ISO/IEC 14776-224

Edition 1.0 2019-10

INTERNATIONAL STANDARD

Information technology – Small computer system interface (SCSI) – Part 224: Fibre Channel Protocol for SCSI, fourth version (FCP-4)

ISO/IEC 14776-224:2019-10(en)

IEC Central Office
3, rue de Varembé
CH-1211 Geneva 20
Switzerland

Tel.: +41 22 919 02 11
info@iec.ch
www.iec.ch

ISO/IEC 14776-224

Edition 1.0 2019-10

INTERNATIONAL STANDARD

Information technology – Small computer system interface (SCSI) – Part 224: Fibre Channel Protocol for SCSI, fourth version (FCP-4)

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 35.200

ISBN 978-2-8322-7520-7

Warning! Make sure that you obtained this publication from an authorized distributor.

Contents

Tables

Figures

INFORMATION TECHNOLOGY –
SMALL COMPUTER SYSTEM INTERFACE (SCSI) –
Part 224: Fibre Channel Protocol, fourth version (FCP-4)

FOREWORD

1) ISO (the International Organization for Standardization) and IEC (the 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. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

2) The formal decisions or agreements of IEC and ISO on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC National Committees and ISO member bodies.

3) IEC, ISO and ISO/IEC publications have the form of recommendations for international use and are accepted by IEC National Committees and ISO member bodies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC, ISO and ISO/IEC publications is accurate, IEC or ISO cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.

4) In order to promote international uniformity, IEC National Committees and ISO member bodies undertake to apply IEC, ISO and ISO/IEC publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any ISO, IEC or ISO/IEC publication and the corresponding national or regional publication should be clearly indicated in the latter.

5) ISO and IEC do not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. ISO or IEC are not responsible for any services carried out by independent certification bodies.

6) All users should ensure that they have the latest edition of this publication.

7) No liability shall attach to IEC or ISO or its directors, employees, servants or agents including individual experts and members of their technical committees and IEC National Committees or ISO member bodies for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication of, use of, or reliance upon, this ISO/IEC publication or any other IEC, ISO or ISO/IEC publications.

8) Attention is drawn to the normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.

9) Attention is drawn to the possibility that some of the elements of this ISO/IEC publication may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 14776-224 was prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

The list of all currently available parts of the ISO/IEC 14776 series, under the general title *Information technology – Small computer system interface (SCSI)*, can be found on the IEC and ISO websites.

The text for this standard is based on the following document:

CDV	Report on voting
JTC1-SC25/2833/CDV	JTC1-SC25/2881/RVC

Full information on the voting for the approval of this standard can be found in the report on voting indicated in the above table.

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2 except as indicated in 3.4.

A bilingual version of this publication may be issued at a later date.

INTRODUCTION

The Small Computer System Interface (SCSI) command set is widely used and applicable to a wide variety of device types. The transmission of SCSI command set information across Fibre Channel links allows the large body of SCSI application and driver software to be successfully used in the high performance Fibre Channel environment.

This standard describes the protocol for transmitting SCSI commands, data, and status using Fibre Channel FC-FS-3 Exchanges and Information Units. Fibre Channel is a high speed serial architecture that allows either optical or electrical connections. The topologies supported by Fibre Channel include point-to-point, fabric switched, and arbitrated loop. All Fibre Channel connections use the same standard frame format and standard hierarchy of transmission units to transmit the Information Units that carry SCSI information.

This standard is divided into the following clauses:

Clause 1 is the scope of this standard.

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

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

Clause 4 provides an overview of the protocol for transmitting SCSI information over Fibre Channel.

Clause 5 describes the FC-FS-3 frame header.

Clause 6 describes the Basic Link Services and Extended Link Services used by the protocol for transmitting SCSI information over Fibre Channel.

Clause 7 describes the Name Server objects defined for FCP-4.

Clause 8 describes the FCP FC-4 Link Service definitions for the protocol for transmitting SCSI information over Fibre Channel.

Clause 9 describes the Information Units used to transfer SCSI commands, data, and status across a Fibre Channel connection.

Clause 10 defines the SCSI mode pages used by the protocol for transmitting SCSI information over Fibre Channel.

Clause 11 defines the timers used for FCP-4 operation and recovery.

Clause 12 defines the link error detection and error recovery procedures for FCP-4.

This standard has the following annexes:

Annex A is a normative description of the relationship between the services defined by SAM-5 and the corresponding functions defined by this standard.

Annex B is an informative annex that provides examples of the protocol for transmitting SCSI information over FCP.

Annex C is an informative annex providing examples of the FCP-4 error recovery mechanisms.

Annex D is an informative annex describing techniques for discovering FCP device capabilities.

Annex E is an informative annex providing examples of the content of ELSs used during FCP-4 recovery operations.

This standard is part of ISO/IEC 14476 (all parts) developed to facilitate the use of the SCSI command sets for many different types of devices across many different types of physical interconnects. The architectural model for the family of standards is ISO/IEC 14776-415, *Information technology - Small computer system interface (SCSI) - Part 415: SCSI architecture model - 5 (SAM-5).*

INFORMATION TECHNOLOGY –
SMALL COMPUTER SYSTEM INTERFACE (SCSI) –
Part 224: Fibre Channel Protocol, fourth version (FCP-4)

1 Scope

This part of ISO/IEC 14776 defines a fourth version of the SCSI Fibre Channel Protocol (FCP). This standard is a mapping protocol for applying the SCSI command set to Fibre Channel. This standard defines how the Fibre Channel services and the defined Information Units (IUs) are used to perform the services defined by the SCSI Architecture Model - 5 (SAM-5). This fourth version includes additions and clarifications to the third version (ISO/IEC 14776-223:2008), removes information that is now contained in other standards, and describes additional error recovery capabilities for the Fibre Channel Protocol.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

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

ISO/IEC 14776-415, *Information technology - Small computer system interface (SCSI) - Part 415: SCSI architecture model - 5 (SAM-5)*

ISO/IEC 14776-454, *Information technology - Small computer system interface (SCSI) - Part 454: SCSI Primary Commands - 4 (SPC-4)*

INCITS 463-2010, *Information Technology - Fibre Channel Generic Services - 6 (FC-GS-6)*

INCITS 470-2011, *Information Technology - Fibre Channel Framing and Signaling Interface - 3 (FC-FS-3)*

INCITS 477-2011, *Information Technology - Fibre Channel - Link Services - 2 (FC-LS-2)*

INCITS 540, *Information Technology - Fibre Channel - Non-Volatile Memory Express (FC-NVMe)*

INCITS 544, *Information technology - Fibre Channel - Single-Byte Command Code Sets Mapping Protocol - 6 (FC-SB-6)*

INCITS TR-49-2012, *INCITS Technical Report For Information Technology - Fibre Channel - Device Attach - 2 (FC-DA-2)*

SFF document SFF-8067, *40-pin SCA-2 Connector w/Bidirectional ESI* [1]

1. SFF specifications are available from the SNIA SFF Technology Affiliate (see http://www.snia.org/sff).

3 Terms, definitions, abbreviations and conventions

3.1 Terms and definitions

3.1.1
acknowledged class
class of service that acknowledges transfers

Note 1 to entry: An example of an acknowledged class is class 2.

Note 2 to entry: See FC-FS-3.

3.1.2
address identifier
address value used to identify the source (S_ID) or destination (D_ID) of a frame

Note 1 to entry: See FC-FS-3.

3.1.3
application client
object that is the source of SCSI commands and task management function requests

Note 1 to entry: See SAM-5.

3.1.4
application client buffer offset
offset in bytes from the beginning of the application client's buffer (i.e., data-in or data-out) to the location for the transfer of the first byte of a data delivery service request

Note 1 to entry: See SAM-5.

3.1.5
command
request describing a unit of work to be performed by a device server

Note 1 to entry: See SAM-5.

3.1.6
command descriptor block
CDB
structure used to communicate a command from an application client to a device server

Note 1 to entry: See SAM-5.

3.1.7
command identifier
numerical identifier of a command

Note 1 to entry: See Annex A and SAM-5.

3.1.8
data buffer size
upper limit on the amount of data (i.e., data-in or data-out) to be transferred by the command

Note 1 to entry: See SAM-5.

3.1.9
Data frame
FC-4 Device_Data frame, FC-4 Video_Data frame, or Link_Data frame

Note 1 to entry: See FC-FS-3.

3.1.10
data overlay
random buffer access capability where data is transmitted using the same application client buffer offset more than one time during the set of delivery actions performed by a single command

Note 1 to entry: See 6.3.3, 6.3.4, and 9.4.

3.1.11
Destination_Identifier
D_ID
address identifier used to indicate the destination of the transmitted frame

Note 1 to entry: See FC-FS-3.

3.1.12
device server
object within the logical unit that processes SCSI commands and enforces the rules for task management

Note 1 to entry: See SAM-5.

3.1.13
discard
remove a frame or Sequence from the destination buffer without making use of the frame or Sequence and without notifying upper layers of the receipt of the frame or Sequence

Note 1 to entry: See FC-FS-3.

3.1.14
Exchange
basic mechanism that transfers information consisting of one or more related non-concurrent Sequences that may flow in the same or opposite directions

Note 1 to entry: The Exchange is identified by an Originator Exchange_ID (OX_ID) and a Responder Exchange_ID (RX_ID).

Note 2 to entry: See FC-FS-3.

3.1.15
FCP device
device containing one or more FCP_Ports operating as an initiator FCP_Port and/or a target FCP_Port

3.1.16
FCP Exchange
SCSI I/O operation for the Fibre Channel FC-2 layer

Note 1 to entry: The SCSI I/O operation for Fibre Channel is contained in a Fibre Channel Exchange.

Note 2 to entry: See FC-FS-3 and 4.1.

3.1.17
FCP I/O operation
SCSI I/O operation for the Fibre Channel FC-4 layer, as defined in this standard

3.1.18
FCP_Port
Nx_Port that supports the SCSI Fibre Channel Protocol

3.1.19
fully qualified Exchange identifier
FQXID
set of addresses and values used to uniquely identify an FCP I/O operation

Note 1 to entry: See 4.16.

3.1.20
image pair
originating and responding processes related by a Process Login operation

Note 1 to entry: For the Fibre Channel Protocol, the image pair is composed of one initiator FCP_Port and one target FCP_Port.

Note 2 to entry: See FC-LS-2.

3.1.21
I_T nexus loss
condition resulting from the events defined by SAM-5 in which the SCSI device performs the I_T nexus loss operations described in SAM-5, SPC-4, and this standard

3.1.22
I_T nexus loss notification
SCSI transport protocol specific event that results in an I_T nexus loss as described in SAM-5

3.1.23
Information Unit
IU
organized collection of data specified by the Fibre Channel Protocol to be transferred as a single Sequence by the Fibre Channel service interface (see FC-FS-3)

3.1.24
initiator FCP_Port
FCP_Port using the Fibre Channel Protocol to perform the SCSI initiator device functions defined by SAM-5

3.1.25
initiator port identifier
value by which a SCSI initiator port is referenced within a SCSI domain (see SAM-5)

Note 1 to entry: In this standard the address identifier of the initiator FCP_Port is an initiator port identifier.

3.1.26
initiator port name
name of a SCSI initiator port

Note 1 to entry: See SAM-5.

3.1.27
interconnect tenancy
period of time that an FCP device owns or may access a shared Fibre Channel link such as an FC-AL-2 loop

Note 1 to entry: See 10.2.1.

3.1.28
logical unit
externally addressable entity within a SCSI target device that implements a SCSI device model and contains a device server

Note 1 to entry: See SAM-5.

3.1.29
logical unit number
LUN
identifier for a logical unit

Note 1 to entry: See SAM-5.

3.1.30
loop initialization primitive
primitive used in Fibre Channel arbitrated loops to start loop initialization

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

3.1.31
management logical unit
logical unit that only performs management functions (e.g., device configuration and discovery of devices)

Note 1 to entry: See SAM-5.

3.1.32
Name_Identifier
64-bit identifier, with a 60-bit value preceded with a 4 bit Network_ Address_Authority Identifier, used to identify entities in Fibre Channel (e.g., N_Port, Node, F_Port, or Fabric)

Note 1 to entry: See FC-FS-3.

3.1.33
Name Server
Fibre Channel service accessed through a well-known address identifier that uses the Common Transfer (CT) protocol as defined in FC-GS-6 to allow a client to determine the address identifier and properties of devices attached to a Fibre Channel switching fabric

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

3.1.34
Node_Name
Name_Identifier associated with a Node

Note 1 to entry: See FC-FS-3.

3.1.35
NL_Port
N_Port that contains arbitrated loop functions associated with the Fibre Channel arbitrated loop topology

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

3.1.36
N_Port
hardware entity that supports the FC-FS-3 FC-2 layer and may act as an Originator, a Responder, or both

Note 1 to entry: See FC-FS-3.

3.1.37
N_Port Login
PLOGI
Fibre Channel Extended Link Service (ELS) that exchanges identification and operation parameters between an originating N_Port and a responding N_Port

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

3.1.38
N_Port_Name
Name_Identifier (see FC-FS-3) associated with an Nx_Port

3.1.39
Nx_Port
end device for Fibre Channel communication

Note 1 to entry: In this standard the term Nx_Port is used to specify behavior of either N_Ports or NL_Ports.

3.1.40
Originator
logical function associated with an N_Port responsible for originating an Exchange

Note 1 to entry: See FC-FS-3.

3.1.41
Originator Exchange_ID
OX_ID
identifier assigned by an Originator to identify an Exchange

Note 1 to entry: See 3.1.48 and FC-FS-3.

3.1.42
private loop
arbitrated loop operating with no attached fabric loop ports

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

3.1.43
public loop
arbitrated loop operating with an attached fabric loop port

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

3.1.44
random buffer access
occurrence of device server data transfer requests that request data transfers to or from segments of the application client's buffer with an arbitrary offset and extent

Note 1 to entry: See SAM-5.

3.1.45
read operation
operation that uses the Data-In action, IU I3 (see 9.1)

3.1.46
request byte count
number of bytes to be moved by a data delivery service request

Note 1 to entry: See SAM-5.

3.1.47
Responder
logical function in an N_Port responsible for supporting the Exchange initiated by the Originator in another N_Port

Note 1 to entry: See FC-FS-3.

3.1.48
Responder Exchange_ID
Responder Exchange Identifier
RX_ID
identifier assigned by a Responder to identify an Exchange and meaningful only to the Responder

Note 1 to entry: See FC-FS-3.

3.1.49
SCSI device
device that contains one or more SCSI ports that are each connected to a service delivery subsystem and supports a SCSI application protocol

Note 1 to entry: See SAM-5.

3.1.50
SCSI I/O operation
operation defined by a command or a task management function

Note 1 to entry: See SAM-5.

3.1.51
SCSI initiator device
SCSI device (see 3.1.49) containing application clients (see 3.1.3) and SCSI initiator ports (see 3.1.52)

Note 1 to entry: SCSI Initiator devices originate device service and task management requests (see SAM-5) to be processed by a SCSI target device (see 3.1.53) and receive device service and task management responses from SCSI target devices.

Note 2 to entry: See SAM-5.

3.1.52
SCSI initiator port
SCSI initiator device (see SAM-5) object that acts as the connection between application clients and a service delivery subsystem through which requests and responses are routed

Note 1 to entry: See SAM-5.

3.1.53
SCSI target device
SCSI device (see 3.1.49) containing logical units (see 3.1.28) and SCSI target ports (see 3.1.54)

Note 1 to entry: SCSI target devices receive device service and task management requests (see SAM-5) for processing and send device service and task management responses to SCSI initiator devices.

Note 2 to entry: See SAM-5.

3.1.54
SCSI target port
SCSI target device (see SAM-5) object that acts as the connection between device servers and task managers and a service delivery subsystem through which requests and responses are routed

Note 1 to entry: See SAM-5.

3.1.55
sense data
data describing command completion information that a device server delivers to an application client in an FCP_RSP IU along with the status or as parameter data in response to a REQUEST SENSE command

Note 1 to entry: See SPC-4.

3.1.56
Sequence
set of one or more Data frames (see 3.1.9) with a common Sequence_ID (SEQ_ID), transmitted unidirectionally from one N_Port to another N_Port with a corresponding response, if applicable, transmitted in response to each Data frame

Note 1 to entry: See FC-FS-3.

3.1.57
Sequence_ID
SEQ-ID
identifier used to identify a Sequence

Note 1 to entry: See FC-FS-3.

3.1.58
Source_Identifier
S_ID
address identifier used to indicate the source port of the transmitted frame

Note 1 to entry: See FC-FS-3.

3.1.59
status
single byte returned by the device server to the application client to indicate the completion and completion state of a command

Note 1 to entry: See SAM-5.

3.1.60
target FCP_Port
FCP_Port using the Fibre Channel Protocol to perform the SCSI target device functions defined by SAM-5

3.1.61
target port identifier
value by which a SCSI target port (see 3.1.54) is referenced within a SCSI domain (see SAM-5)

Note 1 to entry: In this standard the address identifier of the target FCP_Port is a target port identifier.

3.1.62
target port name
name of a SCSI target port

Note 1 to entry: See SAM-5.

3.1.63
task attribute
attribute of a command (see 3.1.5) that specifies the processing relationship of a command with regard to other commands in the task set

EXAMPLE - Examples of task attributes are SIMPLE, ORDERED, HEAD OF QUEUE, ACA.

Note 1 to entry: See SAM-5.

3.1.64
task management function
peer-to-peer confirmed service provided by a task manager that may be invoked by an application client to affect the processing of one or more tasks

Note 1 to entry: See SAM-5.

3.1.65
task retry identifier
identifier that is used to associate a command and the Link Services used to retry any of the command IUs

3.1.66
unacknowledged class
class of service that does not acknowledge transfers

Note 1 to entry: An example of an unacknowledged class is class 3.

Note 2 to entry: See FC-FS-3.

3.1.67
Worldwide_Name
Name_Identifier that is worldwide unique, and represented by a 64-bit value

Note 1 to entry: See FC-FS-3.

3.1.68
word
string of four contiguous bytes occurring on boundaries that are zero modulo 4 from a specified reference

Note 1 to entry: See FC-FS-3.

3.1.69
write operation
operation that uses the Data-Out action, IU T6 (see 9.1)

3.2 Abbreviations

ABTS	Abort Sequence Basic Link Service (see FC-FS-3)
ABTS-LS	ABTS with the PARAMETER field bit 0 set to zero (i.e., Abort Exchange) (see FC-FS-3)
BA_ACC	Basic Link Service Accept (Basic_Accept) (see FC-FS-3)
BA_RJT	Basic Link Service Reject (Basic_Reject) (see FC-FS-3)
BLS	Basic Link Service (see FC-FS-3)
CDB	command descriptor block (see 3.1.6)
CRN	Command Reference Number (see 4.4 and 9.2.2.2)
D_ID	Destination_Identifier (see 3.1.11)
ELS	Extended Link Service (see FC-LS-2)
FC	Fibre Channel (see FC-FS-3)
FC-AL-2	Fibre Channel Arbitrated Loop-2 and Fibre Channel Arbitrated Loop-2/Amd1 (see clause 2)
FC-FS-3	Fibre Channel Framing and Signaling Interface - 3 (see clause 2)
FC-GS-6	Fibre Channel Generic Services - 6 (see clause 2)
FC-LS-2	Fibre Channel - Link Services - 2 (see clause 2)
FCP	Refers to this standard

FCP_ACC	FCP FC-4 Link Service Accept
FCP_LS	FCP FC-4 Link Service
FCP_RJT	FCP FC-4 Link Service Reject (see 8.4)
FCP-4	This standard
FC-4	Fibre Channel Layer 4 mapping layer (see FC-FS-3)
FLOGI	Fabric Login (see FC-LS-2)
FQXID	fully qualified Exchange identifier (see 3.1.19)
ID	identifier
IU	Information Unit (see 3.1.23)
LIFA	Loop Initialization Fabric Assigned (see FC-AL-2)
LIHA	Loop Initialization Hard Assigned (see FC-AL-2)
LIP	Loop Initialization Primitive (see FC-AL-2)
LIPA	Loop Initialization Previously Assigned (see FC-AL-2)
LISA	Loop Initialization Soft Assigned (see FC-AL-2)
LISM	Loop Initialization Select Master (see FC-AL-2)
LOGO	Logout (see FC-LS-2)
LS	Link Service
LS_ACC	Link Service Accept reply Sequence
LS_RJT	Link Service Reject reply Sequence
LUN	logical unit number
NA	Not Applicable
OX_ID	Originator Exchange Identifier (see 3.1.41)
PLOGI	N_Port Login (see FC-LS-2)
PRLI	Process Login (see 6.3 and FC-LS-2)
PRLO	Process Logout (see 6.4 and FC-LS-2)
REC	Read Exchange Concise (see 6.5 and FC-LS-2)
RRQ	Reinstate Recovery Qualifier (see 6.5 and FC-LS-2)
RX_ID	Responder Exchange_ID (see 3.1.48)
SAM-5	SCSI Architecture Model - 5 (see clause 2)
SCSI	Small Computer System Interface, any revision
S_ID	Source_Identifier (see 3.1.58)
SEQ_ID	Sequence_ID
SPC-4	SCSI Primary Commands - 4 (see clause 2)
SRR	Sequence Retransmission Request (see 8.2)
TPRLO	Third Party Process Logout (see FC-LS-2)
ULP	upper layer protocol (see FC-FS-3)

3.3 Keywords

3.3.1
invalid
keyword used to describe an illegal or unsupported bit, byte, word, field or code value

Note 1 to entry: Receipt of an invalid bit, byte, word, field or code value shall be reported as error.

3.3.2
mandatory
keyword indicating an item that is required to be implemented as defined in this standard

3.3.3
may
keyword that indicates flexibility of choice with no implied preference

Note 1 to entry: Equivalent to “may or may not”.

3.3.4
may not
keyword that indicates flexibility of choice with no implied preference

Note 1 to entry: Equivalent to “may or may not”.

3.3.5
obsolete
keyword indicating that an item was defined in a prior SCSI standard but has been removed from this standard

3.3.6
optional
keyword that describes features that are not required to be implemented by this standard

Note 1 to entry: If any optional feature defined by this standard is implemented, then it shall be implemented as defined in this standard.

3.3.7
reserved
keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization

Note 1 to entry: A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard.

Note 2 to entry: Recipients are not required to check reserved bits, bytes, words or fields for zero values.

Note 3 to entry: Receipt of reserved code values in defined fields shall be reported as an error.

3.3.8
restricted
keyword referring to bits, bytes, words, and fields that are set aside for other identified standardization purposes

Note 1 to entry: A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word or field in the context where the restricted designation appears.

3.3.9
shall
keyword indicating a mandatory requirement

Note 1 to entry: Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

Note 2 to entry: This standard prescribes no specific response by a component if it receives information that violates a mandatory behavior.

3.3.10
should
keyword indicating flexibility of choice with a strongly preferred alternative

Note 1 to entry: Equivalent to the phrase “it is strongly recommended”.

3.4 Editorial conventions

Certain words and terms used in this standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in clause 3 or in the text where they first appear.

Names of signals, address frames, primitives and primitive sequences, state machines, SCSI commands, SCSI statuses, SCSI sense keys, and SCSI additional sense codes are in all uppercase (e.g., REQUEST SENSE command).

Names of messages, requests, confirmations, indications, responses, event notifications, timers, SCSI diagnostic pages, SCSI mode pages, and SCSI log pages are in mixed case (e.g., Disconnect-Reconnect mode page).

Names of fields are in small uppercase (e.g., FCP_DL). Normal case is used when the contents of a field are being discussed. Fields containing only one bit are usually referred to as the NAME bit instead of the NAME field.

Normal case is used for words having the normal English meaning.

A binary number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b). Underscores or spaces may be included between characters in binary number representations to increase readability or delineate field boundaries (e.g., 0 0101 1010b or 0_0101_1010b).

A hexadecimal number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 through 9 and/or the upper-case English letters A through F immediately followed by a lower-case h (e.g., FA23h). Underscores or spaces may be included between characters in hexadecimal number representations to increase readability or delineate field boundaries (e.g., B FD8C FA23h or B_FD8C_FA23h).

A decimal number is represented in this standard by any sequence of digits consisting of only the Arabic numerals 0 through 9 not immediately followed by a lower-case b or lower-case h (e.g., 25).

This standard uses the following conventions for representing decimal numbers:

a) the decimal separator (i.e., separating the integer and fractional portions of the number) is a period;
b) the thousands separator (i.e., separating groups of three digits in a portion of the number) is a space; and
c) the thousands separator is used in both the integer portion and the fraction portion of a number.

In all of the figures, tables, and text of this standard, the most significant bit of a binary quantity is shown on the left side. Bit order and byte order are as specified in FC-FS-3.

Table 1 shows some examples of decimal numbers using various numbering conventions.

Table 1 - Numbering conventions

French	English	This standard
0,6	0.6	0.6
3,141 592 65	3.14159265	3.141 592 65
1 000	1,000	1 000
1 323 462,95	1,323,462.95	1 323 462.95

A decimal number represented in this standard with an overline over one or more digits following the decimal point is a number where the overlined digits are infinitely repeating (e.g., $666.\overline{6}$ means 666.666 666... or 666 2/3, and $12.\overline{142\ 857}$ means 12.142 857 142 857... or 12 1/7).

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

In the event of conflicting information the precedence for requirements defined in this standard is:

1) text;
2) tables; and
3) figures.

Notes do not constitute any requirements for implementers.

Notes are numbered consecutively throughout this standard.

4 General

4.1 Structure and concepts

Fibre Channel (FC) is logically a point-to-point serial data channel. The architecture has been designed so that it may be implemented with high performance hardware that requires little real-time software management. The Fibre Channel Physical layer (FC-2 layer) described by FC-FS-3 performs those functions required to transfer data from one Nx_Port to another. In this standard, Nx_Ports capable of supporting Fibre Channel Protocol transactions are collectively referred to as FCP_Ports. The FC-2 layer is a delivery service with information grouping and several defined classes of service.

A switching fabric allows communication among more than two FCP_Ports.

Fibre Channel Arbitrated Loop-2 (FC-AL-2) is an alternative multiple port topology that allows communication between two NL_Ports on an arbitrated loop.

An FC-4 mapping layer uses the services provided by FC-FS-3 to perform the functions defined by the FC-4. The protocol is described in terms of the stream of FC IUs and Exchanges generated by a pair of FCP_Ports that support the FC-4.

The detailed implementation that supports that stream is not defined by this standard. Originator and Responder FCP_Ports are assumed to have a common service interface, for use by all FC-4s, that is similar in characteristics to the service interface defined in FC-FS-3. The requirements for the service interface for SCSI are contained in SAM-5.

This standard defines four kinds of functional management:

a) device management;
b) task management;
c) Process Login and Process Logout management; and
d) link management.

The Fibre Channel Protocol device and task management protocols define the mapping of the SCSI functions (see SAM-5) to the Fibre Channel interface (see FC-FS-3). Link control is performed by standard FC-FS-3 protocols. The task management functions defined by SAM-5 are mapped as described in 4.9 of this standard. The I/O operation defined by SAM-5 is mapped into a Fibre Channel Exchange. A Fibre Channel Exchange carrying information for a SCSI I/O operation is an FCP Exchange. The request and response primitives of an I/O operation are mapped into Information Units (IUs) as shown in table 2.

Table 2 - SCSI and Fibre Channel Protocol functions

SCSI function	FCP equivalent
I/O operation	Exchange
Transport protocol service request and response	Sequence
Send SCSI Command request	Sending an unsolicited command IU (FCP_CMND)
SCSI Command Received indication	Receiving an unsolicited command IU (FCP_CMND)
Receive Data-Out request	Data descriptor IU (FCP_XFER_RDY)
Data-Out Received confirmation	Receipt of solicited data IU (FCP_DATA)
Send Data-In request	Sending solicited data IO (FCP_DATA)
Data-In Delivered confirmation	Dependent on class of service
Send Command Complete response	Sending a command status IU (FCP_ RSP)
Command Complete Received confirmation	Receiving a command status IU (FCP_RSP)
Send Task Management request	Sending the FCP equivalent specified in 4.9.
Task Management Request Received indication	Receiving the FCP equivalent specified in 4.9.
Task Management Function Executed response	Sending the response specified in table 4, table 5, or table 6.
Received Task Management Function Executed response	Receiving the response specified in table 4, table 5, or table 6.

The number of Exchanges that may simultaneously be open between an initiator FCP_Port and a target FCP_Port is defined by the FC-FS-3 implementation. The architectural limit for this value is 65 535. The maximum number of active Sequences that may simultaneously be open between an initiator FCP_Port and a target FCP_Port is restricted by the allowable range of values of the Sequence ID to 256 (see FC-FS-3). To allow task management Exchanges to be originated, a certain number of extra Exchange IDs and at least one extra Sequence_ID should always be available.

4.2 FCP I/O operations

An application client begins an FCP I/O operation when it invokes a Send SCSI Command SCSI transport protocol service request or a Send Task Management Request SCSI transport protocol service request (see SAM-5).

The use of the FCP_CMND IU for Send Task Management Request SCSI transport protocol service requests is specified in 4.9.

The Send SCSI Command SCSI transport protocol service request conveys a single request from the application client to the FCP service delivery subsystem. Each request contains all the information necessary for the processing of one command, including the local storage address and characteristics of data to be transferred by the command. The Fibre Channel Protocol then performs the following actions using FC-FS-3 services to perform the command. The processing of the individual steps of the protocol is consistent with the SCSI architectural model as defined by SAM-5.

The initiator FCP_Port starts an Exchange by transmitting an unsolicited command IU containing the FCP_CMND IU payload, including some command controls, addressing information, and the command descriptor block (CDB). The initiator FCP_Port transmits the FCP_CMND IU payload to start the FCP I/O operation. The Exchange that is started is identified by its fully qualified Exchange identifier (FQXID) during the remainder of the FCP I/O operation and is used only for the IUs associated with that FCP I/O operation. See 4.16.

If the device server has interpreted the command, has determined that a write operation is required, and is prepared to request the data delivery service, then it invokes the Receive Data-Out transport protocol service request and the target FCP_Port transmits a data descriptor IU containing the FCP_XFER_RDY IU payload to the initiator FCP_Port indicating which portion of the data is to be transferred. The initiator FCP_Port then transmits a solicited data IU to the target FCP_Port containing the FCP_DATA IU payload requested by the FCP_XFER_RDY IU. Data delivery requests containing FCP_XFER_RDY IU and returning FCP_DATA IU payloads continue until the data transfer requested by the command is complete. One FCP_DATA IU shall follow each FCP_XFER_RDY IU. If the initiator FCP_Port and target FCP_Port have negotiated to disable the initial FCP_XFER_RDY IU (see 6.3.3), then a first burst shall be transferred (see 10.2.10).

If the device server has interpreted the command and has determined that a read operation is required, then it invokes the Send Data-In transport protocol service request (see SAM-4) and the target FCP_Port transmits a solicited data IU containing the FCP_DATA IU payload to the initiator FCP_Port. Data deliveries containing FCP_DATA IU payloads continue until all data described by the command is transferred.

