

**draft proposed
American National Standard**

**T10
Project 1560-D
Revision 4
September 13, 2005**

Information Technology - Fibre Channel Protocol for SCSI, Third Version (FCP-3)

This is an internal working document of T10, a Technical Committee of Accredited Standards Committee INCITS (International Committee for Information Technology Standards). As such this is not a completed standard and has not been approved. The INCITS T10 Technical Committee may modify this document as a result of comments received during public review and its approval as a standard. Use of the information contained herein is at your own risk.

Permission is granted to members of INCITS, its technical committees, and their associated task groups to reproduce this document for the purposes of INCITS standardization activities without further permission, provided this notice is included. All other rights reserved. Any commercial or for-profit duplication is strictly prohibited.

Technical editor:

David Peterson
McDATA Corporation.
6000 Nathan Lane North
Plymouth, MN 55442
(612) 802-3299
E-mail: david.peterson@mcddata.com

**ISO
LOGO**

**IEC
LOGO**

**Reference number
ISO/IEC xxxxx: 2003
ANSI INCITS.xxx: 2003
Printed September 13, 2005**

POINTS OF CONTACT:

INCITS T10 Chair

John Lohmeyer
LSI Logic
4420 Arrows West Drive
Colorado Springs, CO 80907
Tel: (719) 533-7560
Fax: (719) 533-7183
E-mail: lohmeyer@t10.org

INCITS T10 Vice Chair

George O. Penokie
IBM
3605 Highway 52 N. MS 2C6
Rochester, MN 55901
Tel: (507) 253-5208
Fax: (507) 253-2880
E-mail: gpenokie@us.ibm.com

INCITS Secretariat

INCITS Secretariat
1250 Eye Street, NW Suite 200
Washington, DC 20005

Tel: 202-737-8888
Fax: 202-638-4922
E-mail: INCITS@itic.org

T10 WEb Site <http://www.t10.org>

T10 Reflector To subscribe send email to majordomo@t10.org with 'subscribe' in message body
To unsubscribe send email to majordomo@t10.org with 'unsubscribe' in message body
Internet address for distribution via SCSI reflector: t10@t10.org

Document Distribution

INCITS Online Store
managed by Techstreet
1327 Jones Drive
Ann Arbor, MI 48105

<http://www.techstreet.com/INCITS.html>
Telephone: 1-734-302-7801 or
1-800-699-9277
Facsimile: 1-734-302-7811

or

Global Engineering
15 Inverness Way East
Englewood, CO 80112-570415

<http://global.ihs.com>
Telephone: 1-303-792-2181 or
1-800-854-7179
Facsimile: 1-303-792-2192

ABSTRACT

This standard describes the frame format and protocol definitions required to transfer commands and data between a SCSI (Small Computer System Interface) initiator and target using the Fibre Channel family of standards. The second version added optional retransmission, task ordering, and confirmation capabilities. This third version incorporates bi-directional commands, removes information that is now contained in other standards, and describes additional error recovery capabilities for the Fibre Channel Protocol.

PATENT STATEMENT

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard. As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, notice of one or more claims has been received.

By publication of this standard, no position is taken with respect to the validity of this claim or of any rights in connection therewith. The known patent holder has, however, filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. Details may be obtained from the publisher.

No further patent search is conducted by the developer or the publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Contents

Foreword	xii
Introduction	xiii
1 Scope	1
2 Normative references	1
2.1 Qualification and availability of references	1
2.2 Published standard and technical report references	1
2.3 References under development	1
2.4 Other references	2
3 Definitions, abbreviations and conventions	2
3.1 Definitions	2
3.2 Abbreviations	5
3.3 Keywords	6
3.4 Editorial conventions	7
4 General	9
4.1 Structure and concepts	9
4.2 FCP I/O operations	10
4.3 Bidirectional and unidirectional commands and FCP_RSP IU format	12
4.4 Precise delivery of SCSI commands	12
4.5 Confirmed completion of FCP I/O operations	13
4.6 Retransmission of unsuccessfully transmitted IUs	14
4.7 Task retry identification	14
4.8 Discovery of FCP capabilities	15
4.9 Task management	15
4.10 Clearing effects of task management, FCP, FC-FS-2, FC-LS, and FC-AL-2 actions ..	16
4.11 I_T nexus loss notification events	19
4.12 Transport Reset notification events	19
4.13 Port Login/Logout	19
4.14 Process Login and Process Logout	20
4.15 Link management	20
4.16 FCP addressing and Exchange identification	20
4.17 Use of World Wide Names	20
5 FC-FS-2 frame header	22
5.1 FC-FS-2 frame header overview	22
5.2 FC-FS-2 frame header fields	22
5.2.1 r_ctl field	22
5.2.2 d_id field	22
5.2.3 cs_ctl field	22
5.2.4 s_id field	22
5.2.5 type field	22
5.2.6 f_ctl field	23
5.2.7 seq_id field	23
5.2.8 df_ctl field	23
5.2.9 seq_cnt field	23
5.2.10 ox_id field	23
5.2.11 rx_id field	23
5.2.12 parameter field	23
6 FCP link service definitions	24
6.1 Overview of link service requirements	24

6.2 Overview of Process Login and Process Logout	24
6.3 PRLI	25
6.3.1 Use of PRLI by the Fibre Channel Protocol	25
6.3.2 Process_Associator requirements	25
6.3.3 New or repeated Process Login	25
6.3.4 PRLI request FCP Service Parameter page format	26
6.3.5 PRLI accept FCP Service Parameter page format	29
6.4 PRLO	30
6.5 Read Exchange Concise (REC)	30
7 FC-4 specific Name Server registration and objects	31
7.1 Overview of FC-4 specific objects for the Fibre Channel Protocol	31
7.2 FC-4 TYPEs object	31
7.3 FC-4 Features object	31
8 FC-4 Link Service definitions	32
8.1 FC-4 Link Services for the Fibre Channel Protocol	32
8.2 Sequence Retransmission Request (SRR)	32
8.3 FCP FC-4 Link Service Reject (FCP_RJT)	34
9 FCP Information Unit (IU) usage and formats	36
9.1 FCP Information Unit (IU) usage	36
9.2 FCP_CMND IU	37
9.2.1 Overview and format of FCP_CMND IU	37
9.2.2 FCP_CMND IU field descriptions	38
9.2.2.1 fcp_lun field	38
9.2.2.2 COMMAND REFERENCE NUMBER field	39
9.2.2.3 priority field	39
9.2.2.4 task attribute field	39
9.2.2.5 TASK MANAGEMENT FLAGS FIELD	39
9.2.2.6 additional fcp_cdb length field	41
9.2.2.7 rddata and wrdata bits	42
9.2.2.8 fcp_cdb field	42
9.2.2.9 additional_fcp_cdb field	42
9.2.2.10 fcp_dl field	42
9.2.2.11 fcp_bidirectional_read_dl field	43
9.3 FCP_XFER_RDY IU	43
9.3.1 Overview and format of FCP_XFER_RDY IU	43
9.3.2 fcp_data_ro field	44
9.3.3 fcp_burst_len field	44
9.4 FCP_DATA IU	44
9.4.1 FCP_DATA IU overview	44
9.4.2 FCP_DATA IUs for read and write operations	45
9.4.3 FCP_DATA IUs for bidirectional commands	45
9.4.4 FCP_DATA IU use of fill bytes	46
9.5 FCP_RSP IU	46
9.5.1 Overview and format of FCP_RSP IU	46
9.5.2 retry delay timer field	48
9.5.3 fcp_bidi_rsp bit	48
9.5.4 fcp_bidi_read_resid_under bit	49
9.5.5 fcp_bidi_read_resid_over bit	49
9.5.6 FCP_CONF_REQ BIT	49
9.5.7 fcp_resid_under bit	49

9.5.8 fcp_resid_over bit	49
9.5.9 fcp_sns_len_valid bit	49
9.5.10 fcp_rsp_len_valid bit	49
9.5.11 SCSI status code field	50
9.5.12 fcp_resid field	50
9.5.13 fcp_bidirectional_read_resid field	50
9.5.14 fcp_sns_len field	51
9.5.15 fcp_rsp_len field	51
9.5.16 fcp_rsp_info field	51
9.5.17 fcp_sns_info field	52
9.6 FCP_CONF IU	53
10 SCSI mode parameters for the Fibre Channel Protocol	54
10.1 Overview of mode pages for the Fibre Channel Protocol	54
10.2 Disconnect-Reconnect mode page	54
10.2.1 Overview and format of Disconnect-Reconnect mode page for FCP	54
10.2.2 buffer full ratio field	55
10.2.3 buffer empty ratio field	56
10.2.4 bus inactivity limit field	56
10.2.5 disconnect time limit field	56
10.2.6 connect time limit field	56
10.2.7 maximum burst size field	56
10.2.8 emdp bit	57
10.2.9 FAA, FAB, FAC BITS	57
10.2.10 first burst size field	57
10.3 Fibre Channel Logical Unit Control mode page	58
10.4 Fibre Channel Port Control mode page	58
10.4.1 Overview and format of Fibre Channel Port Control mode page	58
10.4.2 DISABLE TARGET ORIGINATED LOOP INITIALIZATION (DTOLI) bit	59
10.4.3 DISABLE TARGET INITIATED PORT ENABLE (DTIPE) bit	59
10.4.4 ALLOW LOGIN WITHOUT LOOP INITIALIZATION (ALWLI) bit	59
10.4.5 REQUIRE HARD ADDRESS (RHA) bit	60
10.4.6 disable loop master (dlm) bit	60
10.4.7 DISABLE DISCOVERY (DDIS) bit	60
10.4.8 PREVENT LOOP PORT BYPASS (PLPB) bit	60
10.4.9 DISABLE TARGET FABRIC DISCOVERY (DTFD) bit	60
10.4.10 sequence initiative resource recovery timeout value (rr_tovseq_init) field	61
11 Timers for FCP operation and recovery	62
11.1 Summary of timers for the Fibre Channel Protocol	62
11.2 Error_Detect Timeout (E_D_TOV)	63
11.3 Resource Allocation Timeout (R_A_TOV)	63
11.4 Resource Recovery Timeout (RR_TOV)	63
11.5 Read Exchange Concise Timeout Value (REC_TOV)	63
11.6 Upper Level Protocol Timeout (ULP_TOV)	64
12 Link error detection and error recovery procedures	65
12.1 Error detection and error recovery overview	65
12.1.1 Exchange level	65
12.1.2 Sequence level	65
12.2 FCP error detection	65
12.2.1 Overview of FCP-3 error detection	65
12.2.2 FCP-3 error detection using protocol errors for all classes of service	65

12.2.3 Error detection mechanisms for acknowledged classes of service	66
12.3 Exchange level recovery using recovery abort	67
12.3.1 Recovery abort overview	67
12.3.2 Initiator FCP_Port invocation of recovery abort	67
12.3.3 Target FCP_Port response to recovery abort	67
12.3.4 Additional error recovery by initiator FCP_Port	68
12.3.5 Additional error recovery by target FCP_Port	68
12.4 Sequence level error detection and recovery	68
12.4.1 Using information from REC to perform Sequence level recovery	68
12.4.1.1 Polling Exchange state with REC	68
12.4.1.2 Detection of errors while polling with REC	69
12.4.1.3 FCP_CMND IU recovery	69
12.4.1.4 FCP_XFER_RDY IU recovery	69
12.4.1.5 FCP_RSP IU recovery	69
12.4.1.6 FCP_DATA IU recovery - write operations	70
12.4.1.7 FCP_DATA IU recovery - read operations	71
12.4.1.8 FCP_CONF IU recovery	71
12.4.2 Additional error recovery requirements	72
12.4.2.1 Error indicated in ACK	72
12.4.2.2 Missing ACK	72
12.4.2.3 Distinguishing Exchange to be aborted	72
12.5 Second-level error recovery	72
12.5.1 ABTS error recovery	72
12.5.2 REC error recovery	73
12.5.3 SRR error recovery	73
12.6 Responses to FCP type frames before PLOGI or PRLI	73
Annex A	74
A.1 Definition of procedure terms	74
Annex B	76
B.1 Examples of the use of FCP Information Units (IUs)	76
B.1.1 Overview of examples	76
B.1.2 SCSI FCP read operation	76
B.1.3 SCSI FCP write operation	77
B.1.4 SCSI FCP operation with no data transfer or with check condition	77
B.1.5 SCSI FCP read operation with multiple FCP_DATA IUs	78
B.1.6 SCSI FCP write operation with FCP_XFER_RDY disabled	78
B.1.7 SCSI FCP bidirectional command with write before read	79
B.1.8 SCSI FCP bidirectional command with read before write	79
B.1.9 SCSI FCP bidirectional command, write first, write FCP_XFER_RDY disabled ..	80
B.1.10 SCSI FCP bidirectional command with intermixed writes and reads	81
B.1.11 SCSI linked commands	82
B.1.12 SCSI WRITE command with confirmed completion	83
B.1.13 SCSI FCP task management function	83
B.2 FCP write example, frame level	84
B.3 FCP read example, frame level	86
Annex C	88
C.1 Introduction	88
Annex D	122
D.1 FCP Device Discovery Procedure	122
D.1.1 Initiator discovery of Fabric-attached target FCP_Ports	122

D.1.2 Initiator discovery of loop-attached target FCP_Ports	122
D.2 Fabric and Device Authentication	123
D.3 Logical unit authentication	123
Annex E	124
E.1 Formats for recovery link services	124
E.2 Abort Sequence (ABTS) Request	124
E.2.1 Abort Sequence (ABTS) Request fields	124
E.2.2 Basic Accept (BA_ACC) Frame to ABTS	125
E.2.3 Basic Reject (BA_RJT) Frame to ABTS	126
E.3 Reinstate Recovery Qualifier (RRQ)	126
E.3.1 RRQ request format	126

Tables

Table 1 - SCSI and Fibre Channel Protocol functions	10
Table 2 - Discovery of FCP-3 capabilities	15
Table 3 - Task management functions, SAM-3 to FCP-3	16
Table 4 - Clearing effects of link related functions	18
Table 5 - Clearing effects of initiator FCP_Port actions	19
Table 6 - FCP frame header	22
Table 7 - FCP Service Parameter page, PRLI request	26
Table 8 - FCP Service Parameter page, PRLI accept	29
Table 9 - FCP definition of FC-4 Feature bits	31
Table 10 - FCP FC-4 Link Service Requests and Responses	32
Table 11 - SRR Payload	33
Table 12 - SRR Accept Payload	33
Table 13 - FCP_RJT Payload	34
Table 14 - FCP_RJT reason codes	35
Table 15 - FCP_RJT reason code explanations	35
Table 16 - FCP Information Units (IUs) sent to target FCP_Ports	36
Table 17 - FCP Information Units (IUs) sent to initiator FCP_Ports	37
Table 18 - FCP_CMND IU Payload	38
Table 19 - TASK ATTRIBUTE field values	39
Table 20 - task management flags field	40
Table 21 - FCP_XFER_RDY IU payload	43
Table 22 - FCP_RSP IU Payload	48
Table 23 - FCP_RSP_INFO field format	52
Table 24 - RSP_CODE definitions	52
Table 25 - Mode pages for FCP	54
Table 26 - Disconnect-Reconnect mode page (02h)	55
Table 27 - Fibre Channel Logical Unit Control mode page (18h)	58
Table 28 - Fibre Channel Port Control mode page (19h)	59
Table 29 - Values for RR_TOV UNITS	61
Table 30 - Timer summary	62
Table 31 - Initiator FCP_Port REC_TOV Usage	64
Table 32 - Target FCP_Port REC_TOV usage	64
Table A.2 - Procedure terms	75
Table B.1 - FCP read operation, example	76
Table B.2 - FCP write operation, example	77
Table B.3 - FCP operation without data transfer, example	77
Table B.4 - FCP read operation, example	78
Table B.5 - FCP write operation with FCP_XFER_RDY disabled, example	78
Table B.6 - FCP bidirectional command with write before read, example	79
Table B.7 - FCP bidirectional command with read before write, example	79
Table B.8 - FCP bidirectional command, write FCP_XFER_RDY disabled, example	80
Table B.9 - FCP bidirectional command with intermixed writes and reads, example	81
Table B.10 - FCP linked commands, example	82
Table B.11 - FCP write command with confirmed completion	83
Table B.12 - FCP task management function, example	83
Table E.2 - BA_ACC Frame to ABTS	125
Table E.3 - BA_RJT Frame to ABTS	126

Table E.4 - Reinstate Recovery Qualifier	126
--	-----

Figures

Figure B.1 - Example of class 2 FCP write operation	84
Figure B.2 - Example of class 2 FCP_DATA write	85
Figure B.3 - Example of class 2 FCP read operation	86
Figure B.4 - Example of class 2 FCP_DATA read	87
Figure C.1 - Lengthy FCP_CMND or Lost ACK	89
Figure C.2 - FCP_CMND Lost, Unacknowledged Classes	90
Figure C.3 - FCP_CMND Lost, Acknowledged Classes	91
Figure C.4 - FCP_CMND Acknowledgement Lost, Acknowledged Classes	92
Figure C.5 - FCP_XFER_RDY Lost, Unacknowledged Classes	93
Figure C.6 - FCP_XFER_RDY Lost, Acknowledged Classes	94
Figure C.7 - FCP_XFER_RDY Received, ACK Lost, Acknowledged Classes	95
Figure C.8 - FCP_RSP Lost, FCP_CONF not requested, Unacknowledged Classes	96
Figure C.9 - FCP_RSP Lost, FCP_CONF not requested, Acknowledged Classes	97
Figure C.10 - FCP_RSP Lost Read Command, no FCP_CONF, Acknowledged Classes	98
Figure C.11 - FCP_RSP Received, ACK Lost, Acknowledged Classes, Example 1	99
Figure C.12 - FCP_RSP Received, ACK Lost, Acknowledged Classes, Example 2	100
Figure C.13 - Lost Write Data, Last Frame of Sequence, Unacknowledged Classes	101
Figure C.14 - Lost Write Data, Last Frame of Sequence, Acknowledged Classes	102
Figure C.15 - Lost Write Data, Not Last Frame of Sequence, Unacknowledged Classes ...	103
Figure C.16 - Lost Write Data, Not Last Frame of Sequence, Acknowledged Classes	104
Figure C.17 - Lost Read Data, Last Frame of Sequence, Unacknowledged Classes	105
Figure C.18 - Lost Read Data, Last Frame of Sequence, Acknowledged Classes	106
Figure C.19 - Lost Read Data, Not Last Frame of Sequence, Unacknowledged Classes ...	107
Figure C.20 - Lost Read Data, Not Last Frame of Sequence, Acknowledged Classes	108
Figure C.21 - ACK Lost on Read (Acknowledged Classes)	109
Figure C.22 - ACK Lost on Write (Acknowledged Classes)	110
Figure C.23 - FCP_CONF Lost, Unacknowledged Classes	111
Figure C.24 - FCP_CONF Lost, Acknowledged Classes	112
Figure C.25 - ACK lost on FCP_CONF, Acknowledged Classes	113
Figure C.26 - REC or REC Response Lost, Unacknowledged Classes	114
Figure C.27 - REC Lost, Acknowledged Classes	115
Figure C.28 - REC Response Lost, Acknowledged Classes	116
Figure C.29 - Two RECs Lost, Unacknowledged Classes, Abort the original Exchange ...	117
Figure C.30 - SRR Lost, Unacknowledged Classes, Abort original Exchange	118
Figure C.31 - SRR Response Lost, Unacknowledged Classes	119
Figure C.32 - SRR Lost, Acknowledged Classes	120
Figure C.33 - SRR Response Lost, Acknowledged Classes	121

Foreword

(This foreword is not part of BSR INCITS [insert number here].)

The Fibre Channel Protocol for SCSI, Third version (FCP-3) defines a Fibre Channel mapping layer (FC-4) that uses the services defined by INCITS Project 1619-D, *Fibre Channel Framing and Signaling Interface - 2 (FC-FS-2)* to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target. The use of the standard enables the transmission of standard SCSI command formats, the transmission of standard SCSI data and parameter strings, and the receipt of SCSI status and sense information across the Fibre Channel using only the standard Fibre Channel frame and sequence formats. The Fibre Channel Protocol operates with Fibre Channel Classes of Service 1, 2, and 3 and operates across Fibre Channel fabrics and arbitrated loops.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, Information Technology Industry Council (ITI), 1250 Eye Street, NW, Suite 200, Washington, DC 20005.

This standard was processed and approved for submittal to ANSI by International Committee for Information Technology Standards (INCITS). Committee approval of the standard does not necessarily imply that all committee members voted for its approval. At the time it approved this standard, INCITS had the following members:

The INCITS Technical Committee T10 on Lower-Level Interfaces, that developed and reviewed this standard, had the following members:

John B. Lohmeyer, Chair

George O. Penokie, Vice-Chair

Ralph O. Weber, Secretary

(insert T10 member list here)

The INCITS Technical Committee T11 on Device Level Interfaces contributed a great deal of analysis, review, and content to this standard. It had the following members:

Robert Snively, Chair

Claudio DeSanti, Vice-Chair

Neil Wanamaker, Secretary

(insert T11 member list here)

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-2 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 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 Information Units used to transfer SCSI commands, data, and status across a Fibre Channel connection.

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 FC-GS-4 Name Server objects defined for FCP-3.

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

Clause 9 describes the details of the Information Unit formats.

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

Clause 11 defines the timers used for FCP-3 error recovery algorithms.

Clause 12 defines the error recovery algorithms for FCP-3.

The Fibre Channel Protocol for SCSI, Third revision (FCP-3) standard has the following annexes:

Annex A is a normative description of the relationship between the services defined by SAM-3 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 Fibre Channel.

Annex C is an informative annex providing examples of the FCP-3 error recovery mechanisms.

Annex D is an informative annex describing techniques for discovering SCSI device capabilities over Fibre Channel.

Annex E is an informative annex providing examples of the content of ELSs used during FCP-3 recovery operations.

This standard is part of the SCSI family of standards developed by T10 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 INCITS Project 1561D, *Information Technology - SCSI Architecture Model - 3 (SAM-3)*.

American National Standard for Information Technology – **Fibre Channel Protocol for SCSI, Third version (FCP-3)**

1 Scope

This standard defines a third 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 - 3 (SAM-3). This third version includes additions and clarifications to the second version, removes information that is now contained in other standards, and describes additional error recovery capabilities for the Fibre Channel Protocol.

2 Normative references

2.1 Qualification and availability of references

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

Copies of the following standards and technical reports may be obtained from ANSI: Approved ANSI standards, approved and draft international and regional standards (ISO, IEC, CEN/CENELEC, ITUT), and approved and draft foreign standards (including BSI, JIS, and DIN). For further information, contact ANSI Customer Service Department at 212-642-4900 (phone), 212-302-1286 (fax) or via the World Wide Web at <http://www.ansi.org>.

Reference is made to these standards and technical reports by their standard designations.

Additional availability contact information is provided below as needed.

