

INTERNATIONAL
STANDARD

ISO/IEC
14165-141

First edition
2001-06

Information technology –
Fibre Channel –

Part 141:
Fabric Generic Requirements (FC-FG)

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14165-141:2001



Reference number
ISO/IEC 14165-141:2001(E)

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14165-141:2007

INTERNATIONAL STANDARD

ISO/IEC 14165-141

First edition
2001-06

**Information technology –
Fibre Channel –**

**Part 141:
Fabric Generic Requirements (FC-FG)**

© ISO/IEC 2001

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

ISO/IEC Copyright Office • Case postale 56 • CH-1211 Genève 20 • Switzerland



PRICE CODE **M**

For price, see current catalogue

Contents

	Page
FOREWORD	v
INTRODUCTION	vi
1 Scope	1
2 Normative references	1
3 Definitions and conventions	2
3.1 Definitions	2
3.2 Editorial conventions	5
3.3 Abbreviations and acronyms	5
4 Fabric concepts	6
4.1 Fabric and Fabric Elements	6
4.1.1 Sub-Fabric	6
4.1.2 Region	8
4.1.3 Translator	8
4.1.4 Extended Region	8
4.1.5 Zone	8
4.2 Typical Fabric topologies	8
4.2.1 Switch topology	9
4.2.1.1 Dedicated Connections	9
4.2.1.2 Connectionless service	9
4.2.1.3 Connection-oriented service	10
4.2.2 Distributed Fabric Element topology (DFE)	10
4.2.3 Other topologies	10
4.3 Fabric frame	10
4.4 Fabric_Ports	11
4.5 Fabric Service Parameters	11
4.6 Fabric addressing	11
4.6.1 Address identifiers	11
4.6.2 Address space partitioning	12
4.7 Fabric addressable service elements	12
4.7.1 Broadcast Alias_ID	12
4.7.2 Fabric F_Port/Login server	12
4.7.3 Fabric Controller	13
4.7.4 Directory server	13
4.7.5 Time server	13
4.7.6 Management server	13

	Page
4.7.7 Quality of Service Facilitator - Class 4 (QoSf)	13
4.7.8 Alias Server	13
5 Fabric entity requirements and characteristics	13
5.1 General requirements	13
5.2 Link_Control response	14
5.3 Frame validity checking	14
5.4 Connection independence	14
5.5 Class 1 bandwidth & frame jitter	14
5.6 Fabric Controller	14
5.7 Login Server	14
5.8 Service Parameter extent	14
5.9 E_D_TOV, R_A_TOV enforcement	14
5.10 Non-duplication of frames	15
5.11 Phase discontinuities	15
6 Fabric_Port requirements and characteristics	15
6.1 General requirements	15
6.2 Class 1 service - Dedicated Connection	15
6.3 Buffered Class 1 service	16
6.4 Dedicated Simplex service	16
6.5 Class 2 service - Multiplex	16
6.6 Class 3 service - Datagram	17
6.7 Class 4 service - Fractional	17
6.8 Intermix service	17
6.9 Class F service - Fabric signaling	17
6.9.1 Class F Frame formats	17
6.9.2 Class F function	17
6.9.3 Class F rules	18
6.9.4 Class F delimiters	19
6.9.4.1 Class F frame size	19
6.9.4.2 Class F flow control	19
6.9.5 Link Control	19
6.10 Fabric Login	19
7 Initialization and configuration control	20
7.1 Initialization	20

	Page
7.1.1 Power On	20
7.1.2 Link Initialization Protocol	20
7.1.3 Link Attachment Protocol	20
7.1.4 Addressing and Configuration Determination	21
7.1.5 F_Port Activation	21
7.1.6 N_Port Login with Fabric	21
7.2 Configuration Changes	21
8 Fabric inter-operation	22

Annex

A Address Space Partitioning	23
A.1 Address partitioning	23
A.1.1 Port Identifier partition	23
A.1.2 Fabric-Assisted functions	23
A.1.3 Vendor Unique partitions	23

Table

1 Well-known Address Identifiers	12
2 Address Partitioning	24

Figures

1 Document relationship	vii
2 Fabric model	6
3 Fabric with Sub-Fabric illustrations	7
4 Class 2 Sub-Fabrics and an extended region	8
5 Example of Switch topology	9
6 Example of Distributed Fabric Element topology	10
7 The Class F frame format	18

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 141: Fabric Generic Requirements (FC-FG)

FOREWORD

- 1) ISO (International Organization for Standardization) and IEC (International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.
- 2) In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.
- 3) Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any of all such patent rights.

International Standard ISO/IEC 14165-141 was prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

ISO/IEC 14165 consists of the following parts, under the general title *Information technology – Fibre Channel*:

- Part 111: *Physical and Signaling Interface (FC-PH)* ¹⁾
- Part 112: *Physical and Signaling Interface - 2 (FC-PH-2)* ¹⁾
- Part 113: *Physical and Signaling Interface - 3 (FC-PH-3)* ¹⁾
- Part 122: *Arbitrated Loop-2 (FC-AL-2)* ¹⁾
- Part 131: *Switch Fabric Requirements (FC-SW)*
- Part 141: *Fabric Generic Requirements (FC-FG)*
- Part 211: *Mapping to HIPPI-FP (FC-FP)*
- Part 222: *Mapping of Single-Byte Command Code Sets-2 (FC-SB-2)* ¹⁾
- Part 321: *Audio Video (FC-AV)* ¹⁾
- Part 412: *Generic Services (FC-GS-2)* ¹⁾

Annex A is for information only.

¹⁾ Under consideration.

INTRODUCTION

This part of ISO/IEC 14165 defines requirements for Fabrics supporting the Fibre Channel Physical and Signaling Interface (FC-PH) that are independent of specific Fabric topologies. FC-PH, together with second generation (FC-PH-2) and third generation (FC-PH-3), describes the point-to-point physical interface transmission protocol and signaling protocol of a high-performance serial link for support of higher level protocols.

The term Fibre Channel generally refers to: ISO/IEC 14165, part 111, FC-PH; part 112, FC-PH-2; and part 113, FC-PH-3. Fibre Channel provides a general data transport vehicle for Upper Level Protocols (ULPs) such as Intelligent Peripheral Interface (IPI) and Small Computer System Interface (SCSI) command sets, the High-Performance Parallel Interface (HIPPI) data framing, IP (Internet Protocol), ANSI/IEEE 802.2, and others. Proprietary and other command sets may also use and share Fibre Channel, but such use is not defined as part of Fibre Channel.

This part of ISO/IEC 14165 is organized as follows:

- a) FC-0 defines the physical portions of Fibre channel including the fibre, connectors, and optical and electrical parameters for a variety of data rates and physical media. Coax and twisted pair versions are defined for limited distance applications. FC-0 provides the point-to-point physical portion of Fibre Channel. A variety of physical media is supported to address variations in cable plants.
- b) FC-1 defines the transmission protocol which includes the serial encoding, decoding, and error control.
- c) FC-2 defines the signaling protocol which includes the frame structure and byte sequences.
- d) FC-3 defines a set of services which are common across multiple ports of a node.
- e) FC-4 is the highest level in Fibre Channel. It defines the mapping between the lower levels of Fibre Channel and the IPI and SCSI command sets, the HIPPI data framing, IP, and other Upper Level Protocols (ULPs).

Of these levels, FC-0, FC-1, and FC-2 are integrated into the ISO/IEC 14165-111, FC-PH document. The Fibre Channel protocol provides a range of implementation possibilities extending from minimum cost to maximum performance. The transmission medium is isolated from the control protocol so that each implementation may use a technology best suited to the environment of use.

Figure 1 shows the relationship of this part of ISO/IEC 14165 (the highlighted rectangle) with other Fibre Channel documents. Part 112, FC-PH-2 and part 113, FC-PH-3 specify enhanced functions added to part 111, FC-PH. Part 141, FC-FG and part 131, FC-SW are related to Fabric requirements. Part 121, FC-AL specifies the arbitrated loop topology. Part 412, FC-GS-2 is related to Generic fibre Channel Services. Part 221, FC-SB; part 211, FC-FP; part 311, FC-AE, part 321, FC-AV, IPI-3 Disk revision, IPI-3 Tape revision and SCSI-FCP are FC-4 standards. Parts 122, Arbitrated loop-2 (FC-AL-2), and 222, Mapping of single-byte command code sets-2 (FC-SB-2), are under consideration.