If the device server has interpreted the command and has determined that bidirectional transfer is required, then it selects the first FCP_DATA IU to be transferred. The IU may be either a Data-In or a Data-Out transfer. If the device server chooses to request a Data-Out transfer first, it invokes the Receive Data Out transport protocol service and the target FCP_Port transmits a data descriptor IU containing the FCP_XFER_RDY IU payload to the initiator FCP_Port to indicate which portion of the data is to be transferred. The initiator FCP_Port then transmits a solicited data IU containing the FCP_DATA IU payload requested by the FCP_XFER_RDY IU. If the device server chooses to transmit a Data-In transfer first, then it invokes the Send Data-In transport protocol service request (see SAM-4) and the target FCP_Port transmits a solicited data IU containing the FCP_DATA IU payload to the initiator FCP_Port. The device server then selects the next FCP_DATA IU to be transmitted and performs the appropriate procedure to transmit. Data deliveries continue until all data described by the command is transferred. This standard places no restrictions on the order that the device server performs Data-In and Data-Out transfer operations, except that only one Data-In or Data-Out transfer operation is allowed at a time in an Exchange. If the initiator FCP_Port and target FCP_Port have negotiated to disable the initial FCP_XFER_RDY IU (see 6.3.3), then a first burst shall be transferred (see 10.2.10).

After all the data has been transferred, the device server invokes the Send Command Complete transport protocol service response (see SAM-5) and the target FCP_Port transmits a command status IU containing the FCP_RSP IU payload. That payload contains the SCSI status and, if the SCSI status is CHECK CONDITION, the sense data describing the condition. The FCP_RSP IU indicates completion of the command. If no error recovery or confirmed completion is requested, then the FCP_RSP IU is the final Sequence of the Exchange, and the FCP I/O operation and the Exchange are terminated. If an FCP protocol error occurred during processing of the command, the FCP_RSP IU payload carries the FCP response information instead of the SCSI status and sense data.

When the command is completed, the initiator FCP_Port uses returned information to invoke invoke the Command Complete Received transport protocol service confirmation to notify the application client. The returned status indicates whether or not the command was successful. The successful completion of the command indicates that the SCSI target device performed the requested operations with the transferred data and that the information was successfully transferred to or from the SCSI initiator device. Status other than successful completion indicates that either SCSI sense data or warnings about unexpected FCP behaviors are being provided. In this case, the sense data or warning is interpreted to determine whether the desired operation was successfully completed. The device server may request confirmed delivery of the FCP_RSP IU payload as described in 4.5.

The number of FCP I/O operations that may be active at one time depends on the queuing capabilities of the FCP device. If command queueing resources are unavailable in the logical unit when a command is received, then the device server returns TASK SET FULL status or BUSY status in the FCP_RSP IU as specified by SAM-5.

The Fibre Channel Protocol takes full advantage of the multiplexing and shared bandwidth capabilities provided by Fibre Channel classes of service. The protocol is designed to operate with Class 2 or Class

3 service and to provide options for reliable error detection and error recovery independent of the class of service.

The SCSI initiator port function may exist in any FCP_Port and the SCSI target port function may exist in any FCP_Port. For FCP I/O operations between a host and a peripheral subsystem, the host typically takes on the SCSI initiator port role and the peripheral subsystem typically takes on the SCSI target port role. For host to host communications, either one of the communicating pair may take on the SCSI initiator port role. For device to device communications, typically used to implement extended copy and other third-party operations, the SCSI initiator port role is adopted by the managing FCP device.

4.3 Bidirectional and unidirectional commands and FCP_RSP IU format

A device server that supports bidirectional commands should implement both unidirectional and bidirectional commands. Two FCP_RSP IU formats are defined. For commands that set both the RDDATA bit and WRDATA bit to one, the bidirectional FCP_RSP IU payload shall be used for presenting all status and error conditions. For commands that set either the RDDATA bit or WRDATA bit to one or both to zero, the unidirectional FCP_RSP IU payload shall be used for presenting all status and error conditions. The format of the FCP_RSP IU that is returned depends only on the state of the RDDATA bit and WRDATA bit.

If a device server that does not support bidirectional commands receives a command that requests read and write operations by setting both the RDDATA bit and WRDATA bit to one, then the device server may return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN COMMAND INFORMATION UNIT.

4.4 Precise delivery of commands

In applications where SCSI communications between an application client and a device server are stateless, verification of the delivery and processing of commands may not be critical. Any changes in processing sequence caused by link failures or switch latencies may not be important and the recovery and retry mechanisms may be performed while other activities are continued by the application client and the device server.

To assure ordering of commands, an initiator FCP_Port transmits a single command and waits for GOOD status before transmitting the next command (i.e., commands are guaranteed to be processed in order).

This standard defines a second optional mechanism called precise delivery to assure ordering of commands. This may be used by any FCP device, but may be useful for devices performing ordered command queuing where device state is preserved from one command to the next. An application client may determine if a device server supports the precise delivery function by using the MODE SENSE command and MODE SELECT command to examine and set the EPDC bit in the Fibre Channel Logical Unit Control mode page to one (see 10.3).

If a SCSI command requires precise delivery and the device server has the EPDC bit set to one, then the application client provides the Command Reference Number (CRN) argument to the Send SCSI Command transport protocol service. The initiator FCP_Port then places the CRN value in the COMMAND REFERENCE NUMBER field in the FCP_CMND IU.

The CRN is a one byte unsigned integer that starts at the reset value of one and shall be incremented by one for each command requiring precise delivery for that device server. Separate increment counters are maintained for each I_T_L nexus (i.e., each initiator FCP_Port maintains a separate counter for each device server using precise delivery). After the number of precisely delivered commands causes the integer to increment to 255, the integer wraps back to a value of one. The value of zero is reserved and shall be used for those commands that do not require precise delivery and the COMMAND REFERENCE NUMBER field shall be set to zero for task management functions.

The following rules specify how the application client and device server use the CRN to determine that each command requiring precise delivery has been properly received and processed:

a) see table 8 and table 9 for the actions that cause the CRN to be transmitted by the initiator FCP_Port to be set to one and the CRN expected by the device server to be set to one;

b) the CRN shall be equal to one for the first FCP_CMND IU requiring precise delivery between the application client and device server and shall be incremented by one for each subsequent command requiring precise delivery;
c) the CRN shall wrap from 255 to one (i.e., a value of zero in the COMMAND REFERENCE NUMBER field is not valid for an Exchange using precise delivery);
d) the initiator FCP_Port shall not transmit the same CRN again until delivery of the first FCP_CMND IU transmitted with that CRN has been confirmed by receipt of:
   A) an FCP_XFER_RDY IU;
   B) the first Data frame of an FCP_DATA IU;
   C) an FCP_RSP IU;
   D) an ACK; or
   E) a response to an REC ELS;
e) the device server shall not accept a command with a nonzero CRN into the dormant or enabled state until after all commands with a previous CRN have been received by the device server. The commands shall be assumed to be received in the order of increasing CRN, and accounting for a wrap from 255 to one, the highest CRN last. The order of processing of the commands shall be managed by the normal task set management algorithms;
f) the device server shall accept any valid command with a CRN of zero into the dormant or enabled state regardless of whether or not all commands with a nonzero CRN have been received. The processing order of the commands shall be managed by the normal task set management algorithms (see SAM-5); and
g) task management functions shall have the FCP_CMND IU COMMAND REFERENCE NUMBER field set to zero and shall not be tested for precise delivery by the device server.

Any command may use a CRN of zero if precise delivery is not required for that command (e.g., commands such as INQUIRY, TEST UNIT READY, REPORT LUNS, MODE SENSE, and MODE SELECT that are used for booting and initialization may use a CRN of zero).

4.5 Confirmed completion of FCP I/O operations

Some device servers require an acknowledgment of successful delivery of FCP_RSP information. Such an acknowledgment is provided by the optional confirmed completion function. The CONFIRMED COMPLETION ALLOWED bits in the PRLI ELS request FCP Service Parameter page (see 6.3.3) and PRLI ELS accept FCP Service Parameter page (see 6.3.4) are used to negotiate the use of confirmed completion function.

If the CONFIRMED COMPLETION ALLOWED bit is set to one in the PRLI ELS accept FCP Service Parameter page, then the target FCP_Port may request the confirmed completion function by setting the FCP_CONF_REQ bit to one in the FCP_RSP IU. Upon receiving the request in the FCP_RSP IU, the initiator FCP_Port shall transmit an FCP_CONF IU to the target FCP_Port, indicating to the target FCP_Port that the FCP_RSP IU has been received by the initiator FCP_Port.

The confirmed completion function allows the retry of unsuccessful notifications of errors and confirms that the initiator FCP_Port and the target FCP_Port both agree upon the state of a state dependent device. Retry mechanisms for unsuccessful transmission of FCP_RSP IUs and FCP_CONF IUs are defined in this standard.

Target FCP_Ports shall not request confirmed completion for FCP_RSP IUs responding to task management requests.

If confirmed completion is not enabled, then the FCP_CONF IU shall not be requested by the FCP_RSP IU.

Confirmed completion may assist SCSI initiator devices and SCSI target devices in many environments. Examples include:

a) the confirmed completion function may be used to confirm that an initiator FCP_Port has received an FCP_RSP IU reporting a SCSI CHECK CONDITION status, together with accompanying sense data. Upon receiving the FCP_CONF IU, the target FCP_Port may discard its copy of the sense data;
b) the confirmed completion function may be used to confirm that a command has been completed and that the completion information has been successfully transferred to the initiator FCP_Port.

That allows subsequent queued state dependent operations to be performed, since the FCP_CONF IU confirms that the FCP_RSP IU has been received by the initiator FCP_Port; and

c) the confirmed completion function may be used to confirm that an initiator FCP_Port has received the FCP_RSP IU for target FCP_Ports that require state dependent synchronization with initiator FCP_Ports.

4.6 Retransmission of unsuccessfully transmitted IUs

Error detection and IU retransmission algorithms are defined in clause 12.

The Read Exchange Concise (REC) ELS may be used by the initiator FCP_Port to determine the state of an ongoing Exchange. See 6.5.

Support for the REC ELS by both the initiator FCP_Port and target FCP_Port is indicated by the REC SUPPORT bit in the PRLI ELS request FCP Service Parameter page (see 6.3.3) and the PRLI ELS accept FCP Service Parameter page (see 6.3.4).

If the target FCP_Port responds with the REC SUPPORT bit set to one and an error is identified by any of the detection mechanisms defined in clause 12, then the initiator FCP_Port may use the REC ELS to determine the nature of the error.

Target FCP_Ports that do not support the REC SUPPORT bit indicate they do not support the REC ELS by returning a Link Service Reject (LS_RJT) with a Reason Code of "Command not supported" in response to an REC ELS. See 8.4.

If the data retransmission capability is supported by both the initiator FCP_Port and target FCP_Port as indicated by the RETRY bit in the PRLI ELS request FCP Service Parameter page and PRLI ELS accept FCP Service Parameter page (see 6.3.3 and 6.3.4), then the initiator FCP_Port and target FCP_Port shall support the REC ELS and task retry identification (see 4.7).

If an error is identified by any of the mechanisms defined in clause 12, then the initiator FCP_Port may request retransmission using the Sequence Retransmission Request (SRR) FCP_LS request (see 8.2).

4.7 Task retry identification

Task retry identification provides an additional mechanism for relating commands that are being retried to the requests that are sensing the requirement for recovery (i.e., REC ELS request) and performing the recovery (i.e., SRR FCP_LS request).

For example, it is possible that initiator FCP_Ports may re-use OX_ID field values rapidly enough to create an ambiguous situation where the status being preserved in the target FCP_Port for possible retransmission and the new command being presented to the target FCP_Port may have the same OX_ID field values. When recovery of a transmission failure for the new command is attempted, the target FCP_Port instead indicates that the recovery is related to the previous command's status and the initiator FCP_Port is provided status for the completed command. That information is mistakenly interpreted as status for the failed command.

FCP_Ports that agree to perform recovery shall support task retry identification. If the initiator FCP_Port and target FCP_Port agree to support task retry identification, then a task retry identifier shall be provided in the PARAMETER field of each FCP_CMND IU frame. The Link Services associated with retransmission of IUs (i.e., REC ELS and SRR FCP_LS) each contain the same task retry identifier. If the initiator FCP_Port and target FCP_Port do not agree to support task retry identification, then the PARAMETER field shall be set to zero for the FCP_CMND IU frame, REC ELS frame, and SRR FCP_LS frame.

4.8 Discovery of FCP capabilities

A number of Fibre Channel Protocol capabilities require the knowledge and agreement of both the target FCP_Port and the initiator FCP_Port that such capabilities may or shall be used. Table 3 provides references to the discovery process for each of the Fibre Channel Protocol capabilities.

Table 3 - Discovery of FCP-4 capabilities

Capability	Discovery mechanism	Reference
Initiator FCP_Port	Process Login	6.3
Target FCP_Port	Process Login	6.3
Initiator FCP_Port accepts data overlay	Process Login	6.3.3
Target FCP_Port performs data overlay	Disconnect-Reconnect mode page EMDP bit	10.2.8
Initiator FCP_Port transmits FCP_CONF IU	Process Login	6.3
Target FCP_Port requests FCP_CONF IU	Process Login	6.3
Initiator FCP_Port transmits REC ELS request	None required, Process Login allowed	4.6 and 6.3
Target FCP_Port accepts REC ELS request	Process Login[a]	4.6 and 6.3
Initiator FCP_Port transmits SRR FCP_LS request	Process Login	6.3
Target FCP_Port accepts SRR FCP_LS request	Process Login	6.3
Initiator FCP_Port provides CRN	Fibre Channel Logical Unit Control mode page EPDC bit	4.4 and 10.3
Target FCP_Port accepts CRN	Fibre Channel Logical Unit Control mode page EPDC bit	4.4 and 10.3
Task Retry Identification	Process Login	6.3

a) If the target FCP_Port does not support the REC SUPPORT bit in the PRLI ELS, then the target FCP_Port may return an LS_RJT in response to an REC ELS (see 4.6).

4.9 Task management functions

4.9.1 Task management functions overview

An application client requests a task management function (see SAM-5) to control explicitly the processing of one or more FCP I/O operations (see 9.2.2.5).

The task management function mappings are specified in table 4.

Table 4 - Task management functions, SAM-5 to FCP-4

SAM-4 task management function	FCP-4 implementation
ABORT TASK	ABTS-LS (see 4.9.2 and FC-FS-3)
ABORT TASK SET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_ABORT_TASK_SET
CLEAR TASK SET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_CLEAR_TASK_SET
CLEAR ACA	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_CLEAR_ACA
I_T NEXUS RESET	Not supported (see SPC-4)
LOGICAL UNIT RESET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_LOGICAL_UNIT_RESET
QUERY TASK	REC ELS (see 4.9.2 and FC-LS-2)
QUERY TASK SET	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_QUERY_TASK_SET
QUERY ASYNCHRONOUS EVENT	FCP_CMND IU with TASK MANAGEMENT FLAGS field set to FCP_QUERY_ASYNCHRONOUS_EVENT

Task management functions that use the FCP_CMND IU are transmitted as the first IU in a new Exchange. A task management function that uses the FCP_CMND IU ends with an FCP_RSP IU that indicates the completion status of the function. If the addressed logical unit is not supported or is not available (e.g., not connected or not configured), then the target FCP_Port:

a) should end the Exchange with an FCP_RSP IU completion status of 09h (i.e., Task Management function incorrect logical unit number) (see table 29); or
b) may end the Exchange with an FCP_RSP IU completion status of 00h (i.e., Task Management function complete) (see table 29).

The FCP_CDB field in an FCP_CMND IU that performs task management functions is ignored.

FC-FS-3 BLSs and FC-LS-2 ELSs are used to perform the ABORT TASK task management function, to perform the QUERY TASK task management function, to recover Exchange resources, and to re-establish other initial conditions.

Table 5 specifies the SCSI Service Response mappings for an initiator FCP_Port for FCP_CMND delivered task management functions.

Table 5 - SCSI Service Response mapping for FCP_CMND delivered task management functions

SCSI Service Response	Response to the FCP_CMND IU
FUNCTION_COMPLETE	FCP_RSP IU with the RSP_CODE field set to TMF_COMPLETE
FUNCTION_REJECTED	FCP_RSP IU with the RSP_CODE field set to TMF_REJECTED
INCORRECT LOGICAL UNIT NUMBER	FCP_RSP IU with the RSP_CODE field set to TMF_INCORRECT_LUN
SERVICE DELIVERY OR TARGET FAILURE	All other responses, including: a) timeout; and b) FCP_RSP IU with the RSP_CODE field set to TMF_FAILED.

4.9.2 ABORT TASK task management function

The ABORT TASK task management function causes the device server to abort the specified command, if it exists, using ABTS-LS (see FC-FS-3). An example ABTS frame format is shown in Annex E.

The ABORT TASK task management function may not immediately release all Exchange resources, since a Recovery_Qualifier may be established to allow for the management of information that may already have been delivered to the fabric.

In addition to recovering Exchange resources that may have been left unavailable while processing task management functions, ABTS-LS is used to recover Exchange resources left in an undefined state.

Table 6 specifies the SCSI Service Response mappings for an initiator FCP_Port for the ABORT TASK task management function.

Table 6 - SCSI Service Response mapping for ABORT TASK

SCSI Service Response	Response to the ABTS-LS
FUNCTION_COMPLETE	Either: a) BA_ACC; or b) BA_RJT with Reason Code set to 03h (i.e., Logical error) and Reason Code Explanation set to 03h (i.e., Invalid OX_ID-RX_ID combination.
FUNCTION_REJECTED	BA_RJT with Reason Code set to 01h (i.e., Invalid command code) or 09h (i.e., Unable to perform command request).
INCORRECT LOGICAL UNIT NUMBER	No responses map to this SCSI Service Response.
SERVICE DELIVERY OR TARGET FAILURE	All other responses, including timeout.

4.9.3 QUERY TASK task management function

The QUERY TASK task management function maps to the REC ELS. The REC ELS shall be transmitted in a new Exchange.

Table 7 specifies the SCSI Service Response mappings for an initiator FCP_Port for the QUERY TASK task management function.

Table 7 - SCSI Service Response mapping for QUERY TASK

SCSI Service Response	Response to the REC ELS
FUNCTION_COMPLETE	Either: a) LS_ACC with E_STAT (i.e., word 5) bit 29 (i.e., Completion) set to one (i.e., complete); or b) LS_RJT with Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to perform command request) and Reason Code Explanation set to 17h (i.e., Invalid OX_ID-RX_ID combination) or 15h (i.e., Invalid Originator S_ID).
FUNCTION_SUCCEEDED	LS_ACC with E_STAT (i.e., word 5) bit 29 (i.e., Completion) set to zero (i.e., open).
FUNCTION_REJECTED	LS_RJT with Reason Code set to: a) 01h (i.e., Invalid ELS command code); or b) 0Bh (i.e., Command not supported).
INCORRECT LOGICAL UNIT NUMBER	No responses map to this SCSI Service Response.
SERVICE DELIVERY OR TARGET FAILURE	All other responses, including: a) timeout; and b) LS_RJT with Reason Code set to 09h (i.e., Unable to perform command request) and Reason Code Explanation set to 15h (i.e., Invalid Originator S_ID).

4.10 Clearing effects of task management, FCP, FC-FS-3, FC-LS-2, and FC-AL-2 actions

Table 8 and table 9 summarize the clearing effects resulting from Fibre Channel link actions and SCSI operations, respectively. The clearing effects are applicable only to Sequences and Exchanges associated with Fibre Channel Protocol actions. Sequences and Exchanges associated with other actions follow rules specified in FC-FS-3 or other relevant protocol standards. Rows indicating a clearing effect for all initiator FCP_Ports have the specified clearing effect on all initiator FCP_Ports, regardless of the link that attaches the initiator FCP_Port to the target FCP_Port.

Table 8 - Clearing effects of link related actions

Clearing effect	FC link action					
	Target Power Cycle	Reset LIP(y,x)[b]	LOGO ELS[e], PLOGI ELS	PRLI ELS[d], PRLO ELS[e],	TPRLO ELS[c]	ABTS-LS	ABTS (Sequence)
PLOGI ELS parameters set to default values (see FC-LS-2)							
For all logged-in initiator FCP_Ports	Y	Y	N	N	N	N	N
Only for initiator FCP_Port associated with the action	-	-	Y	N	N	N	N
Open FCP Exchanges terminated							
For all initiator FCP_Ports	Y	Y	N	N	Y	N	N
Only for initiator FCP_Port associated with the action	-	-	Y	Y	-	N	N
Only for FCP Exchange associated with ABTS	-	-	-	-	-	Y	-
FCP Sequence associated with ABTS terminated	-	-	-	-	-	-	Y
Login BB_Credit_CNT set to login value (see FC-FS-3)							
For all Logged-In NL_Ports	Y	Y	N	N	N	N	N
For transmitting NL_Port only	-	-	Y	N	N	N	N
Hard address acquisition attempted by NL_Port	Y[a]	Y[a]	N	N	N	N	N
Process Login parameters cleared[f]							
For all logged-in initiator FCP_Ports	Y	Y	N	N	Y	N	N
Only for FCP_Port associated with the action	-	-	Y	Y	-	N	N
CRN set to one							
For all initiator FCP_Ports	Y	Y	N	N	Y	N	N
Only for initiator FCP_Port associated with the action	-	-	Y	Y	-	N	N

Key:
"Y" indicates the clearing effect upon successful completion of the specified action.
"N" indicates the clearing effect is not performed by the specified action.
"-" indicates the clearing effect is not applicable.

a) If the NL_Port has an AL_PA different than its hard address and the NL_Port experiences a power cycle or recognizes LIP(AL_PD,AL_PS), then the NL_Port shall relinquish its current AL_PA and attempt to acquire its hard address.

b) This is also known as LIP(AL_PD,AL_PS). If the destination recognizes a selective hard reset LIP where the AL_PD matches the AL_PA of the receiving NL_Port, then the receiving NL_Port shall perform the behavior described in this column.

c) For a TPRLO ELS, the actions listed shall be performed when the GLOBAL bit is set to one. If the GLOBAL bit is set to zero, then the actions listed under PRLI ELS/PRLO ELS shall be performed for the designated initiator FCP_Port. See FC-FS-3.

d) The target FCP_Port shall clear the object only if ESTABLISH IMAGE PAIR is set to one and if the referenced image pair is FCP type. See 6.2.

e) Logout and Process Logout may be either implicit or explicit. Implicit logout and Process Logout are specified in FC-FS-3.

f) A target FCP_Port should transmit a PRLO ELS to all logged-in initiator FCP_Ports that are logged out as a result of processing a TPRLO ELS with the GLOBAL bit set to one. The PRLO ELS(s) may be transmitted before or after transmitting the LS_ACC for the TPRLO ELS.

Table 9 - Clearing effects of initiator FCP_Port actions

Clearing effect	Initiator FCP_Port action		
	LOGICAL UNIT RESET[b]	CLEAR TASK SET[b]	ABORT TASK SET[b]
PLOGI ELS parameters set to default values (see FC-LS-2)			
For all logged-in initiator FCP_Ports	N	N	N
Only for initiator FCP_Port associated with the action	N	N	N
Open FCP Sequences Terminated			
For all initiator FCP_Ports with open FCP Sequences	Y[a]	Y[a]	N
Only for initiator FCP_Port associated with the action	-	-	Y[a]
Only for FCP Sequences associated with Aborted FCP Exchanges	-	-	-
Login BB_Credit_CNT set to login value (see FC-FS-3)			
For all Logged-In NL_Ports	N	N	N
For transmitting NL_Port only	N	N	N
Hard address acquisition attempted by NL_Port	N	N	N
Process Login parameters cleared			
For all logged-in initiator FCP_Ports	N	N	N
Only for FCP_Port associated with the action	N	N	N
CRN set to one			
For all initiator FCP_Port	Y	Y	N
Only for initiator FCP_Port associated with the action	-	-	Y

Key:
"Y" indicates the clearing effect upon successful completion of the specified action.
"N" indicates the clearing effect is not performed by the specified action.
"-" indicates the clearing effect is not applicable.

a) Exchanges are cleared internally within the target FCP_Port, but open FCP Sequences shall be individually aborted by the initiator FCP_Port using ABTS-LS. This has the effect of aborting the associated FCP Exchange. See 12.3.
b) For multiple-logical unit SCSI target devices, CLEAR TASK SET, ABORT TASK SET, and LOGICAL UNIT RESET affect only the addressed logical unit.

4.11 I_T nexus loss notification events

An FCP_Port shall deliver an I_T nexus loss notification (see SAM-5) as a consequence of the following events:

a) transmitting or receiving a LOGO ELS (explicit or implicit);
b) transmitting or receiving a PRLO ELS (explicit or implicit);
c) receiving a TPRLO ELS;
d) transmitting a TPRLO ELS with a Third Party Originator N_Port_ID (see FC-LS-2) that matches the N_Port_ID of the transmitting FCP_Port; or
e) transmitting a TPRLO ELS with the GLOBAL bit set to one to a target FCP_Port that has an I_T nexus with the transmitting initiator FCP_Port.

4.12 Transport Reset notification events

If the AL_PD matches the AL_PA of the receiving NL_Port, then an NL_Port shall deliver a Transport Reset notification (see SAM-5) for a Reset LIP(y,x) (see FC-AL-2) FC link event.

4.13 Port Login/Logout

The N_Port Login (PLOGI) ELS is optionally used to establish the Fibre Channel operating parameters between any two Fibre Channel ports, including FCP_Ports. Implicit login functions are allowed.

If a target FCP_Port receives a PLOGI ELS request and it finds there are not enough login resources to complete the login, then the target FCP_Port shall respond to the PLOGI ELS with LS_RJT and Reason Code "Unable to perform command request" and Reason Code Explanation "Insufficient resources to support Login" as defined in FC-LS-2. By means outside the scope of this standard, the target FCP_Port may select another initiator FCP_Port and release some login resources by performing an explicit logout of the other initiator FCP_Port, thus freeing resources for a future PLOGI ELS.

4.14 Process Login and Process Logout

The Process Login (PRLI) ELS request is used to establish the FCP operating relationships between two FCP_Ports (see 6.3). The Process Logout (PRLO) ELS request is used to de-establish the FCP operating relationships between two FCP_Ports (see 6.4). Implicit Process Login and Process Logout parameters may be defined for FCP_Ports. Such definitions are outside the scope of this standard.

4.15 Link management

FC-FS-3 allows management protocols above the FC-FS-3 interface to perform link data functions. The standard primitive sequences, link management protocols, BLSs, and ELSs are used as required by FCP devices (see FC-FS-3 and FC-LS-2).

4.16 FCP addressing and Exchange identification

The address of each FCP_Port is defined by its address identifier as described in FC-FS-3.

Each FCP I/O operation is identified by the FCP I/O operation's fully qualified exchange identifier (FQXID). The FQXID is composed of the initiator port identifier, the target port identifier, the OX_ID field value, and the RX_ID field value. Other definitions of FQXID are outside the scope of this standard. The method used to identify FCP I/O operations internal to the application client and the device server is not defined by this standard.

Logical unit numbers are contained in the FCP_LUN field of FCP_CMD IUs. Subsequent identification of the FCP I/O operation and the Exchange that carries the protocol interactions for the FCP I/O operation uses the FQXID.

The target FCP_Port uses the OX_ID field value, and, if it has been assigned, the RX_ID field value to perform error recovery and task management functions. The task retry identifier is used as a supplemental command identifier if task retry identification is supported and enabled.

4.17 Use of Worldwide_Names

As specified in FC-FS-3, each Fibre Channel node shall have a Node_Name that is a Worldwide_Name and each Fibre Channel port shall have an N_Port_Name that is a Worldwide_Name. The Worldwide_Name shall be a unique name using one of the formats defined by FC-FS-3. See Annex A for a description of the mapping of FCP-4 terminology to SAM-5 terminology.

Each target FCP_Port and its associated logical units has knowledge of the N_Port_Name of each initiator FCP_Port through the Fibre Channel login process. As a result, the relationship between the address identifier of the initiator FCP_Port and a persistent reservation for a logical unit may be adjusted (see SPC-4) during those reconfiguration events that may change the address identifier of the initiator FCP_Port.

If a target FCP_Port receives a PRLI ELS or a PLOGI ELS from an initiator FCP_Port with a previously known N_Port_Name, but with a changed initiator port identifier, then the device server shall:

a) assign the new initiator port identifier to the associated I_T nexus;
b) maintain the existing registration on that I_T nexus; and
c) set the reservation holder to the I_T nexus having the same N_Port_Name.

Each logical unit shall be able to present a Worldwide_Name through the INQUIRY command Device Identification VPD page (see SPC-4). For devices compliant with this standard and having a LUN 0, the Worldwide_Name of the logical unit having a LUN of 0 may be the same as the Node_Name of the SCSI target device. The Worldwide_Name for the FCP_Port shall be different from the Worldwide_Name for the node (i.e., the N_Port_Name shall be different than the Node_Name).

Each FCP device should include a SCSI device name in NAA IEEE Registered format (see SPC-4). If the FCP device includes a Platform Name (see FC-GS-6), then the Platform Name should be the same as the SCSI device name.

In the Device Identification VPD page, a device server in an FCP target device that implements a SCSI device name:

a) shall report the SCSI device name in binary NAA format; and
b) should report the SCSI device name in SCSI name string format (e.g., "naa." followed by 16 hexadecimal digits followed by 4 ASCII null characters).

5 FC-FS-3 frame header

5.1 FC-FS-3 frame header overview

The format of the standard FC-FS-3 header as used by the Fibre Channel Protocol is defined in table 10.

Table 10 - FCP frame header

Bits Word	31–24	23–16	15–08	07–00
0	R_CTL	D_ID		
1	CS_CTL	S_ID		
2	TYPE	F_CTL		
3	SEQ_ID	DF_CTL	SEQ_CNT	
4	OX_ID		RX_ID	
5	PARAMETER			

All fields in the FCP frame header use the standard FC-FS-3 definitions. The following explanations of the fields provide information about the use of those fields to implement FCP functionality.

5.2 FC-FS-3 frame header fields

5.2.1 R_CTL field

The R_CTL field is subdivided into a ROUTING field and an INFORMATION field (see FC-FS-3). The ROUTING field shall be set to 0h (i.e., Device_Data) and the INFORMATION field shall be set to the value specified in table 21 and table 22.

5.2.2 D_ID field

The value in the D_ID field is the D_ID of the frame. For FCP FC-4 Device_Data frames, the D_ID transmitted by the Exchange Originator is the address identifier of the target FCP_Port. The D_ID transmitted by the Exchange Responder is the address identifier of the initiator FCP_Port.

5.2.3 CS_CTL field

The values in the CS_CTL field are defined by FC-FS-3 for class specific control information and do not interact with the Fibre Channel Protocol.

5.2.4 S_ID field

The value in the S_ID field is the S_ID of the frame. For FCP FC-4 Device_Data frames, the S_ID transmitted by the Exchange Originator is the address identifier of the initiator FCP_Port. The S_ID transmitted by the Exchange Responder is the address identifier of the target FCP_Port.

5.2.5 TYPE field

The value in the TYPE field shall be 08h (i.e., Fibre Channel Protocol) (see FC-FS-3) for all frames of SCSI FCP Exchanges.

A TYPE field value of 08h is also used for Information Units specified in FC-NVMe and FC-SB-6.

5.2.6 F_CTL field

The bits in the F_CTL field manage the beginning and normal or abnormal termination of Sequences and Exchanges. The bits and definitions shall be as defined by FC-FS-3. See 5.2.12.

5.2.7 SEQ_ID field

The value in the SEQ_ID field identifies each Sequence between a particular Exchange Originator and Exchange Responder with a unique value as defined by FC-FS-3.