2.2 Published standard and technical report references

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

ANSI/INCITS 402-2005, *SCSI Architecture Model - 3 (SAM-3)*

INCITS TR-36-2004, *Fibre Channel - Device Attach (FC-DA)*

2.3 References under development

At the time of publication, the following referenced standards and technical reports were still under development by INCITS T10 and INCITS T11. For information on the current status of the standard or technical report, or regarding availability, contact the INCITS Secretariat, Information Technology Industry Council (ITI) at 202-737-8888 (phone) or by mail at 1250 Eye Street NW, Suite 200, Washington, DC 20005-3922.

INCITS Project 1619-D, *Fibre Channel Framing and Signaling Interface (FC-FS-2)*

INCITS Project 1620-D, *Fibre Channel - Link Services (FC-LS)*

INCITS Project 1505-D, *Fibre Channel Generic Services - 4 (FC-GS-4)*

INCITS Project 1739-DT, *Fibre Channel - Device Attach - 2 (FC-DA-2)*

INCITS Project 1683-D, *SCSI Architecture Model - 4 (SAM-4)*

INCITS Project 1416-D, *SCSI Primary Commands - 3 (SPC-3)*

Copies of these INCITS T10 and T11 draft standards and technical reports are available for purchase from Global Engineering Documents. For further information, contact Global Engineering Documents at 800-854-7179 (phone) or 303-792-2181 (phone) or by mail at 15 Inverness Way East, Englewood, CO 80122-5704. The INCITS T10 draft standards are also available on the web site www.t10.org. The INCITS T11 draft standards and technical reports are also available on the web site www.t11.org.

2.4 Other references

The following references are the product of the SFF committee. For information on the current status and availability of the documents, contact the SFF committee at 408-867-6630 (phone) or by mail at 14426 Black Walnut Court, Saratoga, CA 95070.

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

3 Definitions, abbreviations and conventions

3.1 Definitions

3.1.1 acknowledged class: Any class of service that acknowledges transfers (e.g., Class 1, Class 2, and Class 4) See FC-FS-2.

3.1.2 address identifier: An address value used to identify the source (S_ID) or destination (D_ID) of a frame. See FC-FS-2.

3.1.3 application client: An object that is the source of SCSI commands. See SAM-3.

3.1.4 application client buffer offset: Offset in bytes from the beginning of the application client's buffer (data-in or data-out) to the location for the transfer of the first byte of a data delivery service request. See SAM-3.

3.1.5 autosense data: Sense data (see 3.1.51) that is returned in the FCP_RSP IU payload. See SAM-3.

3.1.6 command: A request describing a unit of work to be performed by a device server. See SAM-3.

3.1.7 command descriptor block: A structure used to communicate a command from an application client to a device server. See SAM-3.

3.1.8 data buffer size: Upper limit on the extent of the data (data-in or data-out) to be transferred by the SCSI command. See SAM-3.

3.1.9 Data-In delivery service: A confirmed service used by the device server to request the transfer of data to the application client. See SAM-3.

3.1.10 Data-Out delivery service: A confirmed service used by the device server to request the transfer of data from the application client. See SAM-3.

3.1.11 Data frame: An FC-4 Device_Data frame, an FC-4 Video_Data frame, or a Link_Data frame. See FC-FS-2.

3.1.12 data overlay: The use of 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. See SAM-3.

3.1.13 Destination_Identifier (D_ID): The address identifier used to indicate the destination of the transmitted frame. See FC-FS-2.

3.1.14 device server: An object within the logical unit that processes SCSI tasks and enforces the rules for

task management. See SAM-3.

3.1.15 discard: The term used in FC-FS-2 to describe removing 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. See FC-FS-2.

3.1.16 Exchange: The basic mechanism that transfers information consisting of one or more related non-concurrent Sequences that may flow in the same or opposite directions. The Exchange is identified by an Originator Exchange_ID (OX_ID) and a Responder Exchange_Identifier (RX_ID). See FC-FS-2.

3.1.17 Execute Command service: A peer-to-peer, confirmed service requested by the application client to perform a SCSI Command. See SAM-3.

3.1.18 FCP Exchange: A SCSI I/O operation for the Fibre Channel FC-2 layer. The SCSI I/O operation for Fibre Channel is contained in a Fibre Channel Exchange. See FC-FS-2 and 4.1.

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

3.1.20 FCP_Port: An N_Port or NL_Port that supports the SCSI Fibre Channel Protocol.

3.1.21 fully qualified exchange identifier: A set of addresses and values used to uniquely identify an FCP I/O operation. See 4.16.

3.1.22 image pair: The originating and responding processes related by a Process Login operation. For the Fibre Channel Protocol, the image pair is composed of one initiator FCP_Port and one target FCP_Port. See FC-LS.

3.1.23 I_T nexus loss: A condition resulting from the events defined by SAM-3 in which the SCSI device performs the operations described in SAM-3 and this standard.

3.1.24 I_T nexus loss event: A SCSI transport protocol specific event that triggers I_T nexus loss as described in SAM-3.

3.1.25 Information Unit (IU): An 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-2).

3.1.26 initiator: A SCSI device containing application clients that originate device service requests and task management functions to be processed by a target SCSI device. In this standard, the word initiator also refers to an FCP_Port using the Fibre Channel Protocol to perform the SCSI initiator functions defined by SAM-3.

3.1.27 initiator port identifier: A value by which a SCSI initiator port is referenced within a domain. See SAM-3.

3.1.28 initiator port name: A SCSI port name of a SCSI initiator port or of a SCSI target/initiator port when operating as a SCSI initiator port. See SAM-3.

3.1.29 interconnect tenancy: The period of time that an FCP device owns or may access a shared Fibre Channel link such as an FC-AL-2 loop. See 10.2.1.

3.1.30 logical unit: A SCSI target device object, containing a device server and task manager, that implements a device model and manages tasks to process commands sent by an application client. See SAM-3.

3.1.31 logical unit number: An encoded 64-bit identifier for a logical unit. See SAM-3.

3.1.32 loop initialization primitive: A primitive used in Fibre Channel arbitrated loops to start loop initialization. See FC-AL-2.

3.1.33 Name_Identifier: A 64-bit identifier, with a 60-bit value preceded with a 4 bit Network_Address_Authority Identifier, used to identify entities in Fibre Channel such as N_Port, Node, F_Port, or Fabric.

See FC-FS-2.

3.1.34 Name Server: A Fibre Channel service accessed through a well-known address identifier that uses the Common Transfer (CT) protocol as defined in FC-GS-4 to allow a client to determine the address identifier and properties of devices attached to a Fibre Channel switching fabric. See FC-GS-4.

3.1.35 Node_Name: A Name_Identifier associated with a Node. See FC-FS-2.

3.1.36 NL_Port: An N_Port that contains arbitrated loop functions associated with the Fibre Channel arbitrated loop topology. See FC-AL-2.

3.1.37 N_Port: A hardware entity that supports the FC-FS-2 FC-2 layer. It may act as an Originator, a Responder, or both. See FC-FS-2.

3.1.38 Originator: The logical function associated with an N_Port responsible for originating an Exchange. See FC-FS-2.

3.1.39 Originator Exchange Identifier: An identifier assigned by an Originator to identify an Exchange. See 3.1.50 and FC-FS-2.

3.1.40 Port Identifier: An address identifier (see 3.1.2) assigned to an N_Port or NL_Port during implicit or explicit Fabric Login (see FC-LS).

3.1.41 Port (N_Port) Login (PLOGI): The Fibre Channel Extended Link Service (ELS) that exchanges identification and operation parameters between an originating N_Port and a responding N_Port. See FC-LS.

3.1.42 Port_Name: A Name_Identifier associated with an N_Port or an NL_Port.

3.1.43 private loop: A loop operating with no attached fabric loop ports. See FC-DA.

3.1.44 public loop: A loop operating with an attached fabric loop port. See FC-DA.

3.1.45 random buffer access: The 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. See SAM-3.

3.1.46 read operation: An operation that uses the Data-In action, IU I3 (see 9.1).

3.1.47 recovery abort: An FC-LS protocol that recovers FCP_Port resources by terminating the Exchange and FCP I/O operation. See 12.3.

3.1.48 request byte count: The number of bytes to be moved by a data delivery service request. See SAM-3.

3.1.49 Responder: The logical function in an N_Port responsible for supporting the Exchange initiated by the Originator in another N_Port. See FC-FS-2.

3.1.50 Responder Exchange Identifier: An identifier assigned by a Responder to identify an Exchange and meaningful only to the Responder. See FC-FS-2.

3.1.51 sense data: Data returned to an application client as a result of an autosense operation or REQUEST SENSE command. See SPC-3.

3.1.52 Sequence: A set of one or more Data frames 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. See FC-FS-2.

3.1.53 Source Identifier (S_ID): The address identifier used to indicate the source port of the transmitted frame. See FC-FS-2.

3.1.54 SCSI device: A device that originates or services SCSI commands. See SAM-3.

3.1.55 SCSI I/O operation: An operation defined by a SCSI command, a series of linked SCSI commands, or

a task management function. See SAM-3.

3.1.56 SCSI initiator port: A SCSI initiator device object that acts as the connection between application clients and the service delivery subsystem through which requests, indications, responses, and confirmations are routed (see SAM-3). In all cases when this term is used it refers to an initiator port or a SCSI target/initiator port operating as a SCSI initiator port. In this standard, the term SCSI initiator port also refers to an FCP_Port using the Fibre Channel protocol to perform the SCSI initiator port functions defined by SAM-3.

3.1.57 SCSI target port: A SCSI target device object that contains a task router and acts as the connection between device servers and task managers and the service delivery subsystem through which indications and responses are routed (see SAM-3). In this standard, the term SCSI target port also refers to an FCP_Port using the Fibre Channel protocol to perform the SCSI target port functions defined by SAM-3.

3.1.58 status: A single byte returned by the device server to the application client to indicate the completion and completion state of a command. See SAM-3.

3.1.59 tag: The application client specified component of a task identifier that uniquely identifies one task among the several tasks coming from an application client to a logical unit. The fully qualified exchange identifier performs the function of the SCSI tag in this standard. See 4.16 and SAM-3.

3.1.60 target port identifier: An address identifier (see 3.1.2) that a SCSI initiator port uses to identify the SCSI target port. See SAM-3.

3.1.61 target port name: A SCSI port name of a SCSI target port or of a SCSI target/initiator port when operating as a SCSI target port. See SAM-3.

3.1.62 task: An object within the logical unit representing the work associated with a command or group of linked commands. See SAM-3.

3.1.63 task attribute: The queuing specification for a task (e.g., SIMPLE, ORDERED, HEAD OF QUEUE, ACA). See SAM-3.

3.1.64 task identifier: The information uniquely identifying a task. See SAM-3.

3.1.65 task management function: A 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. See SAM-3.

3.1.66 unacknowledged class: Any class of service that does not acknowledge transfers (e.g., Class 3). See FC-FS-2.

3.1.67 Worldwide_Name: A Name_Identifier that is worldwide unique, and represented by a 64-bit unsigned binary value. See FC-FS-2.

3.1.68 word: A string of four contiguous bytes occurring on boundaries that are zero modulo 4 from a specified reference. See FC-FS-2.

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

3.2 Abbreviations

ABTS	Abort Sequence (see FC-FS-2).
CRN	Command Reference Number (see 4.4 and 9.2.2.2).
D_ID	Destination_Identifier (see 3.1.13).
ELS	Extended Link Service (see FC-LS).
FC	Fibre Channel (see FC-FS-2).

FC-AL-2	INCITS 332:1999, <i>Fibre Channel Arbitrated Loop - 2</i> (see 2.2).
FC-FS-2	INCITS Project 1619-D, <i>Fibre Channel Framing and Signaling Interface - 2</i> (see 2.3).
FC-GS-4	INCITS Project 1505-D, <i>Fibre Channel Generic Services - 4</i> (see 2.3).
FC-LS	INCITS Project 1620-D, <i>Fibre Channel - Link Services</i> (see 2.3).
FCP	Refers to this standard.
FCP_RJT	FCP FC-4 Link Service Reject link service (see 8.3).
FCP-3	This standard (Fibre Channel Protocol for SCSI, third version).
FC-4	Fibre Channel Layer 4 mapping layer (see FC-FS-2).
FLOGI	Fabric Login ELS (see FC-LS).
FQXID	fully qualified exchange identifier (see 3.1.21).
ID	identifier
IU	Information Unit (see 3.1.25).
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	Log Out ELS (see FC-LS).
NA	Not Applicable
OX_ID	Originator Exchange Identifier (see FC-FS-2).
PLOGI	Port (N_Port) Login ELS (see FC-LS).
PRLI	Process Login ELS (see 6.3 and FC-LS).
PRLO	Process Logout ELS (see 6.4 and FC-LS).
REC	Read Exchange Concise ELS (see 6.5 and FC-LS).
RX_ID	Responder Exchange Identifier (see FC-FS-2).
SAM-3	INCITS Project 1561D, <i>SCSI Architecture Model - 3</i> (see 2.3).
SCSI	Small Computer System Interface, any revision.
S_ID	Source_Identifier (see 3.1.53).
SPC-3	INCITS Project 1416D, <i>SCSI Primary Commands - 3</i> (see 2.3).
SRR	Sequence Retransmission Request FC-4 Link Service (see 8.2).
TPRLO	Third Party Process Logout ELS (see FC-LS).
ULP	upper layer protocol (see FC-FS-2).

3.3 Keywords

3.3.1 expected: A keyword used to describe the behavior of the hardware or software in the design models

assumed by this standard. Other hardware and software design models may also be implemented.

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

3.3.3 ignored: A keyword used to describe a bit, byte, word, field or code value that shall not be examined by the receiving SCSI device. The bit, byte, word, field or code value has no meaning in the specified context.

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

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

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

3.3.7 obsolete: A keyword indicating that an item was defined in a prior SCSI standard but has been removed from this standard.

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

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

3.3.10 restricted: A keyword referring to bits, bytes, words, and fields that are set aside for use in other SCSI standards. A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word or field for the purposes of the requirements defined in this standard.

3.3.11 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard. This standard prescribes no specific response by a component if it receives information that violates a mandatory behavior.

3.3.12 should: A keyword indicating flexibility of choice with a strongly preferred alternative; 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 3.1 or in the body of the standard where they first appear. Names of commands, statuses, task management functions, task attributes, Information Units, sense keys, additional sense codes, and additional sense code qualifiers are in all uppercase (e.g., REQUEST SENSE). Lowercase is used for words having the normal English meaning.

The names of fields are in small uppercase (e.g., ALLOCATION LENGTH). When a field name is a concatenation of acronyms, uppercase letters may be used for readability (e.g., NORMACA). 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. Where fields defined in another standard are referenced in this standard, the capitalization conventions of the originating standard are used.

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

Numbers immediately followed by lower-case b (0101b) are binary values.

Numbers or upper case letters immediately followed by lower-case h (FA23h) are hexadecimal values.

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-2.

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

ISO	American
0,6	0.6
1 000	1,000
1 323 462,9	1,323,462.9

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

If a conflict arises between text, tables, or figures, the order of precedence to resolve the conflicts is text, then tables, and finally figures. Exceptions to this convention are indicated in the appropriate subclauses. Not all tables or figures are fully described in the text. Tables show data format and values. Notes do not constitute any requirements for implementors.

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-2 performs those functions required to transfer data from one N_Port or NL_Port to another. In this standard, N_Ports and NL_Ports capable of supporting Fibre Channel Protocol transactions are collectively referred to as FCP_Ports. The FC-2 layer may be treated as a very powerful 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 ports on the loop or between a port on the loop and a port on a switching fabric attached to the loop.

An FC-4 mapping layer uses the services provided by FC-FS-2 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-2. The requirements for the service interface for SCSI are contained in SAM-3.

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 defined in SAM-3 to the Fibre Channel interface defined by FC-FS-2. Link control is performed by standard FC-FS-2 protocols. The task management functions defined by SAM-3 are mapped as described in 4.9 of this standard. The I/O operation defined by SAM-3 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 1.

Table 1 - SCSI and Fibre Channel Protocol functions

SCSI function	FCP equivalent
I/O operation	Exchange
Protocol Service request and response	Sequence
Send SCSI Command request	Unsolicited command IU (FCP_CMND)
Data delivery request	Data descriptor IU (FCP_XFER_RDY)
Data delivery action	Solicited data IU (FCP_DATA)
Send Command Complete response	Command status IU (FCP_RSP)
REQ/ACK for Command Complete	Confirmation IU (FCP_CONF)

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-2 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, as defined in FC-FS-2. 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-3). The Send SCSI Command SCSI transport protocol service request conveys a single request or a list of linked requests from the application client to the FCP service delivery subsystem. Each request contains all the information necessary for the processing of one SCSI command or task management function, including the local storage address and characteristics of data to be transferred by the SCSI command. The Fibre Channel Protocol then performs the following actions using FC-FS-2 services to perform the SCSI command or task management function. The processing of the individual steps of the protocol is consistent with the SCSI architectural model as defined by SAM-3.

The FCP_Port that is the initiator for the command starts an Exchange by sending an unsolicited command IU containing the FCP_CMND IU payload, including some command controls, addressing information, and the SCSI command descriptor block (CDB). The initiator FCP_Port sends the FCP_CMND IU payload to invoke the Send SCSI Command SCSI transport protocol service request (see SAM-3) and 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.

When the device server has interpreted the command, has determined that a write operation is required, and is prepared to request the data delivery service, it sends a data descriptor IU containing the FCP_XFER_RDY IU payload to the initiator indicating which portion of the data is to be transferred. The initiator FCP_Port then transmits a solicited data IU to the target containing the FCP_DATA IU payload requested by the FCP_XFER_RDY IU. The FCP_XFER_RDY IU and FCP_DATA IU payloads constitute the Receive Data-Out protocol service request and Data-Out Received service confirmation described in SAM-3. Data delivery requests containing FCP_XFER_RDY IU and returning FCP_DATA IU payloads continue until the data transfer

requested by the SCSI command is complete. One FCP_DATA IU shall follow each FCP_XFER_RDY IU. If the initiator and target have negotiated to disable the initial FCP_XFER_RDY IU (see 6.3.4), a first burst shall be transferred (see 10.2.10).

When the device server has interpreted the command and has determined that a read operation is required, the target FCP_Port transmits a solicited data IU to the initiator FCP_Port. The solicited data IU shall contain the FCP_DATA IU payload. The FCP_DATA IU constitutes the Send Data-In protocol service request described in SAM-3. Data deliveries containing FCP_DATA IU payloads continue until all data described by the SCSI command is transferred.

When the device server has interpreted the command and has determined that bidirectional transfer is required, 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 sends a data descriptor IU containing the FCP_XFER_RDY IU payload to the initiator to indicate which portion of the data is to be transferred. The initiator FCP_Port then transmits the solicited data IU to the target FCP_Port. The solicited data IU shall contain the FCP_DATA IU payload requested by the FCP_XFER_RDY IU. The FCP_XFER_RDY IU and FCP_DATA IU payloads constitute the Receive Data-Out protocol service request and Data-Out Received service confirmation described in SAM-3. If the device server chooses to send a Data-In transfer first, the target FCP_Port transmits a solicited data IU to the initiator FCP_Port. The solicited data IU shall contain the FCP_DATA IU payload. The FCP_DATA IU constitutes the Send Data-In protocol service request described in SAM-3. 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 SCSI command is transferred. This standard places no restrictions on the order that the device server performs Data-In and Data-Out transfer operations. If the initiator and target have negotiated to disable the initial FCP_XFER_RDY IU (see 6.3.4), a first burst shall be transferred (see 10.2.10).

After all the data has been transferred, the device server transmits the Send Command Complete protocol service response (described in SAM-3) by requesting the transmission of an IU containing the FCP_RSP IU payload. That payload contains the SCSI status and, if the SCSI status is CHECK CONDITION, the autosense data describing the condition. The FCP_RSP IU indicates completion of the SCSI command. If no command linking, error recovery, or confirmed completion is requested, the FCP_RSP IU is the final sequence of the Exchange. The device server determines whether additional linked commands are to be performed in the FCP I/O operation. If this is the last or only command processed in the FCP I/O operation, 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 autosense data.

When the command is completed, returned information is used to prepare and return the Command Complete Received protocol service confirmation to the application client that requested the operation. The returned status indicates whether or not the command was successful. The successful completion of the command indicates that the SCSI device performed the requested operations with the transferred data and that the information was successfully transferred to or from the initiator. 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 a protocol service indication that confirms delivery of the FCP_RSP IU payload, as described in 4.5.

If the command is linked to another command, the FCP_RSP IU payload shall contain the proper status (i.e., INTERMEDIATE or INTERMEDIATE-CONDITION MET) indicating that another command shall be processed. The target FCP_Port shall present the FCP_RSP using the IU that allows command linking, I5 (see 9.1). The initiator FCP_Port shall continue the same Exchange with an FCP_CMND IU, beginning the next SCSI

command. All SCSI commands linked in the FCP I/O operation except the last are processed in the manner described above. SAM-3 defines the cases that interrupt and terminate a series of linked commands. In those cases, the FCP_RSP IU of the last command in the set of linked commands shall be transmitted using the IU that does not allow command linking, I4 (see 9.1). See 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, the device server returns TASK SET FULL status or BUSY status in the FCP_RSP IU as specified by SAM-3.

The Fibre Channel Protocol takes full advantage of the multiplexing and shared bandwidth capabilities provided by various Fibre Channel classes of service. The protocol is designed to operate with any class of service and to provide options for reliable error detection and error recovery independent of the class of service.

SCSI allows the SCSI initiator port function in any FCP_Port and the SCSI target port function 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 may implement both unidirectional and bidirectional commands. Two FCP_RSP IU formats are defined. For commands that set both the RDDATA and WRDATA bits 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 or WRDATA bit 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 and WRDATA bits.

A device server that does not support bidirectional commands shall use the unidirectional FCP_RSP IU payload for presenting all status and error conditions. If a device server that does not support bidirectional commands receives a command that requests read and write operations by setting both the RDDATA and WRDATA bits to one, 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 SCSI commands

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

SAM-3 defines a mechanism to assure ordering of commands. If the initiator transmits a single command and waits for GOOD status before transmitting the next command, the 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 and MODE SELECT commands to examine and set the enable precise delivery checking (EPDC) bit in the Fibre Channel Logical Unit Control mode page. See 10.3.

If a SCSI command requires precise delivery and the device server has the EPDC bit set to one, the application client provides the Command Reference Number (CRN) argument to the Send SCSI Command 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 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 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 4 and table 5 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 CRN field is not valid for an Exchange using precise delivery);
- d) the initiator 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 an FCP_XFER_RDY IU, the first Data frame of an FCP_DATA IU, an FCP_RSP IU, an ACK, or a response to an REC;
- 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-3; and
- g) task management functions shall have the CRN 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 used for booting and initialization may use a CRN of zero).

4.5 Confirmed completion of FCP I/O operations

Some FCP devices 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 bit in the PRLI request FCP Service Parameter page request (see 6.3.4) and PRLI accept FCP Service Parameter page (see 6.3.5) is used to negotiate the use of confirmed completion function.