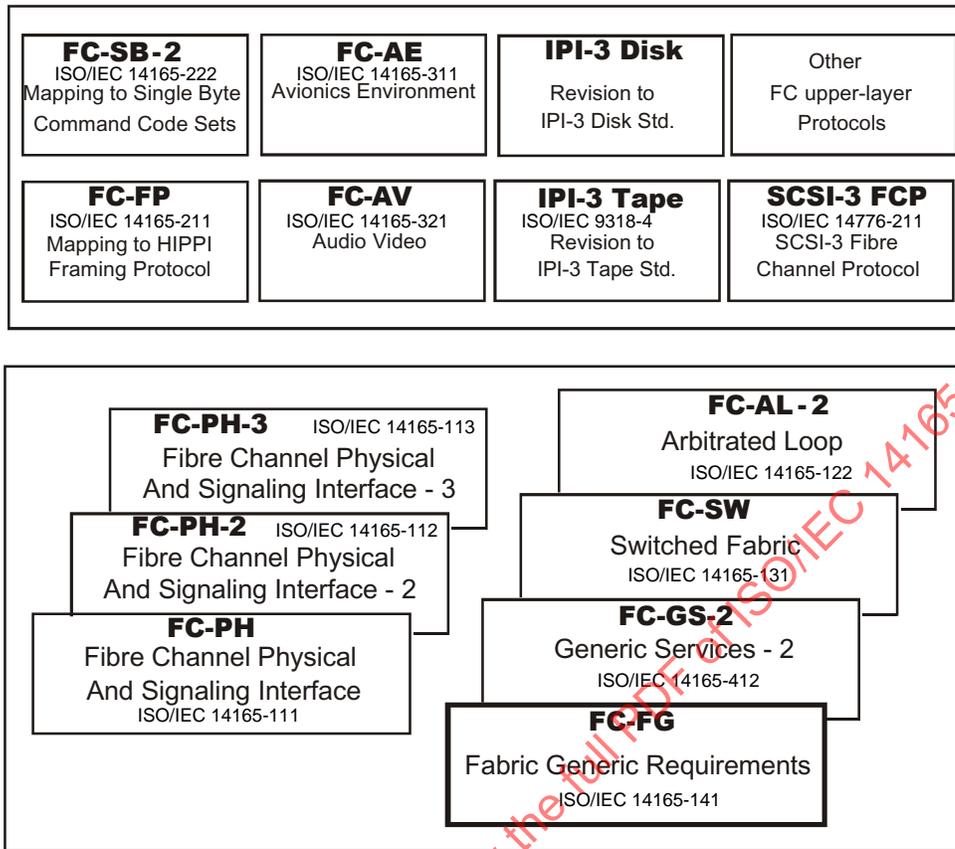


Figure 1 – Document relationship

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14165-141:2007

Information technology — Fibre Channel —

Part 141: Fabric Generic Requirements (FC-FG)

1 Scope

This part of ISO/IEC 14165 describes generic requirements for a communications transport medium called the Fabric, an entity that provides switched interconnect between pairs of user attachment points. Fabrics may be implemented using one or more topologies and this document describes requirements that are generic across all topologies.

A companion document, ISO/IEC 14165-111, *Fibre Channel - Physical and Signaling Interface (FC-PH)*, describes the physical interface, transmission protocol, and signaling protocol of high-performance serial links which attach user nodes to the Fabric. The Fabric serves to extend these serial links between pairs of attachment points. ISO/IEC 14165-111, FC-PH also describes features and behaviors of the Fabric required by user nodes.

2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 14165. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this part of ISO/IEC 14165 are encouraged to investigate the possibility of applying the most recent editions of the standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards (see foreword).

ISO/IEC 14165-111, *Information Technology - Fibre Channel - Physical and Signaling Interface (FC-PH)*

ISO/IEC 14165-112, *Information Technology - Fibre Channel - Physical and Signaling Interface-2 (FC-PH-2)*

ISO/IEC 14165-122, *Information Technology - Fibre Channel - Arbitrated Loop-2 (FC-AL-2)*

ISO/IEC 14165-412, *Information Technology - Fibre Channel - Fabric Generic Services-2 (FC-GS-2)*

3 Definitions and conventions

For the purpose of ISO/IEC 14165-141, FC-FG, the following definitions, conventions, abbreviations, and acronyms apply.

3.1 Definitions

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

3.1.2 alias address identifier (Alias): One or more address identifiers which may be recognized by an N_Port in addition to its N_Port Identifier. An alias address identifier is Fabric unique and may be common to multiple N_Ports.

3.1.3 Arbitrated Loop topology: A configuration that allows multiple ports to be connected serially (see ISO/IEC 14165-122, FC-AL-2).

3.1.4 Area: The second hierarchical level in the three-level addressing hierarchy.

3.1.5 bandwidth: Maximum effective transfer rate for a given set of physical variants such as communication model, Payload size, Fibre speed, and overhead specified by FC-PH (see 4.7 and annex M of ISO/IEC 14165-111, FC-PH).

3.1.6 Broadcast: A simplified Multicast service in which all available destinations are implicitly registered (see 3.1.5 of ISO/IEC 14165-112, FC-PH-2).

3.1.7 Broadcast Alias_ID: The address of the logical entity within the Fabric that provides a Broadcast service.

3.1.8 circuit: A bidirectional path within the Fabric.

3.1.9 Class 1 service: A service which establishes a Dedicated Connection between communicating N_Ports.

3.1.10 Class 2 service: A service which multiplexes frames at frame boundaries to or from one or more N_Ports with acknowledgment provided.

3.1.11 Class 3 service: A service which multiplexes frames at frame boundaries to or

from one or more N_Ports without acknowledgment.

3.1.12 Class 4 service: A service that establishes Virtual Connections to provide fractional bandwidth service between communicating N_Ports. The service multiplexes frames at frame boundaries to or from one or more N_Ports with acknowledgment provided.

3.1.13 Class F service: A service which multiplexes frames at frame boundaries that is used for control and coordination of the internal behavior of the Fabric.

3.1.14 Classes of service: Different types of services provided by the Fabric and used by the communicating N_Ports.

3.1.15 Connection Initiator: The source N_Port which initiates a Class 1 Connection with a destination N_Port through a connect-request and also receives a valid response from the destination N_Port to complete the Connection establishment.

3.1.16 Connection Recipient: The destination N_Port which receives a Class 1 connect-request from the Connection Initiator and accepts establishment of the Connection by transmitting a valid response.

3.1.17 Connectionless service: Communication between two N_Ports performed without a Dedicated Connection.

3.1.18 Dedicated Connection: A communicating circuit guaranteed and retained by the Fabric for two given N_Ports.

3.1.19 destination F_Port: The F_Port which is directly connected through a link to a destination N_Port.

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

3.1.21 destination N_Port: The N_Port to which a frame is targeted.

3.1.22 disconnection: The process of removing a Dedicated Connection between two N_Ports.

3.1.23 Distributed Fabric Element (DFE):

A Fabric Element with a 1-to-n tree-like topology.

3.1.24 Domain: The highest or most significant hierarchical level in the three-level addressing hierarchy.

3.1.25 E_Port: A Fabric Inter-Element Port used to establish Inter-Element Links (IEL).

3.1.26 Element Controller: A logical entity which is that portion of the Fabric Controller that manages a Fabric Element.

3.1.27 Element Name: A Name_Identifier associated with a Fabric Element.

3.1.28 Exclusive Connection: A Class 1 Dedicated Connection without Intermix.

3.1.29 Extended Region: A section of two or more Sub-Fabrics with compatible service parameters forming an extended communication group.

3.1.30 F_Port: The Link_Control_Facility within the Fabric which attaches to an N_Port through a link. An F_Port is addressable by the N_Port attached to it, with a common well-known address identifier (hex 'FFFFFFE').

3.1.31 F_Port Name: A Name_Identifier associated with an F_Port.

3.1.32 Fabric: The entity which interconnects various N_Ports attached to it and is capable of routing frames by using only the D_ID information in a FC-2 frame header.

3.1.33 Fabric Controller: The logical entity responsible for operation of the Fabric.

3.1.34 Fabric Element: A Fabric Element is the smallest unit of a Fabric which meets the definition of a Fabric. A Fabric may consist of one or more Fabric Elements, interconnected E_Port to E_Port, each with its own Fabric controller. To the attached N_Ports, a Fabric consisting of multiple Fabric Elements is indistinguishable from a Fabric consisting of a single Fabric Element.

3.1.35 Fabric frame: A frame whose first delimiter is the Start_of_Frame Fabric (**SOFF**) (see 17.2.4 of ISO/IEC 14165-111, FC-PH).

3.1.36 Fabric Name: A Name_Identifier associated with a Fabric.

3.1.37 Fabric_Port: A generic reference to an E_Port, F_Port, FL_Port, G_Port, or GL_Port.

3.1.38 Fabric topology: The physical configuration of a Fabric.

3.1.39 FL_Port: An F_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology.

3.1.40 Frame delimiter: An Ordered Set used to indicate a frame boundary (see 11.4 of ISO/IEC 14165-111, FC-PH).

3.1.41 G_Port: A generic Fabric_Port that can function either as an E_Port or as an F_Port.

3.1.42 GL_Port: A generic Fabric_Port that can function either as an E_Port or as an FL_Port.

3.1.43 Identified N_Port: An N_Port which has been assigned an N_Port identifier by the initialization procedure.

3.1.44 Inter-Element Link (IEL): A link connecting the E_Port of one Fabric Element to the E_Port of another Fabric Element.