5.2.8 DF_CTL field

The bits in the DF_CTL field indicate any optional headers that may be present. The DF_CTL field shall be set to 00h (i.e., no optional headers) or 40h (i.e., Encapsulating Security Payload).

5.2.9 SEQ_CNT field

The value in the SEQ_CNT field indicates the frame order within the Sequence as defined by FC-FS-3.

5.2.10 OX_ID field

The value in the OX_ID field is the Originator Exchange Identifier and is one of the identifiers contained in the FQXID. The OX_ID field shall be assigned and shall have a value other than FFFFh.

5.2.11 RX_ID field

The value in the RX_ID field is the Responder Exchange Identifier and is one of the identifiers contained in the FQXID. The RX_ID field shall have the unassigned value of FFFFh until the Exchange Responder assigns a different value in its response to the Exchange Originator as specified in FC-FS-3. The Exchange Originator shall use the value assigned by the Exchange Responder for subsequent frames.

5.2.12 PARAMETER field

The PARAMETER field has two definitions for Device_Data frames with the TYPE field set to 08h (i.e., Fibre Channel Protocol).

For a frame with the R_CTL field set to 01h (i.e., solicited data) (i.e., an FCP_DATA IU) (see 9.1 and 9.4), the PARAMETER field shall contain a relative offset. The RELATIVE OFFSET PRESENT bit of the F_CTL field shall be set to one, indicating that the PARAMETER field value is a relative offset. The relative offset is the application client buffer offset as described by SAM-5. The relative offset shall have a value that is a multiple of 4 (i.e., each frame of each FCP_DATA IU shall begin on a word boundary).

For a frame with the R_CTL field set to 02h (i.e., unsolicited control) (i.e., an FCP_CMND IU) (see 9.1 and 9.2), the PARAMETER field value depends on whether task retry identification (see 4.7) is active. If the target FCP_Port and initiator FCP_Port have agreed upon performing task retry identification, then the PARAMETER field shall contain the task retry identifier. If the target FCP_Port and initiator FCP_Port have not agreed upon performing task retry identification, then the PARAMETER field shall be set to zero. In both cases, the RELATIVE OFFSET PRESENT bit of the F_CTL field shall be set to zero.

For a frame with the R_CTL field other than 01h and 02h, the RELATIVE OFFSET PRESENT bit of the F_CTL field shall be set to zero and the PARAMETER field shall contain a value of zero.

For FCP FC-4 Link Service frames, the PARAMETER field is specified in the description of the individual link services.

6 FCP link service definitions

6.1 Overview of link service requirements

The FCP link-level protocol includes the BLSs (see FC-FS-3) and ELSs defined by FC-LS-2. The protocol also includes the PRLI ELS and PRLO ELS specified in FC-LS-2, the PRLI FCP Service Parameter pages specified in 6.3, and the REC ELS with usage as defined in this standard.

Link-level protocols are used to configure the FC environment, including the establishment of configuration information and address information. FCP devices introduced into a configuration or modifications in the addressing or routing of the configuration may require the login and discovery procedures to be performed again. Annex D includes a model for the discovery procedure.

6.2 Overview of Process Login and Process Logout

Process Login allows for a process at one FCP_Port to be related to a corresponding process at another FCP_Port as an image pair. In addition, the PRLI ELS allows one or more FC-4 capabilities to be reported by the initiating FCP_Port to the recipient FCP_Port. The recipient FCP_Port indicates its acceptance or rejection of the capabilities in its response to the PRLI ELS request.

Since implicit login may be established by configuration conventions outside the scope of this standard, Process Login is optional except in the case where an initiator FCP_Port is not using implicit login and is operating in a point-to-point topology. In this case, the initiator FCP_Port shall always transmit an explicit PRLI ELS.

NOTE 1 - The requirement to transmit a PRLI ELS for an initiator FCP_Port that is not using implicit login and operating in a point-to-point topology is to remove a deadlock condition that occurs when the target FCP_Port N_Port_Name is larger than the initiator FCP_Port N_Port_Name. In this case the target FCP_Port PLOGI ELS request is processed, but the target FCP_Port is prohibited from transmitting a PRLI ELS. If the initiator FCP_Port does not transmit a PRLI ELS, then a deadlock occurs.

PRLI ELS requests shall only be initiated by devices having the initiator FCP_Port capability. Devices having only target FCP_Port capability shall not perform a PRLI ELS request.

An initiator FCP_Port shall have successfully completed Process Login with a target FCP_Port that establishes an image pair before any FCP IUs are exchanged. An image pair may also be established by an implicit Process Login performed by methods outside the scope of this standard. An image pair is removed by an implicit Process Logout or explicit PRLO ELS (see 6.4). If an image pair is not established by an initiator FCP_Port to a target FCP_Port, then the initiator FCP_Port and target FCP_Port shall not exchange any FCP IUs. Any FCP IUs received by a target FCP_Port from an Nx_Port that does not have an image pair with that target FCP_Port shall be discarded. In addition, a target FCP_Port that receives an FCP_CMND IU from an Nx_Port that has successfully completed PLOGI ELS, but does not have an image pair with that target FCP_Port, shall discard the FCP_CMND IU and respond with an explicit PRLO ELS (see 12.6). Reasons why the Nx_Port does not have an image pair with the target FCP_Port include:

a) the Nx_Port has not established an image pair with that target FCP_Port;
b) the target FCP_Port performed an implicit Process Logout of the Nx_Port; or
c) the target FCP_Port processed a TPRLO ELS that effected the Nx_Port.

The creation of image pairs behind an FCP_Port has no effect on the Fibre Channel Protocol.

The FC-4 Service Parameter pages for the Fibre Channel Protocol are defined in 6.3.3 and 6.3.4.

Processing of a PRLI ELS or PRLO ELS request performs the clearing actions defined in 4.10.

Process Login has two actions that may be performed, selected by the ESTABLISH IMAGE PAIR bit (see 6.3.3):

a) informative action - service parameter information is exchanged during the Process Login enabling subsequent negotiation for image pair establishment; or
b) binding action - service parameter information is exchanged that establishes an image pair relationship between processes in the communicating Nx_Ports. The relationship does not allow any communication types or paths other than those established by the PRLI ELS.

6.3 PRLI ELS

6.3.1 Use of PRLI ELS by the Fibre Channel Protocol

The PRLI ELS request is transmitted from an Originator FCP_Port to a Responder FCP_Port to identify to the destination the capabilities that the Originator FCP_Port expects to use with the Responder FCP_Port and to determine the capabilities of the Responder (see FC-LS).

If the PRLI ELS is requesting an informative action by setting the ESTABLISH IMAGE PAIR bit to zero, then the PRLI ELS accept reports the capabilities of the Responder to the Originator.

If the PRLI ELS is requesting a binding action by setting the ESTABLISH IMAGE PAIR bit to one, then the PRLI ELS accept reports the capabilities of the Responder to the Originator and establishes an image pair. An image pair shall be established only if the FCP devices have complementary initiator FCP_Port and target FCP_Port capabilities. If both FCP devices have both initiator FCP_Port and target FCP_Port capabilities, then a single image pair allows both initiator FCP_Ports to access the complementary target FCP_Port capabilities of the other device in the pair. Some capabilities require support by both the Originator and Responder before they may be used (see 6.3.3). If the IMAGE PAIR ESTABLISHED bit is set to one in the PRLI ELS accept, then an image pair was successfully established. If the IMAGE PAIR ESTABLISHED bit is set to zero in the PRLI ELS accept, then an image pair was not successfully established.

An accept response code indicating other than REQUEST EXECUTED (see 6.3.4 and FC-LS-2) shall be provided if the PRLI ELS FCP Service Parameter page is incorrect or if a requested image pair is not established.

A Link Service Reject (LS_RJT) indicates that the PRLI ELS request is not supported or is incorrectly formatted.

The PRLI ELS common service parameters and accept response codes are defined in FC-LS. FC-4 service parameters for mappings other than the Fibre Channel Protocol are defined in other FC-4 standards.

6.3.2 New or repeated Process Login

After the completion of any new or repeated binding Process Login, all clearing actions specified in 4.10 shall be performed.

After the completion of any new or repeated informative Process Login, the state of the Originator and Responder remains unchanged.

FCP devices may have default Process Login information provided in a manner outside the scope of this standard. Such devices do not require the processing of a PRLI ELS to perform normal FCP I/O operations. If default Process Login information is sufficient for login (i.e., PLOGI ELS) to perform an implicit Process Login, then the PLOGI ELS shall perform the same clearing actions and establish the same Unit Attention condition that is normally performed and established by Process Login.

6.3.3 PRLI ELS request FCP Service Parameter page format

The FCP Service Parameter page for the PRLI ELS request is shown in table 11.

Table 11 - PRLI ELS request FCP Service Parameter page

Item	Word	Bit(s)
TYPE CODE (08h for this standard)	0	31–24
TYPE CODE EXTENSION	0	23–16
Obsolete	0	15
Obsolete	0	14
ESTABLISH IMAGE PAIR	0	13
Reserved	0	12–0
Obsolete	1	31–0
Obsolete	2	31–0
Service Parameters		
Reserved	3	31–12
ENHANCED DISCOVERY	3	11
REC SUPPORT	3	10
TASK RETRY IDENTIFICATION REQUESTED	3	9
RETRY	3	8
CONFIRMED COMPLETION ALLOWED	3	7
DATA OVERLAY ALLOWED	3	6
INITIATOR FUNCTION	3	5
TARGET FUNCTION	3	4
Obsolete	3	3
Obsolete	3	2
READ FCP_XFER_RDY DISABLED (shall be set to one)	3	1
WRITE FCP_XFER_RDY DISABLED	3	0

Word 0, Bits 31–24: TYPE CODE: The value of 08h in this byte indicates that this Service Parameter page is defined for the Fibre Channel Protocol (see FC-FS-3).

Word 0, Bit 13: ESTABLISH IMAGE PAIR: If the ESTABLISH IMAGE PAIR bit is set to zero, then the PRLI ELS only exchanges service parameters as defined in FC-LS-2. If the ESTABLISH IMAGE PAIR bit is set to one, then the PRLI ELS exchanges service parameters and attempts to establish an image pair as defined in FC-LS-2.

Word 3, Bit 11: ENHANCED DISCOVERY: If the ESTABLISH IMAGE PAIR bit is set to zero, then the ENHANCED DISCOVERY bit shall be ignored. If the ENHANCED DISCOVERY bit is set to one, then the Originator is requesting that an image pair be established only if the initiator FCP_Port has been authorized to access one or more logical units, not including management logical units (see SAM-5), that are addressed through the target FCP_Port. If the initiator FCP_Port is not authorized to access one or more logical units, not including management logical units, then the image pair shall not be established, the PRLI ELS accept response code shall be set to THE EXCHANGE RECIPIENT HAS A PREDEFINED CONFIGURATION THAT PRECLUDES ESTABLISHING THIS IMAGE PAIR (i.e., 0101b), and the IMAGE PAIR ESTABLISHED bit shall be set to zero. If the ENHANCED DISCOVERY bit is set to zero by either the Originator or the Responder to the PRLI ELS, then accessibility of logical units shall not affect establishment of the image pair.

Word 3, Bit 10: REC SUPPORT: If the REC ELS supported (REC SUPPORT) bit is set to one, then the Originator is indicating that it supports the transmission of the REC ELS when it is acting as an initiator FCP_Port.

The capability of the initiator FCP_Port to retransmit unsuccessfully transmitted data is determined by the RETRY bit (i.e., a REC SUPPORT bit set to one does not indicate the initiator FCP_Port supports retransmission of data). If the REC SUPPORT bit is set to zero, then the Originator is providing no information about whether it supports transmission of the REC ELS.

Word 3, Bit 9: TASK RETRY IDENTIFICATION REQUESTED: If the TASK RETRY IDENTIFICATION REQUESTED bit is set to one, the Originator of the PRLI ELS requests that task retry identification (see 4.7) be used. If both the Originator of the PRLI ELS and the Responder to the PRLI ELS request that task retry identification be used, then it shall be used between the initiator FCP_Port and all logical units addressed through the target FCP_Port. The PARAMETER field for each FCP_CMND IU shall be set to a unique non-zero value. The PARAMETER field for any REC ELS request or SRR FCP_LS request for that command shall be set to the same value.

If the TASK RETRY IDENTIFICATION REQUESTED bit is set to zero by either the Originator of or the Responder to the PRLI ELS, then task retry identification shall not be used. The PARAMETER fields for FCP_CMND IUs, for REC ELS requests, and for SRR FCP_LS requests shall be zero.

Word 3, Bit 8: RETRY: If the RETRY bit is set to one, then the Originator or Responder is indicating that its initiator FCP_Port functions support the capability of requesting a retransmission of unsuccessfully transmitted data or that its target FCP_Port functions support the capability of performing a requested retransmission. If the RETRY bit is set to zero, then the Originator or Responder is indicating that it does not support the capability of requesting or performing retransmissions of unsuccessfully transmitted data.

If the process has both initiator FCP_Port and target FCP_Port capabilities, then the RETRY bit shall apply to both. The SRR FCP_LS request may be both transmitted by and accepted by the process.

An initiator FCP_Port and target FCP_Port shall use the retransmission capability only if the RETRY bit is set to one in both the request payload and in the accept payload. If the RETRY bit is set to zero in either the request payload or the accept payload, then the SRR FCP_LS request shall not be transmitted by the initiator FCP_Port. If an SRR FCP_LS request is received by a target FCP_Port that has set the RETRY bit to zero, then the SRR FCP_LS request shall be rejected with an FCP_RJT.

If the image pair is allowed to use the retransmission capability, then overlay of data as defined for retransmission shall be allowed regardless of the state of the DATA OVERLAY ALLOWED bit.

Word 3, Bit 7: CONFIRMED COMPLETION ALLOWED: If the CONFIRMED COMPLETION ALLOWED bit is set to one, then the Originator's or Responder's initiator FCP_Port function has the capability of supporting confirmed completion. If the CONFIRMED COMPLETION ALLOWED bit is set to zero, then the initiator FCP_Port function does not have the capability of supporting confirmed completion. The CONFIRMED COMPLETION ALLOWED bit shall be zero for FCP devices having only target FCP_Port function. If the initiator FCP_Port function supports confirmed completion, then the target FCP_Port may request an FCP_CONF IU by setting the FCP_CONF_REQ bit to one as specified by 4.5. If the initiator FCP_Port function does not have the capability of supporting confirmed completion, then the target FCP_Port shall not set the FCP_CONF_REQ bit to one.

Word 3, Bit 6: DATA OVERLAY ALLOWED: If the DATA OVERLAY ALLOWED bit is set to one, then the Originator or Responder is indicating that its initiator FCP_Port function has the capability of supporting data overlay. If the DATA OVERLAY ALLOWED bit is set to zero, then the initiator FCP_Port function does not have the capability of performing data overlay. The DATA OVERLAY ALLOWED bit shall be zero for FCP devices having only target FCP_Port function. If the initiator FCP_Port function supports data overlay, then the target FCP_Port may perform random buffer access that performs a transfer to or from the same offset in the application client buffer more than once during processing of a command.

Data transmission requested by the initiator FCP_Port during the optional retry procedures defined by this standard is managed by the initiator FCP_Port. Such data retransmissions are not considered data overlays, even if retransmission occurs to the same offset in the application client buffer.

Word 3, Bit 5: INITIATOR FUNCTION: If the INITIATOR FUNCTION bit is set to one, then the Originator or Responder is indicating it has the capability of operating as an initiator FCP_Port. If the INITIATOR FUNCTION bit is set to zero, then the Originator or Responder does not have the capability of operating as an initiator FCP_Port.

Word 3, Bit 4: TARGET FUNCTION: If the TARGET FUNCTION bit is set to one, then the Originator or Responder is indicating that it has the capability of operating as a target FCP_Port. If the TARGET FUNCTION bit is set to zero, then the Originator or Responder does not have the capability of operating as a target FCP_Port. Both the INITIATOR FUNCTION bit and the TARGET FUNCTION bit may be set to one. If neither the INITIATOR FUNCTION bit nor the TARGET FUNCTION bit is set to one, then the service parameters for the FCP Service Parameter page are assumed to be invalid. A Responder receiving such an invalid FCP Service Parameter page shall notify the Originator with a PRLI ELS accept response code of SERVICE PARAMETERS ARE INVALID and the IMAGE PAIR ESTABLISHED bit set to zero. An Originator receiving such an invalid FCP Service Parameter page shall not perform Fibre Channel Protocol operations with the Responder.

Word 3, Bit 1: READ FCP_XFER_RDY DISABLED: The READ FCP_XFER_RDY DISABLED bit shall be set to one. Target FCP_Ports shall not transmit FCP_XFER_RDY on read operations.

Word 3, Bit 0: WRITE FCP_XFER_RDY DISABLED: If the WRITE FCP_XFER_RDY DISABLED bit is set to zero, then FCP_XFER_RDY IUs shall be transmitted by the target FCP_Port to request each of the SCSI write FCP_DATA IUs from the initiator FCP_Port. If the WRITE FCP_XFER_RDY DISABLED bit is set to one, then FCP_XFER_RDY IUs shall not be used before the first FCP_DATA IU to be transferred in the write operation. If both the Originator and Responder choose to disable write FCP_XFER_RDY IUs, then all write operations between the FCP_Ports shall operate without using the FCP_XFER_RDY IU before the first FCP_DATA IU. The FCP_XFER_RDY IU shall be transmitted to request each additional FCP_DATA IU, if any. If either the Originator or the Responder requires the use of FCP_XFER_RDY IUs during SCSI writes, then the Exchange Responder shall transmit an FCP_XFER_RDY IU requesting each FCP_DATA IU, including the first, from the Exchange Originator.

6.3.4 PRLI ELS accept FCP Service Parameter page format

The FCP Service Parameter page for the PRLI ELS accept is shown in table 12.

Table 12 - PRLI ELS accept FCP Service Parameter page

Item	Word	Bit(s)
TYPE CODE (08h for this standard)	0	31–24
TYPE CODE EXTENSION	0	23–16
Obsolete	0	15
Obsolete	0	14
IMAGE PAIR ESTABLISHED	0	13
Reserved	0	12
RESPONSE CODE	0	11–8
Reserved	0	7–0
Obsolete	1	31–0
Obsolete	2	31–0
Service Parameters		
Reserved	3	31–12
ENHANCED DISCOVERY	3	11
REC SUPPORT	3	10
TASK RETRY IDENTIFICATION REQUESTED	3	9
RETRY	3	8
CONFIRMED COMPLETION ALLOWED	3	7
DATA OVERLAY ALLOWED	3	6
INITIATOR FUNCTION	3	5
TARGET FUNCTION	3	4
Obsolete	3	3
Obsolete	3	2
READ FCP_XFER_RDY DISABLED (shall be one)	3	1
WRITE FCP_XFER_RDY DISABLED	3	0

With the following exceptions, the service parameter definitions are identical for the PRLI ELS request (see table 11) and accept FCP Service Parameter pages.

Word 0, Bit 13: IMAGE PAIR ESTABLISHED: The IMAGE PAIR ESTABLISHED bit is defined in FC-LS-2. If the IMAGE PAIR ESTABLISHED bit is set to zero, then the image pair was not established. The RESPONSE CODE field has additional information. If the IMAGE PAIR ESTABLISHED bit is set to one, then the image pair was established.

WORD 0, BITS 11-8: RESPONSE CODE: The RESPONSE CODE field is defined in FC-LS-2. The values of the RESPONSE CODE field indicate whether the image pair was successfully created. If the image pair was not created, then the value of the RESPONSE CODE field indicates why the request failed or was rejected.

Word 3, Bit 11: ENHANCED DISCOVERY: If the ENHANCED DISCOVERY bit is set to one, then the Responder is indicating that it supports enhanced discovery (i.e., an image pair is established only if the initiator FCP_Port is authorized to access logical units, other than management logical units (see SAM-5), that are addressed through the target FCP_Port). If the ENHANCED DISCOVERY bit is set to zero, then the Responder is indicating that it does not support enhanced discovery when it is acting as a target FCP_Port.

Word 3, Bit 10: REC SUPPORT: If the REC ELS supported (REC SUPPORT) bit is set to one, then the Responder is indicating that it supports the receipt of the REC ELS, when it is acting as a target FCP_Port. The capability of the target FCP_Port to retransmit unsuccessfully transmitted data is determined by the RETRY bit (i.e., a REC SUPPORT bit set to one does not indicate the target FCP_Port supports retransmission of data). If the REC SUPPORT bit is set to zero, then the Responder is indicating that it may not support receipt of the REC ELS.

6.4 PRLO ELS

The format for the PRLO ELS request and PRLO ELS accept is specified in FC-LS-2.

The PRLO ELS request is transmitted from an Originator FCP_Port to a Responder FCP_Port to indicate to the Responder that the image pair specified in the FCP Service Parameter pages of the PRLO ELS is being discontinued by the Originator. If the PRLO ELS logs out the image pair between an initiator FCP_Port and a target FCP_Port, then all clearing actions specified in 4.10 shall be performed and an I_T nexus loss notification shall be delivered (see 4.11).

The PRLO ELS accept is returned to the Originator FCP_Port to indicate that the Responder FCP_Port recognizes that the image pair is being discontinued. The accept shall present a response FCP Service Parameter page for the request FCP Service Parameter page. It is not an error to perform Process Logout for an image pair that does not exist.

A Link Service Reject (LS_RJT) indicates that the PRLO ELS request is invalid and not accepted.

After Process Logout, no further Fibre Channel Protocol communication is possible between those Nx_Ports.

The PRLO ELS accept response codes are defined in FC-LS-2.

6.5 Read Exchange Concise (REC) ELS

See FC-LS-2 for a description of the REC ELS. FCP-4 specific usage of the REC ELS is as follows:

a) if task retry identification is active for the Originator and the Responder, the PARAMETER field of the request Sequence shall contain the task retry identifier for the command specified by the OX_ID field value and RX_ID field value;
b) if the destination FCP_Port of the REC ELS request determines that the ORIGINATOR S_ID, OX_ID, or RX_ID fields, or task retry identifier are inconsistent, then it shall respond with an LS_RJT Sequence with the Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to process command request), and the Reason Code Explanation set to 17h (i.e., Invalid OX_ID-RX_ID combination);
c) the REC ELS shall be sent in a new Exchange. The Exchange shall be ended by the response to the REC ELS;
d) if the RX_ID field value in the REC ELS request payload was FFFFh, then the RX_ID field value in the REC ELS LS_ACC payload may be set to the value selected by the Responder when the first frame of the Exchange was received; and
e) the FC4VALUE field in the REC ELS LS_ACC payload shall be set to:
   A) for a write comand, the number of bytes successfully received by the device server. Data that has been retransmitted or overlaid shall be counted only once;
   B) for a read command, the number of bytes transmitted by the target FCP_Port. Data that has been retransmitted or overlaid shall be counted only once;
   C) for a non-data command, 00000000h;

D) for a bidirectional command, 00000000h; and

E) for any other Exchange specified in this standard (e.g., a task management function, an SRR ELS or another REC ELS), 00000000h.

7 FC-4 specific Name Server registration and objects

7.1 Overview of FC-4 specific objects for the Fibre Channel Protocol

The Name Server for a Fibre Channel fabric is defined by FC-GS-6. FCP specific objects are defined in this clause for use by the Name Server. FC-GS-6 provides complete descriptions of the operations that are performed to register objects with a Name Server and to query the Name Server for the value of the objects.

7.2 FC-4 TYPEs object

The FC-4 TYPEs object (see FC-GS-6) indicates a set of supported data structure type values for Device_Data and FC-4 Link_Data frames (see FC-FS-3).

An FCP_Port shall register the FCP TYPE (08h) with the Name Server using the RFT_ID request CT_IU. This registration shall precede registration of the FC-4 TYPE 08h FC-4 Features object.

If additional FC-4 Features (see 7.3) are to be registered, then an FCP_Port shall register the additional FCP Features TYPE (0Ah) with the Name Server using the RFT_ID request CT_IU. This registration shall precede registration of the additional FCP Features TYPE 0Ah FC-4 Features object.

7.3 FC-4 Features object

The FC-4 Features object (see FC-GS-6) defines a 4-bit field for each FC TYPE code. The FC-4 Features object is a 32-word array of 4-bit values. The 4-bit FC-4 Feature field for FCP TYPE 08h is inserted in bits 3 - 0 of word 1. The format of the 4-bit FC-4 Feature field for FCP TYPE 08h is shown in table 13.

Table 13 - FCP TYPE 08h definition of FC-4 Feature bits

Word 1 bit	Description of bit
3	Additional FCP Features
2	Contains one or more peripheral devices that are not peripheral device type 00h (i.e., direct access block device).
1	FCP initiator function supported
0	FCP target function supported

If bit 3 of word 1 is set to one, then additional FCP FC-4 Features may be registered for additional FCP Features TYPE 0Ah (see table 14). If bit 3 of word 1 is set to zero, then no additional FCP FC-4 Features are registered for additional FCP Features TYPE 0Ah.

The 4-bit FC-4 Feature field for additional FCP Features TYPE 0Ah is inserted in bits 3 - 0 of word 6. The format of the 4-bit FC-4 Feature field for additional FCP Features TYPE 0Ah is shown in table 14.

Table 14 - Additional FCP Features TYPE 0Ah definition of FC-4 Feature bits

Word 1 bit	Description of bit
3	Reserved
2	Reserved
1	Reserved
0	Reserved

An FCP_Port shall register the FCP TYPE 08h FC-4 Features object with a Name Server using the RFF_ID request CT_IU. An FCP_Port may register the additional FCP Features TYPE 0Ah FC-4 Features object with a Name Server using the RFF_ID request CT_IU.

The FC-4 Features object may be obtained by any Nx_Port from a Name Server using a GFF_ID request CT_IU, which requests the FC-4 Features object for a specified Port Identifier (see FC-GS-6). The FC-4 Features object is provided in the GFF_ID accept CT_IU.

A list of all the Port Identifiers matching the domain and area addressing and a specified FC-4 Features object may be obtained by any Nx_Port from a Name Server using the GID_FF request CT_IU. The FC-4 Features object is a parameter in the GID_FF request CT_IU.

8 FCP FC-4 Link Service (FCP_LS) definitions

8.1 FCP_LS overview

FC-4 Link Service functionality is specified in FC-LS. For FCP FC-4 Link Services, the FCP frame header fields (see 5.2) shall be set as follows:

a) R_CTL Routing field (word 0, bits 31-28) shall be set to 0011b (i.e., an FC-4 Link_Data frame);
b) the TYPE field shall be set to 08h (i.e., FCP FC-4 Link Service frame); and
c) the R_CTL Information field (word 0, bits 27-24) shall be set to 0010b (i.e., unsolicited control) for request Sequences and 0011b (i.e., solicited control) for response Sequences.

The FCP_LS requests and responses defined in this standard are specified in table 15.

Table 15 - FCP_LS requests and responses

R_CTL (word 0, bits 31-24)	Name	Description	Request/ Response	Reference
14h	SRR	Sequence Retransmission Request	Request	8.2
02h	FCP_ACC	FCP_LS Accept	Response	8.3
01h	FCP_RJT	FCP_LS Reject	Response	8.4

8.2 Sequence Retransmission Request (SRR) FCP_LS request

The SRR FCP_LS request Sequence is transmitted by an initiator FCP_Port to request that a target FCP_Port retransmit information or request retransmission of information for the specified Exchange.

If task retry identification (see 4.7) is active for the Originator and the Responder, then the PARAMETER field of the request Sequence shall contain the task retry identifier for the command specified by the OX_ID and RX_ID field values.

If the target FCP_Port determines that the Originator S_ID, OX_ID, or RX_ID field values, or task retry identifier are inconsistent, then it shall respond with an FCP_RJT (see 8.4) with the Reason Code set to "Unable to perform the command request" and the Reason Code Explanation set to "Invalid OX_ID-RX_ID combination".

If the target FCP_Port is unable to retransmit the Sequence or data at the requested Relative Offset, then the target FCP_Port shall respond with an FCP_RJT Sequence with the Reason Code set to "Unable to perform the command request" the Reason Code Explanation set to "Unable to supply requested data".

If the initiator FCP_Port receives an FCP_RJT response, then the initiator FCP_Port shall terminate the Exchange referenced by the SRR FCP_LS request using ABTS-LS (see 12.3.2).

The SRR FCP_LS request shall be sent in a new Exchange. The Exchange shall be ended by the response to the SRR FCP_LS request.

Sequence Initiative for the Exchange referenced by the SRR FCP_LS request shall be transferred to the target FCP_Port to retransmit the requested Sequence.

For unacknowledged classes, the Sequence Count for a retransmitted FCP_DATA IU may start at zero, even if continuously increasing sequence count is being used. For acknowledged classes, the Sequence Count for a retransmitted FCP_DATA IU shall start at one higher than the last Sequence Count used in the Exchange to prevent it from being within the range of the Recovery_Qualifier.

A target FCP_Port that has agreed during Process Login to support retransmission should not reject requests for retransmission of the requested frames, but if the target FCP_Port does responds to the SRR FCP_LS request with an FCP_RJT, then the target FCP_Port shall return CHECK CONDITION status with the sense key set to HARDWARE ERROR and the additional sense code set to INITIATOR DETECTED ERROR MESSAGE RECEIVED.

SRR FCP_LS requests for Exchanges involving logical units that do not support retransmission on a target FCP_Port that supports retransmission for other logical units shall be rejected with an FCP_RJT with the Reason Code set to “Unable to support command request” and the Reason Code Explanation set to “Unable to supply requested data”.

Addressing:

The S_ID field designates the initiator FCP_Port requesting the information retransmission. The D_ID field designates the target FCP_Port that is to receive the request.

SRR FCP_LS request payload:

The format of the SRR FCP_LS request payload is specified in table 16.

Table 16 - SRR FCP_LS request payload

Word \ Bits	Bits 31-24	Bits 23-16	Bits 15-8	Bits 7-0
0	14000000h			
1	OX_ID		RX_ID	
2	RELATIVE OFFSET			
3	R_CTL FOR IU	Reserved		

The OX_ID field shall be set to the OX_ID of the Exchange for which retransmission is being requested.

The RX_ID field shall be set to the RX_ID of the Exchange for which retransmission is being requested.

The RELATIVE OFFSET field is only valid if the R_CTL FOR IU field is set to 01h (i.e., Solicited Data) or to 05h (i.e., Data Descriptor). The RELATIVE OFFSET field contains the Relative Offset of the lowest byte the initiator FCP_Port has identified as requiring retransmission. The two low-order bits of the RELATIVE OFFSET field shall be set to zero, such that the data to be retransmitted begins on a four-byte boundary.

The amount of data to transfer is implicitly the remainder of that for the Exchange.

The R_CTL FOR IU field encoding is as described in FC-FS-3 (i.e., Data Descriptor for an FCP_XFER_RDY IU, Command Status for an FCP_RSP IU, and Solicited Data for an FCP_DATA IU).

Possible responses to the SRR FCP_LS request:

FCP_ACC

Specifies acceptance of the SRR FCP_LS request.

FCP_RJT

Specifies rejection of the SRR FCP_LS request.

8.3 FCP_LS Accept (FCP_ACC)

FCP_ACC notifies the originator of an FCP_LS request that the FCP_LS request Sequence has been accepted. An FCP_ACC may be a response Sequence to any FCP_LS request.

Addressing:

The D_ID field specifies the source of the FCP_LS request being accepted. The S_ID field specifies the destination of the FCP_LS request being accepted.

FCP_ACC payload:

The payload for the FCP_ACC is specified in table 17.