If the CONFIRMED COMPLETION ALLOWED bit is set to one in the PRLI accept FCP Service Parameter page, 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, the FCP_CONF IU shall not be requested by the FCP_RSP IU.

If command linking is being performed, the target FCP_Port shall not request confirmed completion for an FCP_RSP IU containing INTERMEDIATE or INTERMEDIATE-CONDITION MET status. The target FCP_Port may request confirmed completion:

- a) when providing the FCP_RSP IU for the last command of the set of linked commands; or
- b) when providing the FCP_RSP IU for a command that terminates linking because of an error or CHECK CONDITION status.

Confirmed completion may assist initiators and targets in many environments. Particular 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 autosense data. Upon receiving the FCP_CONF IU, the target FCP_Port may discard its copy of the autosense data;
- b) the confirmed completion function may be used to confirm that a queued SCSI 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. Target FCP_Ports that do not support REC indicate this by performing a Link Service Reject (LS_RJT). See 8.3.

If an error is identified by any of the mechanisms defined in clause 12 and 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 request FCP Service Parameter page and PRLI accept FCP Service Parameter page (see 6.3.4 and 6.3.5):

- a) the initiator FCP_Port may request retransmission using the Sequence Retransmission Request (SRR) FCP FC-4 Link Service Request. See 8.2; and
- b) the initiator FCP_Port and target FCP_Port shall support REC and task retry identification (see 4.7).

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 (REC) and performing the recovery (SRR). The particular case that has been identified as a problem is related to the recovery procedure diagrammed in figure C.7. 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. Many small variations on this scenario may exist.

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, 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 and SRR) each contain the same task retry identifier, unambiguously relating them to the particular command. If the initiator FCP_Port and target FCP_Port do not agree to support task retry identification, the PARAMETER field shall be zero for FCP_CMND IU, REC, and SRR frames.

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 2 provides references to the discovery process for each of the Fibre Channel Protocol capabilities.

Table 2 - Discovery of FCP-3 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	6.3.4
Target FCP_Port performs data overlay	Disconnect-Reconnect mode page EMDP bit	10.2.8
Initiator FCP_Port generates FCP_CONF IU	Process Login	6.3
Target FCP_Port requests FCP_CONF IU	Process Login	6.3
Initiator FCP_Port performs REC	None required, Process Login allowed	4.6
Target FCP_Port accepts REC	LS_RJT if REC not accepted	4.6
Initiator FCP_Port performs SRR	Process Login	6.3
Target FCP_Port accepts SRR	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

4.9 Task management

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

The ABORT TASK task management function is mapped to the FC-FS-2 ABTS basic link service while the other task management functions are mapped into control bits (see table 20) in the FCP_CMND IU. 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) the FCP_CMND IU:

- a) should end with an FCP_RSP IU completion status of 09h (i.e., Task Management function incorrect logical unit number) (see table 24); and
- b) may end with an FCP_RSP IU completion status of 00h (i.e., Task Management function complete) (see table 24).

The FCP_CDB field in FCP_CMND IUs that perform task management functions is ignored.

The QUERY TASK task management function is not supported.

The task management function mappings are specified in table 3.

Table 3 - Task management functions, SAM-3 to FCP-3

SAM-3 function	FCP-3 equivalent
ABORT TASK	FCP recovery abort ^a
ABORT TASK SET	FCP_CMND ABORT TASK SET
CLEAR TASK SET	FCP_CMND CLEAR TASK SET
CLEAR ACA	FCP_CMND CLEAR ACA
LOGICAL UNIT RESET	FCP_CMND LOGICAL UNIT RESET
a) FC-FS-2 basic link services are used to perform the ABORT TASK function.	

FC-FS-2 basic link services and FC-LS extended link services are used to perform the ABORT TASK task management function, to recover Exchange resources, and to re-establish other initial conditions.

The ABORT TASK task management function causes the device server to abort the specified task using the recovery abort protocol, if the task exists. The action is defined in SAM-3. The ABORT TASK task management function is performed by the initiator FCP_Port (i.e., Exchange Originator) using the recovery abort (see 12.3). The specified Exchange shall be terminated by the initiator FCP_Port using the recovery abort. To be compliant with FC-FS-2, 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, recovery abort may be used to recover Exchange resources left in an undefined state by any of the task abort events defined in SAM-3 or by any similar events

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

Tables 4 and 5 summarize the clearing effects resulting from Fibre Channel link actions and SCSI operations, respectively. A 'Y' in the corresponding column of either table indicates the clearing effect upon successful completion of the specified action. 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-2 or other relevant protocol standards. An 'N' in the corresponding column indicates the clearing effect is not performed by the specified action. A '-' in the column indicates that the clearing effect

is not applicable. 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 4 - Clearing effects of link related functions

Clearing effect	FC link action						
	Target Power Cycle	Reset LIP(y,x) ²	LOGO ⁵ , PLOGI	PRLI ⁴ , PRLO ⁵	TPRLO ³	ABTS (Exchange)	ABTS (Sequence)
PLOGI parameters set to default values (see FC-LS) For all logged-in initiator FCP_Ports Only for initiator FCP_Port associated with the action	Y -	Y -	N Y	N N	N N	N N	N N
Open FCP Exchanges terminated For all initiator FCP_Ports Only for initiator FCP_Port associated with the action Only for FCP Exchange associated with ABTS	Y - -	Y - -	N Y -	N Y -	Y - -	N N Y	N N -
FCP Sequence associated with ABTS terminated	-	-	-	-	-	-	Y
Login BB_Credit_CNT set to login value (see FC-FS-2) For all Logged-In NL_Ports For transmitting NL_Port only	Y - -	Y - -	N Y -	N N N	N N N	N N N	N N N
Hard Address Acquisition Attempted	Y ¹	Y ¹	N	N	N	N	N
Process Login parameters cleared ⁶ For all logged-in initiator FCP_Ports Only for FCP_Port associated with the action	Y - -	Y - -	N Y -	N Y -	Y -	N N	N N
CRN set to one For all initiator FCP_Ports Only for initiator FCP_Port associated with the action	Y -	Y -	N Y	N Y	Y -	N N	N N
NOTES: 1 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), the NL_Port shall relinquish its current AL_PA and attempt to acquire its hard address. 2 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, the receiving NL_Port shall perform the behavior described in this column. 3 For a TPRLO, 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/PRLO shall be performed for the designated initiator FCP_Port. See FC-LS. 4 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. 5 Logout and Process Logout may be either implicit or explicit. Implicit logout and Process Logout are specified in FC-LS. 6 A target FCP_Port should send a PRLO to all logged-in initiator FCP_Ports that are logged out as a result of processing a TPRLO with the GLOBAL bit set to one. The PRLO(s) may be sent before or after sending the LS_ACC for the TPRLO.							

Table 5 - Clearing effects of initiator FCP_Port actions

Clearing effect	Initiator FCP_Port action		
	LOGICAL UNIT RESET ²	CLEAR TASK SET ²	ABORT TASK SET ²
PLOGI parameters set to default values (see FC-LS) For all logged-in initiator FCP_Ports Only for initiator FCP_Port associated with the action	N N	N N	N N
Open FCP Sequences Terminated For all initiator FCP_Ports with open FCP Sequences Only for initiator FCP_Port associated with the action Only for FCP Sequences associated with Aborted FCP Exchanges	Y ¹ - -	Y ¹ - -	N Y ¹ -
Login BB_Credit_CNT set to login value (see FC-FS-2) For all Logged-In NL_Ports For transmitting NL_Port only	N N	N N	N N
Hard Address Acquisition Attempted	N	N	N
Process Login parameters cleared For all logged-in initiator FCP_Ports Only for FCP_Port associated with the action	N N	N N	N N
CRN set to one For all initiator FCP_Port Only for initiator FCP_Port associated with the action	Y -	Y -	N Y
NOTES: 1 Exchanges are cleared internally within the target FCP_Port, but open FCP Sequences shall be individually aborted by the initiator FCP_Port via the recovery abort protocol that also has the effect of aborting the associated FCP Exchange. See 12.3. 2 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-3) for the following:

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

4.12 Transport Reset notification events

An NL_Port shall deliver a Transport Reset notification (see SAM-3) for a Reset LIP(y,x) (see FC-AL-2) FC link event if the AL_PD matches the AL_PA of the receiving NL_Port.

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 request and it finds there are not enough login resources to complete the login, the target FCP_Port responds to the PLOGI 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. 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.

4.14 Process Login and Process Logout

The Process Login ELS (PRLI) request is used to establish the FCP operating relationships between two FCP_Ports (see 6.3). The Process Logout ELS (PRLO) 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-2 allows management protocols above the FC-FS-2 interface to perform link data functions. The standard primitive sequences, link management protocols, basic link services, and extended link services are used as required by FCP devices (see FC-FS-2 and FC-LS).

4.16 FCP addressing and Exchange identification

The address of each FCP_Port is defined by its address identifier as described in FC-FS-2. 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.

Addressability of logical units uses the FCP_LUN field provided in the FCP_CMND IU. 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. FCP devices do not use the Process_Associator.

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 task identifier if task retry identification is supported and enabled.

4.17 Use of World Wide Names

As specified in FC-FS-2, each Fibre Channel node and each Fibre Channel port shall have a Worldwide_Name. The Worldwide_Name shall be a unique name using one of the formats defined by FC-FS-2. See Annex A for a description of the mapping of FCP-3 terminology to SAM-3 terminology.

Each target FCP_Port and its associated logical units has knowledge of the Port_Name of each initiator FCP_Port through the Fibre Channel login process. As a result, the relationship between address identifier of the initiator FCP_Port and a persistent reservation for a logical unit may be adjusted (see SPC-3) during those reconfiguration events that may change the address identifier of the initiator FCP_Port. If a target FCP_Port receives a PRLI or a PLOGI from an initiator FCP_Port with a previously known Worldwide_Name, but with a changed initiator port identifier, the device server shall assign the new initiator port identifier to the existing registration and reservation to the initiator FCP_Port having the same Worldwide_Name.

Each logical unit shall be able to present a Worldwide_Name through the INQUIRY command Device Identification VPD page (see SPC-3). 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.

5 FC-FS-2 frame header

5.1 FC-FS-2 frame header overview

The format of the standard FC-FS-2 header as used by the Fibre Channel Protocol is defined in table 6.

Table 6 - FCP frame header

Bits	31– 24	23–16	15–08	07–00
Word				
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-2 definitions. The following explanations of the fields provide information about the use of those fields to implement FCP functionality.

5.2 FC-FS-2 frame header fields

5.2.1 R_CTL field

The values in the R_CTL field identify the frame as part of an FCP I/O operation and identify the information category. All Sequences containing FCP command, data, response, and data descriptor information shall be composed of Device_Data frames.

The information category associated with each IU is defined in table 16 and table 17.

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-2 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 for all frames of SCSI FCP Exchanges.

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-2. 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-2.

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-2.

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. 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 FCP type (i.e., 08h).

For frames of the solicited data category (i.e., FCP_DATA IUs) (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. For the solicited data category (FCP_DATA IUs), the relative offset is the application client buffer offset as described by SAM-3. For solicited data category frames, 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 frames of the unsolicited control category (i.e., FCP_CMND IUs) (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, 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, the PARAMETER field shall contain a value of zero. In both cases, the RELATIVE OFFSET PRESENT bit of the F_CTL field shall be set to zero.

For all other Device_Data frames with the FCP type (i.e., 08h), 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 basic link services (see FC-FS-2) and extended link services defined by FC-LS. The protocol also includes the PRLI and PRLO defined by FC-LS, the PRLI FCP Service Parameter pages defined 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.

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 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 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 send an explicit PRLI.

Note 1 - The requirement to send a PRLI 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. Consider the case where the target FCP_Port WWPN is larger than the initiator FCP_Port WWPN. In this case the target FCP_Port PLOGI request will be processed, but the target FCP_Port is prohibited from sending a PRLI. If the initiator FCP_Port does not send a PRLI, a deadlock occurs.

PRLI 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 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 established by methods outside the scope of this standard. An image pair is removed by an implicit Process Logout or explicit PRLO (see 6.4). If an image pair is not established by an initiator FCP_Port to a target FCP_Port, 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, but does not have an image pair with that target FCP_Port, shall discard the FCP_CMND IU and respond with an explicit PRLO (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 that effected the Nx_Port.

Process_Associators shall not be used in initiator FCP_Ports and target FCP_Ports. If multiple images are required in an initiator FCP_Port, they shall be provided by transparent aliasing of the N_Port Identifier of the initiator FCP_Port. If multiple images are required in a target FCP_Port, they shall be provided by SCSI logical units.

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 6.3.4 and 6.3.5.

Processing of a PRLI or PRLO command 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 6.3.4):

- 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 N_Port or NL_Ports. The relationship does not allow any communication types or paths other than those established by the PRLI.

6.3 PRLI

6.3.1 Use of PRLI by the Fibre Channel Protocol

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

If the PRLI is requesting a binding action by setting the ESTABLISH IMAGE PAIR bit to one, the PRLI 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, 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 6.3.4). The IMAGE PAIR ESTABLISHED bit in the PRLI accept indicates that an image pair was successfully established.

An accept response code indicating other than REQUEST EXECUTED (see 6.3.5 and FC-LS) shall be provided if the PRLI 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 request is not supported or is incorrectly formatted.

The PRLI 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 Process_Associator requirements

Operation of the Process Login depends on the Originator's and Responder's requirements for Process_Associators as specified in FC-LS. Process_Associators are not used in the Fibre Channel Protocol and shall not be used by FCP devices.

The Fibre Channel Protocol assumes that the Originator has knowledge of the capabilities of the Responder. That information may be obtained by performing a PRLI requesting an informative action or by other mechanisms outside the scope of this standard.

6.3.3 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 to perform normal FCP I/O operations. If default

Process Login information is complete enough so that login (i.e., PLOGI) is sufficient to perform an implicit Process Login, then PLOGI shall perform the same clearing actions and establish the same Unit Attention condition that would normally be performed and established by Process Login.

6.3.4 PRLI request FCP Service Parameter page format

The FCP Service Parameter page for the PRLI request is shown in table 7.

Table 7 - FCP Service Parameter page, PRLI request

FCP service parameter	Word	Bit
SCSI FCP (08h)	0	31–24
Reserved for TYPE code extension	0	23–16
ORIGINATOR PROCESS_ASSOCIATOR VALID	0	15
RESPONDER PROCESS_ASSOCIATOR VALID	0	14
ESTABLISH IMAGE PAIR	0	13
Reserved	0	12–0
ORIGINATOR PROCESS_ASSOCIATOR	1	31–0
RESPONDER PROCESS_ASSOCIATOR	2	31–0
Reserved	3	31–10
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

Word 0, Bits 31–24: FCP specific code: The value of 08h in this byte indicates that this Service Parameter page is defined for the Fibre Channel Protocol. (See FC-FS-2.)

Word 0, Bit 15: ORIGINATOR PROCESS_ASSOCIATOR VALID: The ORIGINATOR PROCESS_ASSOCIATOR VALID bit is defined in FC-FS-2. For the Fibre Channel Protocol, the ORIGINATOR PROCESS_ASSOCIATOR VALID bit shall be zero, indicating that the ORIGINATOR PROCESS_ASSOCIATOR is not valid.

Word 0, Bit 14: RESPONDER PROCESS_ASSOCIATOR VALID: The RESPONDER PROCESS_ASSOCIATOR VALID bit is defined in FC-FS-2. For the Fibre Channel Protocol, the RESPONDER PROCESS_ASSOCIATOR VALID bit shall be zero, indicating that the RESPONDER PROCESS_ASSOCIATOR is not valid.

Word 0, Bit 13: ESTABLISH IMAGE PAIR: If the ESTABLISH IMAGE PAIR bit is set to zero, the PRLI only exchanges service parameters as defined in FC-LS.

If the ESTABLISH IMAGE PAIR bit is set to one, the PRLI exchanges service parameters and attempts to establish an image pair as defined in FC-LS.

Word 1: ORIGINATOR PROCESS_ASSOCIATOR: The ORIGINATOR PROCESS_ASSOCIATOR field is the Originator Process_Associator as defined by FC-FS-2.

Word 2: RESPONDER PROCESS_ASSOCIATOR: The RESPONDER PROCESS_ASSOCIATOR field is the Responder Process_Associator as defined by FC-FS-2.

Word 3, Bit 9: TASK RETRY IDENTIFICATION REQUESTED: When the TASK RETRY IDENTIFICATION REQUESTED bit is set to one, the Originator of the PRLI requests that task retry identification (see 4.7) be used. If both the Originator of the PRLI and the Responder to the PRLI request that task retry identification be used, then it shall be used between the initiator FCP_Port and all logical units addressed through that initiator 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 FC-4 Link Service request for that command shall be set to the same value.

When the TASK RETRY IDENTIFICATION REQUESTED bit is set to zero by either the Originator or the Responder to the PRLI, task retry identification shall not be used. The PARAMETER fields for FCP_CMND IUs, for REC ELS requests, and for SRR FCP FC-4 Link Service requests shall be zero.

Word 3, Bit 8: RETRY: When the RETRY bit is set to one, the Originator or Responder is indicating that it supports as an initiator FCP_Port the capability of requesting a retransmission of unsuccessfully transmitted data or as a target FCP_Port the capability of performing a requested retransmission. When the RETRY bit is set to zero, 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, the RETRY bit shall apply to both. SRR 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 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, the SRR shall not be performed by the initiator FCP_Port. If an SRR FCP FC-4 Link Service is received by a target FCP_Port that has set the RETRY bit to zero, the SRR shall be rejected with FCP_RJT.

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

Word 3, Bit 7: CONFIRMED COMPLETION ALLOWED: When the CONFIRMED COMPLETION ALLOWED bit is set to one, the Originator's or Responder's initiator FCP_Port function has the capability of supporting confirmed completion. When the CONFIRMED COMPLETION ALLOWED bit is set to zero, 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 a 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, the target FCP_Port shall not set the FCP_CONF_REQ bit to one.

Word 3, Bit 6: DATA OVERLAY ALLOWED: When the DATA OVERLAY ALLOWED bit is set to one, the Originator or Responder is indicating that its initiator FCP_Port function has the capability of supporting data overlay. When the data overlay allowed bit is set to zero, 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 a 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: When the INITIATOR FUNCTION bit is set to one, the Originator or Responder is indicating it has the capability of operating as an initiator FCP_Port. When the INITIATOR FUNCTION bit is set to zero, the process does not have the capability of operating as an initiator FCP_Port.

Word 3, Bit 4: TARGET FUNCTION: When the TARGET FUNCTION bit is set to one, the Originator or Responder is indicating that it has the capability of operating as a target FCP_Port. When the TARGET FUNCTION bit is set to zero, the process does not have the capability of operating as a target FCP_Port. Both the INITIATOR FUNCTION and the TARGET FUNCTION bits may be set to one. If neither the INITIATOR FUNCTION nor the TARGET FUNCTION bit is set to one, 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 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 send FCP_XFER_RDY on read operations.

Word 3, Bit 0: WRITE FCP_XFER_RDY DISABLED: When the WRITE FCP_XFER_RDY DISABLED bit is set to zero, 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. When the WRITE FCP_XFER_RDY DISABLED bit is set to one, FCP_XFER_RDY IUs shall not be used before the first FCP_DATA IU to be transferred in the SCSI write operation. If both the Originator and Responder choose to disable write FCP_XFER_RDY IUs, then all FCP I/O operations performing SCSI writes 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.

6.3.5 PRLI accept FCP Service Parameter page format

The FCP Service Parameter page for the PRLI accept is shown in table 8.

Table 8 - FCP Service Parameter page, PRLI accept

FCP service parameter	Word	Bit
SCSI FCP (08h)	0	31–24
Reserved for TYPE Code Extension	0	23–16
ORIGINATOR PROCESS_ASSOCIATOR VALID	0	15
RESPONDER PROCESS_ASSOCIATOR VALID	0	14
IMAGE PAIR ESTABLISHED	0	13
Reserved	0	12
ACCEPT RESPONSE CODE	0	11–8
Reserved	0	7–0
Originator Process_Associator	1	31–0
Responder Process_Associator	2	31–0
Reserved	3	31–10
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 request (see table 7) and accept FCP Service Parameter pages.

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

If the IMAGE PAIR ESTABLISHED bit is set to one, the image pair was established.

PRLI ACCEPT RESPONSE CODE: The PRLI ACCEPT RESPONSE CODE field is defined in FC-LS. The values of the PRLI ACCEPT RESPONSE CODE field indicate whether the image pair was successfully created. If the image pair was not created, the value of the PRLI ACCEPT RESPONSE CODE indicates why the request failed or was rejected.

6.4 PRLO

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

The PRLO 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 is being discontinued by the Originator. If the PRLO 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).

For the Fibre Channel Protocol, the PRLO FCP Service Parameter page identifies an image pair where neither the Originator or Responder supports Process_Associators by marking the Originator Process_Associator and Responder Process_Associator as invalid.

The PRLO 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 request is invalid and not accepted.

After Process Logout, no further Fibre Channel Protocol communication is possible between those N_Ports or NL_Ports.

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

6.5 Read Exchange Concise (REC)

See FC-LS for a description of the REC ELS. FCP-3 specific usage of REC 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 task specified by the OX_ID field value and RX_ID field value;
- b) if the destination FCP_Port of the REC 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 a reason code of "Logical error" and a reason code explanation of "Invalid OX_ID-RX_ID combination";
- c) the REC shall be sent in a new Exchange. The Exchange shall be ended by the response to the REC;
- d) if the RX_ID field value in the REC request payload was FFFFh, the RX_ID field value in the REC accept 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 is the number of bytes successfully received by the Device Server for a write or the number of bytes transmitted by the target FCP_Port for a read.

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-4. FCP specific objects are defined in this clause for use by the Name Server. FC-GS-4 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-4) indicates a set of supported data structure type values for Device_Data and FC-4 Link_Data frames (see FC-FS-2). 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 Features object.

7.3 FC-4 Features object

The FC-4 Features object (see FC-GS-4) defines a 4-bit field for each possible FC TYPE code. The object is a 32-word array of 4-bit values. The 4-bit FC-4 Feature field for FCP is inserted in bits 3 - 0 of word 1. The format of the 4-bit FC-4 Feature field for FCP is shown in table 9.

Table 9 - FCP definition of FC-4 Feature bits

Word 1 bit	Description of bit
3	Reserved
2	Reserved
1	FCP initiator function supported
0	FCP target function supported

An FCP_Port shall register its 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 N_Port or NL_Port from a Name Server using a GFF_ID Request CT_IU, which requests the FC-4 Features object for a specified Port Identifier. The 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 N_Port or NL_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 FC-4 Link Service definitions

8.1 FC-4 Link Services for the Fibre Channel Protocol

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 bits 31-28 (Word 0) 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 Category 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 FC-4 Link Service Requests and Responses defined in this standard are shown in table 10.

