2. This transition is made for a Proxy Nx_Port when the pair-match condition is met.

State P1: Exist. In this state, a Native Nx_Port that has been configured in an Inter-Fabric Router Zone is online in its Native Fabric.

Transition P1:P0. This transition is made when a Native Nx_Port goes offline.

State P2: Initializing. With respect to this Proxy Nx_Port, the internal resources to support inter-fabric routing are initializing. Although the Native Fabric does not contain the Proxy Nx_Port in the Name Server, frames are properly translated and forwarded for the Proxy Nx_Port.

Transition P2:P0. This transition is made for a Proxy Nx_Port when the pair-match condition is no longer met.

Transition P2:P3. This transition occurs when the pair-match condition is still met and the Native Nx_Port is in state P2 or state P3 as a Proxy Nx_Port in the Remote Fabric. The Fabric is notified of the Proxy Nx_Ports availability.

State P3: Imported. The Proxy Nx_Port is online in the Fabric.

Transition P3:P0. This transition is made when the pair-match condition is no longer met.

9 Redundant Inter-Fabric Router operation

9.1 Overview

To enable the use of redundant Inter-Fabric Routers, specific information (e.g., Link Service Exchange context, SW_RSCN state) needs to be distributed between Inter-Fabric Routers that are connected to the same Fabric.

In an environment where redundant Inter-Fabric Routers routers are present, a selection process determines which of the Front Domains will operate on behalf of a Translate Domain. This is in the context of the local Fabric that all Front Domains can reach. Following this selection process, one Front Domain is designated as the Owning Front Domain. The Owning Front Domain is responsible for the following:

a) sending the RDI on behalf of a Translate Domain;
b) sending LSUs for Translate Domains owned; and
c) sending SW_RSCNs on behalf of Proxy N_Ports attached to the Translate Domain.

Ownership selection is performed by transmitting an FD_ILS_HLO every IFR_ILS_HLO_INT by all Front Domains. The FD_ILS_HLO contains the reachable remote Fabrics and residual Translate Domain information for non-reachable remote fabrics.

The ownership selection algorithm is as follows:

a) the Front Domain with Translate Domain information is given priority;
b) if no such information is available, then the Front Domain with lowest Domain_ID is given priority; and
c) the Front Domain is able to reach the remote fabric.

9.1.1 Owner selection state machine

The Owner selection state machine is specified in figure 9.

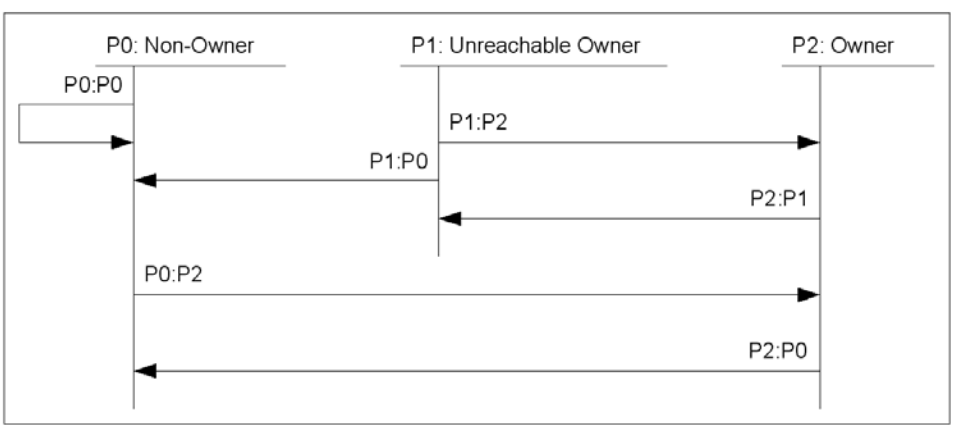

Figure 9 – Owner selection state machine

State P0: Non-Owner. This is the initial state for all Front Domains.

Transition P0:P0. This transition is made when the Front Domain receives another Front Domain's Ownership attainment. The Translate Domain's information is maintained only if the remote Fabric is reachable.

State P1: Unreachable-Owner. In this state, the Front Domain keeps ownership for a remote Fabric that is no longer reachable.

Transition P1:P0. This transition is made when an FD_ILS_HLO is received that indicates a new owner has been selected.

State P2: Owner. The Front Domain has become owner on behalf of a remote Fabric. Ownership attainment is broadcast using the FD_ILS_HLO.

Transition P1:P2. This transition is made when the remote Fabric is reachable again.

Transition P0:P2. This transition is made when the ownership selection algorithm is satisfied.

Transition P2:P1. This transition is made when the remote Fabric is no longer reachable.

Transition P2:P0. This transition is made when an FD_ILS_HLO is received specifying that an owner has been selected and the specified Owner has a lower Domain_ID.

9.1.2 Link Service Exchange context storage processing

Link Service Exchange context saving for embedded N_Ports shall be performed at the ingress Routing Function.

When a Link Service specified in 5.10.5 is received by a Front Domain, Exchange context information is forwarded to all Front Domains connected to the same local Fabric as the device that generated the Link Service, that may be used to reach the destination Fabric. The EF Staging Context IFR_ILS is used for this purpose (see 1.8.2). Each Front Domain then responds with an acknowledgement when the Exchange context is saved. The EF Staging Context ACK IFR_ILS is used for this purpose. When the ingress Front Domain has received EF Staging Context ACKs from all other Front Domains, the Link Service is forwarded to the destination Fabric. To avoid any potential timeout issues, the original Link Service shall be forwarded after 1 second regardless of the even if not all ACKs have been received from other Front Domains.

When any Front Domain attached to the local Fabric receives the LS_ACC for the Link Service, that Front Domain shall release the Exchange context in the other Front Domains by transmitting an EF Staging Context Delete IFR_ILS to all Front Domains. No acknowledgement is required for the EF Staging Context Delete IFR_ILS. All Front Domains shall delete any Exchange context for a particular Link Service after 2 seconds even if an EF Staging Context Delete IFR_ILS has not been received.

10 Timers and constants

10.1 Inter-Fabric Router time-out values

Inter-Fabric Router time-out values are specified in table 30.

Table 30 – Inter-Fabric Router time-out values

Timer	Time-out value (seconds)
IFR_TOV	200
IF_ZONE_TOV	5
IFR_ILS_HLO_INT	5