Table 17 - FCP_ACC payload

Bits Word	Bits 31-24	Bits 23-16	Bits 15-8	Bits 7-0
0	02000000h			

8.4 FCP_LS Reject (FCP_RJT)

FCP_RJT notifies the originator of an FCP_LS request that the FCP_LS request Sequence has been rejected. An FCP_RJT may be a response Sequence to any FCP_LS request.

Addressing:

The D_ID field specifies the source of the FCP_LS request being rejected. The S_ID field specifies the destination of the FCP_LS request being rejected.

FCP_RJT payload:

The first word of the FCP_RJT payload shall contain the FCP_RJT code (i.e., 01000000h). The next four bytes of the FCP_RJT payload shall contain a Reason Code, Reason Code Explanation, and vendor specific information, if any, for rejecting the request. The format of the FCP_RJT payload is shown in table 18.

Table 18 - FCP_RJT payload

Bits Word	Bits 31-24	Bits 23-16	Bits 15-8	Bits 7-0
0	01000000h			
1	Reserved	REASON CODE	REASON CODE EXPLANATION	Vendor specific

The REASON CODE field value are specified in table 19.

Table 19 - FCP_RJT Reason Codes

Encoded Value (Bits 23-16)	Description	Meaning
01h	Invalid FCP_LS command code	The FCP_LS command code in the Sequence being rejected is invalid.
03h	Logical error	The request identified by the FCP_LS command code and payload content is invalid or logically inconsistent for the conditions present.
05h	Logical busy	The Link Service is logically busy and unable to process the request at this time.
07h	Protocol error	An error has been detected that violates the rules of the FC-FS-3 signaling protocol, but that is not specified by other error codes.
09h	Unable to perform command request	The Recipient of a Link Service command is unable to perform the request at this time.
0Bh	Command not supported	The Recipient of a Link Service command does not support the command requested.
FFh	Vendor Specific Error (See Bits 7-0)	The Vendor Specific Error bits may be used by vendors to specify additional reason codes.
other	Reserved	

The first error condition encountered shall be the error reported.

The Reason Code Explanation field values are specified in Table 20.

Table 20 - FCP_RJT Reason Code Explanations

Encoded Value (Bits 15-8)	Description	Applicable FCP_LS requests
00h	No additional explanation	SRR
17h	Invalid OX_ID-RX_ID combination	SRR
2Ah	Unable to supply requested data	SRR

9 FCP Information Unit (IU) usage and formats

9.1 FCP Information Unit (IU) usage

The IUs used by the Fibre Channel Protocol and their characteristics are shown in table 21 for IUs sent to target FCP_Ports, and in table 22 for IUs sent to initiator FCP_Ports. Each IU shall be contained in a single Sequence (see 3.1.56). Each Sequence carrying an FCP IU shall contain only one IU. Examples of typical Fibre Channel Protocol operations using these IUs are included in Annex B.

Table 21 - FCP Information Units (IUs) sent to target FCP_Ports

IU [a]	Description	Data block		F/M/L	SI	M/O
		R_CTL field	Content			
T1	Command / Task management request	06h	FCP_CMND	F	T	M
T2 [b, c]	Command request	06h	FCP_CMND	F	H	O
T6	Data-Out action	01h	FCP_DATA	M	T	M
T12 [d]	Confirm	03h	FCP_CONF	L	T	O

Key:

IU	Information Unit identifier	
Content	Contents (payload) of data block	
F/M/L	First/Middle/Last Sequence of Exchange (FC-FS-3)	
	F	First
	M	Middle
	L	Last
SI	Sequence Initiative: Held or Transferred (FC-FS-3)	
	H	Held
	T	Transferred
M/O	Mandatory/Optional Sequence	
	M	Mandatory
	O	Optional

a T3, T4, T5, T7, T8, T9, T10, and T11 are obsolete.
b T2 is only permitted while transfer ready IUs are disabled (see table 11).
c T2 allows optional Sequence streaming during write operations.
d T12 is only permitted in response to an I5 frame (see table 22).

Table 22 - FCP Information Units (IUs) sent to initiator FCP_Ports

IU [a]	Description	Data block		F/M/L	SI	M/O
		R_CTL field	Content			
I1	Data-Out delivery request	05h	FCP_XFER_RDY (Write)	M	T	M
I3 [b]	Data-In action	01h	FCP_DATA	M	H	M
I4	Command/Task management response	07h	FCP_RSP	L	T	M
I5 [c]	Response (Confirm request)	07h	FCP_RSP	M	T	O
I8	Extended Response	08h	(see FC-NVMe)			
I9	Extended Response (Confirm response)	08h	(see FC-NVMe)			

Key:

- IU — Information Unit identifier
- Content — Contents (payload) of data block
- F/M/L — First/Middle/Last Sequence of Exchange (FC-FS-3)
  - F First
  - M Middle
  - L Last
- SI — Sequence Initiative: Held or Transferred (FC-FS-3)
  - H Held
  - T Transferred
- M/O — Mandatory/Optional Sequence
  - M Mandatory
  - O Optional

a I2, I6, and I7 are obsolete.
b I3 allows optional Sequence streaming to I3, I4, or I5.
c I5 is permitted to request the confirm completion protocol.

9.2 FCP_CMND IU

9.2.1 Overview and format of FCP_CMND IU

The FCP_CMND IU carries either a command or a task management request. If an invalid combination of bits is set in the FCP_CMND IU, then the target FCP_Port shall respond with an FCP_RSP IU with the RSP_CODE field set to INVALID_FIELD (see table 29). The format of the FCP_CMND IU is specified in table 23.

Table 23 - FCP_CMND IU payload

Bit / Byte	7	6	5	4	3	2	1	0
0	FCP_LUN							
7								
8	COMMAND REFERENCE NUMBER							
9	Rsvd	COMMAND PRIORITY				TASK ATTRIBUTE		
10	TASK MANAGEMENT FLAGS							
11	ADDITIONAL FCP_CDB LENGTH = (n-27)/4						RDDATA	WRDATA
12	FCP_CDB							
27								
28	ADDITIONAL FCP_CDB (if any)							
n								
n+1	(MSB)							
n+2	FCP_DL							
n+3								
n+4								(LSB)
n+5	(MSB)							
n+6	FCP_BIDIRECTIONAL_READ_DL (if any)							
n+7								
n+8								(LSB)

9.2.2 FCP_CMND IU field descriptions

9.2.2.1 FCP_LUN field

If the value of byte zero of the FCP_LUN field is equal to FDh and the value of byte one of the FCP_LUN field is equal to 28h, then see FC-NVMe.

If the value of byte zero of the FCP_LUN field is equal to FEh, then see FC-SB-6.

If the value of byte zero of the FCP_LUN field is not equal to FDh or FEh, then the FCP_LUN field contains the LUN of the logical unit in the SCSI target device. See SAM-5.

If the FCP_LUN field contains a valid LUN, then the command or task managmenent function shall be routed to the specified logical unit. If the addressed logical unit does not exist, then the SCSI target device shall follow the rules for responding to selection of an incorrect logical unit as specified in SAM-5.

9.2.2.2 COMMAND REFERENCE NUMBER field

The COMMAND REFERENCE NUMBER field contains the number sent by the initiator FCP_Port to assist in performing precise delivery checking for FCP commands. If precise delivery is enabled, then a nonzero value in the COMMAND REFERENCE NUMBER field shall be treated as a command reference number in determining the receipt and ordering of commands from a particular initiator FCP_Port to the particular logical unit as described in 4.4. If precise delivery is enabled, then a zero value in the COMMAND REFERENCE NUMBER field indicates that command shall not be verified for precise delivery. If precise delivery checking is not enabled, then the COMMAND REFERENCE NUMBER field shall be ignored by the device server. If the FCP_CMND IU specifies a task management function, then the COMMAND REFERENCE NUMBER field is reserved and set to zero and the FCP_CMND IU shall not be verified for precise delivery.

9.2.2.3 COMMAND PRIORITY field

The COMMAND PRIORITY field specifies the relative scheduling importance of a command with the TASK ATTRIBUTE field set to 000b (i.e., SIMPLE) in relation to other commands already in the task set with SIMPLE task attributes (see SAM-5).

9.2.2.4 TASK ATTRIBUTE field

The TASK ATTRIBUTE field contains values that specify the task attribute (see SAM-5) associated with the CDB, as shown in table 24.

Table 24 - TASK ATTRIBUTE field

Task attribute code	Task attribute	Description
000b	SIMPLE	Requests that the task be managed according to the rules for a simple task attribute, and the rules for priority if implemented (see SAM-5).
001b	HEAD OF QUEUE	Requests that the task be managed according to the rules for a head of queue task attribute (see SAM-5).
010b	ORDERED	Requests that the task be managed according to the rules for an ordered task attribute (see SAM-5). Mechanisms to assure delivery of commands to a device server in the correct order are described in 4.4.
011b	Reserved	
100b	ACA	Requests that the task be managed according to the rules for an automatic contingent allegiance task attribute (see SAM-5).
101b	Obsolete	
110b-111b	Reserved	

9.2.2.5 TASK MANAGEMENT FLAGS field

The TASK MANAGEMENT FLAGS field specifies the task management function to be performed.

Task management functions shall be requested by the initiator FCP_Port (Exchange Originator) using a new Exchange. If the TASK MANAGEMENT FLAGS field is set to a nonzero value, the FCP_CDB field, then the FCP_DL field, the TASK ATTRIBUTE field, the RDDATA bit, and the WRDATA bit shall be ignored and the FCP_BIDIRECTIONAL_READ_DL field shall not be included in the FCP_CMND IU payload. If the TASK MANAGEMENT FLAGS field is set to a reserved value, then the target FCP_Port shall return an FCP_RSP IU containing the FCP_CODE field set to INVALID_FIELD.

The clearing actions performed by task management functions are specified in table 9 (see 4.10).

The TASK MANAGEMENT FLAGS field is specified in table 25.

Table 25 - TASK MANAGEMENT FLAGS field

Code	Name	Task management function [a]	Support
00h	None [b]	None [b]	Mandatory
01h	FCP_QUERY_TASK_SET	QUERY TASK SET	Optional
02h	FCP_ABORT_TASK_SET	ABORT TASK SET	Mandatory
04h	FCP_CLEAR_TASK_SET	CLEAR TASK SET	Mandatory
08h	FCP_QUERY_ASYNCHRONOUS_EVENT	QUERY ASYNCHRONOUS EVENT	Optional
10h	FCP_LOGICAL_UNIT_RESET	LOGICAL UNIT RESET	Mandatory
20h	Obsolete	Obsolete	
40h	FCP_CLEAR_ACA	CLEAR ACA	Optional
80h	Obsolete	Obsolete	
All others	Reserved	Reserved	

a) The ABORT TASK and QUERY TASK task management functions are specified in 4.9.
b) Indicates no task management function is performed and the FCP_CDB field shall be used.

The QUERY TASK SET task management function is specified in SAM-4.

The ABORT TASK SET task management function is specified in SAM-4. The ABORT TASK SET task management function resets internal states of the target FCP_Port (see 4.10). Exchange resources may be cleared by an ABTS-LS (see 12.3) transmitted by the initiator FCP_Port that sent the ABORT TASK SET task management function for each command known to the initiator FCP_Port.

The CLEAR TASK SET task management function is specified in SAM-4. The CLEAR TASK SET task management function resets internal states of the target FCP_Port (see 4.10). Exchange resources to be cleared may be cleared by one or more of the following mechanisms:

a) an ABTS-LS (see 12.3) may be transmitted by the initiator FCP_Port that sent the CLEAR TASK SET for each command known to that initiator FCP_Port;
b) a command, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET is ended in the logical unit. The initiator FCP_Port for that command shall determine by a timeout that the command did not finish. If a timeout occurs, then the initiator FCP_Port shall clear the Exchange resources by transmitting an ABTS-LS. (see 12.3); or
c) a command for an initiator FCP_Port other than the initiator FCP_Port that sent the CLEAR TASK SET may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-5.

NOTE 2 - SAM-5 has defined the TASK ABORTED status for commands terminated by a CLEAR TASK SET task management function if the Control mode page (see SPC-4) indicates that the TASK ABORTED status is supported.

The QUERY ASYNCHRNOUS EVENT task management function is specified in SAM-4.

The LOGICAL UNIT RESET task management function is specified in SAM-4. The LOGICAL UNIT RESET task management function resets the internal states of the target FCP_Port and logical unit (see 4.10). Exchange resources to be cleared may be cleared by the following mechanisms:

a) an ABTS-LS (see 12.3) may be transmitted by the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function for each command in the logical unit known to that initiator FCP_Port;
b) a command, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function is ended in the logical unit. The initiator FCP_Port for that command shall determine by a timeout that the command did not finish. If a timeout occurs, then the initiator FCP_Port shall clear the Exchange resources by transmitting an ABTS-LS (see 12.3); or
c) a command for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET task management function may be completed by returning CHECK CONDITION status with the sense key set to UNIT ATTENTION and the additional sense code set as specified in SAM-5.

NOTE 3 - SAM-5 has defined the TASK ABORTED status for commands terminated by a LOGICAL UNIT RESET task management function if the Control mode page (see SPC-4) indicates that the TASK ABORTED status is supported.

The CLEAR ACA task management function is specified in SAM-4.

When the task manager clears the ACA condition, any command within that task set may be completed subject to the rules for task management specified by SAM-5. If there is no ACA condition present, then the CLEAR ACA task management function shall be accepted and the FCP_RSP IU shall contain a RSP_CODE field set to TMF_COMPLETE (see table 29).

The use of the NACA bit in the CDB CONTROL field and the implementation of ACA is described in SAM-5.

Depending on the mode page parameters that have been established (see SPC-4), additional FCP I/O operations may have to be aborted by transmitting an ABTS-LS as part of the process of clearing the ACA condition.

9.2.2.6 ADDITIONAL FCP_CDB LENGTH field

The ADDITIONAL FCP_CDB LENGTH field contains the length in 4-byte words of the ADDITIONAL FCP_CDB field. The ADDITIONAL FCP_CDB LENGTH field shall be set to zero for task management requests.

9.2.2.7 RDDATA bit and WRDATA bit

If the RDDATA bit is set to one, then the initiator FCP_Port expects to receive FCP_DATA IUs from the target FCP_Port (i.e., a read operation).

If the WRDATA bit is set to one, then the initiator FCP_Port expects to transmit FCP_DATA IUs to the target FCP_Port (i.e., a write operation).

If the RDDATA bit and WRDATA bit are both set to one, then the initiator FCP_Port expects both a read operation and a write operation (i.e., a bidirectional command). The FCP_BIDIRECTIONAL_READ_DL field shall be included in the FCP_CMND IU payload. The initiator FCP_Port shall not set both the RDDATA bit and the WRDATA bit to one except for a bidirectional command.

If the RDDATA bit and WRDATA bit are both set to zero, then there shall be no FCP_DATA IUs and the FCP_DL field shall be set to zero.

The device server shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT if the following protocol errors are detected:

a) a read operation has the RDDATA bit set to zero or the WRDATA bit set to one;
b) a write operation has the WRDATA bit set to zero or the RDDATA bit set to one;
c) a bidirectional command has either the RDDATA bit set to zero or the WRDATA bit set to zero; or
d) the RDDATA bit and WRDATA bit are both set to zero and the FCP_DL value is not zero.

NOTE 4 - Device servers compliant to previous versions of this standard may terminate the command and return an FCP_RSP IU with the RSP_CODE field set to INVALID_FIELD (see table 29) for some protocol errors.

9.2.2.8 FCP_CDB field

The FCP_CDB field contains the CDB to be sent to the addressed logical unit. The maximum CDB length is 16 bytes unless the ADDITIONAL FCP_CDB_LENGTH field has specified that there is an ADDITIONAL_FCP_CDB field. The FCP_CDB field shall be ignored if the TASK MANAGEMENT FLAGS field is set to a nonzero value.

The CDB format is defined by SAM-5 and SPC-4 and the contents of the CDB are defined in the SCSI command standards. Bytes between the end of a CDB and the end of the FCP_CDB field or, if applicable, the ADDITIONAL_FCP_CDB field shall be reserved.

9.2.2.9 ADDITIONAL_FCP_CDB field

The ADDITIONAL FCP_CDB field contains any CDB bytes beyond those contained within the 16 byte FCP_CDB field.

The ADDITIONAL FCP_CDB field shall not be present if the TASK MANAGEMENT FLAGS field is set to a nonzero value. The contents of the field shall be those bytes of an extended CDB beyond the first 16 bytes of the CDB as defined in the SCSI command standards.

9.2.2.10 FCP_DL field

For a read command, the FCP_DL field contains a count of the maximum number of all bytes to be transferred to the application client buffer in FCP_DATA IU payloads by the command. The FCP_DL field is the Data-In Buffer Size defined by SAM-5.

For a write command, the FCP_DL field contains a count of the maximum number of all bytes to be transferred from the application client buffer in FCP_DATA IU payloads by the command. The FCP_DL field is the Data-Out Buffer Size defined by SAM-5.

For a bidirectional command, the FCP_DL field contains a count of the maximum number of all bytes to be transferred from the application client buffer in FCP_DATA IU payloads by the command. The FCP_DL field is the Data-Out Buffer Size defined by SAM-5.

An FCP_DL field value of zero indicates that no data transfer is expected regardless of the state of the RDDATA and WRDATA bits and that no FCP_XFER_RDY or FCP_DATA IUs shall be transferred.

9.2.2.11 FCP_BIDIRECTIONAL_READ_DL field

For a bidirectional command, the FCP_BIDIRECTIONAL_READ_DL field contains a count of the maximum number of all bytes to be transferred to the application client buffer in FCP_DATA IU payloads by the command. The FCP_BIDIRECTIONAL_READ_DL field is the Data-In Buffer Size defined by SAM-5.

If the FCP_BIDIRECTIONAL_READ_DL field value is set to zero, then no read operation is expected regardless of the state of the RDDATA bit and no FCP_DATA IUs shall be transferred for read data.

If either the RDDATA bit or the WRDATA bit is set to zero, then the FCP_BIDIRECTIONAL_READ_DL field shall not be included in the FCP_CMND IU payload.

9.3 FCP_XFER_RDY IU

9.3.1 Overview and format of FCP_XFER_RDY IU

The FCP_XFER_RDY IU indicates that the target FCP_Port is prepared to receive part or all of the data for a write operation. The FCP_XFER_RDY IU contains those parameters of the SAM-5 data delivery service required by the initiator FCP_Port, including the length and beginning relative offset of the FCP_DATA IU that is requested. Since the target FCP_Port has established buffering and caching resources based on the requested data, the initiator FCP_Port shall provide the described data in the requested FCP_DATA IU. The initiator FCP_Port shall be ready to transmit any part or all of the number of bytes indicated in the FCP_DL field when requested.

An FCP_XFER_RDY IU shall be transmitted preceding each write FCP_DATA IU when the WRITE FCP_XFER_RDY DISABLED bit is set to zero by Process Login (see 4.14 and 6.3). If the target FCP_Port and initiator FCP_Port have negotiated write FCP_XFER_RDY disabled, then FCP_XFER_RDY IUs shall be transmitted to request each write FCP_DATA IU after the first FCP_DATA IU of the command.

The first FCP_DATA IU is transmitted without a preceding FCP_XFER_RDY IU (see 6.3.3).

The first 8 bytes of the FCP_XFER_RDY IU payload are defined in FC-FS-3 for all IUs of category 5 (i.e., the data descriptor category). The fields defined in FC-FS-3 are given FCP names for use in this standard. The format of the FCP_XFER_RDY IU payload is shown in table 26.

Table 26 - FCP_XFER_RDY IU payload

Bit / Byte	7	6	5	4	3	2	1	0
0	(MSB)							
				FCP_DATA_RO				
3								(LSB)
4	(MSB)							
				FCP_BURST_LEN				
7								(LSB)
8								
				Reserved				
11								

9.3.2 FCP_DATA_RO field

The FCP_DATA_RO field contains a value specifying the relative offset in the PARAMETER field for the first data byte of the requested FCP_DATA IU (see 5.2.12). The FCP_DATA_RO field is the "Offset of the data being transferred" field specified in FC-FS-3.

The FCP_DATA_RO field may be used by the target FCP_Port to request data out of order on writes if allowed by the EMDP bit in the Disconnect-Reconnect mode page (see 10.2.8). This is equivalent to the SAM-5 application client buffer offset.

The FCP_DATA_RO field shall have a value that is a multiple of 4 (i.e., each FCP_DATA IU shall begin on a word boundary).

9.3.3 FCP_BURST_LEN field

The FCP_BURST_LEN field contains a value indicating the amount of buffer space prepared for all bytes to be transferred in the next FCP_DATA IU and requests the transfer of an FCP_DATA IU of that length from the initiator FCP_Port.

The FCP_BURST_LEN field is the "Length of the data being transferred" field specified in FC-FS-3 and the value in the FCP_BURST_LEN field is equivalent to the SCSI data delivery request byte count (see SAM-5).

The value in the FCP_BURST_LEN field shall not exceed the MAXIMUM BURST SIZE field value defined in the Disconnect-Reconnect mode page (see 10.2.7). The sum of the value of the FCP_BURST_LEN field and the value of the FCP_DATA_RO field shall not exceed the value of the FCP_DL field. The value in the FCP_BURST_LEN field shall not be zero.

9.4 FCP_DATA IU

9.4.1 FCP_DATA IU overview

The data associated with a particular FCP I/O operation is transmitted in the same Exchange that sent the FCP_CMND IU requesting the transfer.

SCSI data transfers may be performed by one or more data delivery requests, each one performing a transfer no longer than:

a) the FIRST BURST SIZE field value (see 10.2.10) if the WRITE FCP_XFER_RDY DISABLED bit is set to one; or
b) THE MAXIMUM BURST SIZE field value (see 10.2.7) if a FCP_XFER_RDY IU was received.

If more than one FCP_DATA IU is used to transfer the data, the relative offset value in the PARAMETER field is used to ensure that the SCSI data is reassembled in the proper order (see 5.2.12). If an FCP_XFER_RDY IU is used to describe a data transfer and the first frame of the requested FCP_DATA IU has a relative offset that differs from the value in the FCP_DATA_RO field of the FCP_XFER_RDY IU, then the target FCP_Port shall return an FCP_RSP IU with the RSP_CODE field set to FCP_DATA_RO_MISMATCH (see table 29).

If required by the PRLI FCP Service Parameters, then each Data-Out action FCP_DATA IU shall be preceded by an FCP_XFER_RDY IU containing a standard data descriptor payload that indicates the location and length of the data delivery. If the WRITE FCP_XFER_RDY DISABLED bit is set to one in the PLRI FCP Service Parameter page (see 6.3), then the first FCP_DATA IU shall be transmitted without a preceding FCP_XFER_RDY IU.

If the DATA OVERLAY ALLOWED bit is set to one in the PLRI FCP Service Parameter page (see 6.3) for the initiator FCP_Port, then the target FCP_Port may request that data be overlaid. If the DATA OVERLAY ALLOWED bit is set to zero in the PLRI FCP Service Parameter page (see 6.3) for the initiator FCP_Port, then the target FCP_Port shall not request that data be overlaid. If data overlay is not allowed and the target FCP_Port attempts to overlay data, then the initiator FCP_Port may not be able to guarantee data integrity and may indicate service delivery failure. Data retransmission as part of an error recovery process is not considered data overlay, even if retransmission occurs to the same offset in the application client buffer.

The target FCP_Port may request data bursts in any order if allowed by the EMDP bit in the Disconnect-Reconnect mode page (see 10.2). By the time data transfer has been terminated, all data between the offset of zero and the highest offset shall have been transferred. If error conditions occur that prevent the transfer of data in the middle of a data transfer, the FCP_SNS_INFO shall indicate that only data from the offset of zero up to the first byte of missing data has been transferred. Even if data of a higher offset was successfully transferred, it shall not be considered valid.

FC-FS-3 specifies the mechanisms used to transfer an IU. The mechanisms vary with the class of service being used and the service parameters that are in effect.

9.4.2 FCP_DATA IUs for read and write operations

During any data transfer, the initiator FCP_Port shall have available a buffer of the length specified by the FCP_DL field in the FCP_CMND IU. The buffer contains data to be transferred to the target FCP_Port if the operation is a write operation (i.e., an operation that uses the Data-Out action, IU T6). The buffer receives the data if the operation is a read operation (i.e., an operation that uses the Data-In action, IU I3). The target FCP_Port shall not request or deliver data outside the buffer length defined by FCP_DL.

If the command requested that data beyond the length specified by the FCP_DL field be transferred, then the device server shall set the FCP_RESID_OVER bit (see 9.5.8) to one in the FCP_RSP IU and:

a) process the command normally except that data beyond the FCP_DL count shall not be requested or transferred;
b) transfer no data and return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT; or
c) may transfer data and return CHECK CONDITION status with the sense key set to ABORTED COMMAND and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

During a write operation that is not using FCP_XFER_RDY IUs, the initiator FCP_Port indicates that it has transferred all the required data by transferring Sequence Initiative to the target FCP_Port.

The initiator FCP_Port shall not transfer data outside the buffer length defined by FCP_DL. If the write operation requires a total amount of data less than the amount of data provided by the initiator FCP_Port, then the target FCP_Port shall discard the excess bytes. If an operation results in the transfer of fewer bytes than required by FCP_DL, the FCP_RESID_UNDER bit (see 9.5.7) shall be set to one in the FCP_RSP IU. The command is completed according to the rules specified by the SCSI command standard for that command.

If the amount of data requested or transferred does not match the number of bytes of data calculated from the value of the FCP_DL field and the value of the FCP_RESID field (see 9.5.12), then the error detection and recovery procedure described in clause 12 may be invoked or the FCP I/O operation may be terminated with an Abort Exchange (see 12.3) or other failure indication. The mechanism an initiator FCP_Port uses to determine that the correct amount of data has been returned is vendor specific. Data that has been retransmitted and overlaid shall be counted only once for the purposes of calculating residual values.

9.4.3 FCP_DATA IUs for bidirectional commands

During a bidirectional command, the initiator FCP_port shall always have available a buffer with the length specified by the FCP_DL field in the FCP_CMND IU to transfer data to the target FCP_Port. The target FCP_Port shall not request data outside the buffer length specified by the FCP_DL field.

During a bidirectional command, the initiator FCP_Port shall always have available a buffer with the length specified by the FCP_BIDIRECTIONAL_DL field in the FCP_CMND IU to receive data from the target FCP_Port. The target FCP_Port shall not transfer data outside the buffer length specified by the FCP_BIDIRECTIONAL_DL field.

If a bidirectional command requested that data beyond the length specified by the FCP_DL field be transferred, then the device server shall set the FCP_RESID_OVER bit (see 9.5.8) to one in the FCP_RSP IU and:

a) process the command normally except that data beyond the FCP_DL count shall not be transferred; or
b) transfer no data in either direction and return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.
c) may transfer data in either direction and return CHECK CONDITION status with the sense key set to ABORTED COMMAND and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

If a bidirectional command requests that data beyond the length specified in the FCP_BIDIRECTIONAL_READ_DL field be transferred, the device server shall set the FCP_BIDI_READ_RESID_OVER bit (see 9.5.5) to one in the FCP_RSP IU and:

a) process the command normally except that data beyond the FCP_BIDIRECTIONAL_READ_DL count shall not be transferred; or
b) transfer no data in either direction and return CHECK CONDITION with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN COMMAND INFORMATION UNIT.

If the amount of data requested or transferred does not match the number of bytes of data calculated from the value of the FCP_DL field and the value of the FCP_RESID field (see 9.5.12) for the write operation of a bidirectional command or the number of bytes of data calculated from the value of the FCP_BIDIRECTIONAL_READ_DL field and the value of the FCP_BIDIRECTIONAL_READ_RESID field (see 9.5.13) for the read operation of a bidirectional command, then the FCP I/O operation may be terminated with an ABTS-LS (see 12.3) or other failure indication. The mechanism an initiator FCP_Port uses to determine that the correct amount of data has been returned is vendor specific. Data that has been retransmitted and overlaid shall be counted only once for the purposes of calculating residual values.

9.4.4 FCP_DATA IU use of fill bytes

During transfer of data in response to an FCP_CMND_IU with the RDDATA bit set to one and the WRDATA bit set to zero, all frames of FCP_DATA_IUs except the frame with the highest relative offset within the Data-In Buffer shall have no fill bytes.

During transfer of data in response to an FCP_CMND_IU with the WRDATA bit set to one and the RDDATA bit set to zero, all frames of FCP_DATA_IUs except the frame with the highest relative offset within the Data-Out Buffer shall have no fill bytes.

During transfer of data in response to an FCP_CMND_IU with the WRDATA bit set to one and the RDDATA bit set to one, all frames of FCP_DATA_IUs except the frame with the highest relative offset within the

Data-In Buffer and the frame with the highest relative offset within the Data-Out Buffer shall have no fill bytes.

9.5 FCP_RSP IU

9.5.1 Overview and format of FCP_RSP IU

The FCP_RSP IU provides completion information for FCP I/O operations. The information includes SCSI status, protocol verification, and any applicable sense data. The target FCP_Port shall transmit an FCP_RSP IU for each task management function delivered with an FCP_CMND IU, indicating the completion status of the task management function in the RSP_CODE field (see table 29).

The bits and fields in byte 10 and byte 11 summarize the completion status of an FCP I/O operation and indicate the meaning and validity of other fields in the FCP_RSP IU. Byte 10 and byte 11 shall be zero upon successful completion of an FCP I/O operation, indicating that no other information is present in the FCP_RSP IU. A nonzero value in either byte 10 or byte 11 should cause the application client to examine the fields in FCP_RSP IU to determine whether a failure, a retryable temporary condition, or an expected response occurred.

If data retransmission is enabled and a Sequence error is detected, then a target FCP_Port shall not transmit an FCP_RSP IU with CHECK CONDITION status. See 12.3.5 for additional target FCP_Port error recovery.

If an error is detected by a target FCP_Port while the target FCP_Port has Sequence Initiative for the Exchange associated with the error, then the target FCP_Port should complete any Sequence that has already been started, keep Sequence Initiative, and transmit an FCP_RSP IU with CHECK CONDITION status and sense data that describes the error. If an error is detected by a device server while the target FCP_Port does not have Sequence Initiative for the Exchange associated with the error, then the target FCP_Port shall wait until Sequence Initiative has been returned and then transmit an FCP_RSP IU with CHECK CONDITION status and sense data that describes the error.

If Sequence Initiative is not received within $RR\_TOV_{SEQ\_INIT}$ (see 11.4), then the target FCP_Port may implicitly terminate the affected Exchange.

The format of the FCP_RSP IU is specified in table 27.

Table 27 - FCP_RSP IU payload

Bit / Byte	7	6	5	4	3	2	1	0
0	Reserved							
7								
8	STATUS QUALIFIER							
9								
10	FCP_ BIDI_ RSP	FCP_ BIDI_ READ_ RESID_ UNDER	FCP_ BIDI_ READ_ RESID_ OVER	FCP_ CONF_ REQ	FCP_ RESID_ UNDER	FCP_ RESID_ OVER	FCP_ SNS_ LEN_ VALID	FCP_ RSP_ LEN_ VALID
11	SCSI STATUS CODE							
12	(MSB)	FCP_RESID						
15								(LSB)
16	(MSB)	FCP_SNS_LEN (= n)						
19								(LSB)
20	(MSB)	FCP_RSP_LEN (= m)						
23								(LSB)
24	FCP_RSP_INFO (m bytes long)(if any)(see table 28)							
23+m								
24+m	FCP_SNS_INFO (n bytes long)(if any)							
23+m+n								
24+m+n	(MSB)	FCP_BIDIRECTIONAL_READ_RESID (if any)						
27+m+n								(LSB)

9.5.2 STATUS QUALIFIER field

The STATUS QUALIFIER field contains the status qualifier (see SAM-5).