Table 10 - FCP FC-4 Link Service Requests and Responses

Encoded Value Word 0 of Payload (bits 31-24)	Description	Abbr.	Request/ Response	Reference
14h	Sequence Retransmission Request	SRR	Request	8.2
01h	FCP FC-4 Link Service Reject	FCP_RJT	Response	8.3

8.2 Sequence Retransmission Request (SRR)

The SRR FCP FC-4 Link Service 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, the PARAMETER field of the request Sequence shall contain the task retry identifier for the task 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 8.3) Sequence with a reason code of “Unable to perform the command request” and a reason code explanation of “Invalid OX_ID-RX_ID combination”.

If the target FCP_Port is unable to retransmit the Sequence or data at the requested Relative Offset, the target FCP_Port shall respond with an FCP_RJT Sequence with a reason code of “Unable to perform the command request” and a reason code explanation of “Unable to supply requested data”.

If the initiator FCP_Port receives an FCP_RJT response, the initiator FCP_Port shall terminate the Exchange referenced by the SRR using recovery abort (see 12.3.2).

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

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

For unacknowledged classes, the Sequence Count for a retransmitted FCP_DATA IU shall 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.

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. In the event that the target FCP_Port responds to the SRR with an FCP_RJT, 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**. A target FCP_Port that has agreed during Process Login to support retransmission should not reject requests for retransmission of the requested frames unless unusual conditions make the retransmission impossible. SRR 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 containing a reason code of "Unable to support command request" and a reason code explanation of "Unable to supply requested data".

Payload for SRR FCP FC-4 Link Service request:

The format of the Payload is shown in table 11.

Table 11 - SRR 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 R_CTL FOR IU field encoding is as described in FC-FS-2 (i.e., Data Descriptor (FCP_XFER_RDY IU), Command Status (FCP_RSP IU), Solicited Data (FCP_DATA IU)).

The RELATIVE OFFSET parameter is only valid if the R_CTL FOR IU field is set to 01h for Solicited Data or to 05h for 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 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.

Possible responses to SRR FCP FC-4 Link Service request:

FCP_RJT

Signifies rejection of the SRR request.

SRR Accept

Signifies that the Payload is accepted.

SRR Accept Payload:

The payload for the SRR Accept is shown in table 12.

Table 12 - SRR Accept Payload

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

8.3 FCP FC-4 Link Service Reject (FCP_RJT)

FCP_RJT notifies the originator of an FCP FC-4 Link Service request that the FCP FC-4 Link Service request Sequence has been rejected. A four-byte reason code shall be contained in the Data_Field (see table 13).

An FCP_RJT may be a response Sequence to any FCP FC-4 Link Service request.

Addressing:

The D_ID field designates the source of the FCP FC-4 Link Service request being rejected. The S_ID field designates the destination of the request FCP FC-4 Link Service frame Sequence being rejected.

Payload for FCP_RJT:

The first word of the Payload shall contain the FCP_RJT code (01000000h). The next four bytes of this field shall contain a reason code and reason code explanation for rejecting the request. The format of the FCP_RJT Payload is shown in table 13.

Table 13 - FCP_RJT Payload

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

The reason codes for FCP_RJT are specified in table 14.

Table 14 - FCP_RJT reason codes

Encoded Value (Bits 23-16)	Description	Meaning
01h	Invalid FCP FC-4 Link Service 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	This indicates that an error has been detected that violates the rules of the FC-FS-2 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.

FCP_RJT reason code explanations

Table 15 lists the reason code explanations for FCP FC-4 Link Service commands.

Table 15 - FCP_RJT reason code explanations

Encoded Value (Bits 15-8)	Description	Applicable commands
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 16 for IUs sent to target FCP_Ports, and in table 17 for IUs sent to initiator FCP_Ports. Each IU shall be contained in a single Sequence (see 3.1.52). 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 16 - FCP Information Units (IUs) sent to target FCP_Ports

IU	SCSI primitive	Data block		F/M/L	SI	M/O																										
		CAT	Content																													
T1	Command / Task Mgmt Rqst	6	FCP_CMND	F	T	M																										
T2	Command request	6	FCP_CMND	F	H	O																										
T3	Command request (Linked)	6	FCP_CMND	M	T	O																										
T4	Command request (Linked)	6	FCP_CMND	M	H	O																										
T6	Data-Out action	1	FCP_DATA	M	T	M																										
T12	Confirm	3	none	L	T	O																										
<p>Notes:</p> <p>T5, T7, T8, T9, T10, and T11 are obsolete.</p> <p>T2 and T4 are only permitted when transfer ready IUs are disabled (see table 7).</p> <p>T3 and T4 are only permitted for linked SCSI commands.</p> <p>T2 and T4 allow optional sequence streaming during write operations.</p> <p>T12 is only permitted in response to an I5 frame requesting the confirmed completion protocol.</p> <p>See table 17</p> <p>Key:</p> <table><tr><td>IU</td><td>Information Unit identifier</td></tr><tr><td>CAT</td><td>Information category of Device_Data frames carrying the data block (see FC-FS-2)</td></tr><tr><td>Content</td><td>Contents (payload) of data block</td></tr><tr><td>F/M/L</td><td>First/Middle/Last Sequence of Exchange (FC-FS-2)</td></tr><tr><td></td><td>F First</td></tr><tr><td></td><td>M Middle</td></tr><tr><td></td><td>L Last</td></tr><tr><td>SI</td><td>Sequence Initiative: Held or Transferred (FC-FS-2)</td></tr><tr><td></td><td>H Held</td></tr><tr><td></td><td>T Transferred</td></tr><tr><td>M/O</td><td>Mandatory/Optional Sequence</td></tr><tr><td></td><td>M Mandatory</td></tr><tr><td></td><td>O Optional</td></tr></table>							IU	Information Unit identifier	CAT	Information category of Device_Data frames carrying the data block (see FC-FS-2)	Content	Contents (payload) of data block	F/M/L	First/Middle/Last Sequence of Exchange (FC-FS-2)		F First		M Middle		L Last	SI	Sequence Initiative: Held or Transferred (FC-FS-2)		H Held		T Transferred	M/O	Mandatory/Optional Sequence		M Mandatory		O Optional
IU	Information Unit identifier																															
CAT	Information category of Device_Data frames carrying the data block (see FC-FS-2)																															
Content	Contents (payload) of data block																															
F/M/L	First/Middle/Last Sequence of Exchange (FC-FS-2)																															
	F First																															
	M Middle																															
	L Last																															
SI	Sequence Initiative: Held or Transferred (FC-FS-2)																															
	H Held																															
	T Transferred																															
M/O	Mandatory/Optional Sequence																															
	M Mandatory																															
	O Optional																															

Table 17 - FCP Information Units (IUs) sent to initiator FCP_Ports

IU	SCSI primitive	Data block		F/M/L	SI	M/O
		CAT	Content			
I1	Data delivery request	5	FCP_XFER_RDY (Write)	M	T	M
I3	Data-In action	1	FCP_DATA	M	H	M
I4	Command/Task Mgmt Response	7	FCP_RSP	L	T	M
I5	Response (Linked or confirm request)	7	FCP_RSP	M	T	O
<div>Notes:</div> <div>I2, I6, and I7 are obsolete.</div> <div>I5 is permitted for linked SCSI commands or to request the confirm completion protocol.</div> <div>I3 allows optional sequence streaming to I3, I4, or I5.</div> <div>Key:</div> <div><div>IU</div><div>Information Unit identifier</div><div>CAT</div><div>Information category of Device_Data frames carrying the data block (FC-FS-2)</div><div>Content</div><div>Contents (payload) of data block</div><div>F/M/L</div><div>First/Middle/Last Sequence of Exchange (FC-FS-2)</div><div><div>F</div><div>First</div><div>M</div><div>Middle</div><div>L</div><div>Last</div></div><div>SI</div><div>Sequence Initiative: Held or Transferred (FC-FS-2)</div><div><div>H</div><div>Held</div><div>T</div><div>Transferred</div></div><div>M/O</div><div>Mandatory/Optional Sequence</div><div><div>M</div><div>Mandatory</div><div>O</div><div>Optional</div></div></div>						

9.2 FCP_CMND IU

9.2.1 Overview and format of FCP_CMND IU

The FCP_CMND IU carries either a SCSI Command or a task management request. If an invalid combination of bits is set in the FCP_CMND IU, the target FCP_Port shall respond with an FCP_RSP IU with the RSP_CODE

field set to 02h (i.e., FCP_CMND fields invalid). The FCP_CMND IU shall contain the values and control fields defined in table 18 in its payload.

Table 18 - FCP_CMND IU Payload

Bit Byte	7	6	5	4	3	2	1	0
0	FCP_LUN							
7								
8	COMMAND REFERENCE NUMBER							
9	Reserved	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)	FCP_DL						(LSB)
n+2								
n+3								
n+4								
n+5	(MSB)	FCP_BIDIRECTIONAL_READ_DL (if any)						(LSB)
n+6								
n+7								
n+8								

9.2.2 FCP_CMND IU field descriptions

9.2.2.1 FCP_LUN field

The FCP logical unit number (FCP_LUN) field contains the address of the logical unit (i.e., the logical unit number) in the SCSI target device. See SAM-3.

Each target FCP_Port shall accept an INQUIRY command addressed to LUN 0. If LUNs other than zero are supported by the SCSI target device, LUN 0 shall implement the REPORT LUNS command. See SPC-3.

If the FCP_LUN field contains a valid logical unit address the command or task management function shall be routed to the addressed logical unit. If the addressed logical unit does not exist, the SCSI target device shall follow the rules for selection of incorrect logical units as specified in SAM-3.

9.2.2.2 COMMAND REFERENCE NUMBER field

The COMMAND REFERENCE NUMBER (CRN) 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, a nonzero value in the CRN 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, a zero value in the CRN field indicates that command shall not be verified for precise delivery. If precise delivery checking is not enabled, the COMMAND REFERENCE NUMBER field shall be ignored by the device server. If the FCP_CMND IU specifies a task management function, the CRN field shall be reserved and set to zero and the FCP_CMND IU shall not be verified for precise delivery.

9.2.2.3 PRIORITY field

The PRIORITY field specifies the relative scheduling of this task in relation to other tasks already in the task set for processing by the device server (see SAM-3). If the TASK ATTRIBUTE field contains a value other than SIMPLE, then this field is reserved.

9.2.2.4 TASK ATTRIBUTE field

The TASK ATTRIBUTE field contains values that specify the task attribute (see SAM-3) associated with the CDB, as shown in table 19.

Table 19 - TASK ATTRIBUTE field values

Task attribute code	Task attribute	Description
000b	SIMPLE	Requests that the task be managed according to the rules for a simple task attribute (see SAM-3).
		Requests that the task be managed according to the rules for a simple task attribute and priority (see SAM-3).
001b	HEAD OF QUEUE	Requests that the task be managed according to the rules for a head of queue task attribute (see SAM-3).
010b	ORDERED	Requests that the task be managed according to the rules for an ordered task attribute (see SAM-3). 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-3).
101b	Obsolete	
110b-111b	Reserved	

9.2.2.5 TASK MANAGEMENT FLAGS field

The TASK MANAGEMENT FLAGS field contains flags that request that a task management function be performed. Task management functions shall be requested by the initiator FCP_Port (Exchange Originator) using a new

Exchange. If any task management flag bit is set to one, the FCP_CDB field, the FCP_DL field, the TASK ATTRIBUTE field, the RDDATA bit, and the WRDATA bit shall be ignored. If any bit in the TASK MANAGEMENT FLAGS field is set to one, the FCP_BIDIRECTIONAL_READ_DL field shall not be included in the FCP_CMND IU payload. If more than one task management flag bit is set to one in any FCP_CMND IU, the task management functions shall not be processed and the FCP_RSP IU shall contain the RSP_CODE field set to 02h (i.e., FCP_CMND fields invalid).

The clearing actions performed by task management functions are specified in table 5. The format of the TASK MANAGEMENT FLAGS field is specified in table 20.

Table 20 - TASK MANAGEMENT FLAGS field

bit	task management function ^a
7	Obsolete
6	CLEAR ACA
5	Obsolete
4	LOGICAL UNIT RESET
3	Reserved
2	CLEAR TASK SET
1	ABORT TASK SET
0	Reserved
a) The ABORT TASK management function is specified in 4.9.	

The **CLEAR ACA** bit, when set to one, causes the ACA condition to be cleared. When the task manager clears the ACA condition, any task within that task set may be completed subject to the rules for task management specified by SAM-3. If there is no ACA condition present, the CLEAR ACA task management function shall be accepted and the FCP_RSP IU shall contain a RSP_CODE field set to 00h (i.e., Task Management function complete).

When set to zero, the ACA condition remains unchanged.

The use of the ACA bit in the CDB control field and the implementation of ACA is described in SAM-3.

Depending on the mode page parameters that have been established (see SPC-3), additional FCP I/O operations may have to be aborted by the recovery abort as part of the process of clearing the automatic contingent allegiance.

The CLEAR ACA is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange.

Support of the CLEAR ACA bit is mandatory in the Fibre Channel Protocol if the FCP device sets the NORMACA bit to one in the INQUIRY data. It shall not be sent to a logical unit with a NORMACA bit equal to zero in the INQUIRY data.

The **LOGICAL UNIT RESET** bit, when set to one, performs a LOGICAL UNIT RESET task management function as defined in SAM-3. LOGICAL UNIT RESET aborts all tasks in the task set for the logical unit and

performs a LOGICAL UNIT RESET for all dependent logical units. Support of the LOGICAL UNIT RESET bit is mandatory for the Fibre Channel Protocol.

The LOGICAL UNIT RESET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. LOGICAL UNIT RESET resets the internal states of the target FCP_Port and logical unit as shown in 4.10. Exchange resources to be cleared may be cleared by the following mechanisms:

- a) a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the LOGICAL UNIT RESET for each task in the logical unit known to that initiator FCP_Port;
- b) a task, if any, for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET is ended in the logical unit. The initiator FCP_Port for that task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or
- c) a task for an initiator FCP_Port other than the initiator FCP_Port that sent the LOGICAL UNIT RESET 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-3.

NOTE 3 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a LOGICAL UNIT RESET task management function if the Control mode page indicates that the TASK ABORTED status is supported.

The **CLEAR TASK SET** bit causes all tasks from all initiator FCP_Ports in the specified task set to be aborted as defined in SAM-3. Support of the CLEAR TASK SET bit is mandatory for the Fibre Channel Protocol.

The CLEAR TASK SET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. CLEAR TASK SET resets internal states of the target FCP_Port as shown in 4.10. Exchange resources to be cleared may be cleared by one or more of the following mechanisms:

- a) a recovery abort sequence (see 12.3) may be generated by the initiator FCP_Port that sent the CLEAR TASK SET for each task known to that initiator FCP_Port;
- b) a task, 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 task shall determine by a timeout that the task did not finish. Subsequent retries fail because the task resources have been cleared in the logical unit, so the initiator FCP_Port shall clear the Exchange resources with a recovery abort sequence. See 12.3; or
- c) a task 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-3.

NOTE 4 - SAM-3 has defined the TASK ABORTED status for tasks terminated by a CLEAR TASK SET task management function if the Control mode page indicates that the TASK ABORTED status is supported.

The **ABORT TASK SET** bit set to one requests the ABORT TASK SET task management function to be performed as defined in SAM-3. Support of the ABORT TASK SET bit is mandatory in the Fibre Channel Protocol.

The ABORT TASK SET is transmitted by the initiator FCP_Port (Exchange Originator) using a new Exchange. ABORT TASK SET resets internal states of the target FCP_Port as shown in 4.10. Exchange resources may be cleared by a recovery abort sequence (see 12.3) generated by the initiator FCP_Port that sent the ABORT TASK SET for each task known to the initiator FCP_Port.

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 value of the ADDITIONAL FCP_CDB LENGTH field shall be set to zero for task management requests.

9.2.2.7 RDDATA and WRDATA bits

If the RDDATA bit is set to one, the initiator FCP_Port expects to receive FCP_DATA IUs from the target FCP_Port. This is a SCSI read operation.

If the WRDATA bit is set to one, the initiator FCP_Port expects to transmit FCP_DATA IUs to the target FCP_Port. This is a SCSI write operation.

If the RDDATA bit and WRDATA bit are both set to one, the initiator FCP_Port expects both a SCSI read operation and a SCSI write operation. This is a bidirectional SCSI 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 SCSI command.

If the RDDATA bit and WRDATA bit are both set to zero, 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 SCSI 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 5 - 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 02h (i.e., FCP_CMND fields invalid) (see table 24) 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 shall be ignored if any task management flag is set to one.

The CDB format is defined by SAM-3 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 any task management flag is set to one. 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 SCSI read operation, 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 SCSI command. The FCP_DL field is the Data-In Buffer Size defined by SAM-3.

For a SCSI write operation, 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 SCSI command. The FCP_DL field is the Data-Out Buffer Size defined by SAM-3.

For a bidirectional SCSI 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 SCSI command. The FCP_DL field is the Data-Out Buffer Size defined by SAM-3.

An FCP_DL 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 SCSI 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 SCSI command. The FCP_BIDIRECTIONAL_READ_DL field is the Data-In Buffer Size defined by SAM-3.

An FCP_BIDIRECTIONAL_READ_DL value of zero indicates that no read operation is expected regardless of the state of the RDDATA bit and that no FCP_DATA IUs shall be transferred for read data.

If either RDDATA or WRDATA is set to zero, the FCP_BIDIRECTIONAL_READ_DL field is not 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-3 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 if 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. If the target FCP_Port and initiator FCP_Port have negotiated write FCP_XFER_RDY disabled, 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 6.3.4).

The first 8 bytes of the FCP_XFER_RDY IU payload are defined in FC-FS-2 for all IUs of category 5, the data descriptor category. The fields defined in FC-FS-2 are given FCP names for use in this standard. The format of the FCP_XFER_RDY IU payload is shown in table 21.

Table 21 - FCP_XFER_RDY IU payload

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

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-2.

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 the same as the SAM-3 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-2 and the value in the FCP_BURST_LEN field is the same as the SCSI data delivery request byte count (see SAM-3).

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 FCP_BURST_LEN field and the value of FCP_DATA_RO shall not exceed the value of FCP_DL. 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, the target FCP_Port shall return an FCP_RSP IU with the RSP_CODE field set to 03h (i.e., FCP_DATA parameter mismatch with FCP_DATA RO).

If required by the PRLI FCP Service Parameters, 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), 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, 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, 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, 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-2 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 length FCP_DL. 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 FCP_DL be transferred, 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, the target FCP_Port shall discard the excess bytes. Because there were fewer bytes provided 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 set 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), the error detection and recovery procedure described in clause 12 may be invoked or the FCP I/O operation may be terminated with a recovery abort (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 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 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 FCP_DL be transferred, 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 value 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 data operation of a bidirectional SCSI 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 SCSI command, the FCP I/O operation may be terminated with a recovery abort (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 autosense data. The target FCP_Port shall send 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.

The bits and fields in bytes 10 and 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. Bytes 10 and 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 command linking is being performed, an FCP_RSP IU is provided for each command. For linked commands, INTERMEDIATE status or INTERMEDIATE - CONDITION MET status indicates successful completion of a command with no other information valid if all other fields are zero. If command linking is requested, the use of the INTERMEDIATE or INTERMEDIATE-CONDITION MET status indicates that linking shall be performed. The LINKED COMMAND COMPLETE or LINKED COMMAND COMPLETE (WITH FLAG) Service Response defined by SAM-3 is implicit in the presentation of INTERMEDIATE or INTERMEDIATE-CONDITION MET status in the FCP_RSP IU.

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

If a SCSI device error is detected by a target FCP_Port while the target FCP_Port has Sequence Initiative for the Exchange associated with the error, 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 a SCSI device error is detected by a device server while the target FCP_Port does not have Sequence Initiative for the Exchange associated with the error, 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.

In the event that Sequence Initiative is not received within $RR_TOV_{SEQ_INIT}$ (see 11.4), the target FCP_Port may implicitly terminate the affected Exchange.

The content of the FCP_RSP IU is indicated in table 22.

Table 22 - FCP_RSP IU Payload

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
7								
8	(MSB)							
9	RETRY DELAY TIMER							
	(LSB)							
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)							
15	FCP_RESID							
	(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 23)							
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 RETRY DELAY TIMER field

The RETRY DELAY TIMER field contains the retry delay timer code (see SAM-4).

9.5.3 FCP_BIDI_RSP bit

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

9.5.4 FCP_BIDI_READ_RESID_UNDER bit

If the FCP_BIDI_READ_RESID_UNDER bit is set to one, the FCP_BIDIRECTIONAL_READ_RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred. The application client shall examine the FCP_BIDIRECTIONAL_READ_RESID FIELD field in the context of the command to determine 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, 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. The application client shall examine the FCP_BIDIRECTIONAL_READ_RESID FIELD field in the context of the command to determine whether or not an error condition occurred.

9.5.6 FCP_CONF_REQ bit

If the FCP_CONF_REQ bit is set to one, 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, 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, the FCP_RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred. The application client shall examine the FCP_RESID field in the context of the command to determine whether or not an error condition occurred.

9.5.8 FCP_RESID_OVER bit

If the FCP_RESID_OVER bit is set to one, the FCP_RESID field is valid and contains the count of bytes that could not be transferred because the FCP_DL was not sufficient. The application client shall examine the FCP_RESID field in the context of the command to determine 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, 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. The application client shall examine the FCP_SNS_INFO field to determine whether or not an error condition occurred.

If the FCP_SNS_LEN_VALID bit is set to zero, 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, 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. The application client shall examine the FCP_RSP_INFO field to determine whether or not an error condition occurred. When the FCP_RSP_LEN_VALID bit is set to one, 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 shall indicate the completion status of the task management function.

If the FCP_RSP_LEN_VALID bit is set to zero, the FCP_RSP_LEN field is not valid and shall be treated as if its value were zero. When the FCP_RSP_LEN_VALID bit is set to zero, 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 SCSI command being completed, as defined by SAM-3.

9.5.12 FCP_RESID field

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

For bidirectional SCSI commands, if the FCP_RESID_UNDER bit is set to one, 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 SCSI commands, if the FCP_RESID_OVER bit is set to one, 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 SCSI 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 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 SCSI command.

If the FCP_RESID_UNDER bit is set to one, 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:

$$\text{FCP_RESID} = \text{FCP_DL} - (\text{highest offset of any byte transmitted} + 1)$$

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, refer to 9.4.2 or 9.4.3. The FCP_RESID value is defined as follows:

$$\text{FCP_RESID} = (\text{transfer length required by command}) - \text{FCP_DL}$$

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

NOTE 6 - Some early implementations presented the FCP_RSP IU without the FCP_RESID, FCP_SNS_LEN, and FCP_RSP_LEN fields if the FCP_RESID_UNDER, FCP_RESID_OVER, FCP_SNS_LEN_VALID, and FCP_RSP_LEN_VALID bits 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 SCSI commands.

For bidirectional commands, if the FCP_BIDI_READ_RESID_UNDER bit is set to one, 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, 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 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 SCSI command.