3.1.45 Interject: A service that interleaves Class F frames on an established Class 1 Connection within the Fabric.

3.1.46 Intermix: A service that interleaves Class 2 and Class 3 frames on an established Class 1 Connection.

3.1.47 L_Port: A generic reference to an FL_Port, a GL_Port, or an NL_Port.

3.1.48 link: Two unidirectional fibers transmitting in opposite directions and their associated transmitters and receivers.

3.1.49 Link_Control_Facility: A link hardware facility which attaches to an end of a link and manages transmission and reception of data. It is contained within each Port type.

3.1.50 local F_Port: The F_Port to which an N_Port is directly attached by a link (see remote F_Port).

3.1.51 Logical node: A conceptual N_Port whose implementation is unspecified.

3.1.52 Login: The generic name for a procedure that exchanges link level or end-to-end service parameters.

3.1.53 Login server: The logical entity within the Fabric that processes and responds to Fabric Login frames.

3.1.54 N_Port: A hardware entity which includes a Link_Control_Facility. It may act as an Originator, a Responder, or both.

3.1.55 N_Port Identifier: A Fabric unique address identifier by which an N_Port is uniquely known. The identifier may be assigned by the Fabric during the initialization procedure. The identifier may also be assigned by other procedures not defined in FC-PH. The identifier is used in the S_ID and D_ID fields of a frame.

3.1.56 Name Identifier: A 64 bit identifier, with a 60 bit value preceded with a four bit Network_Address_Authority_Identifier, used to identify physical entities in Fibre Channel such as N_Port, Node, F_Port, or Fabric.

3.1.57 Network Address Authority (NAA): An organization such as ITU or IEEE which administers network addresses.

3.1.58 Network Address Authority Identifier: A four bit identifier defined in FC-PH to indicate a Network_Address_Authority (NAA).

3.1.59 NL_Port: An N_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology.

3.1.60 Node: A collection of one or more N_Ports controlled by a level above FC-2.

3.1.61 Originator: The logical function associated with an N_Port responsible for originating an Exchange.

3.1.62 Payload: Contents of the Data Field of a frame, excluding Optional Headers and fill bytes, if present.

3.1.63 Port: A generic reference to an E_Port, F_Port, G_Port, N_Port or S_Port.

3.1.64 Port Name: A Name_Identifier associated with a Port.

3.1.65 Primitive signal: An ordered set designated to have special meaning (see 16.3 of ISO/IEC 14165-111, FC-PH).

3.1.66 Quality of Service Facilitator (QoSF): A logical entity within the Fabric that allocates Fabric resources where Class 4 operation is available.

3.1.67 Region: A section of a Sub-Fabric with compatible service parameters in which all ports can communicate.

3.1.68 remote F_Port: The F_Port to which the other communicating N_Port is directly attached (see local F_Port).

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

3.1.70 S_Port: A Fabric internal service node that functions both as a Fabric_Port and as an N_Port.

3.1.71 Service parameters: A list of capabilities supported by a Fabric where such capabilities are variable or optional (see clause 23 of ISO/IEC 14165-111, FC-PH).

3.1.72 source F_Port: The F_Port which is directly connected through a link to a source N_Port.

3.1.73 Source Identifier (S_ID): The address identifier used to indicate the source Port of the transmitted frame.

3.1.74 source N_Port: The N_Port from which a frame is transmitted.

3.1.75 Sub-Fabric: The set of ports and services in a Fabric uniquely identified by one data rate and one Class of service.

3.1.76 Switch topology: The physical configuration of a Fabric whose Fabric Elements are switches.

3.1.77 Translator: An agent within a Fabric that performs conversions for data rate and Class of service.

3.1.78 Unidentified N_Port: An N_Port which has not yet had its N_Port identifier assigned by the initialization procedure.

3.1.79 Virtual Connection (VC): A unidirectional path between two communicating N_Ports that permits fractional bandwidth services to be used. Two Virtual Connections are required to form a Class 4 connection.

3.1.80 well-known addresses: A set of address identifiers defined in FC-PH to access global server functions such as a Directory server.

3.1.81 Zone: A non-exclusive administrative partition of a Region or an Extended Region.

3.2 Editorial conventions

In this part of ISO/IEC 14165, 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., Fabric, Class). Any lowercase uses of these words have the normal technical English meanings.

Numbered items in this part of ISO/IEC 14165 do not represent any priority. Priority is explicitly indicated.

In case of any conflict between figure, table, and text, the text takes 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.

The term "shall" is used to indicate a mandatory rule. If such a rule is not followed, the results are unpredictable unless indicated otherwise.

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

3.3 Abbreviations and acronyms

Abbreviations and acronyms applicable to this part of ISO/IEC 14165 are listed. Definitions of several of these items are included in 3.1 "Definitions".

ACC	Accept
ACK	Acknowledgment

Alias	alias address identifier
Credit_CNT	Credit count
D_ID	Destination_identifier
E_D_TOV	Error Detect_Timeout value
E_Port	Fabric Expansion port
F_BSY	F_Port busy
F_Port	Fabric N_Port attachment port
F_RJT	F_Port reject
FL_Port	Fabric Loop Port
G_Port	Generic E_Port/F_Port
GL_Port	Generic Fabric/Loop Port
IEL	Inter-Element Link
IELOGI	Inter-Element Login
L_Port	Generic Loop Port
LA_RJT	Link Application Reject
N_Port	Node Port
NL_Port	Node Loop Port
OLS	Offline Primitive Sequence
P_RJT	N_Port reject
QoS	Quality of Service Facilitator - Class 4
R_A_TOV	Resource_Allocation_Timeout value
RJT	generic F_RJT, P_RJT
S_ID	Source_Identifier
S_Port	Fabric service node
VC	Virtual Connection

4 Fabric concepts

4.1 Fabric and Fabric Elements

The Fabric is a data transport medium that provides switched interconnect between multiple link attachment points called N_Ports. The extent of the Fabric is limited to those Ports that can be addressed by unique values of the 24-bit Port Identifier.

NOTE – In the Fibre Channel context, fabric written with a lower-case 'f' embraces the interconnect of any ports within the 24-bit address space; Fabric written with a capital 'F' describes topologies distinct from Point-to-point topology and Arbitrated Loop topology (see 4.8 of ISO/IEC 14165-111, FC-PH). This document describes the Fabric of FC-PH.

A Fabric may be composed of one or more Fabric Elements as illustrated in figure 2. The link attachment point between the Fabric Element and an N_Port is called the F_Port. The link attachment point between one Fabric Element and another Fabric Element is called the E_Port. The link between Fabric Elements is called an Inter-Element Link (IEL).

The Fabric has characteristics defined in terms of the transport services provided on a bidirectional link between the Link Control Facility within a node N_Port and the Link Control Facility

within the Fabric F_Port. The Fabric provides transport services by routing frames between F_Ports.

Fabric transport services have been divided into multiple Classes of service. These Classes of service are distinguished primarily by whether or not a communication circuit is allocated and retained between the communicating N_Ports, and by the level of delivery integrity provided. Users of a Fabric discover the capabilities available within a particular Fabric through a Login procedure.

This document, ISO/IEC 14165-141, FC-FG, describes more features than any one Fabric or Fabric Element is required to implement. It is expected that Fabric Elements will be optimized for selected environments of use. ISO/IEC 14165-111, FC-PH, defines several implementations of the physical interface to the Fabric. It defines a single logical interface to the Fabric. Neither document intends to prescribe or limit the design of the internal workings of a Fabric.