9.5.3 FCP_BIDI_RSP bit

If the FCP_BIDI_RSP bit is set to one, then the FCP_BIDIRECTIONAL_READ_RESID field is present, and the FCP_BIDI_READ_RESID_OVER bit and FCP_BIDI_READ_RESID_UNDER bit are valid. If the FCP_BIDI_RSP bit is set to zero, then the FCP_BIDIRECTIONAL_READ_RESID field is not present, and the FCP_BIDI_READ_RESID_OVER bit and the FCP_BIDI_READ_RESID_UNDER bit are not valid.

9.5.4 FCP_BIDI_READ_RESID_UNDER bit

If the FCP_BIDI_READ_RESID_UNDER bit is set to one, then the FCP_BIDIRECTIONAL_READ_RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred. In the context of the command, the FCP_BIDIRECTIONAL_READ_RESID field indicates whether or not an error condition occurred.

9.5.5 FCP_BIDI_READ_RESID_OVER bit

If the FCP_BIDI_READ_RESID_OVER bit is set to one, then the FCP_BIDIRECTIONAL_READ_RESID field is valid and contains the count of bytes that were not transferred because the FCP_BIDIRECTIONAL_READ_DL value was not large enough. In the context of the command, the FCP_BIDIRECTIONAL_READ_RESID field indicates whether or not an error condition occurred.

9.5.6 FCP_CONF_REQ bit

If the FCP_CONF_REQ bit is set to one, then the initiator FCP_Port shall transmit an FCP_CONF IU to confirm receipt of the FCP_RSP Sequence. If the FCP_CONF_REQ bit is set to zero, then the initiator FCP_Port shall not transmit an FCP_CONF IU.

9.5.7 FCP_RESID_UNDER bit

If the FCP_RESID_UNDER bit is set to one, then the FCP_RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred. In the context of the command, the FCP_RESID field indicates whether or not an error condition occurred.

9.5.8 FCP_RESID_OVER bit

If the FCP_RESID_OVER bit is set to one, then the FCP_RESID field is valid and contains the count of bytes that were not transferred because the FCP_DL was not sufficient. In the context of the command, the FCP_RESID field indicates whether or not an error condition occurred.

9.5.9 FCP_SNS_LEN_VALID bit

If the FCP_SNS_LEN_VALID bit is set to one, then the FCP_SNS_INFO field contains valid information, the FCP_SNS_LEN field is valid and non-zero and contains the count of bytes in the FCP_SNS_INFO field. In the context of the command, the FCP_SNS_INFO field indicates whether or not an error condition occurred.

If the FCP_SNS_LEN_VALID bit is set to zero, then the FCP_SNS_LEN field is not valid and shall be treated as if its value were zero. See 9.5.14.

9.5.10 FCP_RSP_LEN_VALID bit

If the FCP_RSP_LEN_VALID bit is set to one, then the FCP_RSP_INFO field contains valid information, the FCP_RSP_LEN field is valid and non-zero and contains the count of bytes in the FCP_RSP_INFO field. In the context of the command, the FCP_RSP_INFO field indicates whether or not an error condition occurred.

If the FCP_RSP_LEN_VALID bit is set to one, then the content of the SCSI STATUS CODE field is not reliable and shall be ignored by the application client.

For task management functions transmitted to the logical unit using an FCP_CMND IU, the FCP_RSP_LEN_VALID bit shall be set to one, the FCP_RSP_LEN field shall be set to the specified value, and the information in the RSP_CODE field (see table 29) shall indicate the completion status of the task management function.

If the FCP_RSP_LEN_VALID bit is set to zero, then the FCP_RSP_LEN field is not valid and shall be treated as if its value were zero. If the FCP_RSP_LEN_VALID bit is set to zero, then the FCP_RSP_INFO field shall have a length of zero and shall not be present.

9.5.11 SCSI STATUS CODE field

The SCSI STATUS CODE field contains the status code for the command being completed, as defined by SAM-5.

9.5.12 FCP_RESID field

For read operations and write operations, if the FCP_RESID_UNDER bit is set to one, then the FCP_RESID field contains a count of the number of residual data bytes that were not transferred in FCP_DATA IUs.

For read operations and write operations, if the FCP_RESID_OVER bit is set to one, then the FCP_RESID field contains the excess of the number of bytes required by the command to be transferred over the number of bytes specified by the FCP_DL field.

For bidirectional commands, if the FCP_RESID_UNDER bit is set to one, then the FCP_RESID field contains a count of the number of residual data bytes that were not transferred in the Data-Out FCP_DATA IUs for the command.

For bidirectional commands, if the FCP_RESID_OVER bit is set to one, then the FCP_RESID field contains the excess of the number of bytes required to be transferred in the Data-Out FCP_DATA IUs by the command over the number of bytes specified in the FCP_DL field.

For bidirectional commands, the FCP_BIDIRECTIONAL_READ_RESID field (see 9.5.13) contains the corresponding count for Data-In FCP_DATA IUs.

Upon successful completion of an FCP I/O operation, the residual value is normally zero and the FCP_RESID value is not valid. FCP devices having indeterminate data lengths may have a nonzero residual byte count after completing valid operations. There is no requirement for the device server to verify that the data length implied by the contents of the CDB does not cause an overrun or underrun before beginning the processing of a command.

If the FCP_RESID_UNDER bit is set to one, then a transfer that did not fill the buffer to the expected displacement FCP_DL was performed and the value of FCP_RESID is defined as follows:

FCP_RESID = FCP_DL - (highest offset of any byte transmitted + 1)

A condition of FCP_RESID_UNDER may not be an error for some FCP devices and some commands.

If the FCP_RESID_OVER bit is set to one, then refer to 9.4.2 or 9.4.3. The FCP_RESID value is defined as follows:

FCP_RESID = (transfer length required by command) - FCP_DL

If the FCP_RESID_UNDER and the FCP_RESID_OVER bits are set to zero, then the FCP_RESID field is not meaningful and may have any value. The FCP_RESID field is always included in the FCP_RSP IU.

Some early implementations presented the FCP_RSP IU without the FCP_RESID field, FCP_SNS_LEN field, and FCP_RSP_LEN field if the FCP_RESID_UNDER bit, FCP_RESID_OVER bit, FCP_SNS_LEN_VALID bit, and FCP_RSP_LEN_VALID bit were all set to zero. This non-standard behavior should be tolerated.

9.5.13 FCP_BIDIRECTIONAL_READ_RESID field

The FCP_BIDIRECTIONAL_READ_RESID field is included in the FCP_RSP IU for all bidirectional commands.

For bidirectional commands, if the FCP_BIDI_READ_RESID_UNDER bit is set to one, then the FCP_BIDIRECTIONAL_READ_RESID field contains a count of the number of residual data bytes that were not transferred in Data-In FCP_DATA IUs.

For bidirectional commands, if the FCP_BIDI_READ_RESID_OVER bit is set to one, then the FCP_BIDIRECTIONAL_READ_RESID field contains the excess of the number of bytes required by the command to be transferred in Data-In FCP_DATA IUs over the number of bytes specified by the FCP_BIDIRECTIONAL_READ_DL field.

Upon successful completion of an FCP I/O operation, the residual value is normally zero and the FCP_BIDIRECTIONAL_READ_RESID value is not valid. FCP devices having indeterminate data lengths may have a nonzero residual byte count after completing valid operations. There is no requirement for the device server to verify that the data length implied by the contents of the CDB does not cause an overrun or underrun before beginning the processing of a command.

If the FCP_BIDI_READ_RESID_UNDER bit is set to one, then a transfer that did not fill the buffer to the expected displacement FCP_BIDIRECTIONAL_READ_DL was performed and the value of FCP_BIDIRECTIONAL_READ_RESID is defined as follows:

FCP_BIDIRECTIONAL_READ_RESID = FCP_BIDIRECTIONAL_READ_DL - (highest offset of any byte transmitted + 1)

A condition of FCP_BIDI_READ_RESID_UNDER may not be an error for some FCP devices and some commands.

If the FCP_BIDI_READ_RESID_OVER bit is set to one, then refer to 9.4.3. The FCP_BIDIRECTIONAL_READ_RESID value is defined as follows:

FCP_BIDIRECTIONAL_READ_RESID = (read transfer length required by command) - FCP_BIDIRECTIONAL_READ_DL

If the FCP_BIDI_READ_RESID_UNDER bit and the FCP_RESID_OVER bit are both set to zero, then the FCP_BIDIRECTIONAL_READ_RESID field is not meaningful and may have any value.

9.5.14 FCP_SNS_LEN field

If the FCP_SNS_LEN_VALID bit is one, then the FCP_SNS_LEN field specifies the number of valid bytes of FCP_SNS_INFO.

If the FCP_SNS_LEN_VALID bit is zero, then the FCP_SNS_LEN field is not valid and shall be treated as if its value were zero. The FCP_SNS_INFO field is not present.

The FCP_SNS_LEN field is always included in the FCP_RSP IU.

9.5.15 FCP_RSP_LEN field

If the FCP_RSP_LEN_VALID bit is one, then the FCP_RSP_LEN field specifies the number of valid bytes of FCP_RSP_INFO. The FCP_RSP_LEN field shall be set to 00000004h or 00000008h.

If the FCP_RSP_LEN_VALID bit is zero, then the FCP_RSP_LEN field is not valid and shall be treated as if its value were zero. No FCP_RSP_INFO is provided.

The FCP_RSP_LEN field is always included in the FCP_RSP IU.

9.5.16 FCP_RSP_INFO field

The FCP_RSP_INFO field contains information describing only the protocol failures detected during the processing of an FCP I/O operation. If none of the specified protocol failures have occurred, then the FCP_RSP_INFO field shall not be included in the FCP_RSP IU and the FCP_RSP_LEN_VALID bit shall be zero. The FCP_RSP_INFO does not contain link error information, since FC-FS-3 provides the mechanisms for presenting such errors. The FCP_RSP_INFO field does not contain SCSI logical unit error information, since that is contained in the FCP_SNS_INFO field as described in 9.5.17. The FCP_RSP_INFO field shall contain valid information if the target FCP_Port detects any of the conditions indicated by an FCP RSP_CODE. The format of the FCP_RSP_INFO field is specified in table 28.

Table 28 - FCP_RSP_INFO field format

Byte \ Bit	7	6	5	4	3	2	1	0
0 – 2	ADDITIONAL_RSP_INFO							
3	RSP_CODE							
4 – 7	Reserved (if any)							

The ADDITIONAL_RSP_INFO FIELD contains additional response information for certain task management functions (e.g., QUERY ASYNCHRONOUS EVENT) as defined in SAM-5. If the task management function does not specify additional response information or the logical unit does not support additional response information, then the target FCP_Port shall set the ADDITIONAL_RSP_INFO field to 000000h.

The RSP_CODE field is specified in table 29.

Table 29 - RSP_CODE field

Code	Name	Description
00h	TMF_COMPLETE	Task Management function complete
01h	DATA_LENGTH_MISMATCH	FCP_DATA length different than FCP_BURST_LEN
02h	INVALID_FIELD	FCP_CMND fields invalid
03h	FCP_DATA_RO_MISMATCH	FCP_DATA parameter mismatch with FCP_DATA_RO
04h [a]	TMF_REJECTED	Task Management function rejected
05h [a]	TMF_FAILED	Task Management function failed
08h [a]	TMF_SUCCEEDED	Task Management function succeeded
09h [a]	TMF_INCORRECT_LUN	Task Management function incorrect logical unit number
06h – 07h 0Ah – FFh	Reserved	
a) Only valid when responding to task management functions.		

The completion status of the task management function is indicated by the RSP_CODE field. If the Exchange is aborted before the FCP_RSP IU is returned, then the completion status is unknown. If the RSP_CODE field is set to TMF_FAILED, then the state of the logical unit is unknown.

Activities started by a task management function may continue after the FCP_RSP IU for the task management has been delivered.

9.5.17 FCP_SNS_INFO field

The FCP_SNS_INFO field contains the sense data specified by SPC-4. The proper FCP_SNS_INFO shall be presented when the SCSI STATUS CODE field is set to CHECK CONDITION (see SAM-5). If no conditions requiring the presentation of SCSI sense data have occurred, then the FCP_SNS_INFO field shall not be included in the FCP_RSP IU and the FCP_SNS_LEN_VALID bit shall be set to zero.

9.6 FCP_CONF IU

The FCP_CONF IU has no payload. It is used as described in 4.5 for an initiator FCP_Port to confirm the receipt of the FCP_RSP IU from a target FCP_Port. The frame shall be transmitted by an initiator FCP_Port if the confirmed completion protocol is supported by both the target FCP_Port and the initiator FCP_Port and when the confirmation has been requested by the FCP_CONF_REQ bit in the FCP_RSP IU.

10 SCSI mode parameters for the Fibre Channel Protocol

10.1 Overview of mode pages for the Fibre Channel Protocol

This clause describes the mode pages used with the MODE SELECT and MODE SENSE commands to control and report the behavior of the Fibre Channel Protocol. All mode parameters not defined in this standard shall control the behavior of the FCP devices as specified in the appropriate command set standard. The mode pages are addressed to the device server of a logical unit. The logical unit shall provide the appropriate control parameters, if any, to the state machine implementing the connection to the Fibre Channel loop or link in a vendor-specific manner. The mode pages associated with Fibre Channel Protocol operation are listed in table 30.

Table 30 - Mode pages for FCP

Page code	Subpage code	Description	Reference
02h	00h	Disconnect-Reconnect mode page	10.2
18h	00h	Fibre Channel Logical Unit Control mode page	10.3
19h	00h	Fibre Channel Port Control mode page	10.4

10.2 Disconnect-Reconnect mode page

10.2.1 Overview and format of Disconnect-Reconnect mode page for FCP

The Disconnect-Reconnect mode page (see table 31) allows the application client to modify the behavior of the target FCP_Port. This subclause specifies the parameters defined by SPC-4 that are used by FCP devices and defines how FCP devices interpret the parameters. The application client communicates with the device server to determine what values are most appropriate for a device server. The device server communicates the parameter values in this mode page to the target FCP_Port. This communication is internal to the SCSI target device and FCP device and is outside the scope of this standard. If a field or bit contains a value that is not supported by the FCP device, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

An interconnect tenancy is the period of time when an FCP device owns or may access a shared Fibre Channel interconnect. For arbitrated loops (see FC-AL-2), a tenancy typically begins when an FCP device successfully opens the connection and ends when the FCP device releases the connection for use by other device pairs. Data and other information transfers take place during interconnect tenancies.

Point-to-point or fabric-attached Class 2 or Class 3 links and many other configurations do not have a concept of interconnect tenancy and may perform transfers at any time.

Table 31 - Disconnect-Reconnect mode page (02h)

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	Reserved	PAGE CODE (02h)					
1	PAGE LENGTH (0Eh)							
2	BUFFER FULL RATIO							
3	BUFFER EMPTY RATIO							
4	(MSB)	BUS INACTIVITY LIMIT						
5								(LSB)
6	(MSB)	DISCONNECT TIME LIMIT						
7								(LSB)
8	(MSB)	CONNECT TIME LIMIT						
9								(LSB)
10	(MSB)	MAXIMUM BURST SIZE						
11								(LSB)
12	EMDP	FAA	FAB	FAC	RESTRICTED	RESTRICTED		
13	Reserved							
14	(MSB)	FIRST BURST SIZE						
15								(LSB)

10.2.2 BUFFER FULL RATIO field

The BUFFER FULL RATIO field indicates to the device server, during read operations, how full the buffer should be prior to requesting an interconnect tenancy. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in SPC-4. FCP devices attached to links that do not have the concept of interconnect tenancy shall round the ratio to zero and transmit data in a vendor specific manner.

The value contained in the BUFFER FULL RATIO field is defined by SPC-4.

10.2.3 BUFFER EMPTY RATIO field

The BUFFER EMPTY RATIO field indicates to the device server, during write operations, how empty the buffer should be prior to transmitting an FCP_XFER_RDY IU that requests the initiator FCP_Port to transmit data. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in SPC-4.

The value contained in the BUFFER EMPTY RATIO field is defined by SPC-4.

10.2.4 BUS INACTIVITY LIMIT field

The BUS INACTIVITY LIMIT field indicates the maximum time that the target FCP_Port is permitted to maintain an interconnect tenancy without data or information transfer, measured in transmission word increments. If the bus inactivity limit is exceeded or if the bus is inactive and the target FCP_Port holding the bus detects that the limit is going to be exceeded, then the device server shall end the interconnect tenancy. This value may be rounded as defined in SPC-4. A value of zero indicates that there is no bus inactivity limit.

Because of the low overheads associated with initiating and closing bus tenancy on Fibre Channel links, device servers should end tenancies immediately upon completing the required transfers.

The BUS INACTIVITY LIMIT field is not applicable for FCP devices attached to links that do not have the concept of interconnect tenancy.

10.2.5 DISCONNECT TIME LIMIT field

The DISCONNECT TIME LIMIT field indicates the minimum delay between interconnect tenancies measured in increments of 128 transmission words. Target FCP_Ports in configurations having the concept of interconnect tenancy shall delay at least this time interval after each interconnect tenancy before beginning arbitration. The device server may round this value to any value it prefers. A value of zero indicates that the disconnect time limit does not apply.

The DISCONNECT TIME LIMIT field is not applicable for FCP devices attached to links that do not have the concept of interconnect tenancy.

10.2.6 CONNECT TIME LIMIT field

The CONNECT TIME LIMIT field indicates the maximum duration of a single interconnect tenancy, measured in increments of 128 transmission words. If the connect time limit is exceeded, then the device server shall conclude the interconnect tenancy, within the restrictions placed on it by the applicable Fibre Channel configuration. The device server may round this value to any value it prefers. A value of zero indicates that there is no connect time limit.

The CONNECT TIME LIMIT field is not applicable for FCP devices attached to links that do not have the concept of interconnect tenancy.

10.2.7 MAXIMUM BURST SIZE field

The MAXIMUM BURST SIZE field indicates the maximum size of all bytes in an FCP_DATA IU that the target FCP_Port shall transfer to the initiator FCP_Port in a single Data-In FCP_DATA IU or request from the initiator FCP_Port in an FCP_XFER_RDY IU. This parameter does not affect how much data is transferred in a single interconnect tenancy. This value is expressed in increments of 512 bytes (e.g., a value of 1 means 512 bytes, two means 1024 bytes, etc.). The device server may round this value down as defined in SPC-4. A value of zero indicates there is no limit on the amount of data transferred per data transfer operation. The no limit option (i.e., the zero value) shall be implemented by all FCP devices. The initiator FCP_Port and target FCP_Port may use the value of this parameter to adjust internal maximum buffering requirements.

10.2.8 EMDP bit

The enable modify data pointers (EMDP) bit indicates whether or not the target FCP_Port may use the random buffer access capability to reorder FCP_DATA IUs for a single command. If the EMDP bit is set to zero, then the target FCP_Port shall generate continuously increasing relative offset values for each FCP_DATA IU for a single command. If the EMDP bit is set to one, then the target FCP_Port may transfer the FCP_DATA IUs for a single command in any order. If the EMDP bit is set to zero, then data overlay is prohibited even if it is allowed by the DATA OVERLAY ALLOWED bit in Process Login (see 4.14 and 6.3). The EMDP bit does not affect the order of frames within a Sequence. The enable modify data pointers function is optional for all FCP devices.

For bidirectional commands, the EMDP bit applies independently to the read operation and write operation. If the EMDP bit is set to zero, then the target FCP_Port shall generate continuously increasing relative offset values for the read operation and the write operation, but there is no read operation to write operation or write operation to read operation ordering requirement.

10.2.9 FAA bit, FAB bit, and FAC bit

The fairness access (FA) bits, FAA, FAB, and FAC, indicate whether a target FCP_Port attached to an arbitrated loop (see FC-AL-2) shall use the access fairness algorithm when beginning the interconnect tenancy.

An FA bit set to one indicates that the target FCP_Port shall use the access fairness algorithm for the specified frames. An FA bit set to zero indicates that the target FCP_Port may choose to not use the access fairness algorithm.

The FAA bit controls arbitration when the target FCP_Port has one or more FCP_DATA IU frames to transmit to an initiator FCP_Port.

The FAB bit controls arbitration when the target FCP_Port has one or more FCP_XFER_RDY IU frames to transmit to an initiator FCP_Port.

The FAC bit controls arbitration when the target FCP_Port has an FCP_RSP IU frame to transmit to an initiator FCP_Port.

If the target FCP_Port intends to transmit multiple frame types, then it may choose to not use the access fairness algorithm if any applicable FA bit is set to zero. FCP devices attached to links that do not have the concept of interconnect tenancy shall ignore the FA bits. The FA bits are optional for all FCP devices.

10.2.10 FIRST BURST SIZE field

When the WRITE FCP_XFER_RDY DISABLED bit is negotiated as being set to one in the PRLI ELS FCP Service Parameter page (see 6.3), the FIRST BURST SIZE field indicates the maximum amount of all bytes that shall be transmitted in the first FCP_DATA IU sent from the initiator FCP_Port to the target FCP_Port. If all data is transmitted in the first IU, then no subsequent FCP_XFER_RDY IUs shall be transmitted by the target FCP_Port. If the maximum amount of data has been transmitted, but more data remains to be transferred, then the target FCP_Port shall request that data with subsequent FCP_XFER_RDY IUs.

When the WRITE FCP_XFER_RDY DISABLED bit is negotiated as being set to zero in the PRLI ELS FCP Service Parameter page (see 6.3), the FIRST BURST SIZE field is ignored and permission to transmit data from the initiator FCP_Port to the target FCP_Port is managed using FCP_XFER_RDY IUs. For data transmissions from the target FCP_Port to the initiator FCP_Port, the FIRST BURST SIZE field is ignored.

The FIRST BURST SIZE field value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes). A value of zero indicates that there is no first burst size limit. The FIRST BURST SIZE field shall be implemented by all FCP devices that support the WRITE FCP_XFER_RDY DISABLED bit being set to one. The application client and device server may use the value of this field to adjust internal maximum buffering requirements.

10.3 Fibre Channel Logical Unit Control mode page

The Fibre Channel Logical Unit Control mode page (see table 32) contains those parameters that select FCP logical unit operation options. The implementation of any parameter and its associated functions is optional. The mode page follows the MODE SENSE command and MODE SELECT command rules specified by SPC-4.

Table 32 - Fibre Channel Logical Unit Control mode page (18h)

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (18h)					
1	PAGE LENGTH (06h)							
2	Reserved				PROTOCOL IDENTIFIER (FCP = 0h)			
3	Reserved							EPDC
4	Reserved							
5	Reserved							
6	Reserved							
7	Reserved							

If the enable precise delivery checking (EPDC) bit is set to one, then the logical unit shall use the precise delivery function defined by this standard (see 4.4). If the EPDC bit is set to zero, then the logical unit shall not use the precise delivery function and shall ignore the contents of the COMMAND REFERENCE NUMBER field in the FCP_CMND IU (see 9.2.2.2). The EPDC bit is valid for all types of link connections. If the precise delivery function is not supported and the Fibre Channel Logical Unit Control mode page is supported by the logical unit, then the EPDC bit shall be masked as not changeable and shall follow the MODE SENSE command and MODE SELECT command rules specified by SPC-4.

If the Fibre Channel Logical Unit Control mode page is not supported by a logical unit, then the application client shall assume that the precise delivery function is not supported by that logical unit.

10.4 Fibre Channel Port Control mode page

10.4.1 Overview and format of Fibre Channel Port Control mode page

The Fibre Channel Port Control mode page contains those parameters that select FCP_Port operation options. The mode page policy shall be per target port (see SPC-4). The mode page shall be implemented by LUN 0 and may be implemented by logical units other than LUN 0. The implementation of any bit and its associated functions is optional. The mode page follows the MODE SENSE command and MODE SELECT command rules specified by SPC-4.

NOTE 5 - Some of the bit values defined by the Fibre Channel Port Control mode page require the FCP_Port to violate one or more of the Fibre Channel standards. The non-standard behaviors have been identified as useful for certain specialized operating environments.

The format of the Fibre Channel Port Control mode page is specified in table 33.

Table 33 - Fibre Channel Port Control mode page (19h)

Bit / Byte	7	6	5	4	3	2	1	0
0	PS	SPF (0b)	PAGE CODE (19h)					
1	PAGE LENGTH (06h)							
2	Reserved				PROTOCOL IDENTIFIER (FCP = 0h)			
3	DTFD	PLPB	DDIS	DLM	RHA	ALWI	DTIPE	DTOLI
4	Reserved							
5	Reserved							
6	Reserved					RR_TOV UNITS		
7	SEQUENCE INITIATIVE RESOURCE RECOVERY TIMEOUT VALUE (RR_TOV$_{SEQ_INIT}$)							

10.4.2 Disable target originated loop initialization (DTOLI) bit

If the disable target originated loop initialization (DTOLI) bit is set to one, then the target FCP_Port shall not generate a LIP following insertion into an arbitrated loop (see FC-AL-2). The target FCP_Port shall respond to a LIP when it is received. If the DTOLI bit is set to zero, then the target FCP_Port shall generate LIP(F7, xx) after it enables a port into an arbitrated loop. If the target FCP_Port is attached to an arbitrated loop and detects loop failure at its input, then it shall follow the error initialization process defined by FC-AL-2 regardless of the state of the DTOLI bit. Target FCP_Ports not attached to an arbitrated loop shall ignore the DTOLI bit.

10.4.3 Disable target initiated port enable (DTIPE) bit

If the disable target initiated port enable (DTIPE) bit is set to one, then the target FCP_Port shall wait for an initiator FCP_Port to transmit the Loop Port Enable (LPE) primitive sequence before inserting itself into an arbitrated loop (see FC-AL-2). The target FCP_Port shall wait in a participating state with the Port Bypass circuit, if any, set to bypass the target FCP_Port. The target FCP_Port uses the hard address available in the SCA-2 connector (see SFF-8067) or in device address jumpers to determine whether LPE primitive sequences are addressed to it. An LPE primitive sequence addressed to the broadcast address shall also cause the target FCP_Port to insert itself into the loop. If the DTIPE bit is set to zero, then the target FCP_Port shall enable itself onto the loop according to the rules specified in FC-AL-2. Target FCP_Ports not attached to an arbitrated loop shall ignore the DTIPE bit.

10.4.4 Allow login without loop initialization (ALWLI) bit

If the allow login without loop initialization (ALWLI) bit is set to one, then the target FCP_Port shall use the hard address available in the SCA-2 connector (see SFF-8067) or in device address jumpers, enter the monitoring state in participating mode, and accept logins without using the loop initialization procedure (see FC-AL-2). If the ALWLI bit is set to zero, then the target FCP_Port shall perform the normal loop initialization procedure before entering the monitoring mode and accepting a login ELS. Target FCP_Ports not attached to an arbitrated loop shall ignore the ALWLI bit.

10.4.5 Require hard address (RHA) bit

If the REQUIRE HARD ADDRESS (RHA) bit is set to one, then the target FCP_Port shall only attempt to obtain its hard address available in the SCA-2 connector (see SFF-8067) or device address jumpers during loop initialization. The target FCP_Port shall not attempt to obtain an address during the LISA phase of initialization (see FC-AL-2). If there is a conflict for the hard address selection during loop initialization or the target FCP_Port does not have a valid hard address available, then the target FCP_Port shall enter the nonparticipating state. If the target FCP_Port detects loop initialization while in the nonparticipating state, then it shall again attempt to get its hard address. If the hard address has not changed from the address obtained in a previous successful loop initialization, then the target FCP_Port shall attempt to obtain the address in the LIFA phase if a valid fabric login exists or LIPA phase of loop initialization. If the hard address has changed, the target FCP_Port shall attempt to obtain the new address in the LIHA phase.

If the RHA bit is set to zero, then the target FCP_Port shall follow the normal initialization procedure, including the possibility of obtaining a soft address during the loop initialization process.

Target FCP_Ports not attached to an arbitrated loop shall ignore the RHA bit.

10.4.6 Disable loop master (DLM) bit

If the disable loop master (DLM) bit is set to one, then the target FCP_Port shall not participate in loop master arbitration and shall not become loop master. The target FCP_Port shall only repeat LISM frames it receives. If the DLM bit is set to zero, then the target FCP_Port may participate in loop master arbitration in the normal manner and, if successful, may become loop master during the loop initialization process. Target FCP_Ports not attached to an arbitrated loop shall ignore the DLM bit.

10.4.7 Disable discovery (DDIS) bit

If the disable discovery (DDIS) bit is set to one, a target FCP_Port without a valid fabric login shall not require receipt of Address or Port Discovery (i.e., ADISC ELS or PDISC ELSs) following loop initialization as described in FC-DA-2. The logical units shall resume processing commands on completion of loop initialization. If the DDIS bit is set to zero, then the target FCP_Port shall wait to complete target discovery, as defined by FC-DA-2, before allowing processing of commands to resume.

Target FCP_Ports not attached to an arbitrated loop shall ignore the DDIS bit. A target FCP_Port with a valid fabric login shall ignore the DDIS bit.

10.4.8 Prevent loop port bypass (PLPB) bit

If the prevent loop port bypass (PLPB) bit is set to one, then a target FCP_Port shall ignore any Loop Port Bypass (LPB) and Loop Port Enable (LPE) primitive sequences. The loop port shall always remain participating. If the PLPB bit is set to zero, then the target FCP_Port allows the Loop Port Bypass (LPB) and Loop Port Enable (PBE) primitive sequences to control the port bypass circuit and participation on the loop as specified by FC-AL-2. Target FCP_Ports not attached to an arbitrated loop shall ignore the PLPB bit.

The DTIPE and PLPB bits shall not both be set to one at the same time. If an invalid bit combination is sent by the application client, the device server shall terminate the command with CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN THE PARAMETER LIST.

10.4.9 Disable target fabric discovery (DTFD) bit

If the disable target fabric discovery (DTFD) bit is set to one, then a target FCP_Port shall not recognize the presence of a fabric loop port on the loop. The target FCP_Port shall perform only the private loop functions defined for target FCP_Ports defined by FC-DA-2. If the DTFD bit is set to zero, the target FCP_Port shall discover a fabric loop port if present on the loop and perform the public loop functions defined for target FCP_Ports as described in FC-DA-2. Target FCP_Ports not attached to an arbitrated loop shall ignore the DTFD bit.

10.4.10 RR_TOV UNITS field

The RR_TOV UNITS field indicates the units for the $RR\_TOV_{SEQ\_INIT}$ field value, according to table 34.

Table 34 - Values for RR_TOV UNITS

Value	Units of measure for $RR\_TOV_{SEQ\_INIT}$
000b	No timer is specified

Table 34 - Values for RR_TOV UNITS

Value	Units of measure for $RR\_TOV_{SEQ\_INIT}$
001b	0.001 s
011b	0.1 s
101b	10 s
All other values	Reserved

10.4.11 SEQUENCE INITIATIVE RESOURCE RECOVERY TIMEOUT VALUE ($RR\_TOV_{SEQ\_INIT}$) field

The $RR\_TOV_{SEQ\_INIT}$ (see 11.4) timer is defined by the $RR\_TOV_{SEQ\_INIT}$ field and the RR_TOV UNITS field.