If the FCP_BIDI_READ_RESID_UNDER bit is set to one, 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:

$$\text{FCP_BIDIRECTIONAL_READ_RESID} = \text{FCP_BIDIRECTIONAL_READ_DL} - (\text{highest offset of any byte transmitted} + 1)$$

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, refer to 9.4.3. The FCP_BIDIRECTIONAL_READ_RESID value is defined as follows:

$$\text{FCP_BIDIRECTIONAL_READ_RESID} = (\text{read transfer length required by command}) - \text{FCP_BIDIRECTIONAL_READ_DL}$$

If the FCP_BIDI_READ_RESID_UNDER and the FCP_RESID_OVER bits are both set to zero, 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, the FCP_SNS_LEN field specifies the number of valid bytes of FCP_SNS_INFO.

If the FCP_SNS_LEN_VALID bit is zero, the FCP_SNS_LEN field is not valid and shall be treated as if its value were zero. No FCP_SNS_INFO is provided.

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, the FCP_RSP_LEN field specifies the number of valid bytes of FCP_RSP_INFO. The number shall be 00000004h, or 00000008h.

If the FCP_RSP_LEN_VALID bit is zero, 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, 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-2 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 23.

Table 23 - FCP_RSP_INFO field format

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

The valid RSP_CODE values are specified in table 24.

Table 24 - RSP_CODE definitions

Value	RSP_CODE definition
00h	Task Management function complete
01h	FCP_DATA length different than FCP_BURST_LEN
02h	FCP_CMND fields invalid
03h	FCP_DATA parameter mismatch with FCP_DATA_RO
04h ^a	Task Management function rejected
05h ^a	Task Management function failed
09h ^a	Task Management function incorrect logical unit number
06h – 08h 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. If the Exchange is aborted before the FCP_RSP IU is returned, the completion status is unknown. If the RSP_CODE is set to 05h (i.e., Task Management function failed), 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 autosense data specified by SPC-3. The proper FCP_SNS_INFO shall be presented when the SCSI status byte of CHECK CONDITION is presented as specified by SAM-3. If no conditions requiring the presentation of SCSI sense data have occurred, the FCP_SNS_INFO field shall not be

included in the FCP_RSP IU and the FCP_SNS_LEN_VALID bit shall be zero. FCP devices shall perform autosense.

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 when 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 25.

Table 25 - Mode pages for FCP

Page code	Description	Reference
02h	Disconnect-Reconnect mode page	10.2
18h	Fibre Channel Logical Unit Control mode page	10.3
19h	Fibre Channel Port Control mode page	10.4
3Fh	Return all mode pages (valid only for the MODE SENSE command)	SPC-3

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 26) allows the application client to modify the behavior of the service delivery subsystem. This subclause specifies the parameters defined by SPC-3 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, normally the Fibre Channel interface circuitry. 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, the device server shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

Table 26 - 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						(LSB)
5								
6	(MSB)	DISCONNECT TIME LIMIT						(LSB)
7								
8	(MSB)	CONNECT TIME LIMIT						(LSB)
9								
10	(MSB)	MAXIMUM BURST SIZE						(LSB)
11								
12	EMDP	FAA	FAB	FAC	RESTRICTED	RESTRICTED		
13	Reserved							
14	(MSB)	FIRST BURST SIZE						(LSB)
15								

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) and Fibre Channel Class 1 connections, 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.

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-3. 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-3.

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 send data. Device servers that do not implement the requested ratio should round down to the nearest implemented ratio as defined in SPC-3.

The value contained in the BUFFER EMPTY RATIO field is defined by SPC-3.

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, the device server shall end the interconnect tenancy. This value may be rounded as defined in SPC-3. A value of zero indicates that there is no bus inactivity limit.

NOTE 7 - 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 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-3. A value of zero indicates there is no limit on the amount of data transferred per data transfer operation. This 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 SCSI command. If the EMDP bit is set to zero, the target FCP_Port shall generate continuously increasing relative offset values for each FCP_DATA IU for a single SCSI command. If the EMDP bit is set to one, the target FCP_Port may transfer the FCP_DATA IUs for a single SCSI command in any order. If the EMDP bit is set to zero, data overlay is prohibited even if it is allowed by the state of the PRLI FCP Service Parameter page DATA OVERLAY ALLOWED bit. 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, 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, FAB, FAC bits

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 send 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 send to an initiator FCP_Port.

The FAC bit controls arbitration when the target FCP_Port has an FCP_RSP IU frame to send to an initiator FCP_Port. If the target FCP_Port intends to send multiple frame types, 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 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, 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, 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 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 parameter 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 27) 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 and MODE SELECT command rules specified by SPC-3.

Table 27 - 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							

An ENABLE PRECISE DELIVERY CHECKING (EPDC) bit of one indicates that the logical unit shall use the precise delivery function defined by this standard (see 4.4). When the EPDC bit is set to zero, the logical unit shall not use the precise delivery function and shall ignore the contents of the CRN 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, the EPDC bit shall be masked as not changeable and shall follow the MODE SENSE and MODE SELECT command rules specified by SPC-3.

If the Fibre Channel Logical Unit Control mode page is not supported by a logical unit, the initiator 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-3). The mode page shall be implemented by logical unit 0 and may be implemented by logical units other than logical unit 0. The implementation of any bit and its associated functions is optional. The mode page follows the MODE SENSE and MODE SELECT command rules specified by SPC-3.

Some of the bits 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 shown in shown in table 28.

Table 28 - 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) shall not generate a LIP following insertion into the loop. The target FCP_Port shall respond to a LIP when it is received. If the DTOLI bit is set to zero, the target FCP_Port attached to an arbitrated loop shall generate LIP(F7,xx) after it enables a port into a loop. If the target FCP_Port is attached to an arbitrated loop and detects loop failure at its input, 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) shall wait for an initiator FCP_Port to send 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, the target FCP_Port shall enable itself onto the loop in 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) 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, 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) 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, the target FCP_Port shall enter the nonparticipating state. If the target FCP_Port detects loop initialization while in the nonparticipating state, the target FCP_Port 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, the target FCP_Port shall attempt to obtain the address in the LIPA 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, the target FCP_Port follows 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) 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, 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 FLOGI attached to an arbitrated loop (see FC-AL-2) shall not require receipt of Address or Port Discovery (i.e., ADISC or PDISC ELSS) following loop initialization as described in FC-DA. The logical units shall resume processing tasks on completion of loop initialization. If the DDIS bit is set to zero, the target FCP_Port shall wait to complete target discovery, as defined by FC-DA, before allowing processing of tasks to resume.

Target FCP_Ports not attached to an arbitrated loop shall ignore the DDIS bit. A target FCP_Port with a valid FLOGI 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, a target FCP_Port attached to an arbitrated loop (see FC-AL-2) 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, 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 return 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, a target FCP_Port attached by an arbitrated loop (see FC-AL-2) 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. If the DTFD bit is set to zero, the target FCP_Port attached by an arbitrated loop 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. Target FCP_Ports not attached to an arbitrated loop shall ignore the DTFD bit.

10.4.10 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 UNITS field indicates the units for the RR_TOV_{SEQ_INIT} field value, according to table 29.

Table 29 - Values for RR_TOV UNITS

Byte 6			Units of measure for RR_TOV _{SEQ_INIT}
bit 2	bit 1	bit 0	
0	0	0	No timer is specified
0	0	1	0.001 s
0	1	1	0.1 s
1	0	1	10 s
All other values			Reserved

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, 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-3 recovery procedures. In addition, the clause defines those timers used only by this standard.

Table 30 - Timer summary

Timer	Implementation Mandatory (M) or Optional (O)		Description	Default Value	Notes	Ref
	Initiator	Target				
E_D_TOV	M	O	Error_Detect_Timeout Value	2 s	2,3	11.2
R_A_TOV _{SEQ_QUAL}	M	O	Resource_Allocation Timeout Value	Private loop = 0 s Public loop = 10 s	1,2	11.3
R_A_TOV _{ELS}	M	M		Private loop = 2 s Public loop = 10 s	1,2	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 0: 2 s If RETRY bit is set to 1: ≥ REC_TOV + 2xR_A_TOV _{ELS} + 1 s		11.4
REC_TOV	M	M	REC Timeout Value	≥ E_D_TOV + 1 s (minimum)	4	11.5
ULP_TOV	M		Upper Level Protocol Timeout Value	If RETRY bit is set to 0: ≥ Operation-specific timer + E_D_TOV + 1 s If RETRY bit is set to 1: ≥ Operation-specific timer + 2xRR_TOV _{SEQ_INIT}		11.6

NOTES:

- 1 R_A_TOV is defined by FC-FS-2.
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, the same value is applied for both uses. Other Fibre Channel standards may specify different default values for R_A_TOV for different topologies.
- 2 Target FCP_Ports that support Class 2 delivery service shall implement this timer.
- 3 E_D_TOV default timeout values are defined by FC-FS-2, 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 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-2 before fabric login and the value obtained in the FLOGI ACC after fabric login.
- 4 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-2, FC-AL-2, FC-LS, and FC-DA). For FCP-3 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_Qualifier is composed of the S_ID, D_ID, OX_ID, RX_ID, and SEQ_ID fields. 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, the queue for a given Sequence shall be empty before an ABTS for that Sequence is sent, or the act of sending 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 FLOGI, Public Loop devices shall use the value of R_A_TOV specified by the Fabric in the FLOGI 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 a specific initiator 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). $RR_TOV_{SEQ_INIT}$ 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 either of these two conditions is not recovered successfully before expiration of RR_TOV, a target FCP_Port may implicitly or explicitly perform a LOGO 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 10.4.10).

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 REC 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 31 describes REC_TOV usage pertaining to the initiator FCP_Port.

Table 31 - Initiator FCP_Port REC_TOV Usage

Timer starts or restarts after:	Timer stops without sending REC 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 was sent for an FCP_CMND IU and an ACC was received indicating the command is in progress (i.e. REC polling interval).	FCP_RSP IU is received for the FCP_CMND IU or the Exchange is aborted.

Table 32 describes REC_TOV usage pertaining to the target FCP_Port.

Table 32 - Target FCP_Port REC_TOV usage

Timer starts after:	Timer stops without sending REC 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-2. See 12.2 for a list of these mechanisms.

To recover from these errors, all FCP compliant initiator FCP_Ports shall be capable of invoking the recovery abort function to terminate a failing Exchange and to recover the associated resources as described in 12.3. All FCP compliant target FCP_Ports shall be capable of processing the requested recovery abort 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-3 device may accept or reject error detection inquiries.

12.1.2 Sequence level

Sequence level error recovery as described in 12.4 shall not be used for bidirectional SCSI commands.

To recover from errors, FCP-3 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 FC-4 Link Service by setting the RETRY bit to one in the PRLI before performing the retransmission of individual IUs (see 6.3.4). An FCP-3 device that has agreed to perform retransmission shall use and accept the REC ELS and SRR FCP FC-4 Link Service as defined by this standard to perform the retransmission.

Even after agreeing to perform retransmission, the initiator FCP_Port may choose to request a recovery abort and the target FCP_Port shall be able to accept and perform the recovery abort.

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-2 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-2, 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-3 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-3 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 SCSI commands:

- a) a bidirectional SCSI command completed with the write data count smaller than FCP_DL and the FCP_RESID_UNDER bit is set to zero;
- b) a bidirectional SCSI 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 SCSI 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 SCSI 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 further 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-2.)

The Exchange Originator may further 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 further 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, no error detection or recovery is required.

12.3 Exchange level recovery using recovery abort

12.3.1 Recovery abort overview

The recovery abort is an FC-FS-2 protocol that recovers FCP_Port resources associated with an Exchange that is being terminated, either because of a task management request or because of an error.

Recovery abort may be used whether or not the FCP devices have agreed to Sequence level error recovery.

All initiator FCP_Ports shall be capable of invoking the recovery abort protocol to terminate failing commands for later retry (see 9.2.2.5). All target FCP_Ports shall be capable of accepting and completing the recovery abort protocol.

12.3.2 Initiator FCP_Port invocation of recovery abort

The initiator FCP_Port terminating the Exchange sends an ABTS sequence to the D_ID of the target FCP_Port of the Exchange being terminated. The ABTS sequence shall have the PARAMETER field set to ABORT EXCHANGE. The ABTS sequence shall be generated using the OX_ID and RX_ID field values of the Exchange to be aborted. FC-FS-2 allows ABTS to be generated by an FCP_Port regardless of whether or not it has Sequence Initiative. Following the transmission of ABTS, any Device_Data Frames received for this Exchange shall be discarded until the BA_ACC with "Last Sequence of Exchange" bit set to one is received from the target FCP_Port.

Recovery abort may not take effect immediately (e.g., if ABTS 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 to the ABTS). The initiator FCP_Port shall be capable of receiving this data and providing BB_Credit in order for the target FCP_Port to send the BA_ACC.

After the processing of a task management function that clears tasks, recovery abort 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, and after R_A_TOV_SEQ_QUAL has elapsed, the initiator FCP_Port shall transmit RRQ.

If a BA_ACC, BA_RJT, LOGO, or PRLO is not received from the target FCP_Port within 2 times R_A_TOV_ELS, second level error recovery as described in 12.5 shall be performed.

12.3.3 Target FCP_Port response to recovery abort

When an ABTS (Abort Exchange) 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 and return LOGO if the N_Port or NL_Port issuing the ABTS is not currently logged in (i.e., no PLOGI);
- b) the target FCP_Port shall return BA_RJT with Last Sequence of Exchange bit set to one if the received ABTS 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 Last Sequence of Exchange bit set to one.

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 a Reinstate Recovery Qualifier (RRQ) ELS request has been received.

If the RX_ID field is FFFFh, target FCP_Ports shall qualify the FQXID of the ABTS 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, they shall be reclaimed at the time the response to the RRQ is sent.

When an RRQ is received at the target FCP_Port, it shall return one of the following responses:

- a) the target FCP_Port shall discard the RRQ and return LOGO if the N_Port or NL_Port issuing the RRQ is not currently logged in (i.e., no PLOGI);
- b) the target FCP_Port shall return LS_RJT with Last Sequence of Exchange bit set to one if the received RRQ contains an RX_ID field value, other than FFFFh, that is unknown to the target FCP_Port. The reason code shall be "Logical error" with a reason code explanation set to "Invalid OX_ID-RX_ID combination"; or
- c) the target FCP_Port shall return ACC with 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 level recovery.

If ULP_TOV times out and the Exchange is not complete, 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 using the recovery abort protocol (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 level recovery.

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 bits in the first frame of the Exchange (see FC-FS-2).

For acknowledged classes of service, if a target FCP_Port detects a Sequence error, it may abort the sequence by sending an ABTS with the PARAMETER field to ABORT SEQUENCE (see FC-FS-2). If a Recovery Qualifier range is returned in the BA_ACC for the ABTS the target FCP_Port shall send a RRQ ELS 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 send a LOGO 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); 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 to perform Sequence level recovery

12.4.1.1 Polling Exchange state with REC

REC is periodically transmitted by the initiator FCP_Port to poll each outstanding Exchange to determine if a SCSI task 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 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, 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

If an Exchange Originator receiving an acknowledged service Sequence detects a Sequence error, it shall send an ACK Frame with the F_CTL field Abort Sequence Condition bits set to "Abort Sequence, Perform ABTS" before issuing the REC (see FC-FS-2). The REC 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 is an LS_RJT with a reason code of "command not supported", the initiator FCP_Port shall assume the target FCP_Port is an FCP device not supporting error detection using REC. The initiator FCP_Port shall perform recovery using recovery abort as documented in 12.3.

If an ACC, LS_RJT, LOGO, or PRLO is not received from the target FCP_Port within 2 times R_A_TOV_{ELS}, second level error recovery as described in 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 level recovery.

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

If the ACC for the REC 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), the command is in process and no recovery is needed at this time. At a minimum interval of REC_TOV, the REC 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 level recovery.

If the ACC for an REC 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, the initiator FCP_Port shall issue an SRR in a new Exchange to request retransmission of the FCP_XFER_RDY IU. The target FCP_Port shall first transmit the ACC for the SRR 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 level recovery.

An error in transmitting an FCP_RSP IU is detected if:

- a) the 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 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 issue an SRR FC-4 Link Service frame in a new Exchange to request retransmission of the FCP_RSP IU. The target FCP_Port shall first transmit the ACC for the SRR, 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 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 for the SRR 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 retry the FCP_XFER_RDY IU or, if the actual data transfer length for the command was zero, retry 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.

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, FCP_RSP IU information shall be:

- a) discarded $RR_TOV_{SEQ_INIT}$ after the FCP_RSP IU was transmitted to the initiator FCP_Port; or,
- b) discarded after a new Exchange with the same OX_ID and S_ID field values and task retry identifier is received.

If retransmission is not enabled between the initiator FCP_Port and target FCP_Port, 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 10.4.10).

If task retry identification has been agreed to by both the initiator FCP_Port and target FCP_Port, 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 through 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 level recovery.

If the ACC for an REC 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 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 send a SRR FC-4 Link Service frame 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 ACC for the SRR, the target FCP_Port transmits an FCP_XFER_RDY IU in a new Sequence with the Relative Offset parameter specified by 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 through 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 level recovery.

If the ACC for the REC 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 send a SRR FC-4 Link Service frame 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 to that of the next data requested. If the initiator FCP_Port is unable to determine the Relative Offset of the next data requested, the initiator FCP_Port shall set the RELATIVE OFFSET field to zero. The target FCP_Port shall first transmit the ACC for the SRR, then shall retransmit the requested data specified by the SRR in a new Sequence, and then complete the Exchange in the normal manner, including transmitting or retransmitting the FCP_RSP IU. If the target FCP_Port responds to the SRR with an FCP_RJT and an FCP_RSP IU has not yet been sent or is again requested, the device server shall send 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 REC 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 through 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 level recovery.

This recovery procedure is used by target devices 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 REC 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 reply to the REC from the target FCP_Port shall be a LS_RJT with the reason code of "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 is received, the REC reply shall be an 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, send another REC. 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, 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 through 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, the Sequence Initiator shall send an ABTS for the Sequence. After R_A_TOV times out, an RRQ shall be sent by the Sequence Initiator.

12.4.2.2 Missing ACK

FC-FS-2 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, the target FCP_Port shall abort the sequence by sending an ABTS request with the PARAMETER field set to ABORT SEQUENCE. If a Recovery Qualifier range is returned in the BA_ACC for the ABTS the target FCP_Port shall send an RRQ 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-2.

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-2; 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 send 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.

Recovery abort shall be invoked 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}, the initiator FCP_Port may send the ABTS again, attempt other retry operations allowed by FC-FS-2, or explicitly logout the target FCP_Port. If those retry

operations attempted are unsuccessful, the initiator FCP_Port shall explicitly logout (i.e., use FC-LS Logout, LOGO) the target FCP_Port. All outstanding Exchanges with that target FCP_Port are terminated at the initiator FCP_Port.

12.5.2 REC error recovery

If a response to an REC is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port shall:

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

If the response to the second REC is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port should send an ABTS(Exchange) for the REC followed by an RRQ if a BA_ACC is received for the ABTS;

Other retry mechanisms after the second REC fails are optional and, if implemented, shall comply with FC-FS-2.

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

See figure C.26 through figure C.29.

12.5.3 SRR error recovery

If a response to an SRR is not received within 2 times $R_A_TOV_{ELS}$, the initiator FCP_Port shall send an ABTS(Exchange) for the SRR followed by an RRQ if a BA_ACC is received for the ABTS. The initiator FCP_Port shall then perform an ABTS(Exchange) for the original Exchange.

See figure C.30 through figure C.33.

12.6 Responses to FCP type frames before PLOGI or PRLI

If a target FCP_Port receives an FCP_CMND IU from an FCP_Port that is not successfully logged on to the target FCP_Port using either an implicit or explicit login (i.e. PLOGI), it shall discard the FCP_CMND IU and, in a new Exchange, send a LOGO 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 sending the LOGO.

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, it shall discard the FCP_CMND IU and send PRLO 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, the FCP device shall discard and ignore the content of the frame. If login is not completed, the FCP device may transmit a LOGO request to the source of the unexpected frame. If login is completed, but Process Login is not completed, the FCP device may transmit a PRLO request to the source of the unexpected frame.

Annex A

(normative)

FCP-3 mapping to SAM-3 (Fibre Channel Protocol Service mapping to SCSI Architectural Model - 3)

A.1 Definition of procedure terms

FCP-3 services are provided to the application client by the initiator FCP_Port to request and manage tasks as described by the SAM-3 standard. SAM-3 further defines how the target FCP_Port enables the device server to receive and process the tasks addressed to a logical unit. 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 objects and identifiers used in this standard to the equivalent remote procedure call terms and definitions used in the SCSI Architecture Model-3 standard.

Table A.1 - FCP-3 procedure terms mapped to terms from SAM-3 standard

FCP-3 standard procedure terms	Equivalent SAM-3 terms
address identifier of initiator FCP_Port	initiator port identifier
address identifier of target FCP_Port	target port identifier
Port_Name of initiator FCP_Port	initiator port name
Port_Name of target FCP_Port	target port name
fully qualified exchange identifier	I_T_L_Q Nexus
fully qualified exchange identifier + logical unit number	I_T_L_Q Nexus

See table A.2 for the definitions of the terms used by this standard and the equivalent SAM-3 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:

- the originating object of the term;
- the object that is the final destination of the term; and
- 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-3	SAM-3	DS → targ → init
data buffer size	SAM-3	SAM-3	AC → init
command descriptor block	SAM-3	SAM-3/cmd (note 1)	AC → init → targ → DS
Data-In Buffer	SAM-3	cmd (note 2)	DS → targ → init → AC
Data-Out Buffer	SAM-3	cmd (note 2)	AC → init → targ → DS
device server buffer	SAM-3	cmd (note 2)	DS → targ → init
initiator SCSI ID	SAM-3	this standard	DS → targ or TM → targ
link control function	this standard	this standard	AC → init → targ
logical unit number	SAM-3	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ → init
request byte count	SAM-3	SAM-3	DS → targ
service response	SAM-3	this standard	DS → targ → init → AC or targ → DS
status	SAM-3	SAM-3	DS → targ → init → AC
task tag	SAM-3	this standard	AC → init → targ → DS or AC → init → targ → TM or DS → targ → init
target port identifier	SAM-3	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-3	this standard	AC → init → targ → DS
Key: AC = application client, cmd = SCSI command standards, DS = device server, init = initiator, SAM-3 = SCSI Architecture Model-3 standard, TM = task manager, targ = target			
Notes 1) The portions not defined in SAM-3 are defined in the SCSI command standards (e.g. SPC-3). 2) Parameter lists are defined within one of the SCSI command standards (e.g., SPC-3). 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 SCSI FCP read operation

A typical SCSI FCP read operation with a single data IU is shown in table B.1.

Table B.1 - FCP read operation, 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 SCSI FCP write operation

A typical SCSI FCP write operation with three data IUs and using FCP_XFER_RDY is shown in table B.2.

Table B.2 - FCP write operation, 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 SCSI FCP operation with no data transfer or with check condition

A typical SCSI FCP operation terminating without data transfer, either because of an error or because the SCSI command does not require any data transfer, is shown in table B.3.