4.1.1 Sub-Fabric

Within a Fabric, capabilities for multiple data rates and multiple Classes of service may coexist. This typically occurs where unlike Fabric Elements are connected by Inter-Element Links

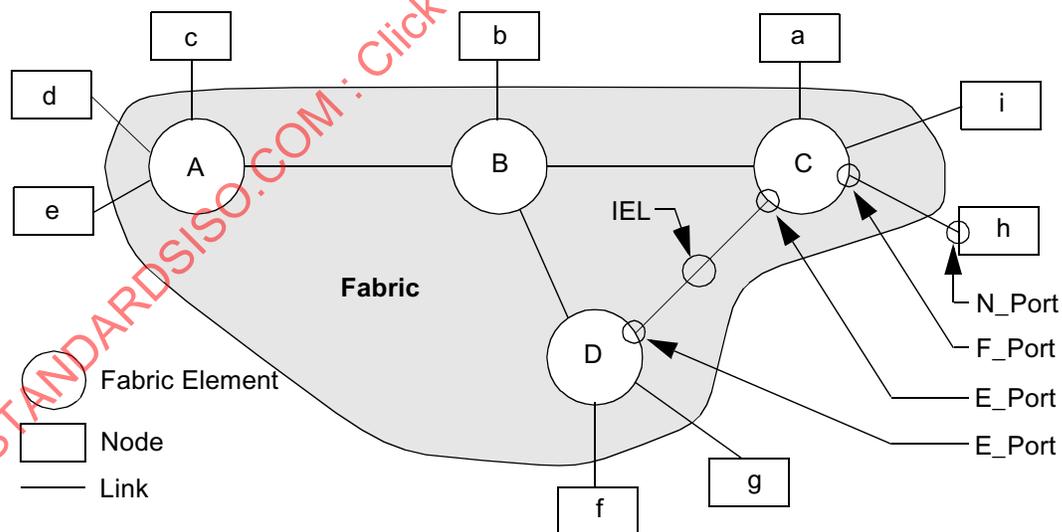


Figure 2 – Fabric model

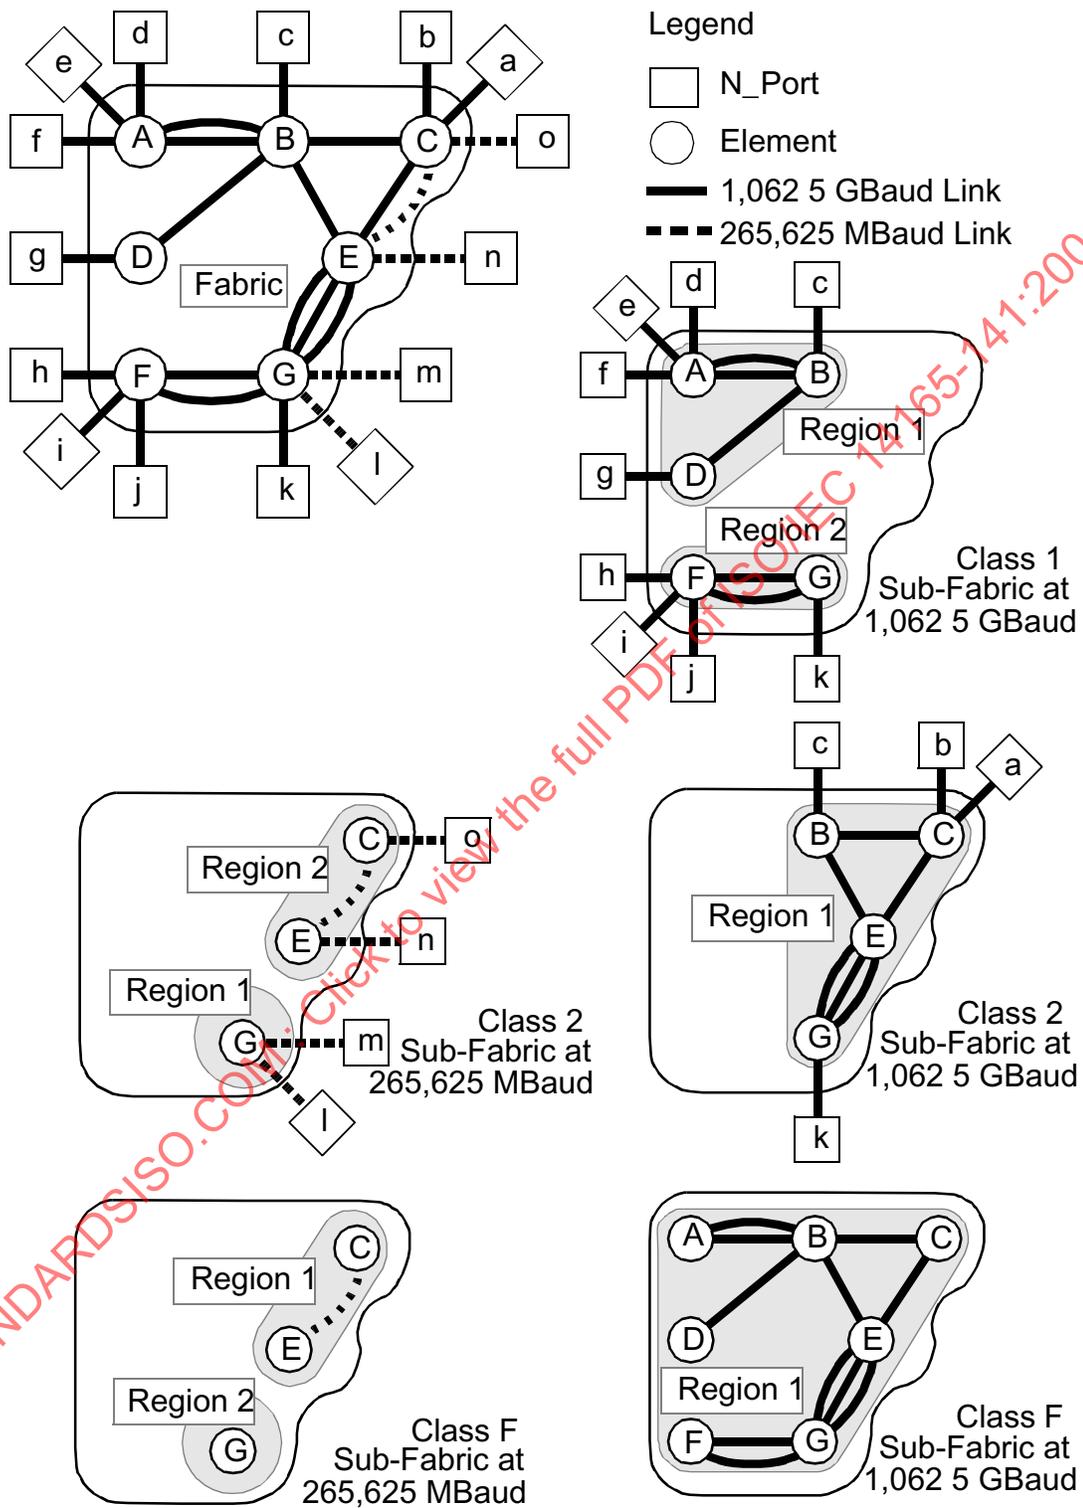


Figure 3 – Fabric with Sub-Fabric illustrations

(see figure 3). Each possible combination of data rate and Class of service that is available within the Fabric defines a logically independent section of the Fabric called a Sub-Fabric.

4.1.2 Region

A Sub-Fabric may be divided into sections that have incompatible variations in service parameters which prevent ports in one section from communicating directly with ports in another section. Each section with compatible service parameters within a Sub-Fabric is a Region (see figures 3 and 4).

4.1.3 Translator

A Fabric may provide a service that permits Sub-Fabrics or Regions to communicate across data rate and Class of service boundaries. The logical entity providing this service is a Translator.

4.1.4 Extended Region

When a single Region in one Sub-Fabric provides access to more than one Region in a second Sub-Fabric, the Fabric may provide a Translator to transparently bridge the disjoint Regions in the second Sub-Fabric using links of the first Sub-Fabric. The bridged Regions of the second Sub-Fabric are said to form an Extended Region (see figure 4).

4.1.5 Zone

For many environments, it is not appropriate for certain Nodes to communicate with other Nodes. There may exist system and application boundaries that must be guarded. The Fabric may support subdividing Regions and Extended Regions into independent partitions for purely administrative purposes. These administered partitions are called Zones.

4.2 Typical Fabric topologies

There are no specific requirements for Fabric topology. Topologies favored by early implementors are illustrated in this clause.

There is also no requirement that a Fabric be implemented with a single, homogeneous topology. It is anticipated, for example, that the Arbitrated Loop topology may attach to the Switch topology.