The $RR\_TOV_{SEQ\_INIT}$ field indicates the number of time units specified by the RR_TOV UNITS field that shall be used by the timer that performs the $RR\_TOV_{SEQ\_INIT}$ timeout function. If no timer is specified, then the $RR\_TOV_{SEQ\_INIT}$ value shall be ignored by the device server and a vendor specific default value shall be used for $RR\_TOV_{SEQ\_INIT}$.

11 Timers for FCP operation and recovery

11.1 Summary of timers for the Fibre Channel Protocol

This clause indicates the use of timers defined by other standards in performing the FCP-4 recovery procedures and those timers used only by this standard. Table 35 shows these timers.

Table 35 - Timer summary

Timer	Implementation Mandatory (M) or Optional (O)		Description	Default Value	Ref
	Initiator FCP_Port	Target FCP_Port			
E_D_TOV[b,c]	M	O	Error_Detect_Timeout Value	2 s	11.2
R_A_TOV[a,b]$_{SEQ\_QUAL}$	M	O	Resource_Allocation Timeout Value	Private loop = 0 s Public loop = 10 s	11.3
R_A_TOV[a,b]$_{ELS}$	M	M		Private loop = 2 s Public loop = 10 s	11.3
RR_TOV$_{AUTH}$		M	Resource Recovery Timeout Value	2 s	11.4
RR_TOV$_{SEQ_INIT}$		M		If RETRY bit is set to zero: 2 s If RETRY bit is set to one: $\geq$ REC_TOV + 2 $\times$ R_A_TOV$_{ELS}$ + 1 s	11.4
REC_TOV[d]	M	M	REC ELS Timeout Value	$\geq$ E_D_TOV + 1 s (minimum)	11.5
ULP_TOV	M		Upper Level Protocol Timeout Value	If RETRY bit is set to zero: $\geq$ Operation-specific timer + E_D_TOV + 1 s If RETRY bit is set to one: $\geq$ Operation-specific timer + 2 $\times$ RR_TOV$_{SEQ_INIT}$	11.6

a) R_A_TOV is defined by FC-FS-3. This standard defines the default R_A_TOV for Sequence Qualifiers as zero for private loops and 10 s for public loops. This standard defines the default R_A_TOV for ELS responses as 2 s for private loops and 10 s for public loops. If ELSs are used to set R_A_TOV, then the same value is applied for both uses. Other Fibre Channel standards may specify different default values for R_A_TOV for different topologies.

b) Target FCP_Ports that support Class 2 delivery service shall implement this timer.

c) E_D_TOV default timeout values are defined by FC-FS-3 and FC-DA. ELSs are provided to set values other than the default value. This standard defines the default value required by the recovery protocol, deriving the value as follows:

A) Public loop devices compliant with FC-DA use an E_D_TOV value of 2 s before fabric login and the value obtained in the FLOGI ELS LS_ACC after fabric login;

B) Private loop devices compliant with FC-DA use the default E_D_TOV value of 2 s; or

C) Devices attached through a fabric or point-to-point connection use the default E_D_TOV value specified by FC-FS-3 before fabric login and the value obtained in the FLOGI ELS LS_ACC after fabric login.

d) REC_TOV is required by the target FCP_Port for FCP_CONF IU error detection.

11.2 Error_Detect Timeout (E_D_TOV)

E_D_TOV is a general error detect timeout value (see FC-FS-3, FC-AL-2, FC-LS-2, and FC-DA-2). For FCP-4 Sequence recovery, it is used to time the following:

a) the maximum time permitted for a Sequence Initiator between the transmission of consecutive Data frames within a single Sequence;
b) the minimum time that a Sequence Recipient shall wait for the reception of the next frame within a single Sequence before recognizing a Sequence timeout; and
c) the minimum time a Sequence Initiator shall wait for an ACK response before it considers the ACK to be missing and begins recovery actions.

Target FCP_Ports that support Class 2 shall implement this timer for the purpose of timing out missing ACKs.

Loop attached Class 2 devices may require a complete fairness cycle plus the fabric and link delay times before an ACK is received.

11.3 Resource Allocation Timeout (R_A_TOV)

R_A_TOV has two separate components, labeled R_A_TOV$_{SEQ_QUAL}$ and R_A_TOV$_{ELS}$.

R_A_TOV$_{SEQ_QUAL}$ is used to define the minimum amount of time that a Sequence Initiator shall wait before reusing the Sequence_Qualifier associated with an aborted Sequence. The Sequence_Qualifer is composed of the S_ID field, D_ID field, OX_ID field, RX_ID field, and SEQ_ID field.This value is also the minimum amount of time that a Sequence Initiator shall wait following receipt of the BA_ACC reply Sequence to an ABTS before transmitting a Reinstate Recovery Qualifier (RRQ) ELS.

Using a value of zero for this timeout value assumes that a Sequence Initiator does not transmit any Frames for a Sequence after an ABTS is sent for that Sequence. If a design uses a queuing mechanism for the transmission of Sequences, then the queue for a given Sequence shall be empty before an ABTS for that Sequence is sent, or the act of transmitting the ABTS purges the queue.

A value of two times R_A_TOV$_{ELS}$ is used to determine the minimum time that the Originator of an Extended Link Service or FC-4 Link Service request shall wait for the response to that request.

After completion of fabric login, public loop devices shall use the value of R_A_TOV specified by the Fabric in the FLOGI ELS LS_ACC.

11.4 Resource Recovery Timeout (RR_TOV)

RR_TOV has two separate components, labeled RR_TOV$_{AUTH}$ and RR_TOV$_{SEQ_INIT}$.

RR_TOV$_{AUTH}$ is the minimum time a target FCP_Port shall wait for an initiator FCP_Port to perform Exchange authentication following the completion of the Loop Initialization Protocol (i.e., the receipt of CLS while in the OPEN-INIT state) (see FC-DA-2). If Exchange authentication is not performed within RR_TOV$_{AUTH}$ of completion of the Loop Initialization protocol, then then a target FCP_Port may implicitly or explicitly perform a logout with that initiator FCP_Port and reclaim the resources associated with those Exchanges (see 12.4.1.5).

RR_TOV$_{SEQ_INIT}$ is the minimum time a target FCP_Port shall wait for an initiator FCP_Port response following transfer of Sequence Initiative from the target FCP_Port to the initiator FCP_Port (e.g., following transmission of the FCP_XFER_RDY IU during a write command). If the initiator FCP_Port does not send a response within RR_TOVseq_init of the transfer of Sequence Initiative, then a target FCP_Port may implicitly or explicitly perform a logout with that initiator FCP_Port and reclaim the resources associated with those Exchanges (see 12.4.1.5).

The value of RR_TOV$_{SEQ_INIT}$ may be set using the Fibre Channel Port Control mode page (see 10.4.11).

11.5 Read Exchange Concise Timeout Value (REC_TOV)

REC_TOV is used by the initiator FCP_Port to provide a minimum polling interval for the REC ELS and by the target FCP_Port for FCP_CONF IU error detection. The REC_TOV timer shall be implemented such that at least one REC_TOV period passes between transmission of a command and the first polling for Exchange status with the REC ELS. Table 36 describes REC_TOV usage pertaining to the initiator FCP_Port.

Table 36 - Initiator FCP_Port REC_TOV usage

Timer starts or restarts after:	Timer stops without transmitting REC ELS after:
FCP_CMND IU has been sent.	FCP_RSP IU is received for the FCP_CMND IU or the Exchange is aborted.
FCP_DATA IU Sequence has been sent by the initiator FCP_Port (optional).	FCP_RSP IU is received for the FCP_CMND IU or the Exchange is aborted.
REC ELS was sent for an FCP_CMND IU and an LS_ACC was received indicating the command is in progress (i.e., REC ELS polling interval).	FCP_RSP IU is received for the FCP_CMND IU or the Exchange is aborted.

Table 37 describes REC_TOV usage pertaining to the target FCP_Port.

Table 37 - Target FCP_Port REC_TOV usage

Timer starts after:	Timer stops without transmitting REC ELS after:
FCP_RSP IU requesting an FCP_CONF IU has been sent.	FCP_CONF IU is received or the Exchange is aborted.

11.6 Upper Level Protocol Timeout (ULP_TOV)

ULP_TOV is an operation-specific timer maintained by the Upper Level Protocol. ULP_TOV is used to time the completion of Exchanges associated with ULP operations. Since the amount of time required varies depending upon the operation, the value assigned for this timer is determined by the operation being timed. Some operations may require extended periods of time to complete.

ULP timers take into account response time increments caused by command queuing and multi-initiator FCP_Port congestion.

12 Link error detection and error recovery procedures

12.1 Error detection and error recovery overview

12.1.1 Exchange level

This standard provides several mechanisms for FCP devices to identify protocol errors caused by frames and responses that have been corrupted and discarded in accordance with the requirements of FC-FS-3. See 12.2 for a list of these mechanisms.

To recover from these errors, all FCP compliant initiator FCP_Ports shall be capable of transmitting an Abort Exchange (i.e., ABTS-LS) to terminate a failing Exchange and to recover the associated resources (see 12.3). All FCP compliant target FCP_Ports shall be capable of processing an ABTS-LS to finish clearing the Exchange and to recover the associated resources. The failed command may then be reissued by higher level programs according to protocols beyond the scope of this standard.

This standard allows the use of the REC ELS to monitor the progress of active Exchanges. An FCP-4 device may accept or reject error detection inquiries.

12.1.2 Sequence level

Sequence level error recovery (see 12.4) shall not be used for bidirectional commands.

To recover from errors, FCP-4 compliant devices may perform retransmission procedures that allow the commands to be completed without requiring higher level programs to perform command retries. Such recovery may be useful for SCSI logical units that depend critically on command ordering and maintaining records of internal device state. The initiator FCP_Port and the target FCP_Port shall agree to perform retransmission using the SRR FCP_LS request by setting the RETRY bit to one in the process login before performing the retransmission of individual IUs (see 6.3.3). An FCP-4 device that has agreed to perform retransmission shall use and accept the REC ELS and SRR FCP_LS request as defined by this standard to perform the retransmission.

Even after agreeing to perform retransmission, the initiator FCP_Port may choose to transmit an ABTS-LS and the target FCP_Port shall be able to accept and process the ABTS-LS.

While the basic error detection and error recovery procedures are class independent, acknowledged classes of services may use the acknowledgement mechanism as an additional error detection feature and may use mechanisms defined in FC-FS-3 to assist in the recovery process.

This clause defines the error detection and recovery mechanisms for fabrics that guarantee in-order frame delivery. However, if continuously increasing sequence count is used and if support for recovery qualifiers is fully implemented as defined in FC-FS-3, then the same recovery mechanisms may be used for fabrics that do not guarantee in-order frame delivery, as shown in the examples in Annex C.

Examples of error detection and error recovery are provided in Annex C.

12.2 FCP error detection

12.2.1 Overview of FCP-4 error detection

The subclauses of 12.2 describe the initial events that indicate an error may have occurred. The error may be recovered at the Exchange level or at the Sequence level.

12.2.2 FCP-4 error detection using protocol errors for all classes of service

The Exchange Originator (i.e., initiator FCP_Port) shall detect any of the following errors:

a) a Sequence error is detected in a Sequence transmitted from a target FCP_Port to an initiator FCP_Port;
b) a read command completed with the data count smaller than FCP_DL and the FCP_RESID_UNDER bit is set to zero;
c) a read command completed with the data count smaller than FCP_DL, the FCP_RESID_UNDER bit is set to one, and the data count plus FCP_RESID is not equal to FCP_DL; or
d) an ABTS is received.

The Exchange Originator (i.e., initiator FCP_Port) shall detect the following errors for bidirectional commands:

a) a bidirectional command completed with the write data count smaller than FCP_DL and the FCP_RESID_UNDER bit is set to zero;
b) a bidirectional command completed with the read data count smaller than FCP_BIDIRECTION_READ_DL and the FCP_BIDI_READ_RESID_UNDER bit is set to zero;
c) a bidirectional command completed with the write data count smaller than FCP_DL, the FCP_RESID_UNDER bit is set to one, and the write data count plus FCP_RESID is not equal to FCP_DL; and
d) a bidirectional command completed with the read data count smaller than FCP_BIDIRECTIONAL_READ_DL, the FCP_BIDI_READ_RESID_UNDER bit is set to one, and the read data count plus FCP_BIDIRECTIONAL_READ_RESID is not equal to FCP_BIDIRECTIONAL_READ_DL.

The Exchange Originator may also identify and recover additional errors as described in 12.4.

The Exchange Responder shall also initiate error detection and recovery after a Sequence error is detected in a Sequence transmitted from an initiator FCP_Port to a target FCP_Port (see 12.3.5).

The Exchange Responder (i.e., target FCP_Port) may detect that REC_TOV times out and an expected FCP_CONF IU has not been received. The Exchange Responder may then use the methods described in 12.4 to determine the presence of an error, regardless of whether Exchange level or Sequence level error recovery is to be used.

12.2.3 Error detection mechanisms for acknowledged classes of service

Acknowledged classes of service provide the additional error detection mechanisms described below.

The Exchange Originator (i.e., initiator FCP_Port) shall detect any of the following errors:

a) after E_D_TOV times out and no ACK has been received for the FCP_CMND IU;
b) after E_D_TOV times out and no ACK_1 has been received for an FCP_DATA IU frame or no ACK_0 has been received for an FCP_DATA IU (see example in figure C.22);
c) after E_D_TOV times out and no ACK has been received for the FCP_CONF IU; or
d) an ACK with the F_CTL field Abort Sequence Condition bits set to Abort Sequence, Perform ABTS is received. (See FC-FS-3.)

The Exchange Originator may also identify and recover additional errors as described in 12.4.

The Exchange Responder (i.e., target FCP_Port) shall detect any of the following errors:

a) after E_D_TOV times out and no ACK has been received for the FCP_XFER_RDY IU (see example in figure C.6);
b) after E_D_TOV times out and no ACK_1 has been received for an FCP_DATA IU frame or no ACK_0 has been received for an FCP_DATA IU (see example in figure C.21); or
c) after E_D_TOV times out and no ACK has been received for the FCP_RSP IU.

The Exchange Responder may also identify and recover additional errors as described in 12.4.

If an ABTS is transmitted by a Sequence Initiator because it had detected a missing ACK and the BA_ACC response to the ABTS indicates the Sequence was correctly received by the Sequence Recipient, then no error detection or recovery is required.

12.3 Exchange level recovery using ABTS-LS

12.3.1 ABTS-LS overview

ABTS-LS is an FC-FS-3 protocol that recovers FCP_Port resources associated with an Exchange that is being terminated, either because of a task management request or an error.

ABTS-LS may be transmitted whether or not the FCP devices have agreed to Sequence level error recovery.

All initiator FCP_Ports shall be capable of transmitting ABTS-LS to terminate failing commands for later retry (see 9.2.2.5). All target FCP_Ports shall be capable of accepting and processing the ABTS-LS.

12.3.2 Initiator FCP_Port Exchange termination

The initiator FCP_Port terminating the Exchange transmits an ABTS-LS to the D_ID of the target FCP_Port of the Exchange being terminated. The ABTS-LS shall be generated using the OX_ID field and RX_ID field values of the Exchange to be aborted. FC-FS-3 allows ABTS-LS to be transmitted by an FCP_Port regardless of whether or not it has Sequence Initiative. Following the transmission of ABTS-LS, any Device_Data Frames received for this Exchange shall be discarded until the BA_ACC with the F_CTL field Last_Sequence bit set to one (i.e., last Sequence of the Exchange) is received from the target FCP_Port.

Exchange termination may not take effect immediately (e.g., if ABTS-LS is sent following transmission of a READ command, the initiator FCP_Port may receive some or all of the requested read data before receiving the BA_ACC for the ABTS-LS). The initiator FCP_Port shall be capable of receiving this data and providing BB_Credit in order for the target FCP_Port to transmit the BA_ACC.

After the processing of any task management function that clears commands (e.g., ABORT TASK SET, CLEAR TASK SET, and LOGICAL UNIT RESET), Exchange termination shall be invoked for all Exchanges not successfully terminated with an FCP_RSP IU status set to COMMAND CLEARED. (See 9.2.2.5).

Following receipt of the BA_ACC in response to an ABTS-LS, and after $R\_A\_TOV_{SEQ\_QUAL}$ has elapsed, the initiator FCP_Port shall transmit an RRQ ELS request (see FC-LS-2).

If a BA_ACC, BA_RJT, LOGO ELS, or PRLO ELS is not received from the target FCP_Port within 2 times $R\_A\_TOV_{ELS}$, then second level error recovery (see 12.5) shall be performed.

12.3.3 Target FCP_Port response to Exchange termination

When an ABTS-LS is received at the target FCP_Port, it shall abort the designated Exchange and return one of the following responses:

a) the target FCP_Port shall discard the ABTS-LS and transmit a LOGO ELS if the Nx_Port issuing the ABTS-LS is not currently logged in (i.e., no N_Port Login exists);
b) the target FCP_Port shall return BA_RJT with the F_CTL field Last_Sequence bit set to one (i.e., last Sequence of the Exchange) if the received ABTS-LS contains an assigned RX_ID field value and a FQXID that is unknown to the target FCP_Port; or
c) the target FCP_Port shall return BA_ACC with the F_CTL field Last_Sequence bit set to one (i.e., last Sequence of the Exchange).

Upon transmission of any of the above responses, the target FCP_Port may reclaim any resources associated with the designated Exchange after $R\_A\_TOV_{SEQ\_QUAL}$ has elapsed or an RRQ ELS request has been received.

If the RX_ID field is set to FFFFh, then target FCP_Ports shall qualify the FQXID of the ABTS-LS based only upon the combined values of the D_ID field, S_ID field, and the OX_ID field, not the RX_ID field.

If the Exchange resources were not reclaimed upon responding to the ABTS-LS, then they shall be reclaimed at the time the response to the RRQ ELS request is transmitted.

If an RRQ ELS request is received at the target FCP_Port, then it shall return one of the following responses:

a) the target FCP_Port shall discard the RRQ ELS request and transmit a LOGO ELS if the Nx_Port that transmitted the RRQ ELS request is not currently logged in (i.e., no N_Port Login);
b) if the received RRQ ELS contains an RX_ID field value, other than FFFFh, that is unknown to the target FCP_Port, then the target FCP_Port shall return LS_RJT with:
   A) the F_CTL field Last_Sequence bit set to one (i.e., last Sequence of the Exchange);
   B) the Reason Code set to 03h (i.e., Logical error) or 09h (i.e., Unable to process command request); and
   C) the Reason Code Explanation set to 17h (i.e., Invalid OX_ID-RX_ID combination); or
c) the target FCP_Port shall return LS_ACC with the Last Sequence of Exchange bit set to one.

12.3.4 Additional error recovery by initiator FCP_Port

This procedure may be used whether or not the FCP devices have agreed to Sequence retransmission.

If ULP_TOV times out and the Exchange is not complete, then the application client shall clear the Exchange resources using the ABORT TASK task management request or the initiator FCP_Port shall clear the Exchange resources by transmitting an ABTS-LS (see 4.9).

12.3.5 Additional error recovery by target FCP_Port

This procedure may be used whether or not the FCP devices have agreed to Sequence retransmission.

If a target FCP_Port detects a Sequence error, it shall discard the Sequence(s) based on the Exchange error policy specified by the F_CTL field Abort Sequence Condition (see FC-FS-3) bits in the first frame of the Exchange.

For acknowledged classes of service, if a target FCP_Port detects a Sequence error, then it may abort the Sequence by transmitting an ABTS with the PARAMETER field bit 0 set to one (i.e., Abort Sequence). If a Recovery Qualifier range is returned in the BA_ACC for the ABTS, then the target FCP_Port shall transmit an RRQ ELS request after $R\_A\_TOV_{SEQ\_QUAL}$ times out after receipt of the BA_ACC.

For unacknowledged classes of service, the target FCP_Port shall not attempt recovery for Sequence errors. The target FCP_Port shall depend on initiator FCP_Port timeouts for recovery.

Target FCP_Ports shall implement RR_TOV as described in 11.4 to facilitate recovery of resources allocated to an initiator FCP_Port that is no longer responding. The target FCP_Port may transmit a LOGO ELS to the initiator FCP_Port and terminate all open Exchanges for that initiator FCP_Port upon detection of the following:

a) the initiator FCP_Port has failed to perform target FCP_Port Exchange authentication within RR_TOV_{AUTH} (see FC-DA-2); or
b) $RR\_TOV_{SEQ\_INIT}$ times out without the initiator FCP_Port transmitting any expected Sequence for any open Exchange at this target FCP_Port (e.g., FCP write Data-In response to an FCP_XFER_RDY IU).

12.4 Sequence level error detection and recovery

12.4.1 Using information from REC ELS to perform Sequence retransmission

12.4.1.1 Polling Exchange state with REC ELS

An REC ELS is periodically transmitted by the initiator FCP_Port to poll each outstanding Exchange to determine if a command is progressing properly and if any Sequences have been received incorrectly. Timing of polling with the REC ELS is controlled by REC_TOV. REC_TOV is normally selected to be long enough that processing the transfers of Sequence Initiative in the Exchange and completing the Exchange occur before REC_TOV times out. If REC_TOV times out, then an REC ELS is performed. The information returned in the REC LS_ACC payload is compared with the expected state information known by the initiator FCP_Port and target FCP_Port. If the information is inconsistent, indicating that a link error occurred, then error recovery actions may be performed to complete the Exchange. Optional error detection and recovery procedures for acknowledged and unacknowledged classes of service are specified in 12.4.1.2, 12.4.1.3, 12.4.1.4, 12.4.1.5, 12.4.1.6, 12.4.1.7, and 12.4.1.8.

12.4.1.2 Detection of errors while polling with REC ELS

If an Exchange Originator receiving an acknowledged service Sequence detects a Sequence error, then it shall transmit an ACK frame with the F_CTL field Abort Sequence Condition bits set to “Abort Sequence, Perform ABTS” (see FC-FS-3) before issuing the REC ELS. The REC ELS for the Exchange containing the FCP_CMND IU shall be issued in a new Exchange.

If the response to the new Exchange issuing the REC ELS is an LS_RJT with the Reason Code set to “command not supported”, then the initiator FCP_Port shall assume the target FCP_Port is an FCP device not supporting error detection using the REC ELS. The initiator FCP_Port shall perform recovery by transmitting an ABTS-LS (see 12.3).

If an LS_ACC, LS_RJT, LOGO ELS, or PRLO ELS is not received from the target FCP_Port within 2 times $R\_A\_TOV_{ELS}$, then second level error recovery (see 12.5) shall be performed.

12.4.1.3 FCP_CMND IU recovery

This procedure may be used whether or not the FCP devices have agreed to Sequence retransmission.

If the FCP_CMND IU was not received by the target FCP_Port (e.g., the initiator FCP_Port receives an LS_RJT for the REC ELS with the Reason code set to Logical error and the Reason Code Explanation set to Invalid OX_ID-RX_ID combination), then retransmit the FCP_CMND IU using a new Exchange. If the precise delivery function is enabled, then the CRN value shall remain the same in the retransmitted FCP_CMND IU.

If the LS_ACC for the REC ELS indicates that the FCP_CMND IU was received by the target FCP_Port and that no reply Sequence has been sent (i.e., by indicating that the initiator FCP_Port does not hold Sequence Initiative, and that the Exchange is not complete), then the command is in process and no recovery is needed at this time. At a minimum interval of REC_TOV, the REC ELS shall be retransmitted to more quickly determine if a reply Sequence has been lost.

For examples of such recoveries, see figure C.1 and figure C.2.

12.4.1.4 FCP_XFER_RDY IU recovery

This procedure shall be used only by FCP devices that have agreed to Sequence retransmission.

If the LS_ACC for an REC ELS indicates that an FCP_XFER_RDY IU was sent by the target FCP_Port (i.e., by indicating that the initiator FCP_Port holds Sequence Initiative, that all bytes were not transferred, and that the Exchange is not complete), but not received by the initiator FCP_Port, then the initiator FCP_Port shall issue an SRR FCP_LS request in a new Exchange to request retransmission of the FCP_XFER_RDY IU. To avoid race conditions between the LS_ACC for an REC ELS and an FCP_XFER_RDY IU, the initiator FCP_Port should wait REC_TOV after receiving the LS_ACC for an REC ELS before transmitting an SRR FCP_LS request to recover for a lost FCP_XFER_RDY IU. If the initiator FCP_Port receives an FCP_XFER_RDY IU for an Exchange after transmitting an REC ELS or within REC_TOV time after receiving an LS_ACC for an REC ELS, then it shall continue normal processing of the FCP I/O operation for that Exchange and ignore the contents of the LS_ACC.

The target FCP_Port shall first transmit the FCP_ACC for the SRR FCP_LS request and then shall retransmit the FCP_XFER_RDY IU in a new Sequence containing the same Relative Offset as the originally transmitted FCP_XFER_RDY IU. After the FCP_XFER_RDY IU is successfully received, the FCP I/O operation continues normally.

For examples of this type of recovery, see figure C.5 and figure C.6.

12.4.1.5 FCP_RSP IU recovery

This procedure shall be used only by FCP devices that have agreed to Sequence retransmission.

An error in transmitting an FCP_RSP IU is detected if:

a) the LS_ACC for the REC ELS indicates that an FCP_RSP IU was sent by the target FCP_Port and no FCP_CONF IU was requested (i.e., E_STAT indicates that the Exchange is complete), but the initiator FCP_Port has not yet received the FCP_RSP IU; or
b) the LS_ACC for the REC ELS indicates that an FCP_RSP IU Sequence was sent by the target FCP_Port and an FCP_CONF IU was requested (i.e., E_STAT indicates that the Exchange is not complete, that the initiator FCP_Port has Sequence Initiative, and that, if the data transfer was from the initiator FCP_Port to the target FCP_Port, the data transfer indicates that all of the bytes expected to be transferred by the command have been transferred.)

When an error in transmitting an FCP_RSP IU is detected, the initiator FCP_Port shall transmit an SRR FCP_LS request in a new Exchange to request retransmission of the FCP_RSP IU. The target FCP_Port shall first transmit the LS_ACC for the SRR FCP_LS request, then shall retransmit the FCP_RSP IU in a new Sequence.

An Exchange carrying a command that was terminated by a CHECK CONDITION requesting an FCP_CONF IU prior to transferring data may have the same REC ELS values as an Exchange carrying a command having an FCP_XFER_RDY IU not received by the initiator FCP_Port. For a command transferring data from the initiator FCP_Port to the target FCP_Port with a non-zero FCP_DL, the parameters in the SRR FCP_LS request shall indicate that an FCP_XFER_RDY IU is expected from the target FCP_Port. The target FCP_Port is aware of the actual present state of the transfer and response and shall either retransmit the FCP_XFER_RDY IU or, if the actual data transfer length for the command was zero, retransmit the FCP_RSP IU.

For non-tagged command queuing operations, the target FCP_Port shall retain the Exchange information until:

a) the next FCP_CMND IU has been received for that LUN from the same initiator FCP_Port;
b) an FCP_CONF IU is received for the Exchange; or
c) after $RR\_TOV_{SEQ\_INIT}$ times out.

NOTE 6 - Non-tagged command queuing is specified in previous versions of the SCSI architecture model.

For tagged command queuing operations, the target FCP_Port shall retain Exchange information until:

a) an FCP_CONF IU is received for the Exchange; or
b) after $RR\_TOV_{SEQ\_INIT}$ times out.

The Exchange information retained shall include data transfer information, data descriptors, and FCP_RSP IU information.

If retransmission is enabled between the initiator FCP_Port and target FCP_Port, then FCP_RSP IU information shall be:

a) discarded $RR\_TOV_{SEQ\_INIT}$ after the FCP_RSP IU was transmitted to the initiator FCP_Port;
b) discarded after a new FCP_CMND IU is received with:
    A) the same OX_ID field value;
    B) the same S_ID field value; and
    C) the same task retry identifier is received; or
c) discarded after an FCP_CONF IU is received.

If retransmission is not enabled between the initiator FCP_Port and target FCP_Port, then FCP_RSP information may be discarded immediately after the FCP_RSP IU has been transmitted to the initiator FCP_Port.

The value of $RR\_TOV_{SEQ\_INIT}$ is set using the Fibre Channel Port Control mode page (see 10.4.11).

If task retry identification has been agreed to by both the initiator FCP_Port and target FCP_Port, then the same task retry value identifier value shall not be used within $RR\_TOV_{SEQ\_INIT}$.

Examples of FCP_RSP IU recoveries are provided in figure C.8, figure C.9, figure C.10, figure C.11, and figure C.12.

12.4.1.6 FCP_DATA IU recovery - write operations

This procedure shall be used only by FCP devices that have agreed to Sequence retransmission.

If the LS_ACC for an REC ELS indicates that an FCP_DATA IU was sent by the initiator FCP_Port, but not received by the target FCP_Port (i.e., the data received count in the REC ELS response is smaller than what the initiator FCP_Port sent, and the target FCP_Port indicates it does not hold Sequence Initiative), then the initiator FCP_Port shall transmit an SRR FCP_LS request in a new Exchange to request retransmission of an FCP_XFER_RDY IU to request the missing data. The target FCP_Port discards the Sequence in error, but does not initiate any recovery action for Class 3 (see 12.3.5). After first transmitting the FCP_ACC for the SRR FCP_LS request, the target FCP_Port transmits an FCP_XFER_RDY IU in a new Sequence with the FCP_DATA_RO field in the FCP_XFER_RDY IU set to the value of the RELATIVE OFSET field in the SRR. The initiator FCP_Port responds with the requested data.

The FCP_DATA IU shall be retransmitted in a new Sequence. For acknowledged classes, the SEQ_CNT field value shall be one greater than that used to transmit the last Sequence, usually the ABTS. For unacknowledged classes, the SEQ_CNT field value may start at zero, even if continuously increasing sequence count is being used.

Examples of data recovery during write operations are provided in figure C.13, figure C.14, figure C.15, and figure C.16.

12.4.1.7 FCP_DATA IU recovery - read operations

This procedure shall be used only by FCP devices that have agreed to Sequence retransmission.

If the LS_ACC for the REC ELS indicates that data was sent by the target FCP_Port but not successfully received by the initiator FCP_Port (i.e., by indicating a data sent count greater than the initiator FCP_Port has successfully received), then the initiator FCP_Port shall transmit an SRR FCP_LS request in a new Exchange to request retransmission of the FCP_DATA IU that was not successfully received. The initiator FCP_Port shall set the RELATIVE OFFSET field in the SRR FCP_LS request to that of the next data requested. If the initiator FCP_Port is unable to determine the Relative Offset of the next data requested, then the initiator FCP_Port shall set the RELATIVE OFFSET field to zero. The target FCP_Port shall:

1) transmit the FCP_ACC for the SRR FCP_LS request;
2) retransmit the requested data specified by the SRR FCP_LS request in a new Sequence; and
3) complete the Exchange in the normal manner, including transmitting or retransmitting the FCP_RSP IU.

If the target FCP_Port responds to the SRR FCP_LS request with an FCP_RJT and an FCP_RSP IU has not yet been sent or is again requested, then the device server shall transmit an FCP_RSP IU with CHECK CONDITION status and sense data containing a sense key of HARDWARE ERROR and an additional sense code of INITIATOR DETECTED ERROR MESSAGE RECEIVED.