Table B.3 - FCP operation 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 SCSI FCP read operation with multiple FCP_DATA IUs

A typical SCSI read operation with multiple FCP_DATA IUs is shown in table B.4.

Table B.4 - FCP read operation, 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 SCSI FCP write operation with FCP_XFR_RDY disabled

A typical SCSI write operation 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 - FCP write operation 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 SCSI FCP bidirectional command with write before read

A typical SCSI FCP 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 - FCP 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 buffer]
	← 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 SCSI FCP bidirectional command with read before write

A typical SCSI FCP 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 - FCP 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 buffer]
	← 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 SCSI FCP bidirectional command, write first, write FCP_XFER_RDY disabled

A SCSI FCP 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 - FCP bidirectional command, 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 SCSI FCP bidirectional command with intermixed writes and reads

A SCSI FCP 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 - FCP 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 SCSI linked commands

A SCSI WRITE command linked after a SCSI READ command is shown in table B.10. The WRITE command is using the FCP_XFER_RDY IU. INTERMEDIATE Status in the FCP_RSP, together with the link control bits present in the CDB of the FCP_CMND indicate that the second operation is linked to the first.

Table B.10 - FCP linked commands, example

Initiator FCP_Port function	IU	Target FCP_Port function
Command request (READ)	T1, FCP_CMND →	
		[Prepare data transfer]
	← I3, FCP_DATA	Data-In action
		[Prepare response message]
	← I5, FCP_RSP	Response (INTERMEDIATE or INTERMEDIATE CONDITION MET status)
[Perform command linking]		
Command request (WRITE)	T3, FCP_CMND →	
		[Prepare data transfer buffer]
	← I1, FCP_XFER_RDY	Data delivery request
Data-Out Action	T6, FCP_DATA →	
		[Prepare response message]
	← I4, FCP_RSP	Response
[Indicate command completion]		

B.1.12 SCSI WRITE command with confirmed completion

A SCSI WRITE command with confirmed completion is shown in table B.11.

Table B.11 - FCP write command with confirmed completion

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.1.13 SCSI FCP task management function

An example of a SCSI Task Management function is shown in table B.12. Additional link services may be required in some cases to complete the activities initiated by the Task Management function.

Table B.12 - FCP 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.2 FCP write example, frame level

A chart of the Sequences and frames typically transmitted to perform an FCP write is shown in figure B.1. All frames of a Sequence have a frame level FC-FS-2 acknowledgment returned automatically as part of the link control.

Figure B.1 - Example of class 2 FCP write operation

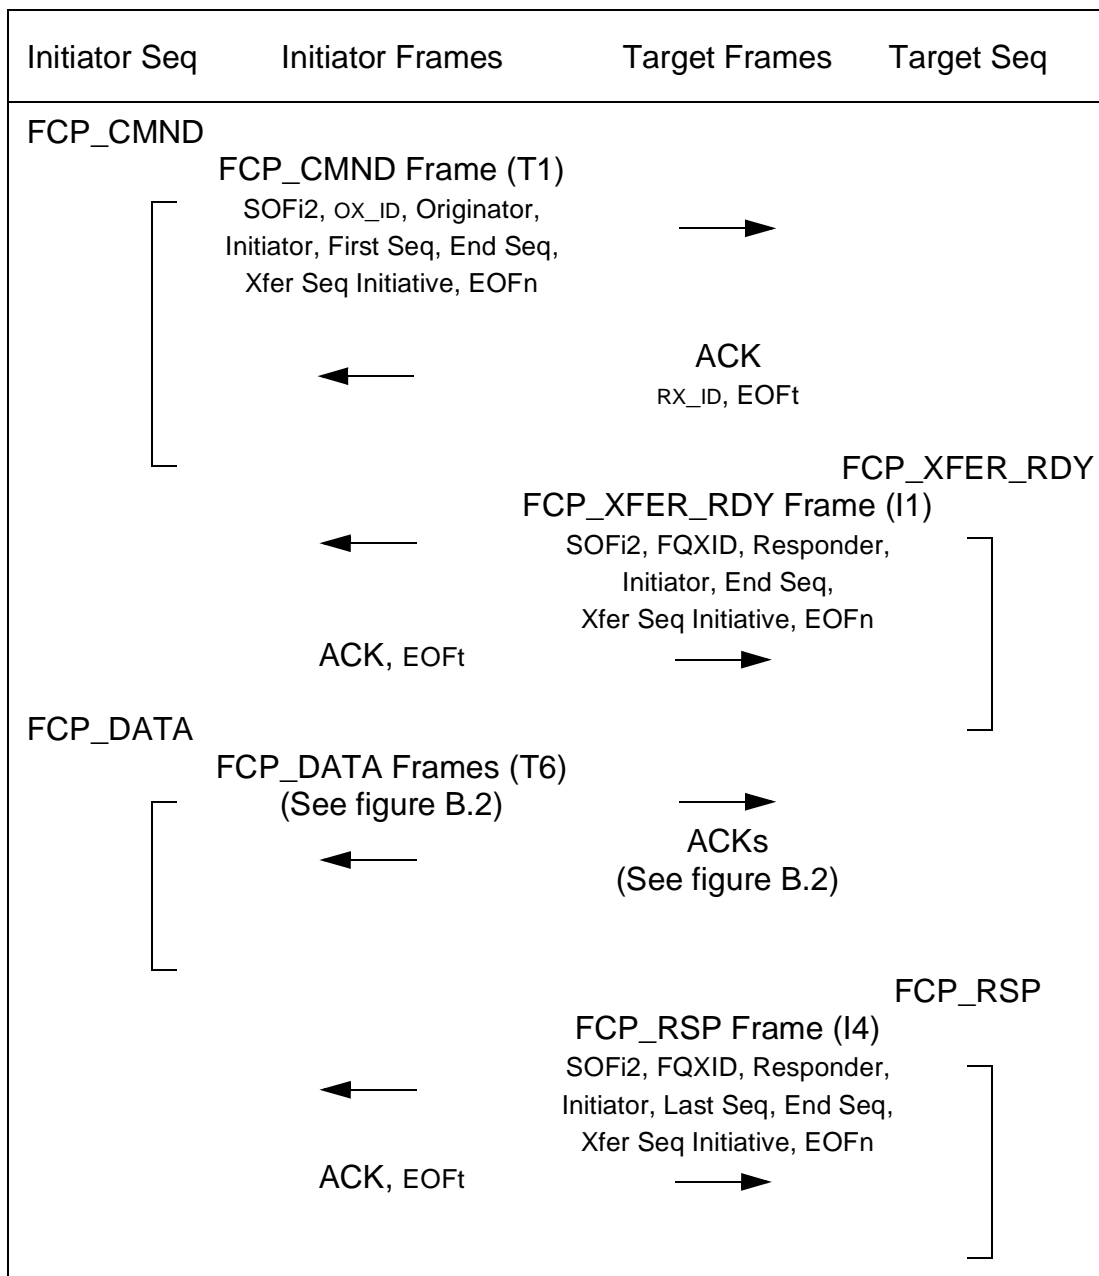
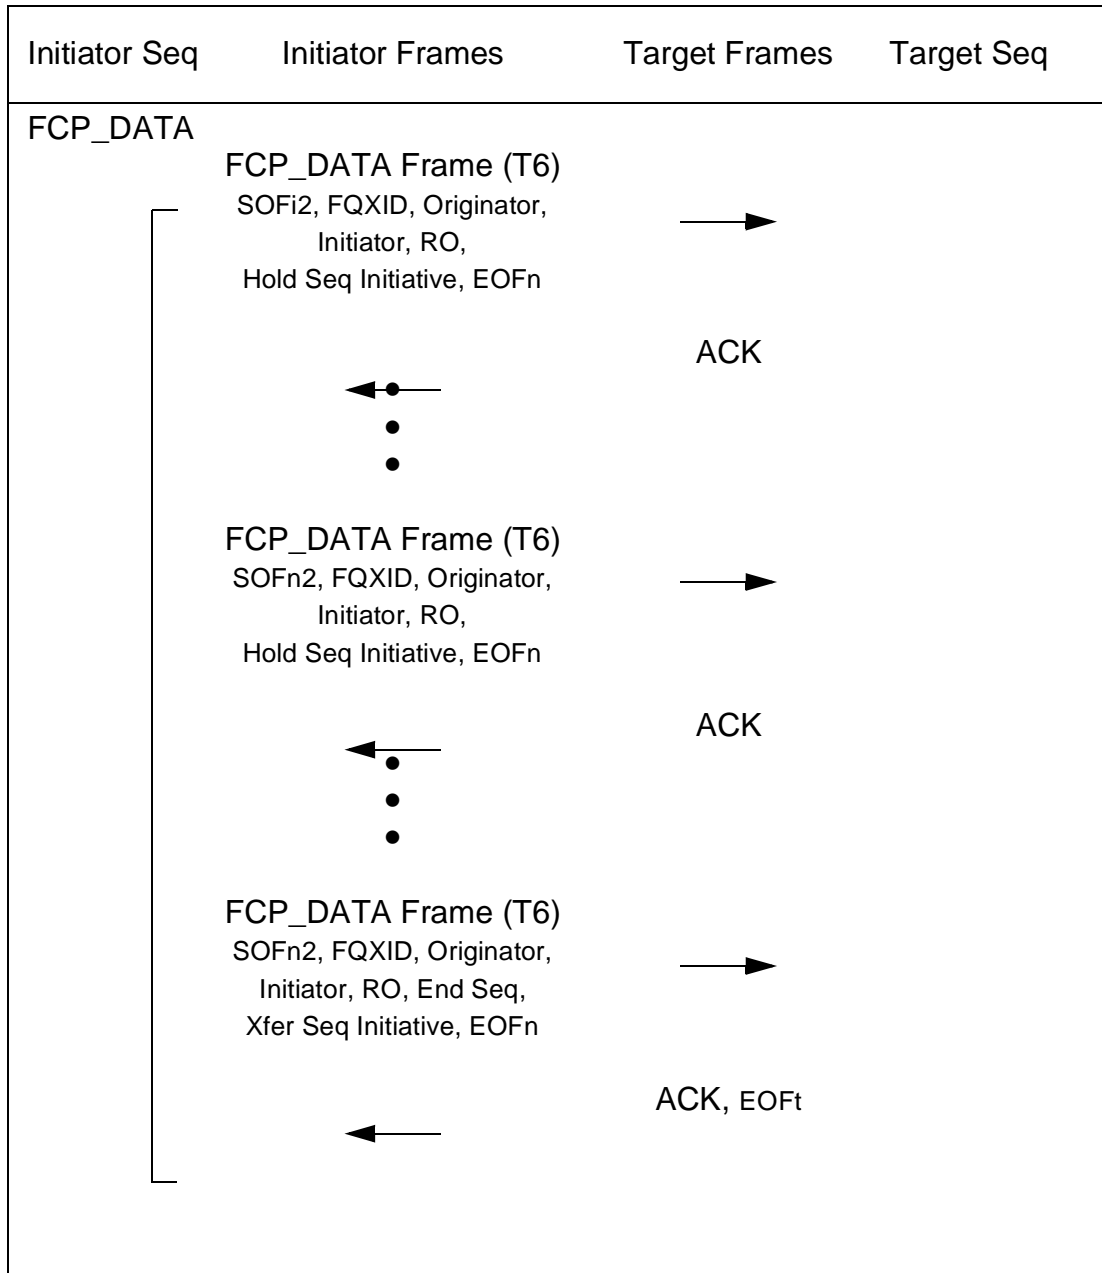


Figure B.2 - Example of class 2 FCP_DATA write



B.3 FCP read example, frame level

A chart of the Sequences typically transmitted to perform an FCP read is shown in figure B.3. All frames of a Sequence have a frame level FC-FS-2 acknowledgment returned automatically as part of the link control.

Figure B.3 - Example of class 2 FCP read operation

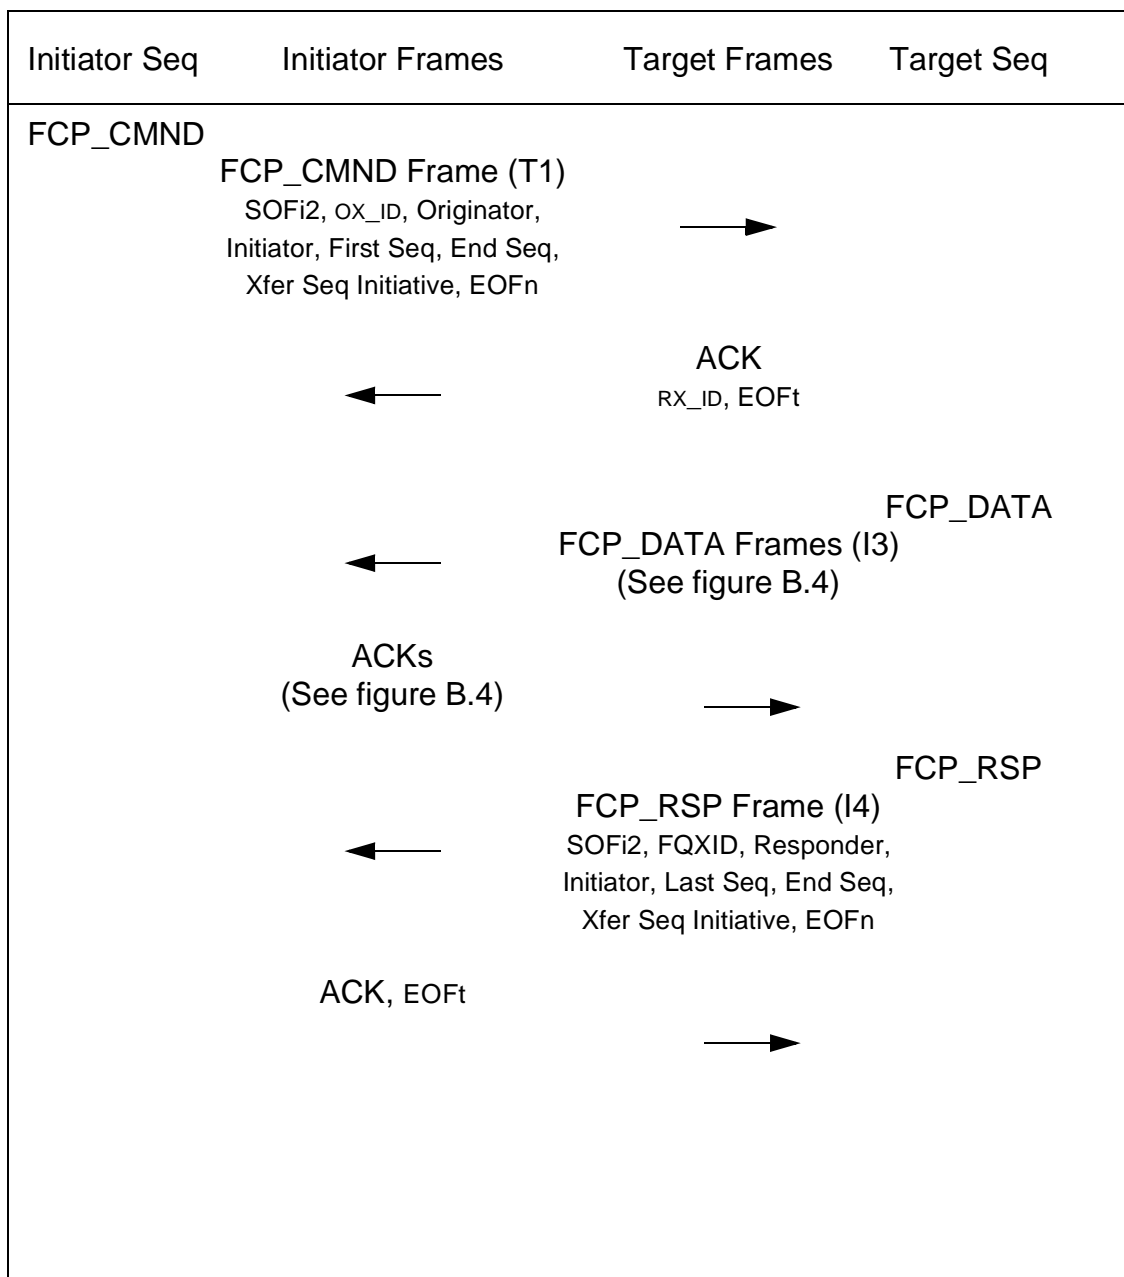
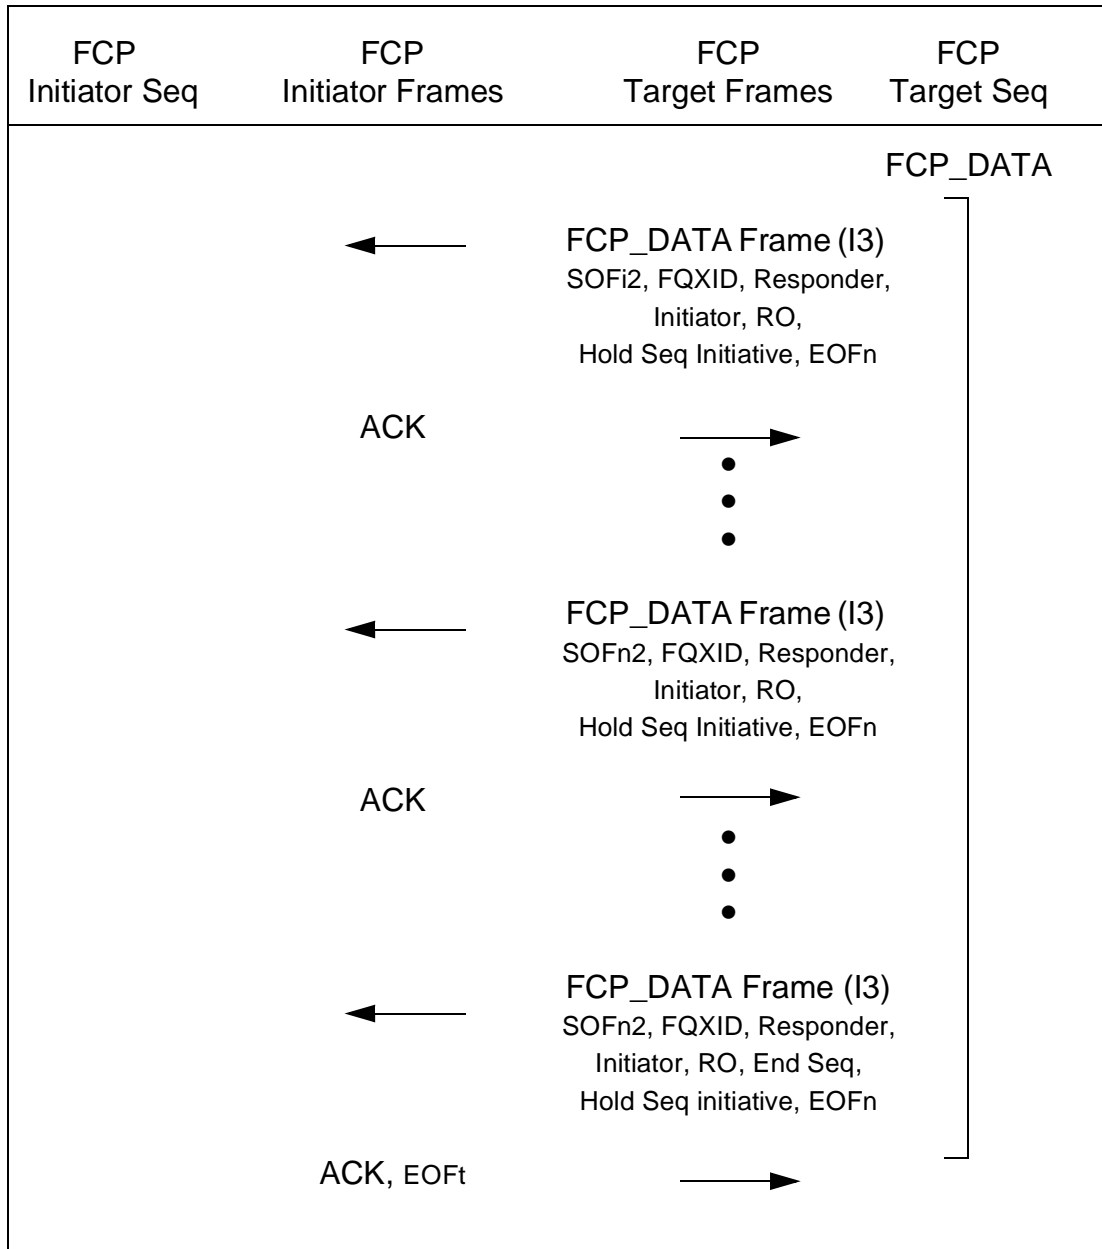


Figure B.4 - Example of class 2 FCP_DATA read

Annex C

(informative)

Error detection and recovery action examples

C.1 Introduction

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-2 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 Link Service, or ELS
	IU or ELS received is processed to transmit IU, FC-4 Link Service, 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.