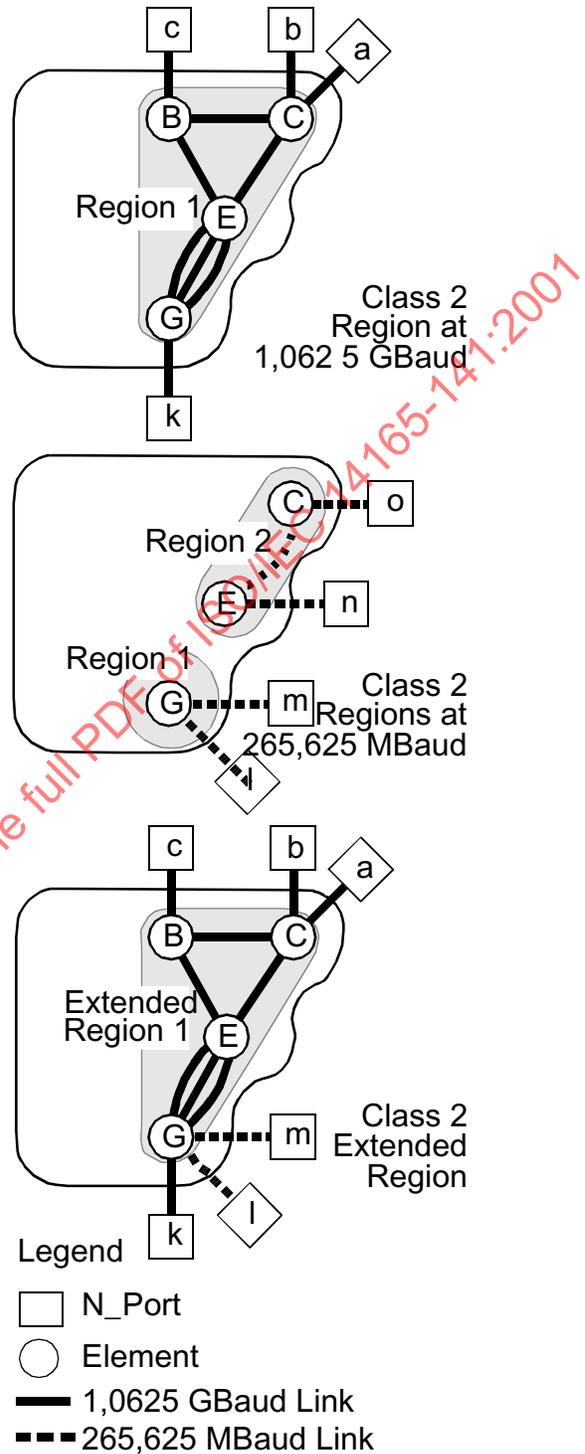


Figure 4 – Class 2 Sub-Fabrics and an extended region

4.2.1 Switch topology

The Switch topology may consist of one or more Fabric Elements as shown in figure 5. Each Fabric Element is a Switch Element.

A Switch topology provides connections between pairs of F_Ports, and thereby between pairs of N_Ports. The connections may be Cir-

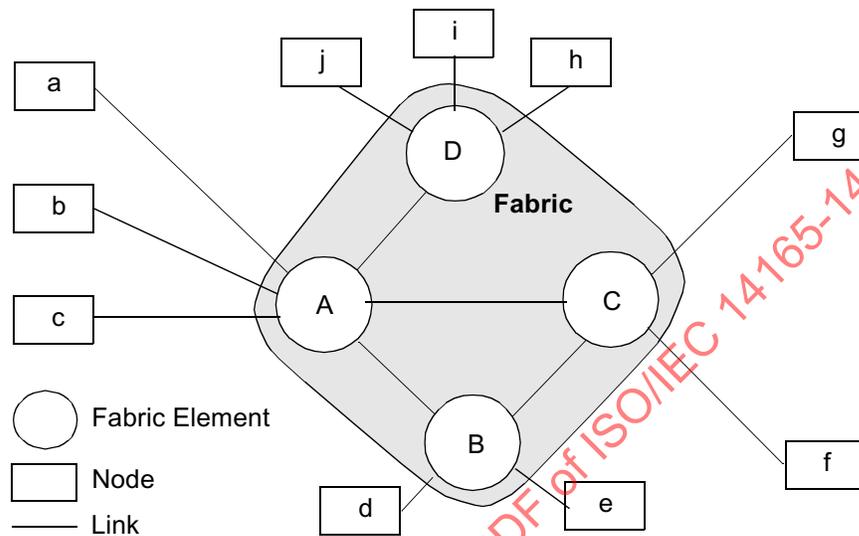


Figure 5 – Example of Switch topology

cuit Switched (Dedicated Connections), or Frame Switched (Connection-oriented and Connectionless services).

The Class of service required by the N_Port and provided by the Fabric is determined by the choice of frame delimiters (see 17.2.2 of ISO/IEC 14165-111, FC-PH, and clause 18 of ISO/IEC 14165-112, FC-PH-2).

4.2.1.1 Dedicated Connections

A Switch Fabric provides the capability to interconnect the links that are attached to it. Multiple Dedicated Connections may exist simultaneously within the Fabric. The interconnection of F_Ports established by the Fabric does not affect the existing interconnection of any other F_Ports, nor does it affect the ability of the Fabric to remove those connections, nor does it affect the ability of any other F_Ports to handle connectionless operations. The maximum number of Dedicated Connections supported is equal to the integer value of one-half the number of attached N_Ports. Fabrics subject to blocking may support less than the maximum number of connections.

When a Dedicated Connection is established, F_Ports and their respective point-to-point links are interconnected within the Fabric, so that the links appear as one continuous link for the duration of the connection. When frames are received by one of the connected F_Ports, the frames are normally passed from one F_Port to the other or others, through the Fabric, for transmission. A Dedicated Connection is retained until a removal request is received.

Class 1, Buffered Class 1, and Dedicated Simplex are services based on Dedicated Connections.

4.2.1.2 Connectionless service

Connectionless service is characterized by the absence of a Dedicated Connection. Frames are multiplexed, on frame boundaries, between an F_Port and any other F_Port and thereby between the N_Ports attached to them.

Consecutive frames received from an attached N_Port by a source F_Port may be directed to the same or different destination N_Ports. Likewise, frames received by a destination F_Port for its attached N_Port may be received from

the same or different source N_Ports. The Fabric treats each frame individually without reference to the frame that precedes it or follows it through the Fabric. Once a frame has been delivered to the destination N_Port by the Fabric, the Fabric may have no memory of the routing taken through the Fabric by the frame. Subsequent frames delivered from the same source N_Port to the same destination N_Port may be routed differently through the Fabric. Some Fabrics ensure that frames sent from a source N_Port to a destination N_Port arrive at the destination in the same order as sent by the source, and some do not.

With Connectionless service, congestion may occur within the Fabric. The Fabric manages this congestion through buffer-to-buffer flow control.

Class 2 and Class 3 are services based on Connectionless service.

4.2.1.3 Connection-oriented service

Connection-oriented service is a virtual connection service that is commonly implemented upon a Connectionless service. The Connection-oriented circuit may provide a fractional allocation of the resources available on the path between connected N_Ports.

Class 4 - Fractional is a Connection-oriented service.

4.2.2 Distributed Fabric Element topology (DFE)

The Distributed Fabric Element is a 1-to-n switch element that may be combined with like Distributed Fabric Elements to form a tree-like topology. An example of the topology is given as figure 6.

4.2.3 Other topologies

Additional Fabric topologies may be defined and implemented. It is not the intent of this document to restrict topologies and implementations to those previously defined.

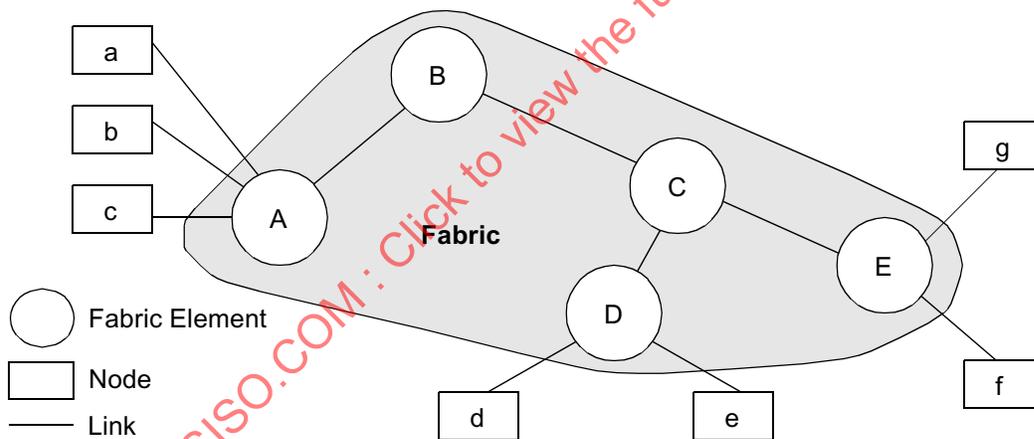


Figure 6 – Example of Distributed Fabric Element topology

4.3 Fabric frame

A Fabric frame is distinguished by a unique Start of Frame Delimiter, the Start_of_Frame Fabric (**SOFF**). It is used for communication within the Fabric itself. Fabric frames are not expected to appear outside of the Fabric, and if they do, they are discarded and ignored.

ISO/IEC 14165-111, FC-PH, requires that N_Ports shall transmit a minimum of six Primitive Signals following each frame, but that receivers need only see two primitive signals ahead of a frame. The difference of four Primitive Signals between the transmitted data stream and the received data stream provides the Fabric with reserved bandwidth that may be used for clock

skew management and for the delivery of Fabric frames.