The FCP_DATA IU shall be retransmitted in a new Sequence. For acknowledged classes, the SEQ_CNT field value shall be one greater than that used to transmit the last Sequence, usually the ABTS. For unacknowledged classes, the SEQ_CNT field value may start at zero, even if continuously increasing sequence count is being used.

It is the responsibility of the initiator FCP_Port to determine the action (e.g., retry, allow ULP timeout, or return status to ULP) based on the information determined by the REC ELS and other internal states. The target FCP_Port does not initiate error recovery for Class 3 (see 12.3.5).

Examples of data recovery during read operations are provided in figure C.17, figure C.18, figure C.19, and figure C.20.

12.4.1.8 FCP_CONF IU recovery

This procedure may be used whether or not an FCP device has agreed to Sequence retransmission.

This recovery procedure is used by target FCP_Ports using all service classes.

Target FCP_Ports that implement confirmed completion shall set the RX_ID field value to a unique value other than FFFFh for each Exchange to enable unambiguous recovery.

If the target FCP_Port has requested that the initiator FCP_Port transmit an FCP_CONF IU by setting the FCP_CONF_REQ in the FCP_RSP IU, then the target FCP_Port may periodically poll the initiator FCP_Port by transmitting an REC ELS to the initiator FCP_Port to determine if the FCP_CONF has been transmitted. Timing of polling with the REC ELS is controlled by REC_TOV.

If the initiator FCP_Port has sent the FCP_CONF IU, the response to the REC ELS from the target FCP_Port shall be a LS_RJT with the Reason Code set to Logical error and Reason Code Explanation set to Invalid OX_ID-RX_ID combination. The target FCP_Port shall assume that the FCP_CONF IU was sent and release the Exchange.

If the initiator FCP_Port has received the FCP_RSP IU with the FCP_CONF_REQ bit set to one and has not sent the FCP_CONF IU before the REC ELS is received, then the REC ELS response shall be an LS_ACC indicating the Exchange is still open. In this case the target FCP_Port shall wait REC_TOV and, if the FCP_CONF IU has not been received, transmit another REC ELS. The target FCP_Port shall repeat this process until the FCP_CONF IU is received, a new FCP_CMND IU is received with the same OX_ID field value as the Exchange waiting for the FCP_CONF IU, or until the Exchange is aborted.

If another FCP_CMND IU is received by the target FCP_Port with the same OX_ID field value as an Exchange waiting for an FCP_CONF IU and with the RX_ID field value unassigned, then the target FCP_Port shall assume that the FCP_CONF IU was sent and release the Exchange.

Examples of recovery of FCP_CONF IUs are provided in figure C.23, figure C.24, and figure C.25.

12.4.2 Additional error recovery requirements

12.4.2.1 Error indicated in ACK

If an ACK is received with the F_CTL field Abort Sequence Condition bits set to Abort Sequence, Perform ABTS, then the Sequence Initiator shall transmit an ABTS for the Sequence. After R_A_TOV times out, an RRQ ELS request shall be transmitted by the Sequence Initiator.

12.4.2.2 Missing ACK

FC-FS-3 requires that an ABTS(Sequence) be transmitted by a Sequence Initiator detecting a missing ACK. If no ACK has been received within E_D_TOV, then the target FCP_Port shall abort the Sequence by transmitting an ABTS request with the PARAMETER field bit 0 set to one (i.e., Abort Sequence). If a Recovery Qualifier range is returned in the BA_ACC for the ABTS, then the target FCP_Port shall transmit an RRQ ELS request at least $R\_A\_TOV_{SEQ\_QUAL}$ after receipt of the BA_ACC. Adjustment of subsequent sequence counts may be required as specified by FC-FS-3.

12.4.2.3 Distinguishing Exchange to be aborted

When OX_ID field values are reused within R_A_TOV and RX_ID field values are not used, and if there is a missing ACK to an FCP_RSP IU, a target FCP_Port may attempt to abort a more recent Exchange that used the same OX_ID field value. To prevent that, a target FCP_Port using acknowledged service behavior and performing error recovery shall:

a) set the RX_ID field to a value other than FFFFh to distinguish outstanding Exchanges as described in FC-FS-3; or
b) always request FCP_CONF IU.

If a Sequence error is detected for an FCP_DATA IU performing a Data-Out action, the target FCP_Port shall transmit an ACK frame with the Abort Sequence Condition bits set to “Abort Perform ABTS”.

Examples of data recovery for acknowledged services are shown in Annex C.

ABTS-LS shall be transmitted for Exchanges that were not successfully recovered by the specified error recovery procedures.

12.5 Second-level error recovery

12.5.1 ABTS error recovery

If a response to an ABTS is not received within 2 times $R\_A\_TOV_{ELS}$, then the initiator FCP_Port may transmit the ABTS again, attempt other retry operations allowed by FC-FS-3, or explicitly logout the target FCP_Port. If those retry operations attempted are unsuccessful, then the initiator FCP_Port shall explicitly logout (i.e., transmit a LOGO ELS) the target FCP_Port. All outstanding Exchanges with that target FCP_Port are terminated at the initiator FCP_Port.

12.5.2 REC ELS request error recovery

If a response to an REC ELS is not received within 2 times $R\_A\_TOV_{ELS}$, then the initiator FCP_Port shall:

1) transmit an ABTS-LS for the REC ELS request followed by an RRQ ELS request if a BA_ACC is received for the ABTS-LS; and
2) transmit another REC ELS request in a new Exchange.

If the response to the second REC ELS is not received within 2 times $R\_A\_TOV_{ELS}$, then the initiator FCP_Port should transmit an ABTS-LS for the second REC ELS followed by an RRQ ELS if a BA_ACC is received for the ABTS-LS.

Other retry mechanisms after the second REC ELS request fails are optional and, if implemented, shall comply with FC-FS-3.

ABTS-LS may be required to clear resources associated with the original failing Exchange if the retry mechanisms are not successful.

See figure C.26, figure C.27, figure C.28 and figure C.29.

12.5.3 SRR FCP_LS request error recovery

If a response to an SRR FCP_LS request is not received within 2 times $R\_A\_TOV_{ELS}$, then the initiator FCP_Port shall transmit an ABTS-LS for the SRR FCP_LS request followed by an RRQ ELS if a BA_ACC is received for the ABTS-LS. The initiator FCP_Port shall then transmit an ABTS-LS for the original Exchange.

See figure C.30, figure C.31, figure C.32, and figure C.33.

12.6 Responses to FCP type frames before port login or process login

If a target FCP_Port receives an FCP_CMND IU from an FCP_Port that is not successfully logged into the target FCP_Port using either an implicit or explicit login, then it shall discard the FCP_CMND IU and, in a new Exchange, transmit a LOGO ELS request to that FCP_Port. No Exchange is created in the target FCP_Port for the discarded request, and the Originator of the discarded request terminates the Exchange associated with the discarded request and any other open Exchanges for the target FCP_Port transmitting the LOGO ELS.

If a target FCP_Port receives an FCP_CMND IU from an FCP_Port that has not successfully completed either implicit or explicit Process Login with the target FCP_Port, then it shall discard the FCP_CMND IU and transmit a PRLO ELS to the initiator FCP_Port. No Exchange is created in the recipient FCP_Port for the discarded request, and the Originator of the discarded request terminates the Exchange associated with the discarded request.

If an FCP device receives a frame of category 0001b or 0011b (i.e., solicited data or solicited control) and the FCP device has not performed successful implicit or explicit login and Process Login with the source of the frame, then the FCP device shall discard and ignore the content of the frame. If login is not completed, then the FCP device may transmit a LOGO ELS request to the source of the unexpected frame. If login is completed, but Process Login is not completed, then the FCP device may transmit a PRLO ELS request to the source of the unexpected frame.

Annex A
(normative)

SAM-5 mapping to FCP-4

FCP-4 services are provided to the application client by the initiator FCP_Port to request and manage commands and task management functions as described by the SAM-5 standard. SAM-5 further defines how the target FCP_Port enables the device server to receive and process the commands addressed to a logical unit and enables the task router and task manager(s) to receive and process task management functions. The Fibre Channel protocol is described in terms of the services provided by the initiator FCP_Port and target FCP_Port.

See table A.1 for the mapping of SAM-5 terms to the equivalent FCP-4 objects and identifiers.

Table A.1 - Mapping of SAM-5 terms to FCP-4 objects and identifiers

SAM-5 terms	Equivalent FCP-4 terms
initiator port identifier	address identifier of initiator FCP_Port
target port identifier	address identifier of target FCP_Port
initiator port name	N_Port_Name of initiator FCP_Port
target port name	N_Port_Name of target FCP_Port
command identifier[a]	OX_ID
I_T nexus	address identifier of initiator FCP_Port + address identifier of target FCP_Port
I_T_L_Q nexus[b]	fully qualified exchange identifier + FCP_LUN

a) If retransmission is enabled, the task retry identifier is also used to construct the command identifier.
b) If retransmission is enabled, the task retry identifier is also used to construct the I_T_L_Qnexus.

See table A.2 for the definitions of the terms used by this standard and the equivalent SAM-5 names of the terms, the name of the standard where the procedure terms are defined, the standard where the binary contents of the terms are defined, and the routing of the terms. The routing shows:

a) the originating object of the term;
b) the object that is the final destination of the term; and
c) the objects that the term moves through to reach the final destination object.

Table A.2 - Procedure terms

FCP terms	Standard where term defined	Standard where binary contents of term defined	Term routing
application client buffer offset	SAM-5	SAM-5	DS → targ → init
data buffer size	SAM-5	SAM-5	AC → init
command descriptor block	SAM-5	SAM-5/cmd[a]	AC → init → targ → DS
command priority	SAM-5	SAM-5/cmd[a]	AC → init → targ → DS
Data-In Buffer	SAM-5	cmd[b]	DS → targ → init → AC
Data-Out Buffer	SAM-5	cmd[b]	AC → init → targ → DS

Key: AC = application client, cmd = SCSI command standards, DS = device server, init = SCSI initiator port, TM = task manager, targ = SCSI target port

a) The portions not defined in SAM-5 are defined in the SCSI command standards (e.g. SPC-4).
b) Parameter lists are defined within one of the SCSI command standards (e.g., SPC-4). SCSI standards do not define non-parameter list information.

Table A.2 - Procedure terms (Continued)

FCP terms	Standard where term defined	Standard where binary contents of term defined	Term routing
device server buffer	SAM-5	cmd[b]	DS → targ → init
link control function	this standard	this standard	AC → init → targ
logical unit number	SAM-5	this standard	AC → init → targ → DS **or** AC → init → targ → TM **or** DS → targ → init
request byte count	SAM-5	SAM-5	DS → targ
service response	SAM-5	this standard	DS → targ → init → AC **or** targ → DS
status	SAM-5	SAM-5	DS → targ → init → AC
status qualifier	SAM-5	SAM-5	DS → targ → init → AC
command identifier	SAM-5	this standard	AC → init → targ → DS **or** AC → init → targ → TM **or** DS → targ → init
target port identifier	SAM-5	this standard	AC → init → targ → DS **or** AC → init → targ → TM **or** DS → targ
target port identifier + initiator port identifier	this standard	this standard	targ → DS **or** targ → TM
task attribute	SAM-5	this standard	AC → init → targ → DS

Key: AC = application client, cmd = SCSI command standards, DS = device server, init = SCSI initiator port, TM = task manager, targ = SCSI target port

a) The portions not defined in SAM-5 are defined in the SCSI command standards (e.g. SPC-4).
b) Parameter lists are defined within one of the SCSI command standards (e.g., SPC-4). SCSI standards do not define non-parameter list information.

Annex B
(informative)

FCP examples

B.1 Examples of the use of FCP Information Units (IUs)

B.1.1 Overview of examples

This annex provides examples of the use of FCP IUs. The functions enclosed in square brackets summarize actions that are not specified by this standard, but are typically performed by SCSI initiators or targets. Sequence streaming may be performed between any two IUs that do not transfer Sequence Initiative.

B.1.2 Read command

A read command with a single data IU is shown in table B.1.

Table B.1 - Read command example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare data transfer]
	← I3, FCP_DATA	Data-In action
		[Prepare response message]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.3 Write command

A write command with three data IUs and using FCP_XFER_RDY is shown in table B.2.

Table B.2 - Write command example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare Data-Out transfer buffer]
	← I1, FCP_XFER_RDY	First data delivery request
First Data-Out Action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Second data delivery request
Second Data-Out Action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Last data delivery request
Last Data-Out Action	T6, FCP_DATA →	
		[Prepare response message]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.4 Command with no data transfer or with check condition

A non-data command or a command terminating without data transfer is shown in table B.3.

Table B.3 - Command without data transfer example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[perform command]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.5 Read command with multiple FCP_DATA IUs

A read command with multiple FCP_DATA IUs is shown in table B.4.

Table B.4 - Read command with multiple FCP_DATA IUs example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare data transfer]
	← I3, FCP_DATA	Data-In action
	← I3, FCP_DATA	Data-In action
	← I3, FCP_DATA	Data-In action
		[Prepare response message]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.6 Write command with FCP_XFER_RDY disabled

A write command performed with FCP_XFER_RDY disabled is shown in table B.5. Only the first transfer is performed without a requesting FCP_XFER_RDY.

Table B.5 - Write command with FCP_XFER_RDY disabled example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T2, FCP_CMND →	
Data-Out Action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Second data delivery request
Data-Out Action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Last data delivery request
Data-Out Action	T6, FCP_DATA →	
		[Prepare response message]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.7 Bidirectional command with write before read

A bidirectional command with a single data IU transferred in each direction is shown in table B.6. The command in the example accepts write data before returning read data.

Table B.6 - Bidirectional command with write before read example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare Data-Out transfer buf-fer]
	← I1, FCP_XFER_RDY	Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
		[Prepare Data-In transfer]
	← I3, FCP_DATA	Data-In action
		[Prepare response message]
	← I4, FCP_RSP	Response
[indicate command completion]		

B.1.8 Bidirectional command with read before write

A bidirectional command with a single data IU transferred in each direction is shown in table B.7. The command in the example accepts write data before returning read data.

Table B.7 - Bidirectional command with read before write example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare Data-In transfer]
	← I3, FCP_DATA	Data-In action
		[Prepare Data-Out transfer buf-fer]
	← I1, FCP_XFER_RDY	Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
		[Prepare response message]
	← I4, FCP_RSP	Response
[indicate command completion]		

B.1.9 Bidirectional command with write first and write FCP_XFER_RDY disabled

A bidirectional command with three write data IUs and one read data IU is shown in table B.8. The command in the example accepts write data before returning read data. The initial write FCP_XFER_RDY IU has been disabled during Process Login.

Table B.8 - Bidirectional command with write first and write FCP_XFER_RDY disabled example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
Data-Out action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Second Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Last Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
		[Prepare Data-In transfer]
	← I3, FCP_DATA	Data-In action
		[Prepare response message]
	← I4, FCP_RSP	Response
[indicate command completion]		

B.1.10 Bidirectional command with intermixed writes and reads

A bidirectional command with three data IUs transferred in each direction is shown in table B.9. The command in the example accepts some write data before returning read data, but intermixes writes and reads thereafter.

Table B.9 - Bidirectional command with intermixed writes and reads example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request	T1, FCP_CMND →	
		[Prepare Data-Out transfer buffer]
	← I1, FCP_XFER_RDY	First Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
		[Prepare Data-In transfer]
	← I3, FCP_DATA	First Data-In action
	← I1, FCP_XFER_RDY	Second Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
	← I1, FCP_XFER_RDY	Last Data-Out delivery request
Data-Out action	T6, FCP_DATA →	
	← I3, FCP_DATA	Second Data-In action
	← I3, FCP_DATA	Last Data-In action
		[Prepare response message]
	← I4, FCP_RSP	Response
[indicate command completion]		

B.1.11 Write command with confirmed completion

A write command with confirmed completion is shown in table B.10.

Table B.10 - Write command with confirmed completion example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request (WRITE)	T1, FCP_CMND →	
		[Prepare data transfer]
	← I1, FCP_XFER_RDY	Data delivery request
Data-Out action	T6, FCP_DATA →	
		[Prepare response message]
	← I5, FCP_RSP	Response, with FCP_CONF_REQ
[indicate command completion]		
Confirm completion	T12, FCP_CONF →	[Accept confirmation]

B.1.12 Task management function

An example task management function is shown in table B.11. Additional link services may be required in some cases to complete the activities initiated by the task management function.

Table B.11 - Task management function example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request, no CDB	T1, FCP_CMND →	
		[Do Task Management]
	← I4, FCP_RSP	Response
[Indicate task management complete]		

B.1.13 Class 2 write command example, frame level

A chart of the Sequences and frames typically transmitted to perform a write command using Class 2 is shown in figure B.1. All frames of a Sequence have a frame level FC-FS-3 acknowledgment returned automatically as part of the link control.

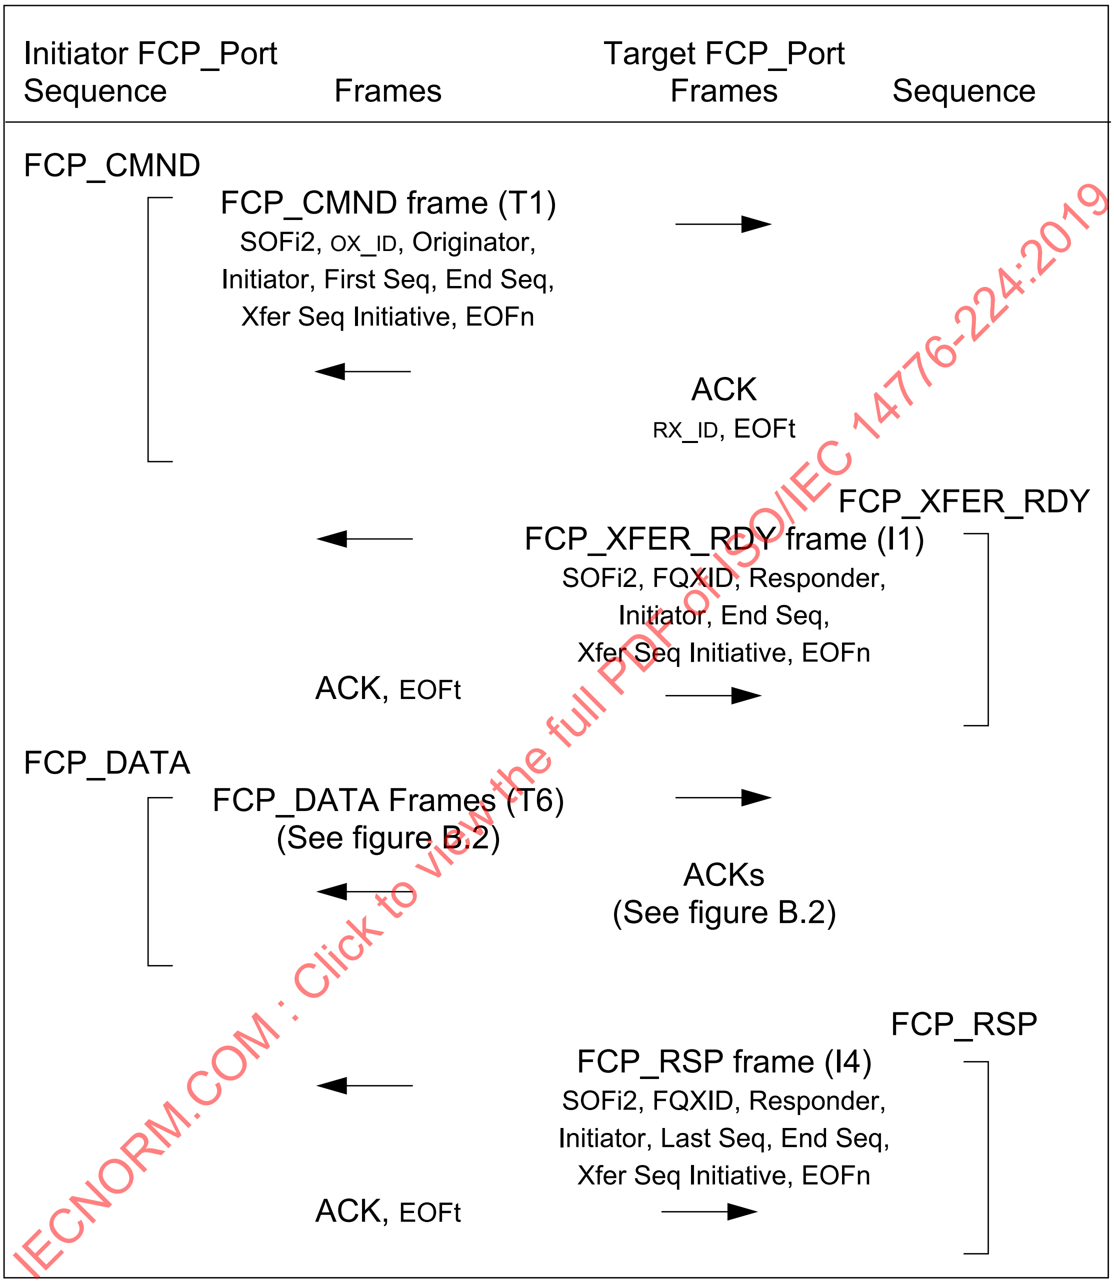

Figure B.1 - Class 2 write command example

A chart of the Sequences and frames typically transmitted to transfer write data using Class 2 is shown in figure B.2

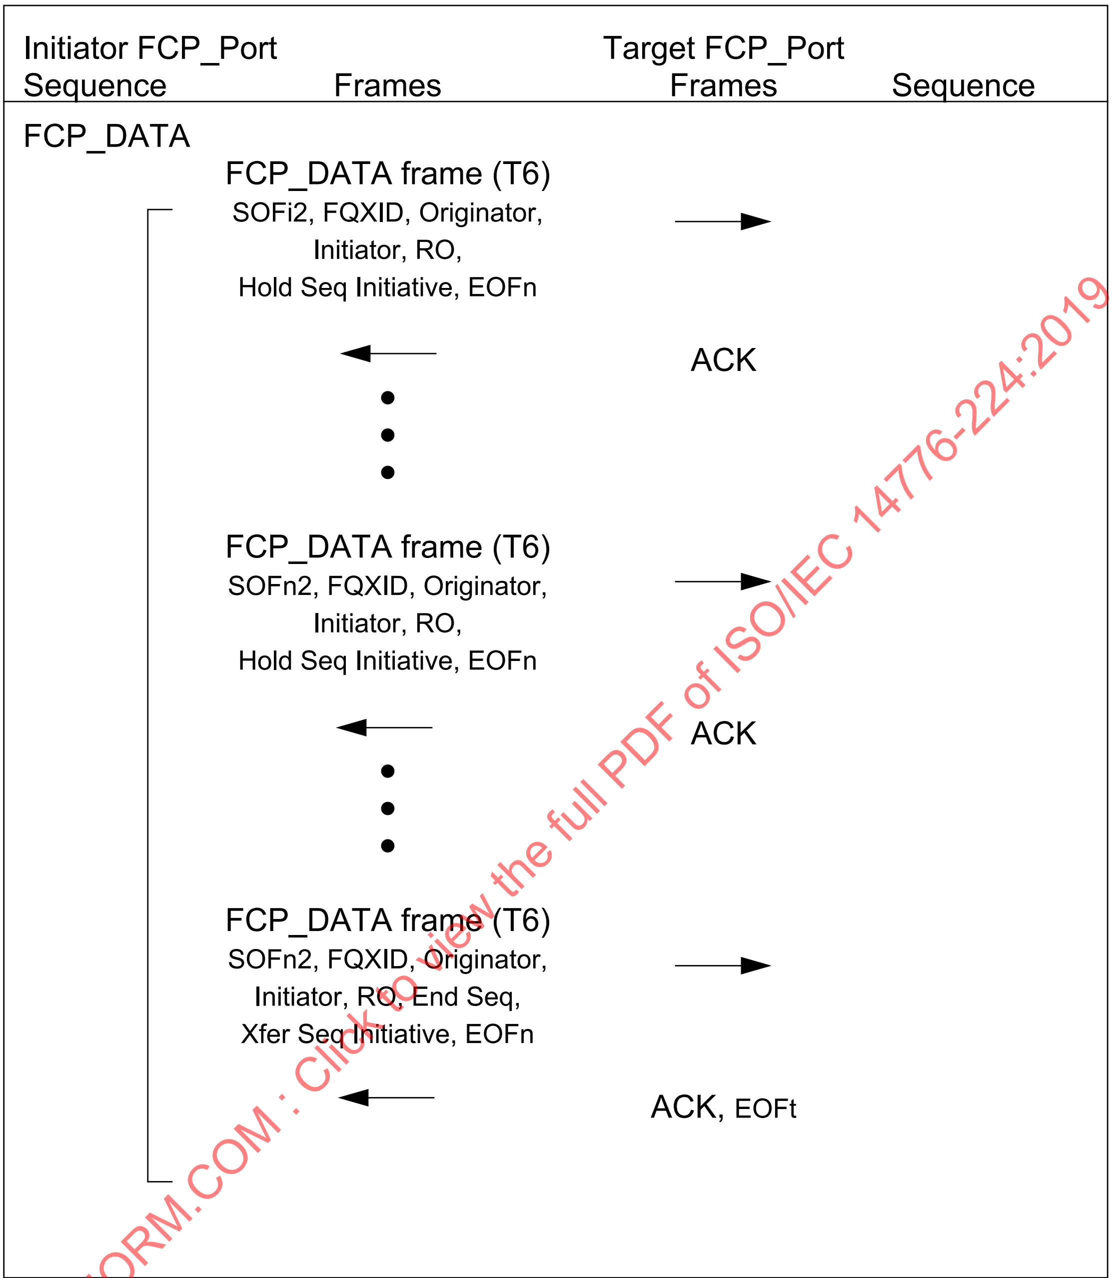

Figure B.2 - Class 2 write data example

B.1.14 Class 2 read command example, frame level

A chart of the Sequences typically transmitted to perform a read command using Class 2 is shown in figure B.3. All frames of a Sequence have a frame level FC-FS-3 acknowledgment returned automatically as part of the link control.

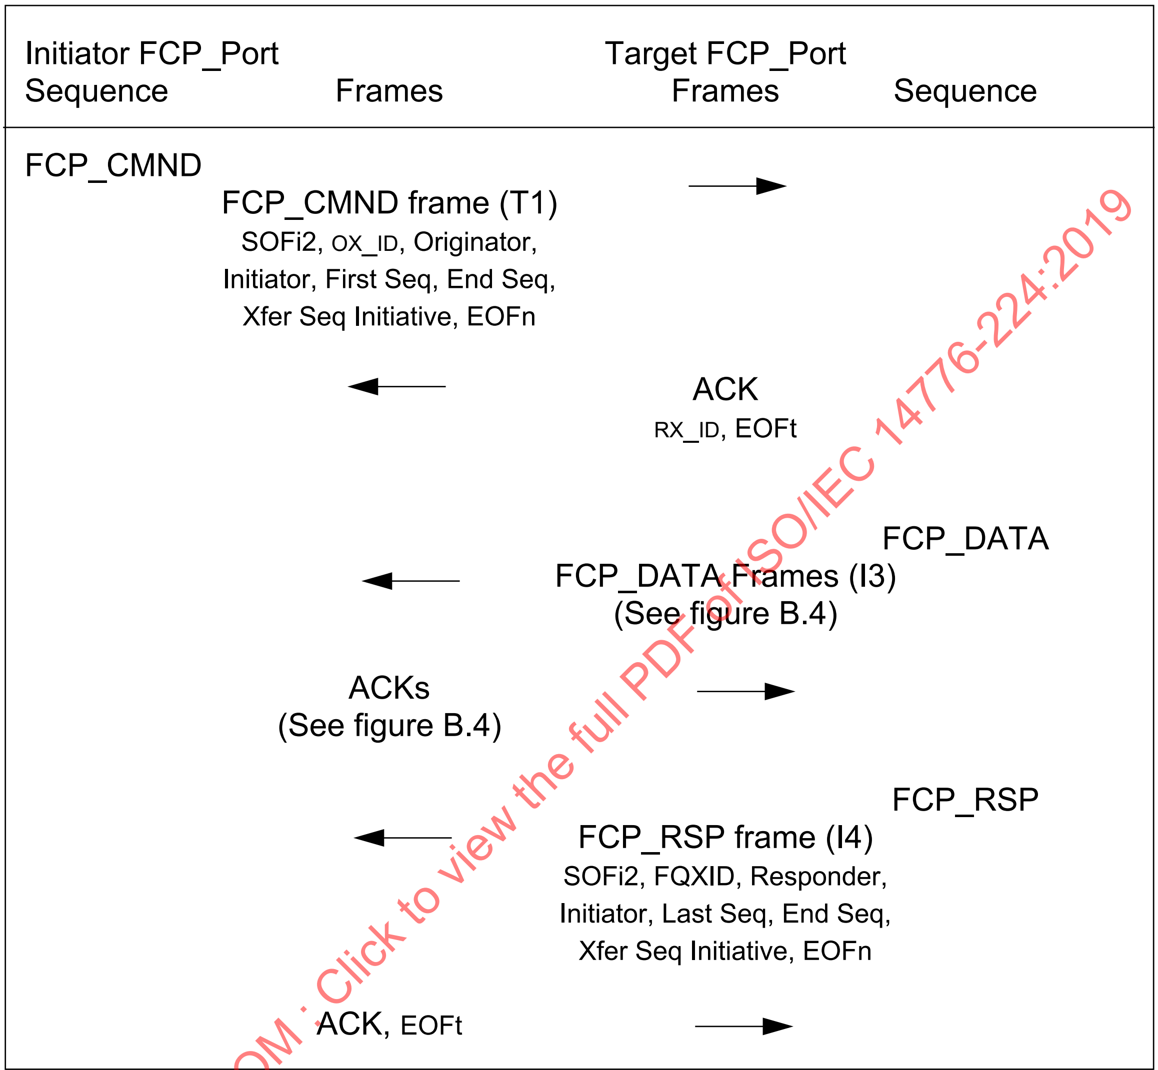

Figure B.3 - Class 2 read command example

A chart of the Sequences and frames typically transmitted to transfer read data using Class 2 is shown in figure B.4

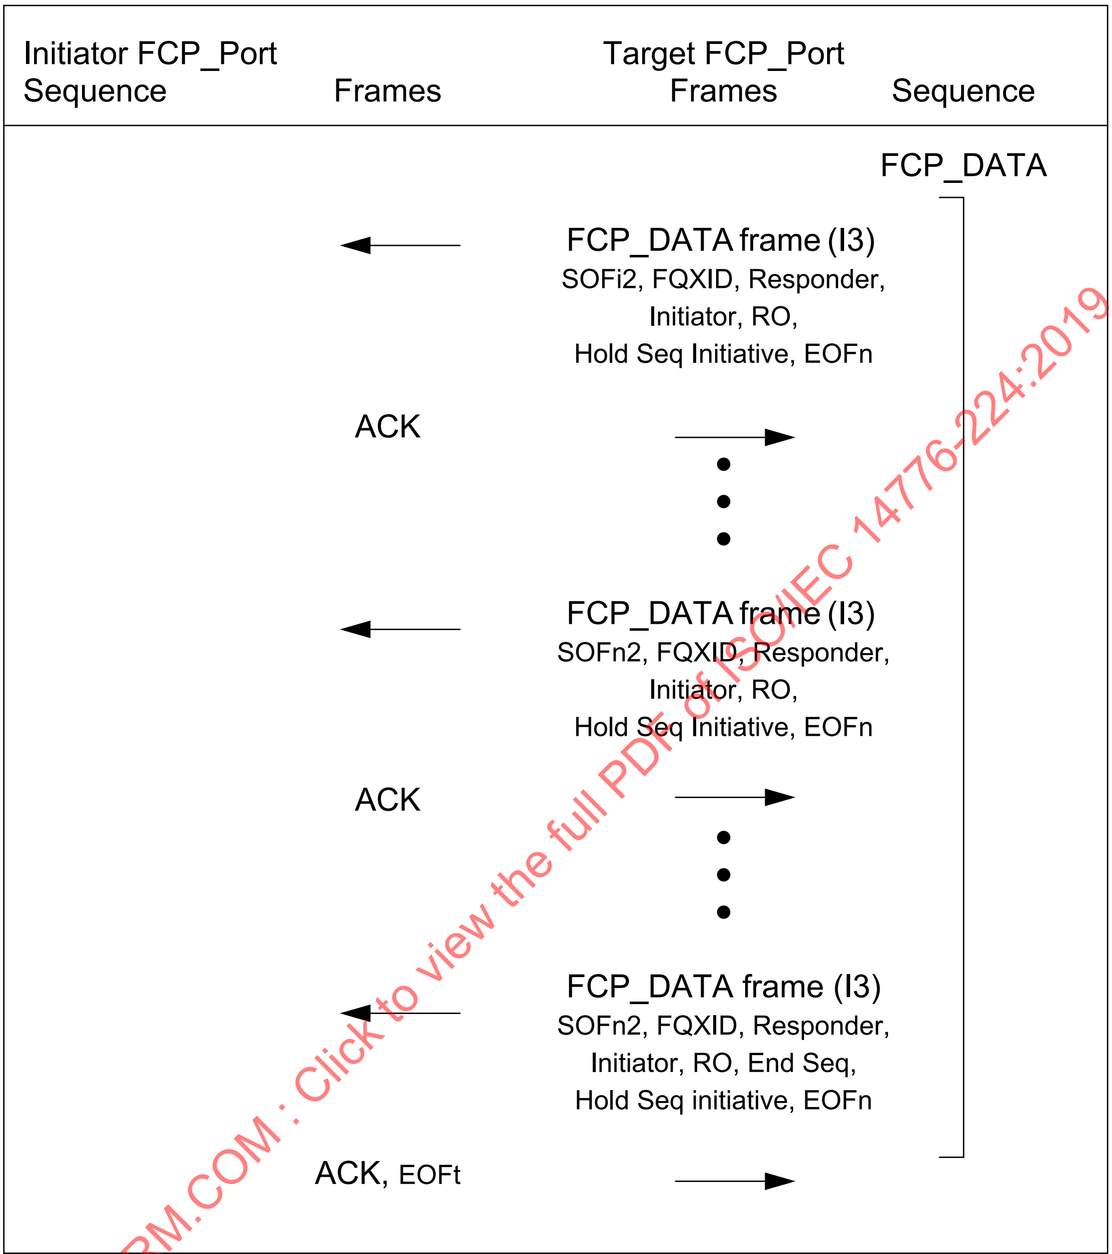

Figure B.4 - Class 2 read data example

Annex C
(informative)

Error detection and recovery examples

This informative annex diagrams various error detection and recovery procedures for FCP devices conforming to this standard. These examples include cases where recovery mechanisms specified by FC-FS-3 interact with recovery mechanisms specified by this standard. The conventions for the diagrams are shown in table C.1.