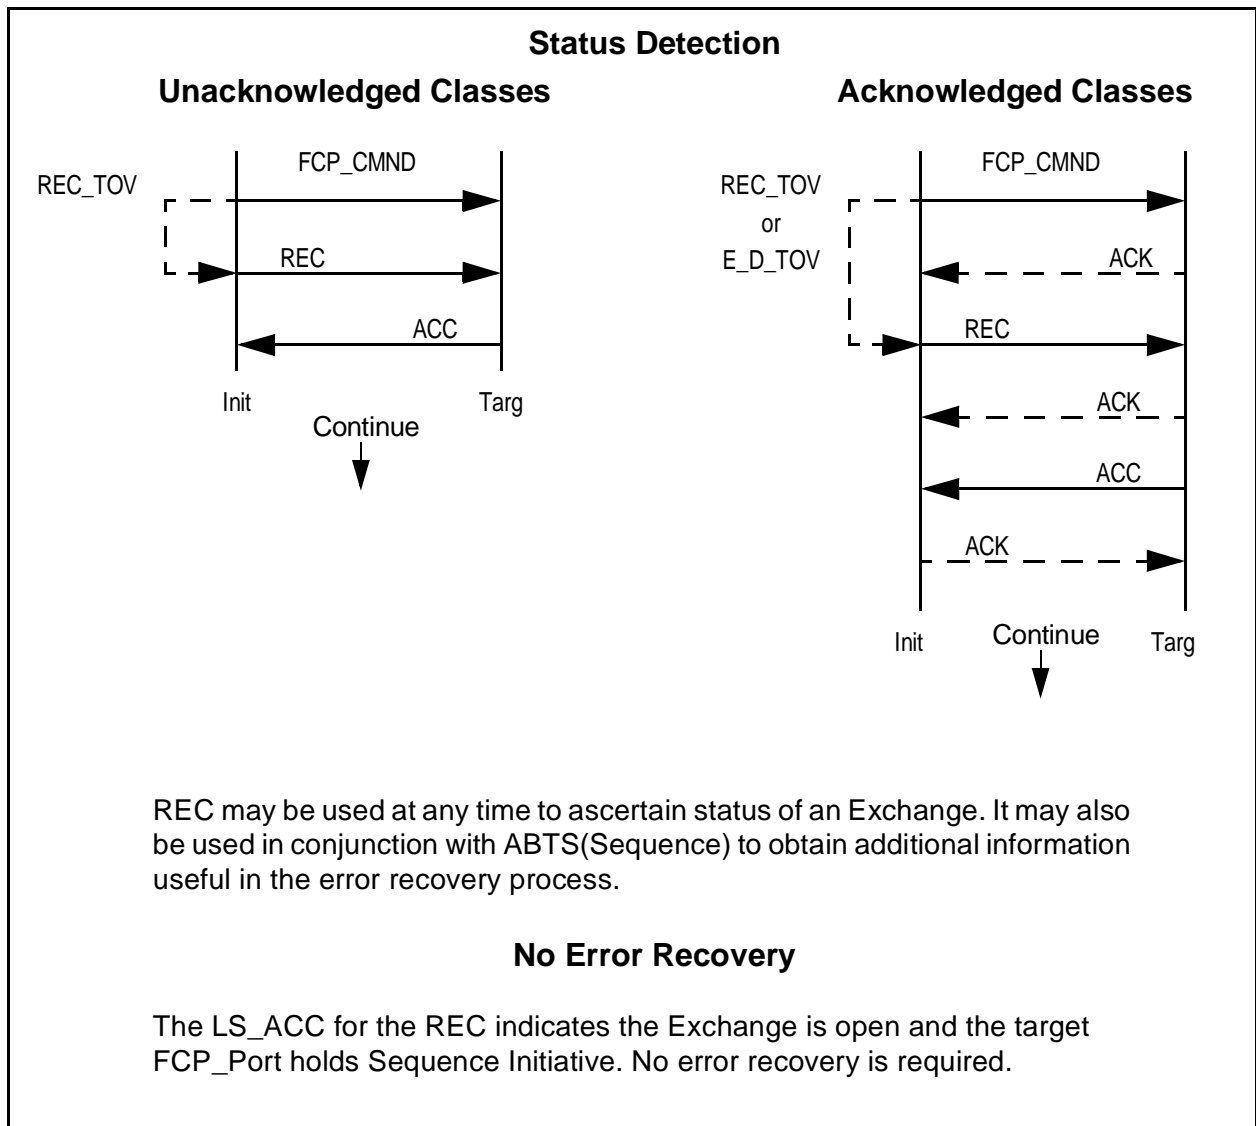
Figure C.1 - Lengthy FCP_CMND or Lost ACK

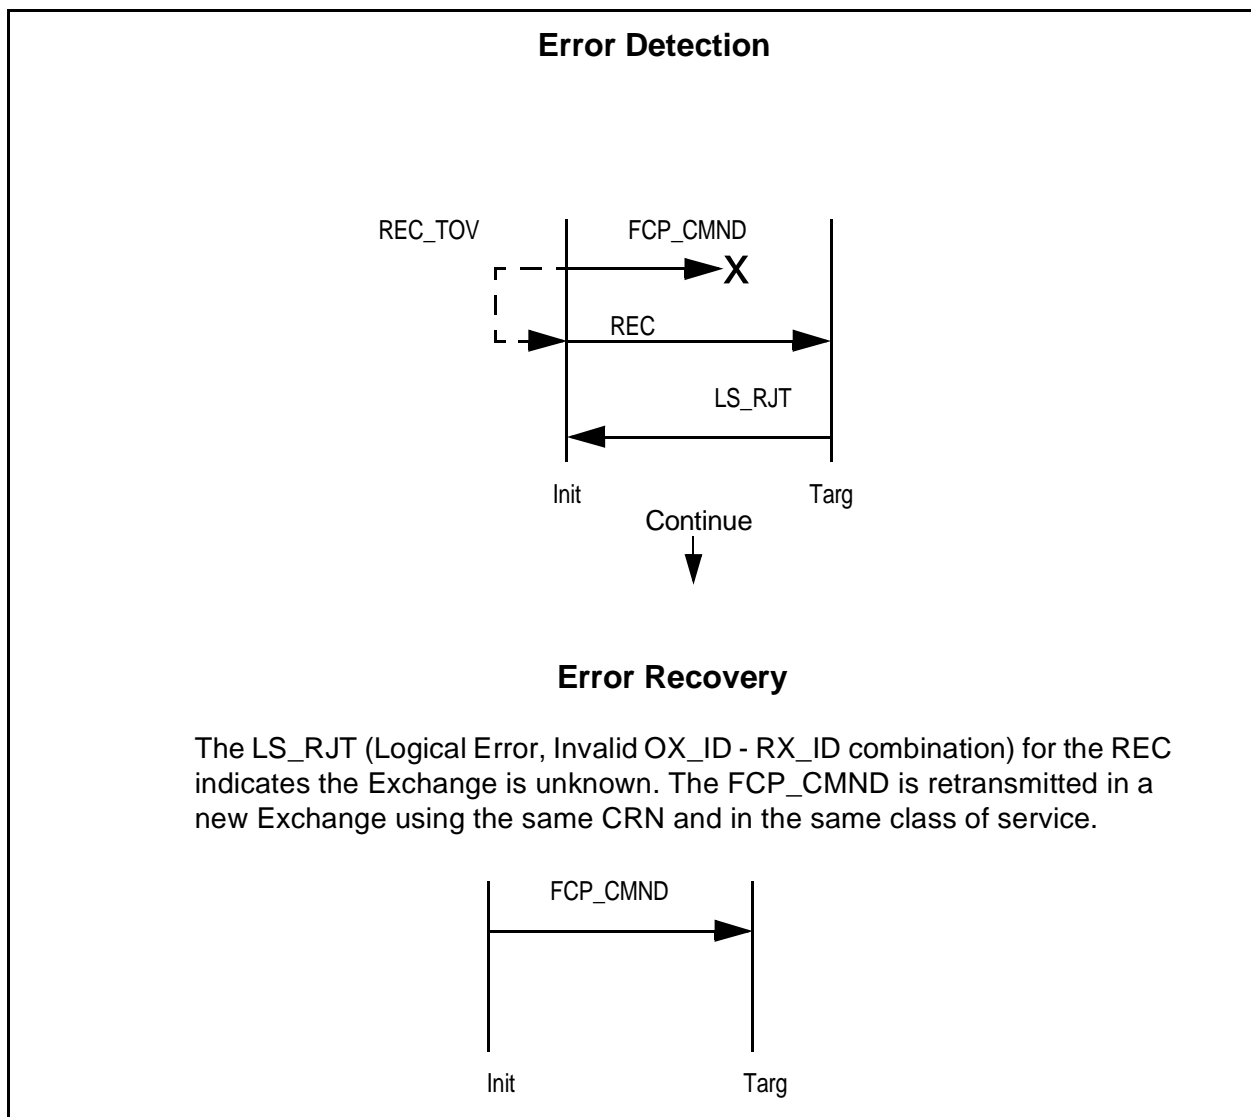
Figure C.2 - FCP_CMND Lost, Unacknowledged Classes

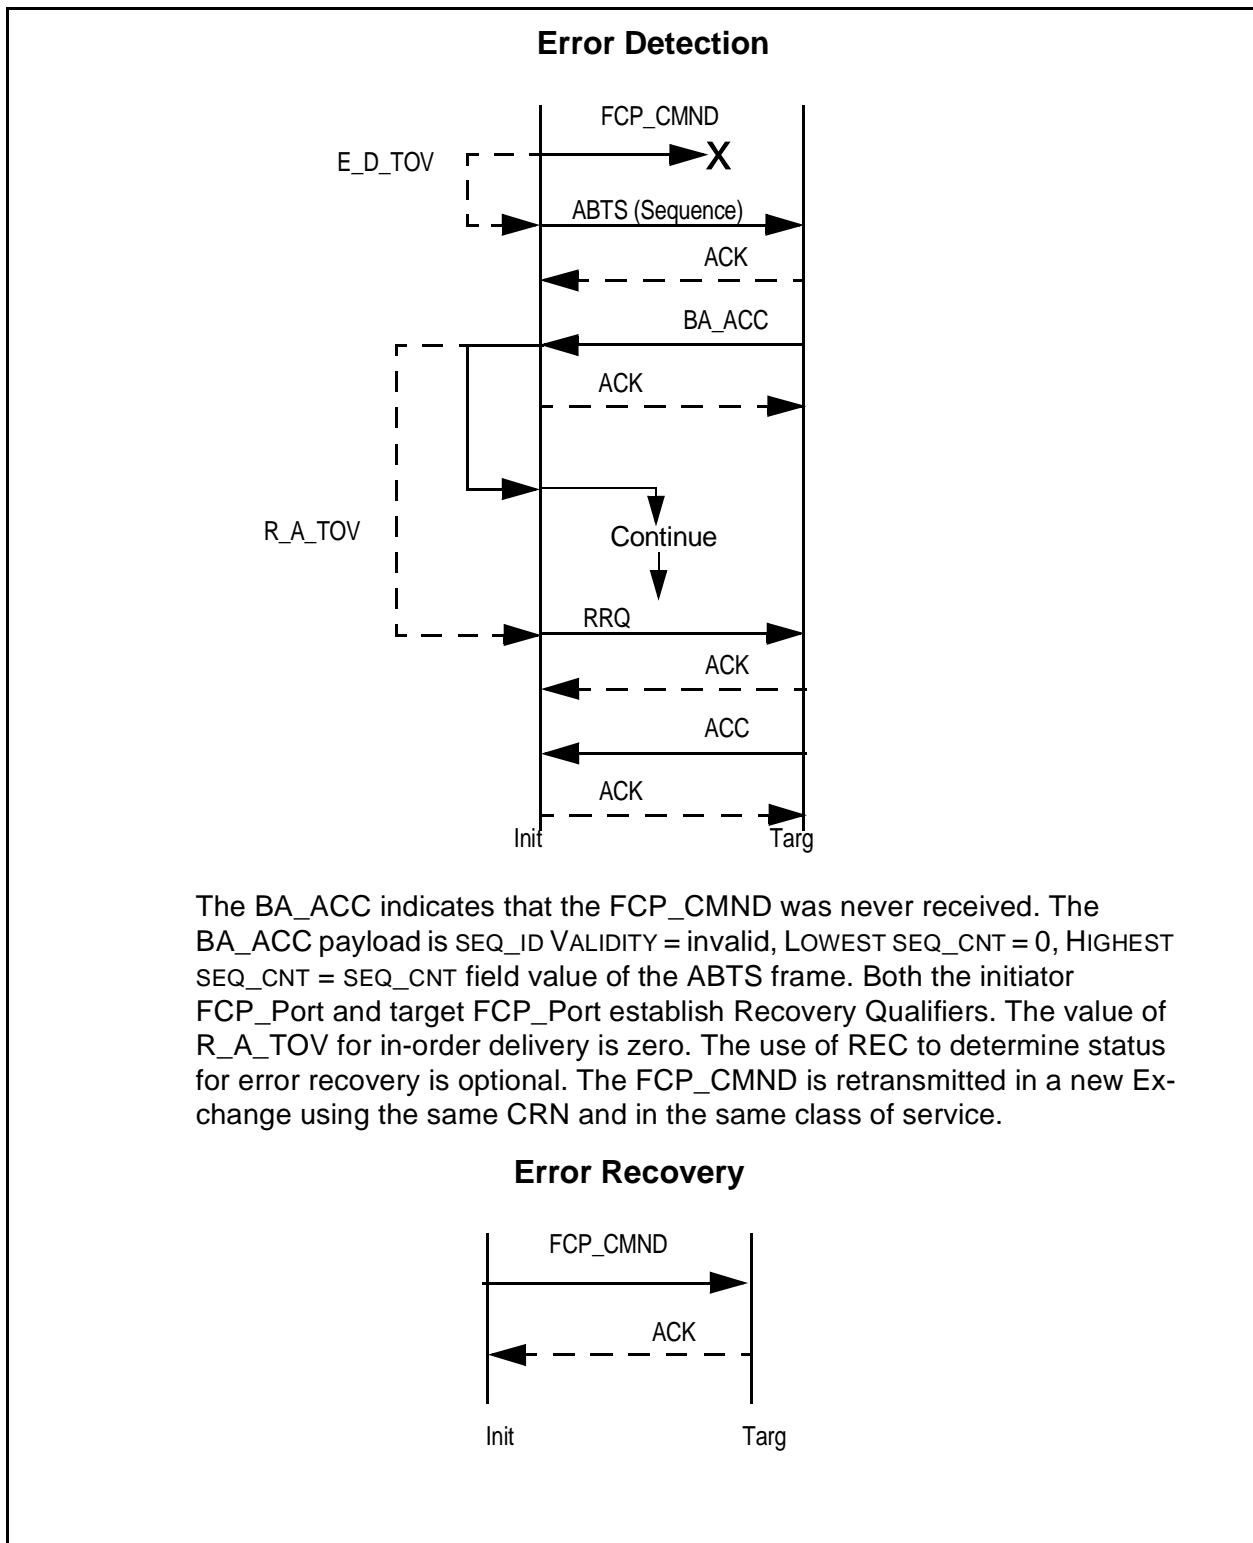
Figure C.3 - FCP_CMND Lost, Acknowledged Classes

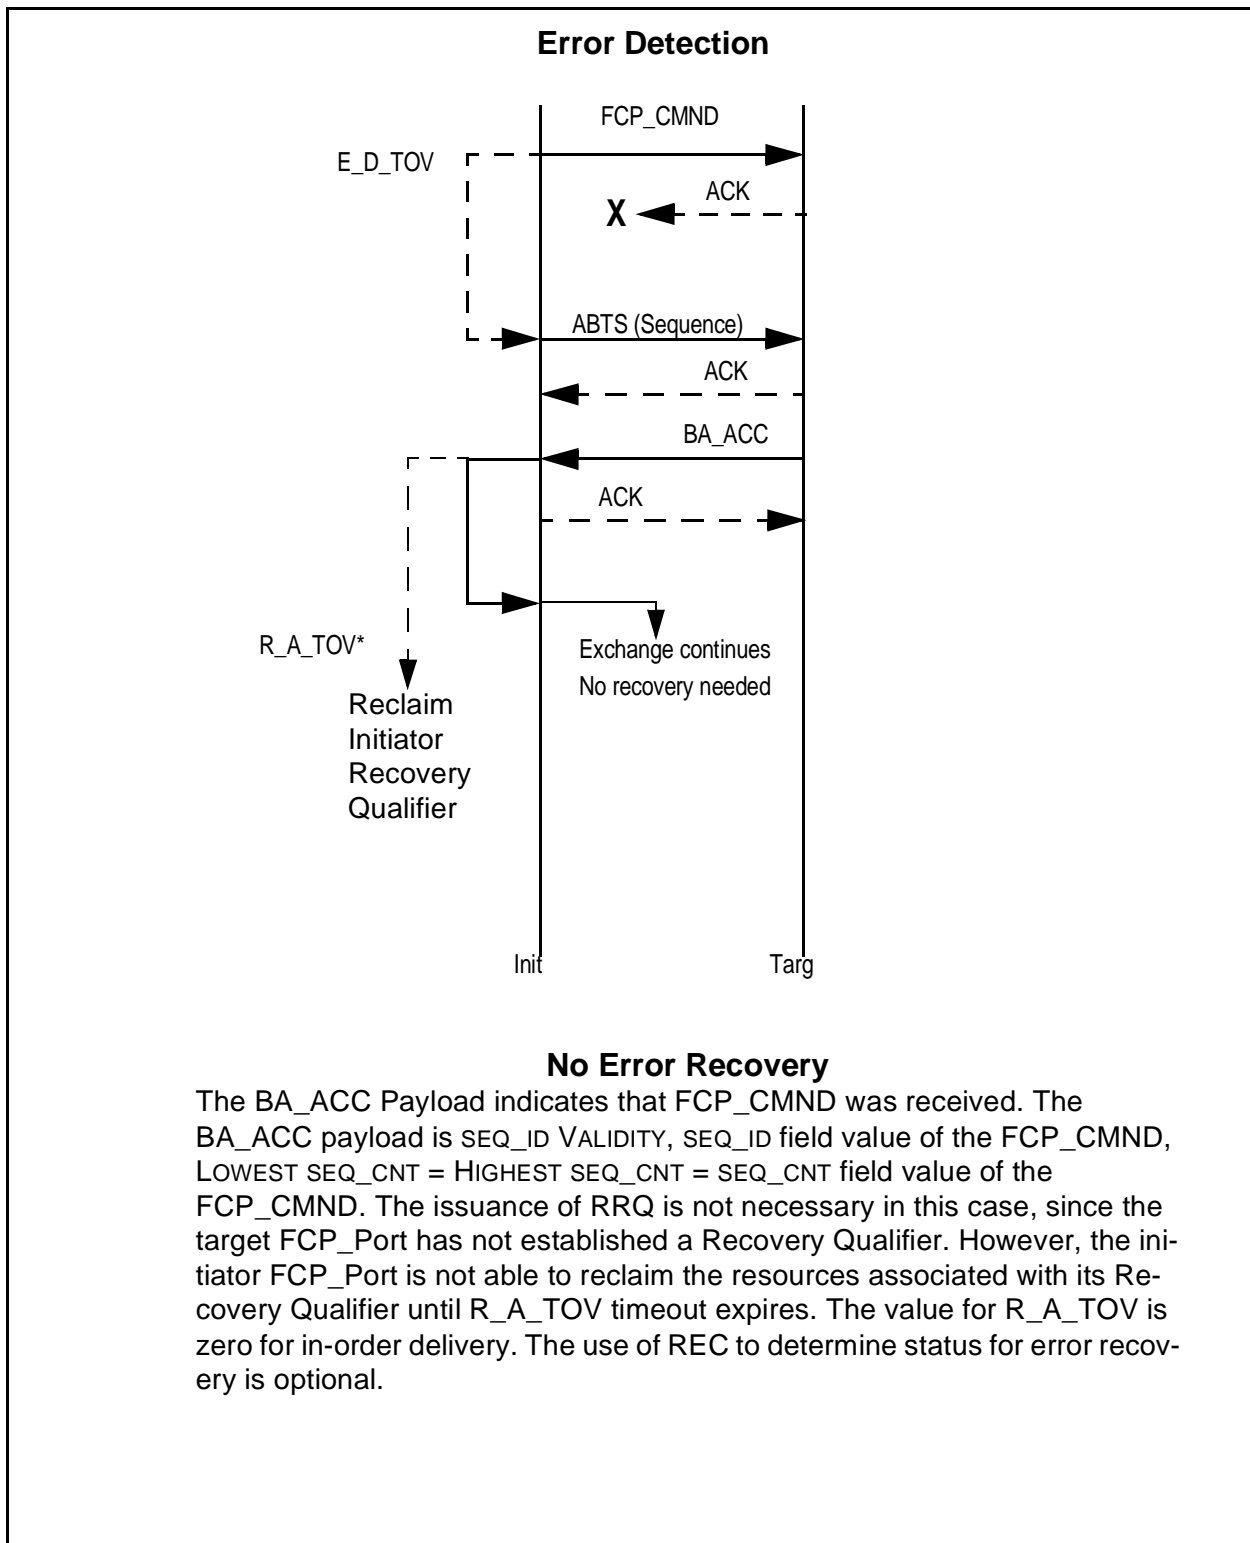
Figure C.4 - FCP_CMND Acknowledgement Lost, Acknowledged Classes