4.4 Fabric_Ports

Fabric_Ports may connect through a link to an N_Port or through a link to the Fabric_Port of another Fabric Element. Six possible configurations of a Fabric_Port are identified:

1. The Fabric_Port that connects through a link to an N_Port is an F_Port that behaves as prescribed in ISO/IEC 14165-111, FC-PH.
2. The Fabric_Port that connects through a link to another Fabric_Port is called an E_Port, an inter-element expansion port. Provision for one or more E_Ports is optional. An E_Port may use FC-PH compliant media, interface, and signaling protocols, or it may use other interconnect not specified by this document.
3. The Fabric may optionally provide a G_Port, a generic Fabric_Port, which is capable of behaving either as an E_Port or as an F_Port. The G_Port determines through Login at initialization whether it has been configured as an E_Port or as an F_Port and thereafter operates as required by the port configuration.
4. The Fabric_Port that may connect either to an N_Port or to an Arbitrated Loop is an FL_Port.
5. The Fabric_Port that may connect either to an N_Port, to an E_Port, or to an Arbitrated Loop is a GL_Port.
6. A logical node within the Fabric, capable of communicating either with other Fabric_Ports or with N_Ports is an S_Port.

4.5 Fabric Service Parameters

During Fabric Login, the F_Port delivers to its associated N_Port a single set of Service Parameters common to all the destinations accessible to that N_Port. Fabrics may or may not be homogeneous. They may be composed of Fabric Elements supporting different Classes of service, different media rates, different data field sizes, different transit delays, and so on. The Fabric has to determine a common set of

Service Parameters available to the N_Port by methods described in the topology specific references.

Changes within the Fabric that affect Service Parameters may require that the N_Port re-Login. The Fabric initiates this process by issuing a Link Reset (see 16.4.4 of ISO/IEC 14165-111, FC-PH).

4.6 Fabric addressing

4.6.1 Address identifiers

Address identifiers are three bytes in length. The Frame Header contains two three-byte fields for address identifiers, the Destination Identification (D_ID) field and the Source Identification (S_ID) field.

Each N_Port has a Fabric unique identifier, the N_Port Identifier, by which it is known. An N_Port may have one or more alias address identifiers as well. The source and destination N_Port Identifiers and alias address identifiers are used to route frames within the Fabric.

The F_Port which is directly connected through a Link to an N_Port has the reserved address identifier 'FFFFFF'. This reserved address identifier is used primarily as the Destination Identifier (D_ID) for Fabric Login.

Certain other address identifiers have been designated well-known addresses (see table 33 of ISO/IEC 14165-111, FC-PH and table 33 of ISO/IEC 14165-112, FC-PH-2) as summarized in table 1.

Table 1 – Well-known Address Identifiers

Hexadecimal value	Description
FFFFF0 to FFFFF7	Reserved
FFFFF8	Alias Server
FFFFF9	Quality of Service Facilitator - Class 4 (QoS4)
FFFFFA	Management Server
FFFFFB	Time Server
FFFFFC	Directory Server
FFFFFD	Fabric Controller
FFFFFE	Fabric F_Port
FFFFFF	Broadcast Alias_ID

The Fabric may or may not provide address assignment. Where the Fabric does provide address assignment, the N_Ports take their address identifiers from the Fabric. The procedures for acquiring address identifiers are described in clause 23 of ISO/IEC 14165-111, FC-PH.

Where the Fabric does provide address assignment, the Fabric first assigns unique address identifiers to each of its F_Ports as part of the initialization procedure (see clause 7). The N_Ports attached to each of the F_Ports then inherit the address identifiers of their associated F_Ports. The F_Port entity is not separately addressable from the N_Ports except that the N_Port can address the F_Port to which it is linked using the reserved F_Port address.

4.6.2 Address space partitioning

All addresses not otherwise reserved or assigned are available for use as N_Port Identifiers or alias address identifiers. While partitioning of the unused address identifiers is not required by this document, it is encouraged as a framework for inter-operation and for anticipated Fibre Channel enhancements. Annex A of this standard describes the recommended address space partitioning.

Within the recommended address space partitioning is a hierarchical division of the address space reserved for Port Identifiers into logical or administrative entities described as Domain, Area, and Port. The suggestion is that, for addressing purposes, Ports are members of Areas, Areas are members of Domains, and Domains are partitions of the Fabric.

4.7 Fabric addressable service elements

The Fabric provides an address mapping between certain well-known addresses (see 4.6.1) and the service elements, or servers, that respond to those addresses. These Fabric addressable service elements have the appearance of Nodes attached to the Fabric, even though they may be an integral part of the Fabric, and are assigned N_Port Identifiers.

Servers appear as single entities to the external N_Ports, regardless of the actual implementation within the Fabric.

4.7.1 Broadcast Alias_ID

The Broadcast Alias_ID is the address of the logical entity within the Fabric that provides a Broadcast service (see 31.5 of ISO/IEC 14165-112, FC-PH-2). Class 3 frames delivered to the well-known address hexadecimal value 'FFFFFF', the Broadcast Alias_ID, are replicated and delivered by the Fabric to all N_Ports able, at the time of delivery, to receive such frames.

4.7.2 Fabric F_Port/Login server

The F_Port is the Link_Control_Facility within the Fabric that attaches to an N_Port through a link. It is addressed by the well-known address identifier hexadecimal value 'FFFFFE'.

The F_Port provides access to the Fabric for Fabric Login. The Login procedure is used by N_Ports and by other Fabric Elements to discover the operating characteristics associated with a Fabric or Fabric Element. The Login server is the logical entity within the Fabric that receives and responds to Fabric Login frames. The Login server also assigns, confirms or re-assigns the N_Port Identifier of the N_Port that initiates the Login. All Fabrics provide a Login server function.

When the F_Port receives a Fabric Login (FLOGI) frame with a D_ID of hexadecimal value 'FFFFFF' it routes the frame to the Fabric Login

server. The Login server provides appropriate responses as described in clauses 21 and 23 of ISO/IEC 14165-111, FC-PH.

4.7.3 Fabric Controller

Every Fabric contains one or more Fabric Controllers. The Fabric Controller is the logical entity responsible for general operation of the Fabric. Fabric control may include such functions as:

- Execution of the Fabric initialization procedure (see clause 7).
- Parsing and routing of frames directed to well-known addresses.
- Setup and tear down of Dedicated Connections.
- General frame routing.
- Generation of F_BSY and F_RJT link responses.

In addition, the Fabric Controller responds to the Read Connection Status Link Service request Sequence as prescribed by 21.4.9 of ISO/IEC 14165-111, FC-PH.

Fabric control may be distributed among Fabric Elements, in which case, the logical entity controlling the Fabric Element is called the Element Controller.

The Fabric Controller has the well-known address identifier hexadecimal value 'FFFFFFD'.

4.7.4 Directory server

At a minimum, the Directory server maintains tables that correlate N_Port Name_Identifiers with N_Port address identifiers. These tables may be used by Nodes to discover one form of identification given knowledge of the other form of identification (see ISO/IEC 14165-412, FC-GS-2).

The Directory server may also provide Directory Services which catalog a much broader list of port characteristics.

The Directory server is optional. When provided, the Directory server has the well-known address identifier hexadecimal value 'FFFFFFC'.

4.7.5 Time server

The Time server is optional, and where provided gives the time values required to manage expiration timers (see clause 11 of ISO/IEC 14165-412, FC-GS-2).

The Time server has the well-known address identifier hexadecimal value 'FFFFFFB'.

4.7.6 Management server

The Management server is optional, and where provided is used to collect and report information on link usage, errors, link quality and the like.

The Management server has the well-known address identifier hexadecimal value 'FFFFFFA'.

4.7.7 Quality of Service Facilitator - Class 4 (QoS)

The Quality of Service Facilitator (QoS) is a function provided within Fabrics that offer Class 4 service (see clause 34 of ISO/IEC 14165-112, FC-PH-2).

The Quality of Service Facilitator has the well-known address identifier hexadecimal value 'FFFFFF9'.

4.7.8 Alias Server

The Alias server is optional, and where provided may be used to issue Hunt Group Identifiers, Multicast Group Identifiers (see clause 32 of ISO/IEC 14165-112, FC-PH-2), and such other Identifiers as may be defined within the Fibre Channel document set (see figure 1).

The Alias server has the well-known address identifier hexadecimal value 'FFFFFF8'.

5 Fabric entity requirements and characteristics

5.1 General requirements

The internal design and behavior of the Fabric is largely unspecified. Equivalent internal mechanisms and/or functions are allowed by ISO/IEC 14165-111, FC-PH, and this document, ISO/IEC 14165-141, FC-FG, so long as the Fabric and each Fabric_Port behaves according to the standards.

5.2 Link_Control response

If a frame is accepted by an F_Port and is found subsequently to be undeliverable to a destination N_Port, the Fabric is required in Class 1 and Class 2 service to issue an appropriate Busy (F_BSY) or Reject (F_RJT) response with a valid reason code. In Class 1 operation, the Fabric shall only issue the Busy or Reject for frames delimited by **SOFC1**.

5.3 Frame validity checking

The Fabric may or may not verify the validity of a frame as the frame passes through the Fabric. The Fabric is required only to parse those fields required for routing. If the Fabric does check frame validity within the meaning of clause 17 of ISO/IEC 14165-111, FC-PH, and if an error is detected within the frame, the frame may be forwarded with a modified **EOF** delimiter to indicate that an error has been detected.