Table C.1 - Diagram drawing conventions

Drawing convention	Meaning
	Acknowledged or unacknowledged frame.
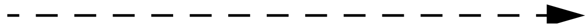	Acknowledgement frame.
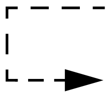	Timeout value exceeded, caused transmission of IU, FC-4 LS, or ELS.
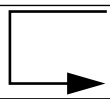	IU or ELS received is processed to transmit IU, FC-4 LS, or ELS.
X	Frame lost or dropped.
CI Continue 	Error detection complete. Operation continues with specified error recovery if continuously increasing Sequence count prerequisites are met.
Continue 	Error detection complete. Operation continues with specified error recovery if continuously increasing Sequence count prerequisites are not met.
cnt	SEQ_CNT
Initiator	initiator FCP_Port
param	relative offset/task retry identifier
RO	FCP_DATA_RO
seq	SEQ_ID
Target	target FCP_Port

An example of a lengthy FCP_CMND or lost ACK is shown in figure C.1.

Status detection

Unacknowledged classes | **Acknowledged classes**

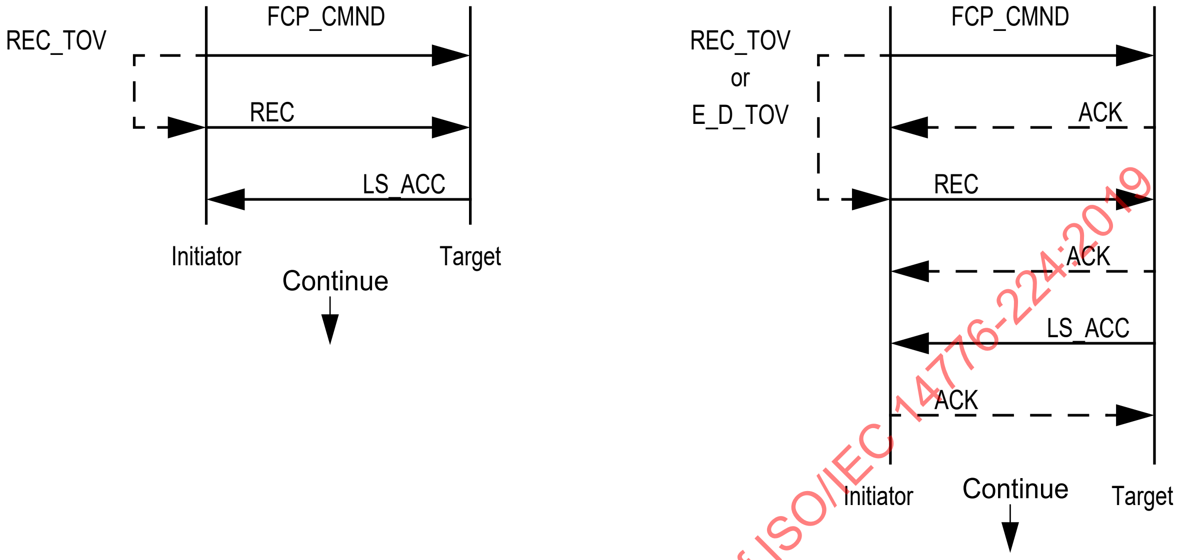

The REC ELS may be used at any time to ascertain status of an Exchange. It may also be used in conjunction with an ABTS (Sequence) to obtain additional information useful in the error recovery process.

No error recovery

The LS_ACC for the REC ELS indicates the Exchange is open and the target FCP_Port holds Sequence Initiative. No error recovery is required.

Figure C.1 - Lengthy FCP_CMND or lost ACK

An example of a lost FCP_CMND using unacknowledged classes is shown in figure C.2.

Error detection

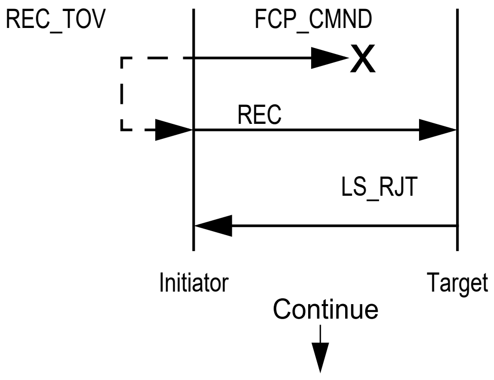

Error recovery

The LS_RJT (Logical error, Invalid OX_ID - RX_ID combination) for the REC ELS indicates the Exchange is unknown. The FCP_CMND is retransmitted in a new Exchange using the same CRN and in the same class of service.

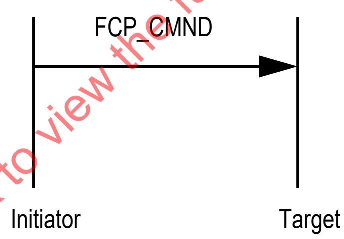

Figure C.2 - FCP_CMND lost, unacknowledged classes

An example of a lost FCP_CMND using acknowledged classes is shown in figure C.3.

Error detection

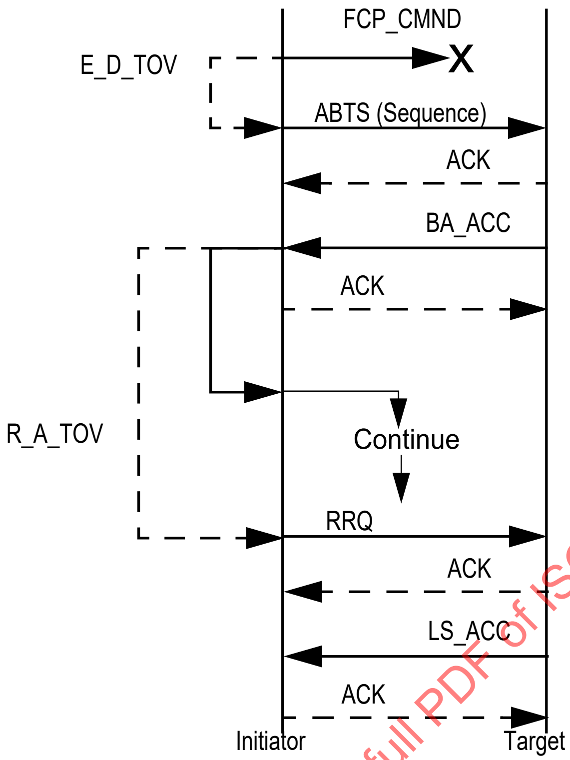

The BA_ACC indicates that the FCP_CMND was never received. The BA_ACC payload is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = SEQ_CNT field value of the ABTS (Sequence) frame. Both the initiator FCP_Port and target FCP_Port establish Recovery Qualifiers. The value of R_A_TOV for in-order delivery is zero. The use of the REC ELS to determine status for error recovery is optional. The FCP_CMND is retransmitted in a new Exchange using the same CRN and in the same class of service.

Error recovery

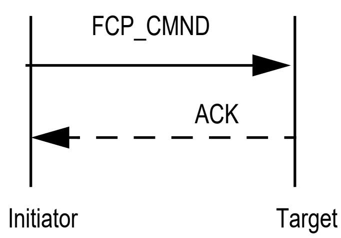

Figure C.3 - FCP_CMND lost, acknowledged classes

An example of a lost FCP_CMND ACK using acknowledged classes is shown in figure C.4.

Error detection

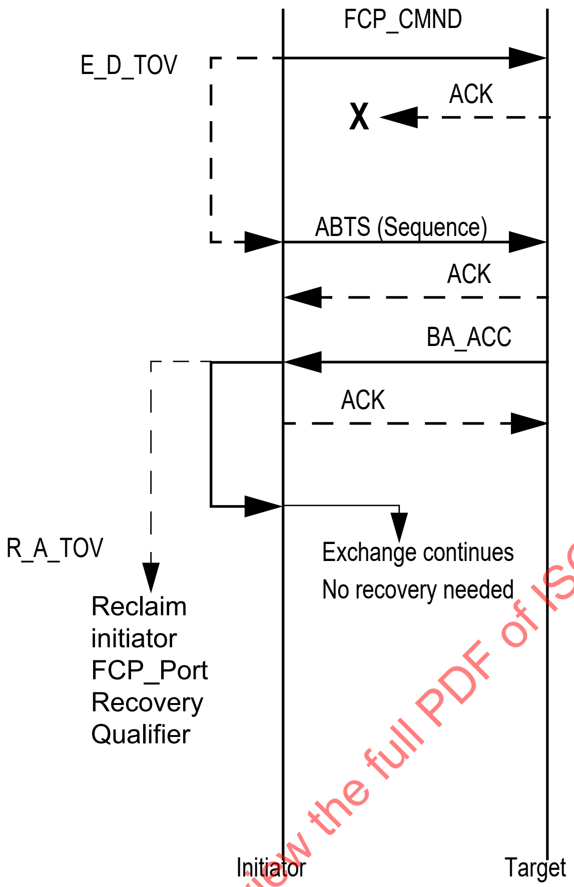

No error recovery

The BA_ACC payload indicates that FCP_CMND was received. The BA_ACC payload is SEQ_ID VALIDITY, SEQ_ID field value of the FCP_CMND, LOWEST SEQ_CNT = HIGHEST SEQ_CNT = SEQ_CNT field value of the FCP_CMND. The issuance of the RRQ ELS is not necessary in this case, since the target FCP_Port has not established a Recovery Qualifier. However, the initiator FCP_Port is not able to reclaim the resources associated with its Recovery Qualifier until R_A_TOV timeout expires. The value for R_A_TOV is zero for in-order delivery. The use of the REC ELS to determine status for error recovery is optional.

Figure C.4 - FCP_CMND ACK lost, acknowledged classes

An example of a lost FCP_XFER_RDY using unacknowledged classes is shown in figure C.5.

Error detection

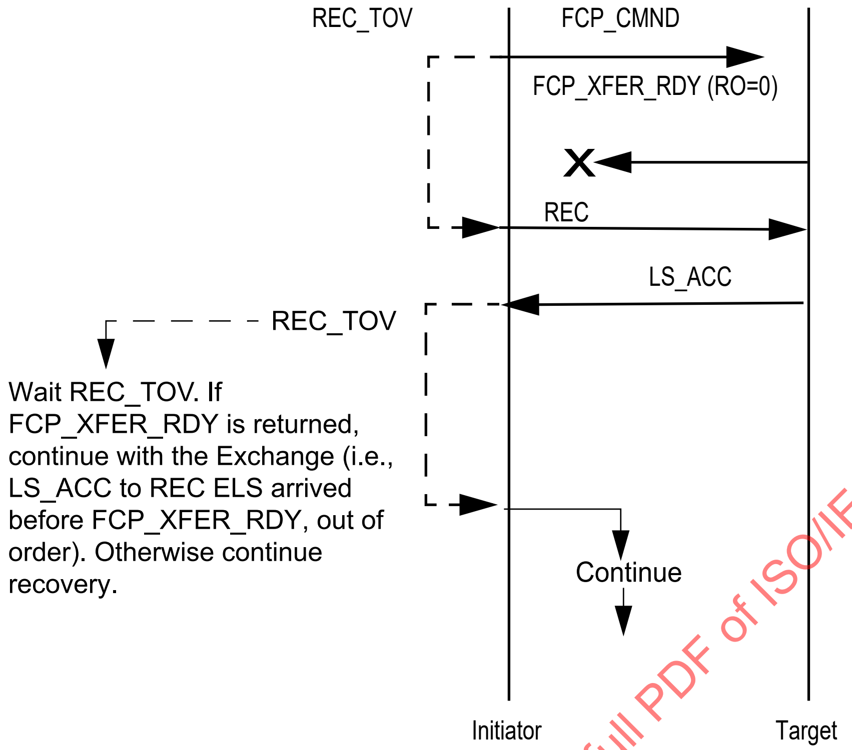

Error recovery

The LS_ACC for the REC ELS indicates the initiator FCP_Port holds Sequence Initiative and the Exchange is open. The initiator FCP_Port transmits an SRR FCP_LS request specifying the FCP_XFER_RDY be resent. The target FCP_Port retransmits the FCP_XFER_RDY using the same Relative Offset. A new SEQ_ID is used for retransmission of the FCP_XFER_RDY. For in-order delivery, the value of REC_TOV is zero.

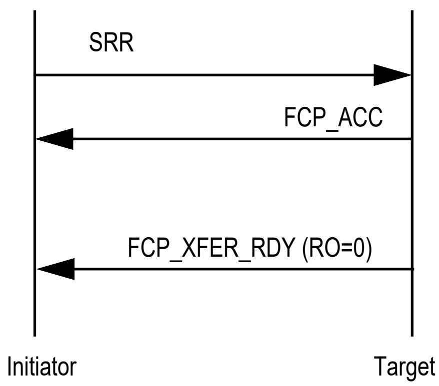

Figure C.5 - FCP_XFER_RDY lost, unacknowledged classes

An example of a lost FCP_XFER_RDY using acknowledged classes is shown in figure C.6.

Error detection

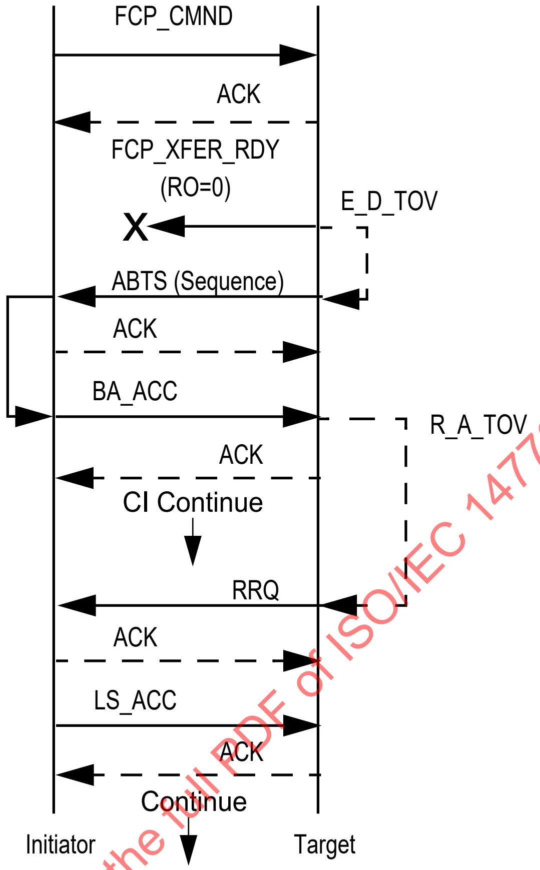

Error recovery

For acknowledged classes, the BA_ACC indicates that the FCP_XFER_RDY was never received by the initiator FCP_Port. The BA_ACC payload is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = SEQ_CNT field value in ABTS frame = 1. Both target and initiator FCP_Ports establish Recovery Qualifiers. The value for R_A_TOV for in-order delivery is zero. A new SEQ_ID is used in the retransmission of FCP_XFER_RDY. The use of an REC ELS to determine status for error recovery is optional.

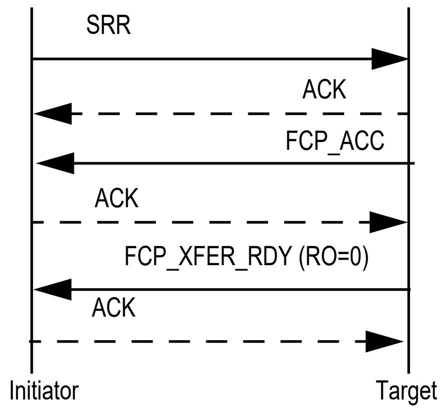

Figure C.6 - FCP_XFER_RDY lost, acknowledged classes

An example of an FCP_XFER_RDY received and subsequent ACK lost using acknowledged classes is shown in figure C.7.

Error detection

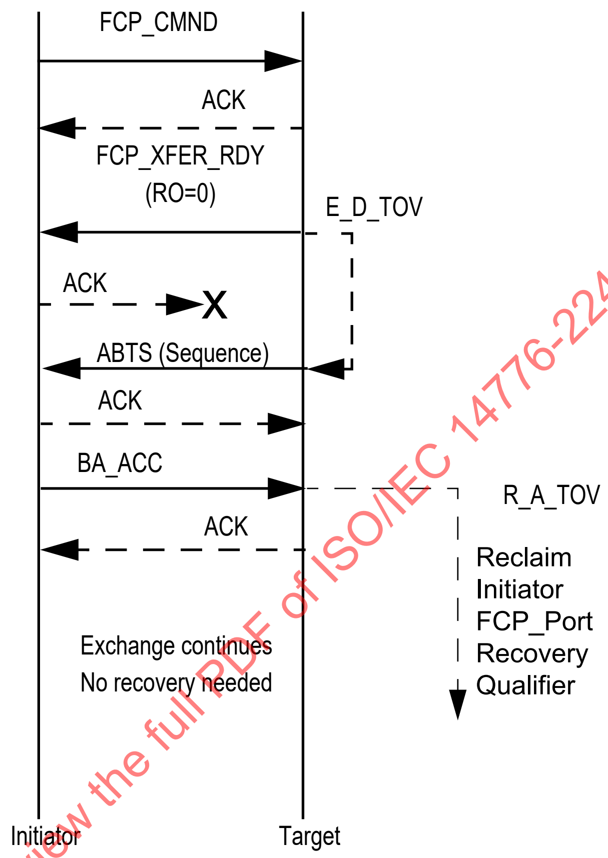

No error recovery

The BA_ACC indicates that the FCP_XFER_RDY was received by the initiator FCP_Port. The BA_ACC payload is SEQ_ID VALIDITY, SEQ_ID = SEQ_ID field of the FCP_XFER_RDY frame, LOWEST SEQ_CNT = HIGHEST SEQ_CNT = SEQ_CNT field value of the ABTS frame. No error recovery is required. There is no need for the target FCP_Port to issue the RRQ ELS since no Recovery Qualifier was established by the initiator FCP_Port in this case. It still lets R_A_TOV* expire before reclaiming the resources associated with its Recovery Qualifier. The value of R_A_TOV for in-order delivery is zero. The use of an REC ELS to determine status for error recovery is optional.

Figure C.7 - FCP_XFER_RDY received, ACK lost, acknowledged classes

An example of a lost FCP_RSP with no FCP_CONF requested using unacknowledged classes is shown in figure C.8.

Error detection

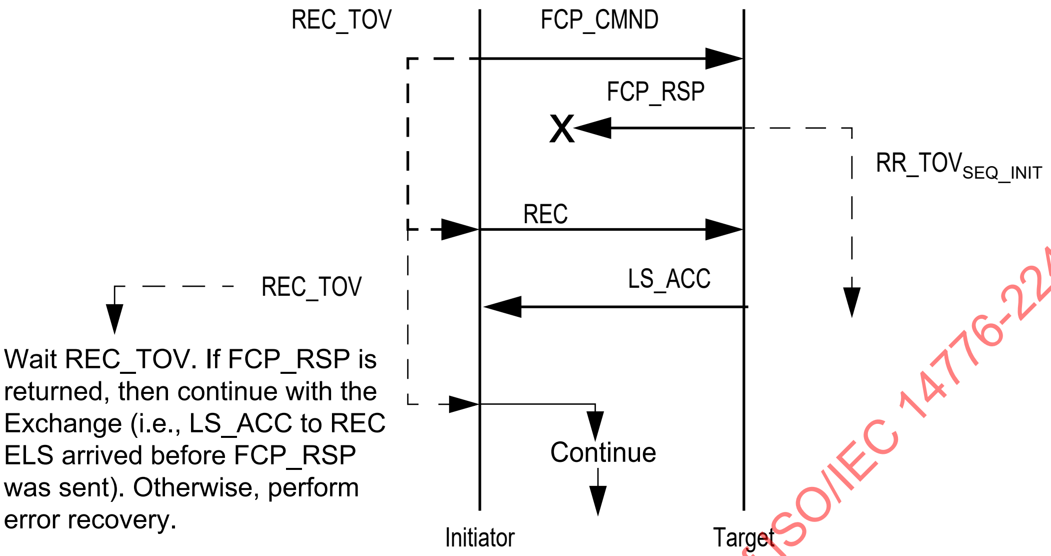

Error recovery

The LS_ACC for the REC ELS indicates the initiator FCP_Port holds Sequence Initiative and the Exchange is complete. The initiator FCP_Port transmits an SRR FCP_LS request specifying the FCP_RSP be retransmitted. The target FCP_Port retransmits the FCP_RSP. For in-order delivery, the value of REC_TOV is zero. For Sequence retransmission, the target FCP_Port keeps the context of this Exchange as defined in 12.4.1.5. For Exchange level recovery, the context of the Exchange may be discarded by the target FCP_Port after FCP_RSP is transmitted.

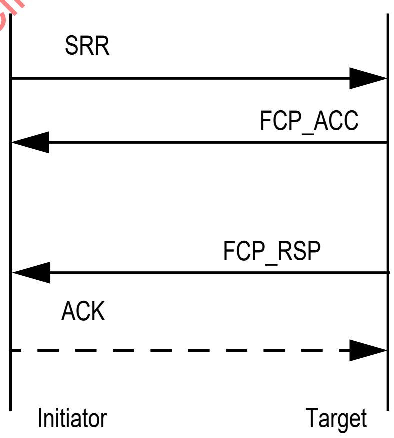

Figure C.8 - FCP_RSP lost, FCP_CONF not requested, unacknowledged classes

An example of a lost FCP_RSP with no FCP_CONF requested using acknowledged classes is shown in figure C.9.

Error detection

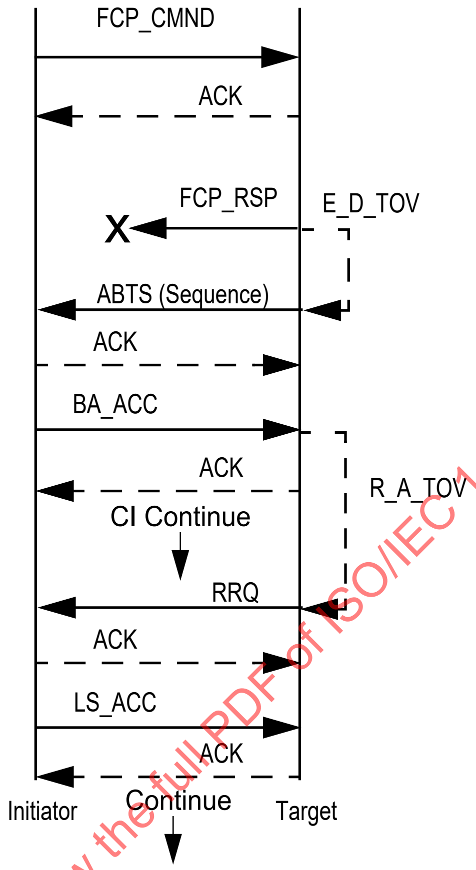

Error recovery

BA_ACC to ABTS indicates that the FCP_RSP was never received by the initiator FCP_Port . The payload is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = SEQ_CNT field value in ABTS frame = 1. The initiator FCP_Port and target FCP_Port establish Recovery Qualifiers. The value of R_A_TOV for in-order delivery is zero. A new SEQ_ID is used in the retransmission of FCP_RSP. The use of an REC ELS to determine status for error recovery is optional.

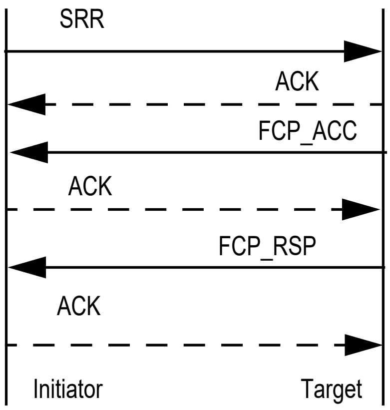

Figure C.9 - FCP_RSP lost, FCP_CONF not requested, acknowledged classes

An example of a lost FCP_RSP during a read command with no FCP_CONF requested using acknowledged classes is shown in figure C.10.

Error detection

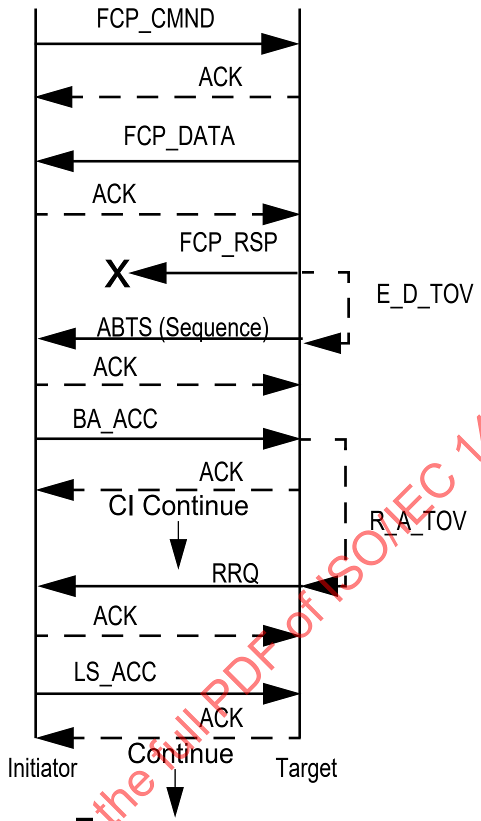

Error recovery

BA_ACC to ABTS indicates that the FCP_RSP was never received by the initiator FCP_Port. The payload is SEQ_ID VALIDITY = valid, SEQ_ID = SEQ_ID field of FCP_DATA Sequence, LOWEST SEQ_CNT = x, HIGHEST SEQ_CNT = SEQ_CNT field value in ABTS frame. The initiator FCP_Port and target FCP_Port establish Recovery Qualifiers. The value of R_A_TOV for in-order delivery is zero. A new SEQ_ID is used in the retransmission of FCP_RSP. The use of an REC ELS to determine status for error recovery is optional.

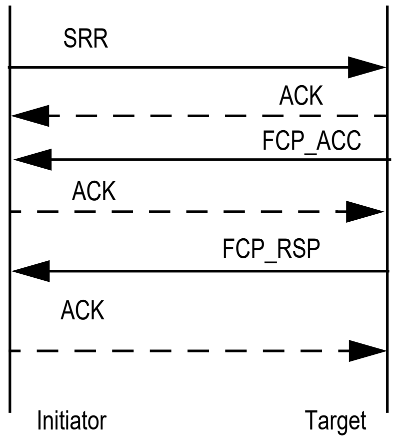

Figure C.10 - FCP_RSP lost read command, no FCP_CONF, acknowledged classes

An example of a received FCP_RSP and lost ACK using acknowledged classes is shown in figure C.11.

Error detection

Acknowledged classes

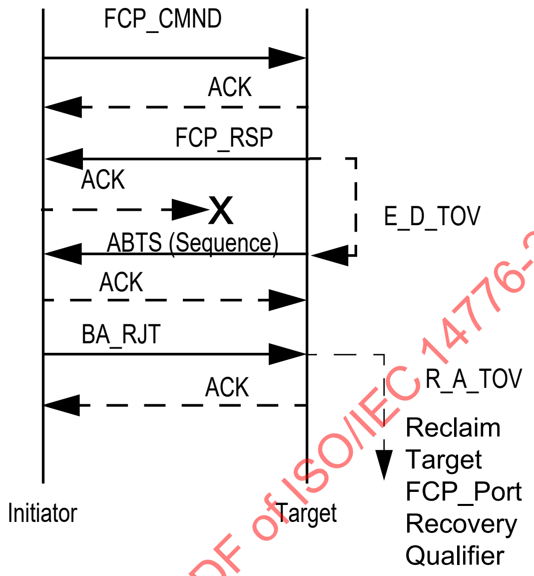

No error recovery

The BA_RJT for the ABTS with Reason Code of Logical error and the Reason Code Explanation of Invalid OX_ID - RX_ID combination indicates the Exchange is unknown and therefore complete. No error recovery is required. The target FCP_Port establishes a Recovery Qualifier. The associated resources are not reused for a period of R_A_TOV. For in-order delivery, the value of R_A_TOV is zero. The target FCP_Port does not need to issue RRQ ELS as no Recovery Qualifier was established by the initiator FCP_Port.

Figure C.11 - FCP_RSP received, ACK lost, acknowledged classes, example 1