The Fabric shall not do anything which will cause the indication of a transmission error in a received frame to be lost, nor do anything which permits part of a path of a frame through the Fabric to be unprotected by an error check. This may be accomplished, among other ways, by passing the CRC contained in the frame from the source N_Port through to the destination N_Port unmodified, and by passing invalid transmission codes through to the destination N_Port.

5.4 Connection independence

The interconnection of F_Ports established by the Fabric shall not affect the existing interconnection of any other F_Ports, nor shall it affect the ability of the Fabric to remove those connections, nor shall it affect the ability of any other F_Ports to handle connectionless operations.

5.5 Class 1 bandwidth & frame jitter

FC-PH requires (see 17.1 of ISO/IEC 14165-111, FC-PH) that at an N_Port transmitter there shall be a minimum of six Primitive Signals (Idles and R_RDY) between frames and that a minimum of two primitive signals shall be guaranteed to precede the start (SOF) of each frame received by a destination N_Port. The surplus of Idles on N_Port transmit is provided to assist the Fabric in clock skew management

and to provide a bandwidth allocation for Fabric frames (delimited by SOFf) and for Intermix.

At the destination F_Port, the Fabric may delay a Class 1 frame, by inserting an Intermix frame, no more than one maximum frame time. Once a frame is inserted, the Fabric shall wait before issuing a subsequent Intermix frame until the number of surplus Primitive Signals received from the source N_Port matches the size of the frame previously transmitted as Intermix.

5.6 Fabric Controller

The Fabric shall contain one or more Fabric controllers. The Fabric Controller is a logical entity at the well-known address hexadecimal value 'FFFFFFD'.

The Fabric Controller shall respond to the Read Connection Status Link Service request sequence as required by 21.4.9 of ISO/IEC 14165-111, FC-PH.

5.7 Login Server

The Fabric shall provide a Login Server function. The Login Server is a logical entity at the well-known address hexadecimal value 'FFFFFFE'. The Login Server provides appropriate responses to requests for Fabric Login originating at attached N_Ports (see clause 23 of ISO/IEC 14165-111, FC-PH), or NL_Ports (see ISO/IEC 14165-122, FC-AL-2).

5.8 Service Parameter extent

A Fabric may be composed of multiple Fabric Elements with different service characteristics (e.g., differing support for Class of Service, different data rates, etc.) During Fabric Login, the Fabric delivers to the source N_Port a single set of Service Parameters which expresses the most limiting characteristics of all destination F_Ports accessible by that source N_Port. The method used internally by the Fabric to discover a common set of Service Parameters is not defined by this document.

5.9 E_D_TOV, R_A_TOV enforcement

The Error_Detect_Timeout Value (E_D_TOV) and the Resource_Allocation_Timeout Value (R_A_TOV) are specified in clause 29 of ISO/IEC 14165-111, FC-PH. When an N_Port performs Fabric Login, the Common Service Parameters provided by the F_Port specify values for E_D_TOV and R_A_TOV. The value of

R_A_TOV shall be greater than or equal to E_D_TOV plus twice the maximum time that a frame may be delayed within a Fabric and still be delivered. The Fabric shall guarantee that a frame received by the Fabric from a source N_Port shall be delivered to a destination N_Port within the R_A_TOV, or that the frame shall never be delivered.

The Fabric shall ensure delivery within the maximum delivery time by requiring each Fabric Element to time out frames stored in receive buffers within the Fabric. Individual Elements may use different timeout values. The maximum Fabric delivery time is the cumulative timeout value for elements along the path or paths joining the source and destination N_Ports.

When the Fabric encounters a Data frame requiring an F_BSY or F_RJT Link_Control frame response from the Fabric, the Fabric shall generate and transmit the F_BSY or F_RJT at a time not later than E_D_TOV after receipt of the Data frame. No Link_Control frame shall be returned in response to the Data frame, if the facility within the Fabric generating the F_BSY or F_RJT Link_Control frame is unable to transmit the Link_Control frame in time.

5.10 Non-duplication of frames

As the Fabric receives frames from a source N_Port, the Fabric shall do one of three things with those frames attempting (a) before (b) and (b) before (c):

- a) Deliver one and only one copy of a frame to a destination N_Port within R_A_TOV.
- b) Discard the frame and, where indicated by Class of service, provide an appropriate F_BSY or F_RJT response within E_D_TOV.
- c) Discard the frame without notice.

5.11 Phase discontinuities

The Fabric is permitted to switch the Fibre Channel serial bit stream and while switching the serial bit stream may produce a phase discontinuity (see 5.3 of ISO/IEC 14165-111, FC-PH). Such phase discontinuities shall not occur within a frame between the Start_of_Frame and End_of_Frame delimiters.

NOTE – Implementations producing such a phase discontinuity are discouraged.

6 Fabric_Port requirements and characteristics

6.1 General requirements

Fabric_Port is a generic name for F_Ports, FL_Ports, E_Ports, G_Ports, and GL_Ports. An F_Port is the Link_Control_Facility within a Fabric that attaches to an N_port. An E_Port is the Link_Control_Facility within a Fabric Element that attaches to another Fabric Element. A G_Port is capable of operating as either an E_Port or an F_Port depending on whether a Fabric_Port or an N_Port is discovered at the other end of the Link during Link initialization. FL_Ports and GL_Ports are F_Ports and G_Ports, respectively, capable of Loop attachment.

An F_Port is a source or destination point of the Fabric. An E_Port is an entry point to a Fabric Element or an exit point from a Fabric Element where two or more Fabric Elements are combined to create a Fabric. A Fabric may use one or more Fabric Elements to achieve its function.

The requirements for F_Port behavior are described throughout ISO/IEC 14165-111, FC-PH. In the event of a conflict between the requirements of this document and the requirements of ISO/IEC 14165-111, FC-PH, the FC-PH standard takes precedence.

An F_Port is not an N_Port addressable entity, separate from the N_Port to which it attaches. The Source_ID (S_ID) and Destination_ID (D_ID) fields within the frame header identify N_Ports. Certain features within the Fabric, such as the Fabric Controller, have been assigned well-known N_Port Identifiers.

F_Ports may, however, be addressed from within the Fabric by entities generating Fabric frames. A Fabric frame addressed to an N_Port Identifier shall be delivered to the F_Port associated with that N_Port.

An F_Port shall support one or more Classes of service. The services provided by an F_Port are discovered during Fabric Login.

6.2 Class 1 service - Dedicated Connection

A Class 1 Dedicated Connection normally guarantees simultaneous, bidirectional flow. However, ISO/IEC 14165-111, FC-PH allows an optional behavior, called Unidirectional Transmit, in which one data path provides full-bandwidth capability and the other data path is used

only for the return of frame acknowledgments (ACKs).

A Dedicated Connection guarantees full-bandwidth transfers and delivery of frames at the destination F_Port in the order of receipt at the source F_Port. The Fabric guarantees, within the error rate, that all frames received from the source N_Port shall be delivered to the destination N_Port.

If two N_Ports are in a Class 1 Dedicated Connection, a Class 1 request from another source to either of the connected N_Ports is handled according to optional behaviors negotiated during Fabric Login (see for example “Camp-On”, clause 35 of ISO/IEC 14165-112, FC-PH-2 and “Stacked Connect Request”, clause 36 of ISO/IEC 14165-112, FC-PH-2).

If the Fabric is unable to establish a Dedicated Connection, it shall return a Busy or a Reject indication with a valid reason code.

The Fabric shall provide buffer-to-buffer flow control for frames delimited by the Start_of_Frame Connect Class 1 (**SOFC1**).

While a Dedicated Connection exists, the Fabric is not involved in buffer-to-buffer flow control. The Fabric may, however, utilize speed matching buffers and may interleave Fabric frames and/or intermixed Class 2 and Class 3 frames with the Class 1 data flow, so long as the Class 1 bandwidth requirement is not violated.

The entry and exit F_Ports, working in conjunction with the Fabric, are required to follow the procedures for establishing and removing Class 1 Dedicated Connections as presented in clause 27 of ISO/IEC 14165-111, FC-PH.

Primitive sequences transmitted from an N_Port to its associated F_Port shall remove an existing or pending Class 1 Dedicated Connection. The affected F_Port shall respond with the appropriate Primitive Sequence protocol. The Fabric shall notify the other F_Port participating in the Dedicated Connection to initiate the Link Reset Protocol with its respective N_Port.

If the Fabric is able to determine that a Dedicated Connection has been broken somewhere within the Fabric, the Fabric may initiate the Link Reset Protocol at the entry and exit F_Ports of the Dedicated Connection.

6.3 Buffered Class 1 service

Buffered Class 1 service is an extension of Class 1 service - Dedicated Connection (see 6.2) in which the flow of all frames is regulated by buffer-to-buffer credit rules (see clause 37 of ISO/IEC 14165-112, FC-PH-2). The purpose of such a connection is to allow ports with differing data rates to communicate in an otherwise Dedicated Connection like environment

Buffered Class 1 service is similar in all respects to Class 1 service - Dedicated Connection - except that the usable bandwidth of the connection is limited to no more than the data rate of the slowest connected port.

6.4 Dedicated Simplex service

Dedicated Simplex service is a variation of Class 1 service - Dedicated Connection (see 6.2) in which the connection is established in one direction only (see clause 33 of ISO/IEC 14165-112, FC-PH-2).

A Dedicated Simplex connection guarantees full-bandwidth transfers and delivery of frames in one direction only. Frame delivery acknowledgments (ACKs) return to the Connection Initiator using the Connectionless service Class 2 (see 6.5).

In Dedicated Simplex service, it is possible for the outbound fiber of a port to have a dedicated path to one destination port at the same time as the inbound fiber has a dedicated path from another source port. In Dedicated Simplex service, the ports are required to support Intermix (see 6.8) so that the ACK response frames may share bandwidth with the Dedicated Simplex connections.

6.5 Class 2 service - Multiplex

Class 2 Multiplex is a Connectionless service that multiplexes frames at the frame boundary between a source F_Port and a destination F_Port of the Fabric. Some Fabrics ensure that frames sent from a source N_Port to a destination N_Port arrive at the destination N_Port in the same order as sent by the source, and some do not.

Class 2 service assures the sender of notification of frame delivery or failure to deliver frames so long as the frames remain error free in transit.

The F_Port participates in buffer-to-buffer flow control (see 26.5 of ISO/IEC 14165-111, FC-PH) and provides R_RDY primitive signalling for Class 2 frames (see 16.3.2 of ISO/IEC 14165-111, FC-PH).

6.6 Class 3 service - Datagram

Class 3 Datagram is a Connectionless service that multiplexes frames at the frame boundary between a source F_Port and a destination F_Port of the Fabric. Some Fabrics ensure that frames sent from a source N_Port to a destination N_Port arrive at the destination N_Port in the same order as sent by the source, and some do not.

Class 3 service provides best-effort delivery with no Busy or Reject indications.

The F_Port participates in buffer-to-buffer flow control (see 26.5 of ISO/IEC 14165-111, FC-PH) and provides R_RDY primitive signalling for Class 3 frames (see 16.3.2 of ISO/IEC 14165-111, FC-PH).

6.7 Class 4 service - Fractional

Class 4 service is a connection-oriented service that provides fractional allocation of the resources of the connection path between communicating N_Ports (see clause 34 of ISO/IEC 14165-112, FC-PH-2).

The Class 4 circuit is bidirectional with one Virtual Circuit (VC) operational in each direction. Each VC may have different Quality of Service (QoS) parameters. The QoS parameters include guaranteed bandwidth and latency.

An N_Port may have up to 254 coexistent Class 4 circuits with the same or different N_Ports. The Virtual Circuit Identifier (VC_ID) is entered in the Class Specific Control Field (CS_CTL) of the frame header (see 18.2.4 of ISO/IEC 14165-112, FC-PH-2).

Class 4 operation is separated into two parts, circuit setup and circuit activation. During the setup process, QoS parameters for both VCs are negotiated between the circuit initiator N_Port, who requests QoS parameters, and the Fabric and the circuit recipient N_Port who guarantee a certain QoS.

Once a Class 4 circuit is setup, it may be activated and deactivated one or more times. Class 4 circuit activation is accomplished separately

on each VC when either N_Port issues a frame delimited by the activate Class 4 Start_of_Frame delimiter (**SOFc4**). Class 4 circuits are deactivated by the End_of_Frame Disconnect_Terminate (**EOFdt**).

Class 4 circuit setup is removed by issuing a frame delimited by the End_of_Frame Remove_Terminate (**EOFrt**).

6.8 Intermix service

The requirements for Class 1 Intermix are given in 22.4 of ISO/IEC 14165-111, FC-PH.

6.9 Class F service - Fabric signaling

Class F service is a connectionless service used for control and coordination of the internal behavior of the Fabric. Class F service is defined in this specification for use by Fabric entities communicating across inter-element links (IEL). It is patterned after Class 2 service with the exceptions that the Start-of-Frame delimiter is the Start-of-Frame Fabric (**SOFF**) and most options are declined.

6.9.1 Class F Frame formats

Class F frames shall use a Fibre Channel compliant Frame Header (see clause 18 of ISO/IEC 14165-111, FC-PH) except when the Routing Control field indicates vendor unique. If the Routing Control field (Word 0, Bits 31-24) contains the hexadecimal value 'FF', indicating vendor unique, then only Word 0 and Word 1 of the Frame Header are defined.

The Class F frame format is illustrated in figure 7. Class F frames shall be delimited by the Start_of_Frame Fabric (**SOFF**) and the End_of_Frame Normal (**EOFn**) and shall include the Cyclic Redundancy Check (CRC) defined by FC-PH. Class F frames shall not include the optional headers described in FC-PH.

6.9.2 Class F function

Class F service is a Connectionless service that multiplexes frames at the frame boundary between a Fabric source entity and a Fabric destination entity. Class F service assures the sender of notification of frame delivery or failure to deliver frames so long as the frames remain error free in transit.

Fabric frames move buffer to buffer and are regulated by R_RDY flow control using the same buffer-to-buffer credit mechanism pre-

scribed for Class 2, Class 3 and Class 1/SOFC1 frames.

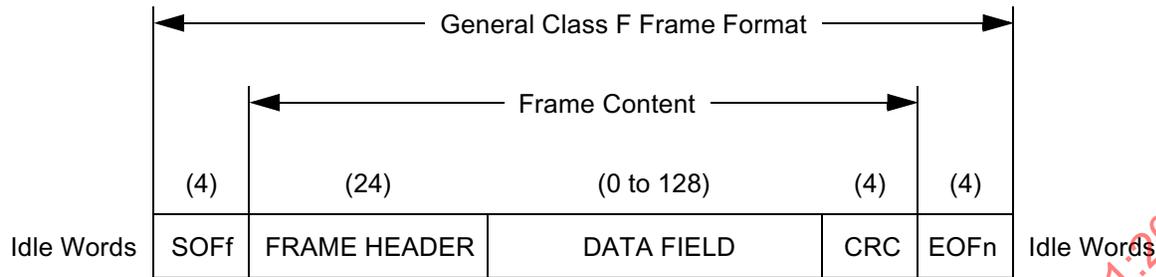


Figure 7 – The Class F frame format

6.9.3 Class F rules

To provide Class F service, the transmitting and receiving E_Ports shall obey the following rules:

- a) Except for some Inter-Element Link Service Protocols, an E_Port supporting Class F service is required to have logged in with the associated Fabric Elements (IELOGI) and the E_Ports (IELOGI) with which it intends to communicate, either explicitly or implicitly. **(Login)**.
- b) The Fabric Element routes Class F frames through Connectionless Sub-Fabric, without establishing a Dedicated Connection between communicating E_Ports. To obtain Class F service from the Fabric Element, the E_Port shall use the Fabric frame **(SOFf)** delimiter. **(Connectionless service)**
- c) An E_Port is allowed to send consecutive frames to one or more destinations. This enables an E_Port to demultiplex multiple Sequences to a single or multiple destinations concurrently. **(demultiplexing)**
- d) A given E_Port may receive consecutive frames from different sources. Each source is allowed to send consecutive frames for one or more Sequences. **(multiplexing)**
- e) A destination E_Port shall provide an acknowledgment to the source for each valid frame received. The destination E_Port shall use ACK_1 for the acknowledgment. If the Fabric Element is unable to deliver the ACK_1 frame, the Fabric Element shall return an F_BSY or F_RJT. **(Acknowledgment)**
- f) The Sequence Initiator shall increment the SEQ_CNT field of each successive frame transmitted within a Sequence. However, the Fabric Elements may not guarantee delivery to the destination in the same order of transmission. **(non-sequential delivery)**
- g) An E_Port may originate multiple Exchanges and initiate multiple Sequences with one or more E_Ports. The E_Port originating an Exchange shall assign an X_ID unique to the Originator called OX_ID and the Responder of the Exchange shall assign an X_ID unique to the responder called RX_ID. The value of OX_ID or RX_ID is unique to a given E_Port. The Sequence Initiator shall assign a SEQ_ID, for each Sequence it initiates, which is unique to the Sequence Initiator and the respective Sequence Recipient pair while the Sequence is Open. (concurrent Exchanges and Sequences)
- h) Each E_Port exercises buffer-to-buffer flow control with the E_Port to which it is directly attached. End-to-end flow control is performed by communicating E_Ports. ACK_1 frames are used to perform end-to-end flow control and R_RDY is used for buffer-to-buffer flow control. **(dual flow control)**