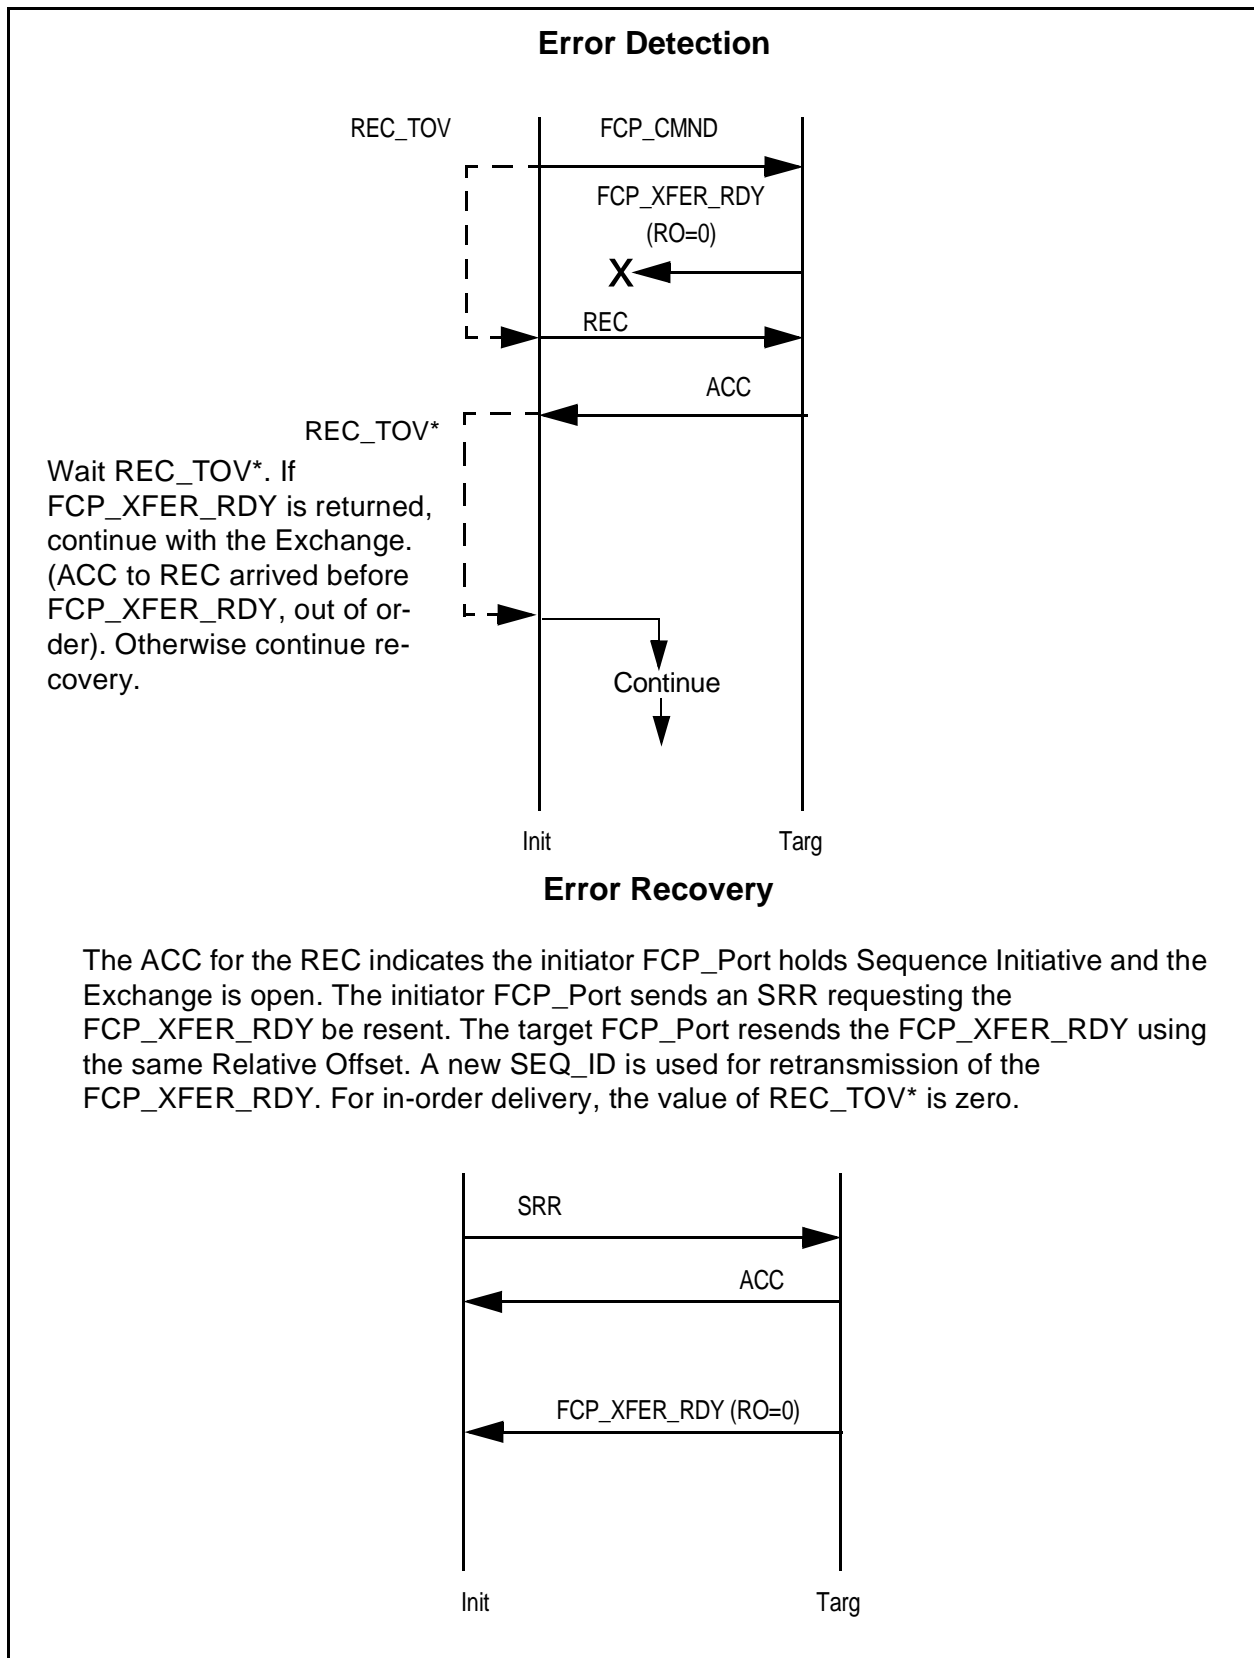
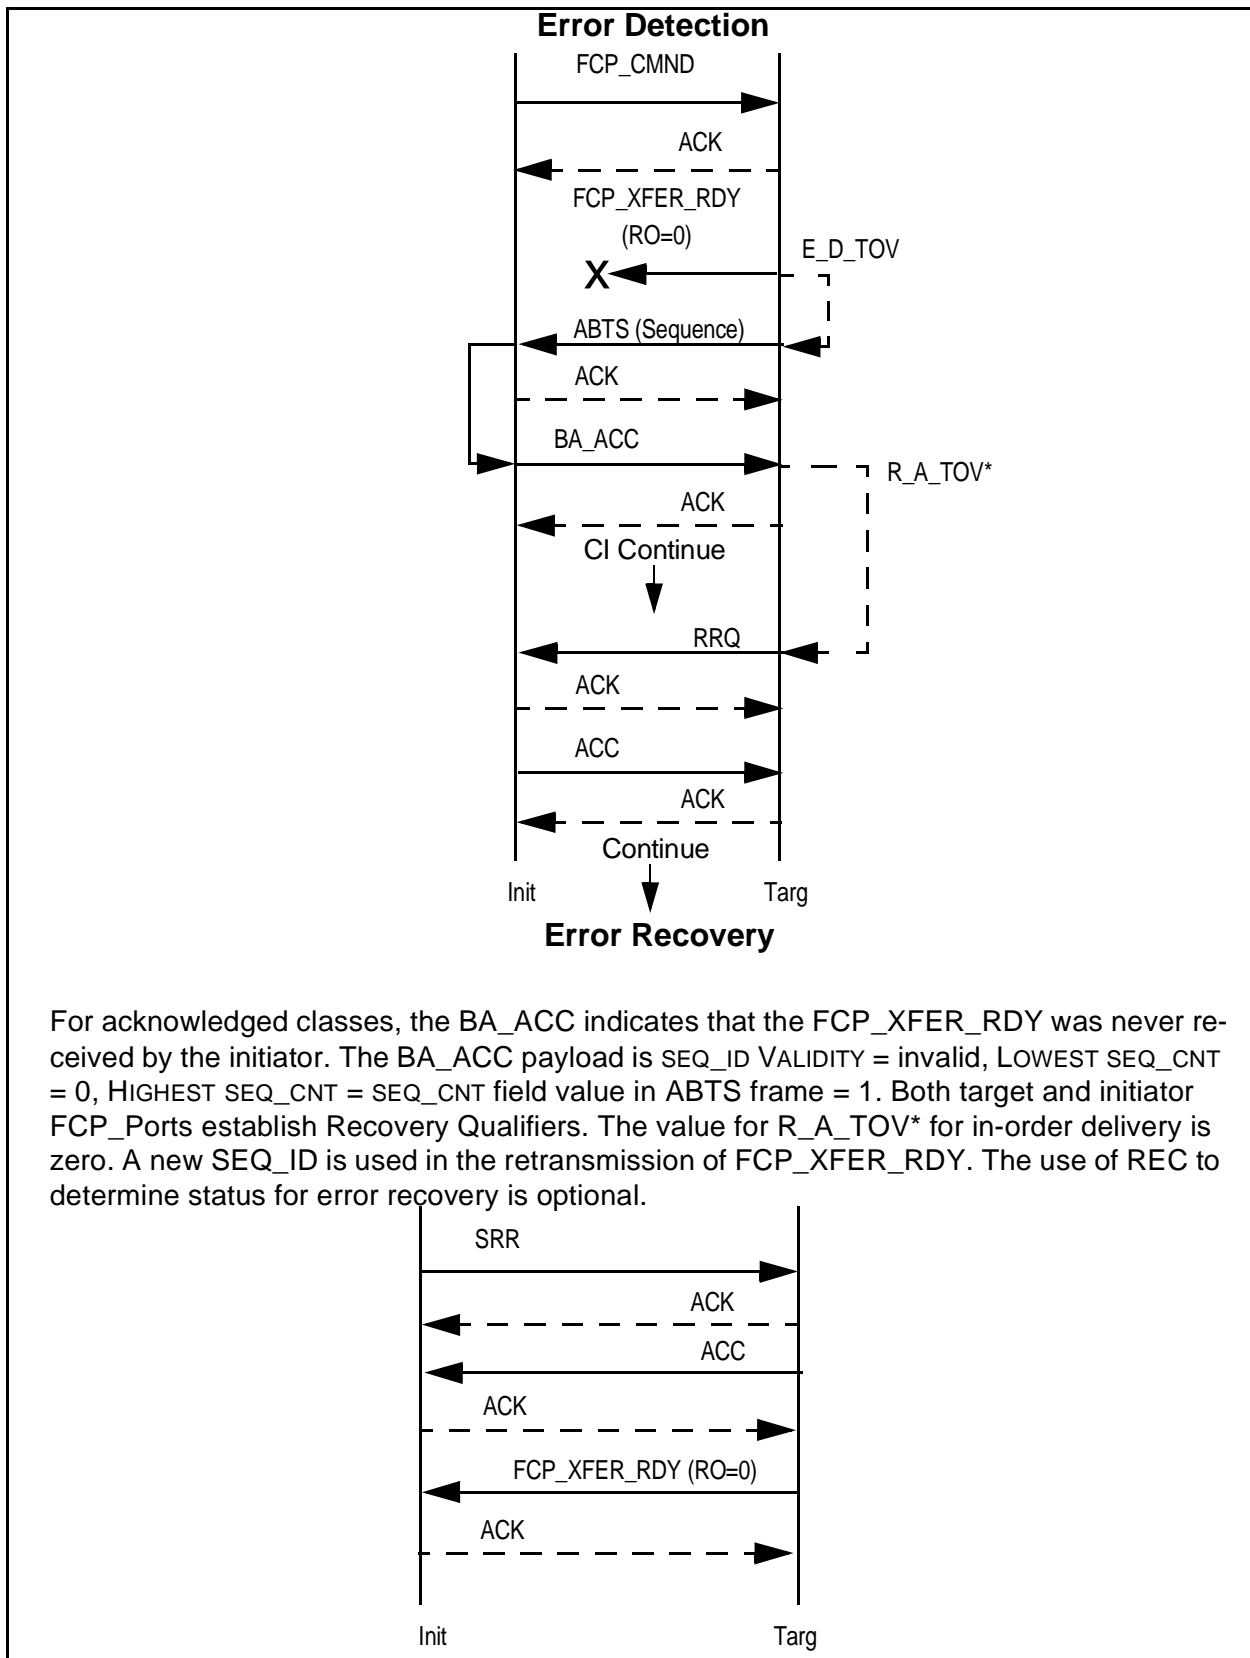
Figure C.5 - FCP_XFER_RDY Lost, Unacknowledged Classes

Figure C.6 - FCP_XFER_RDY Lost, Acknowledged Classes

For acknowledged classes, the BA_ACC indicates that the FCP_XFER_RDY was never received by the initiator. 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 REC to determine status for error recovery is optional.

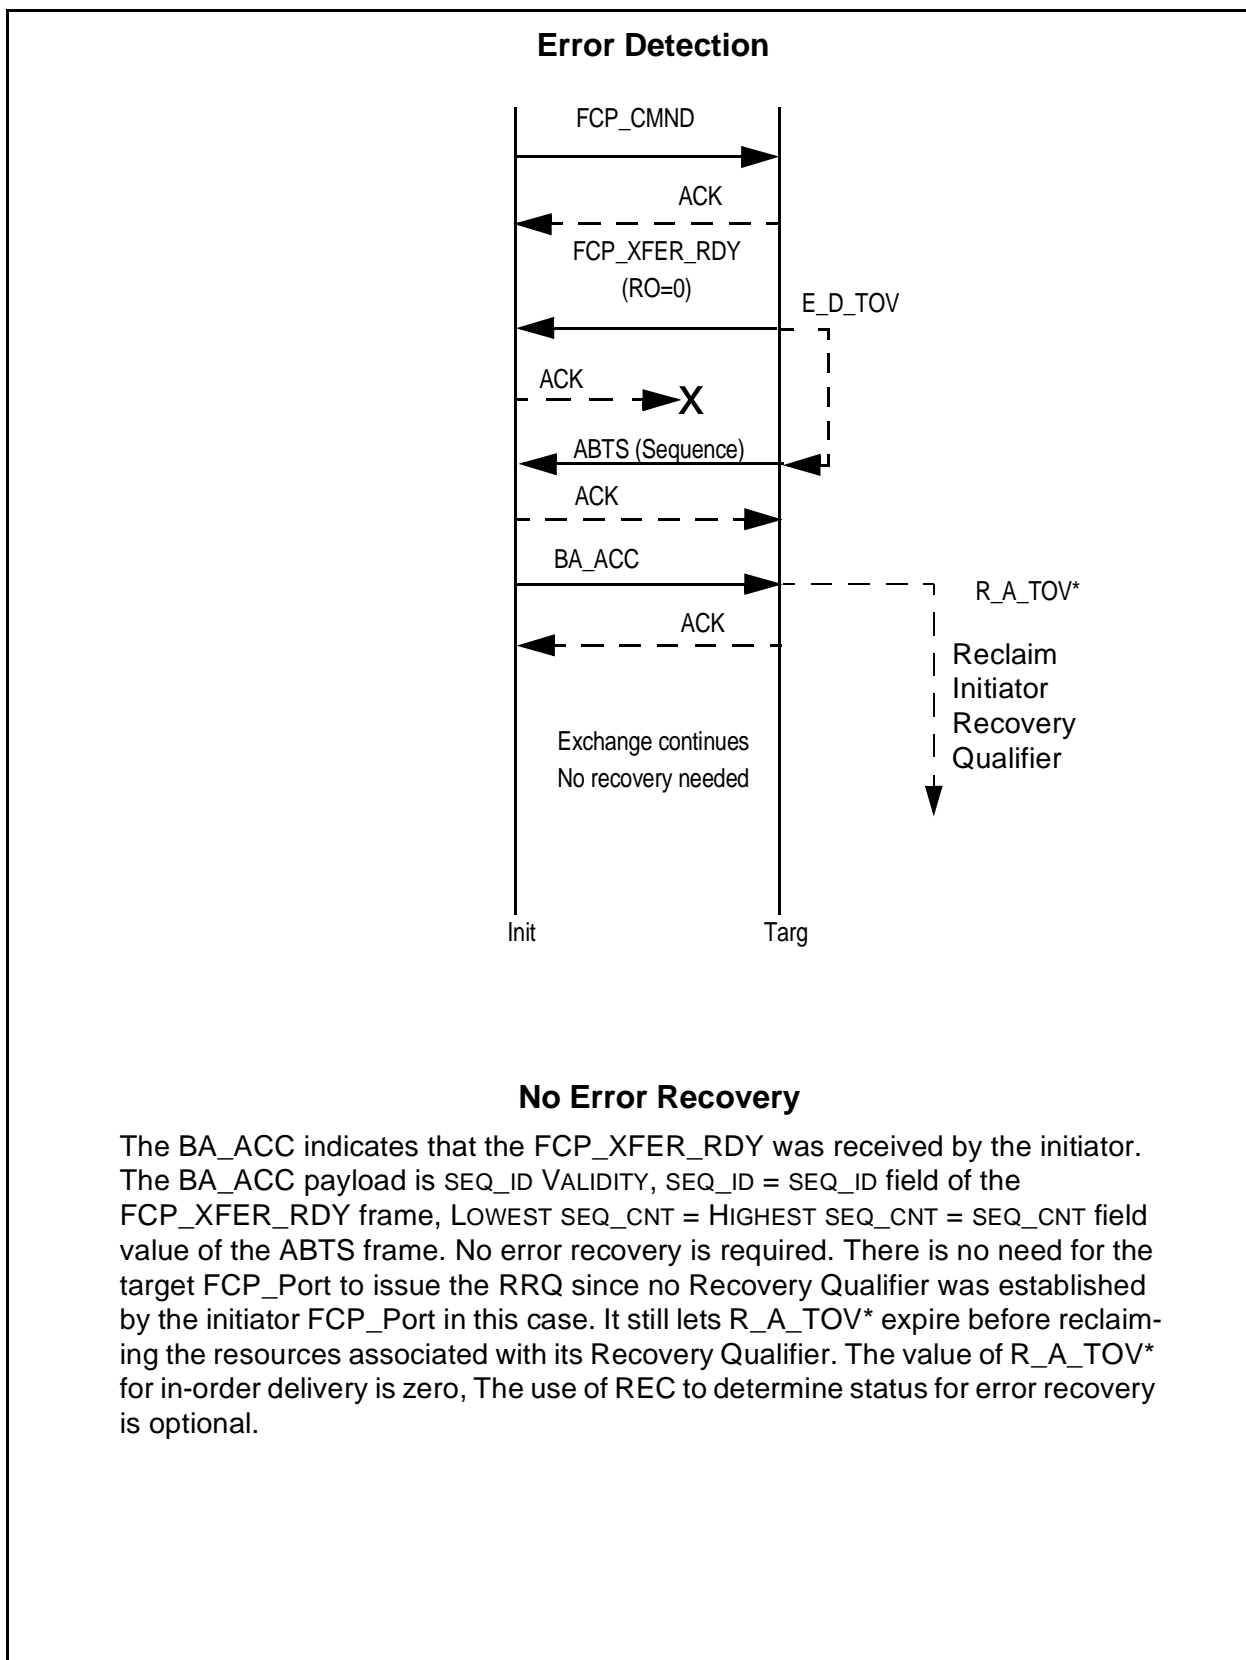
Figure C.7 - FCP_XFER_RDY Received, ACK Lost, Acknowledged Classes

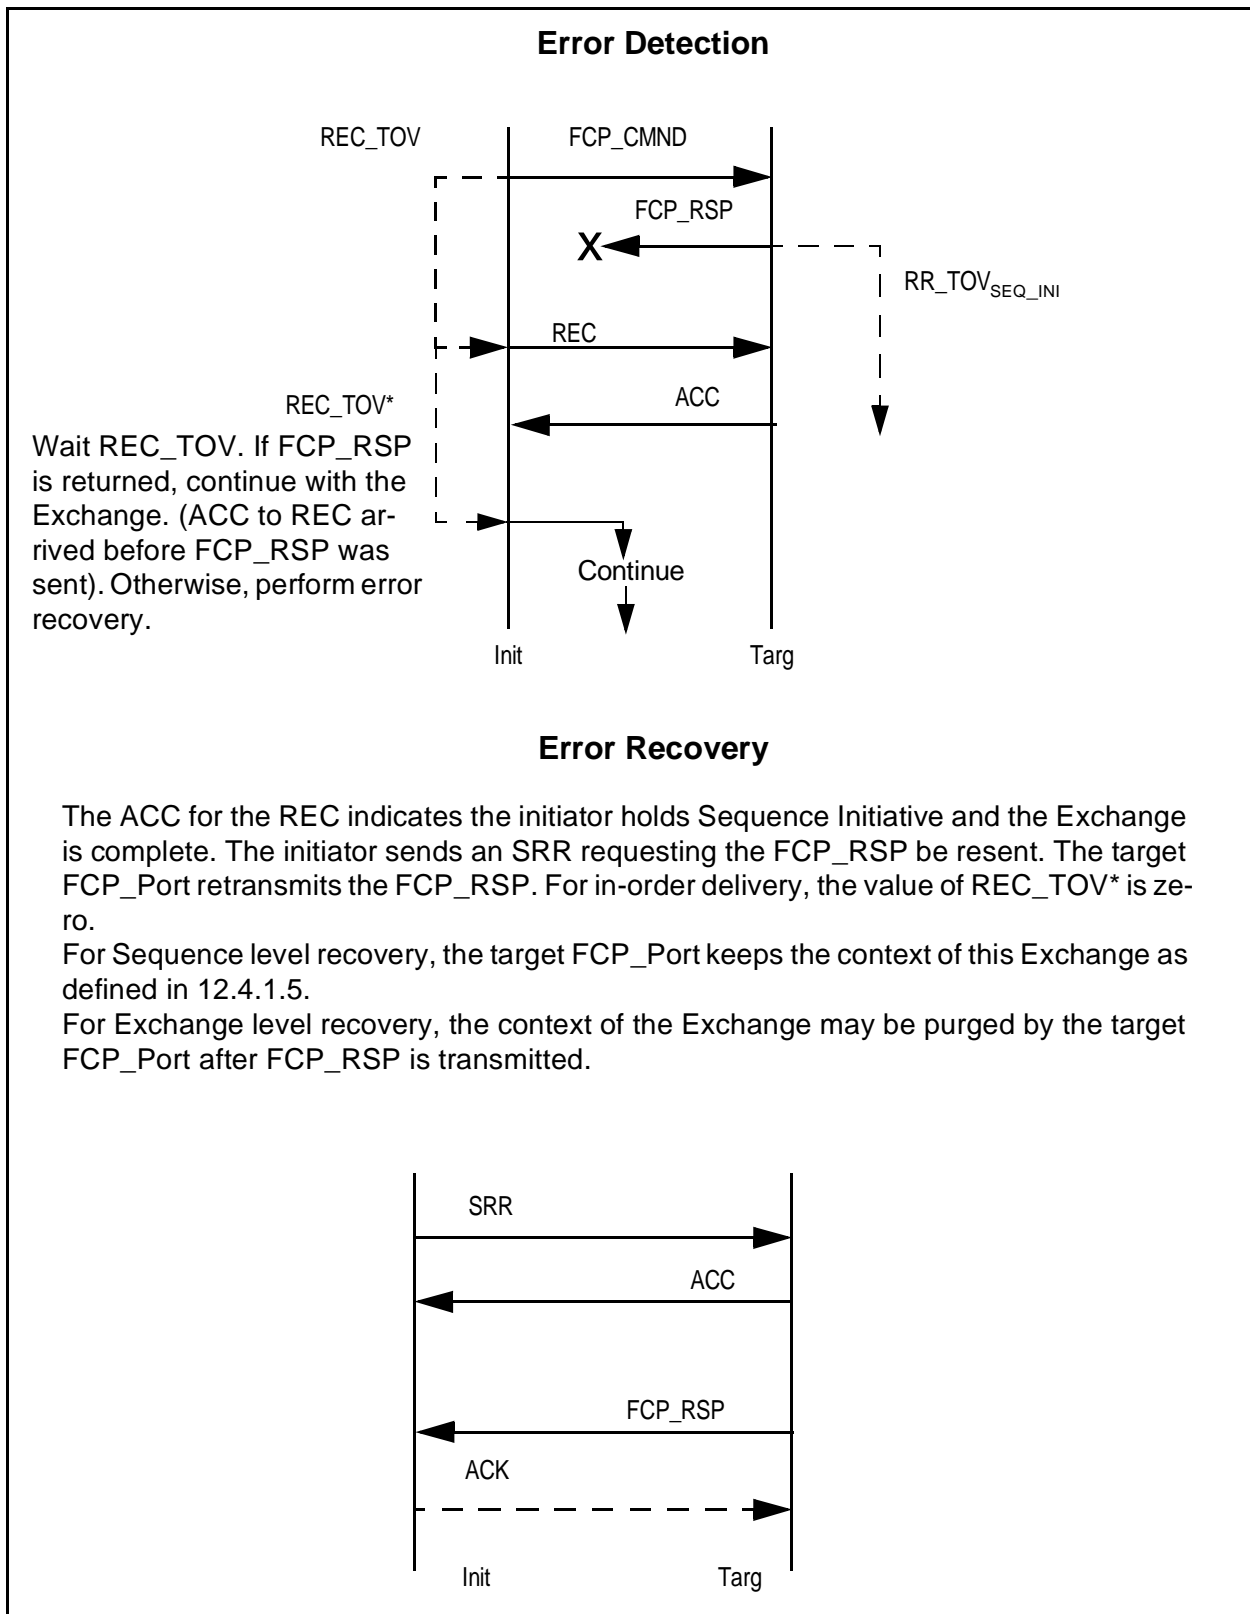
Figure C.8 - FCP_RSP Lost, FCP_CONF not requested, Unacknowledged Classes

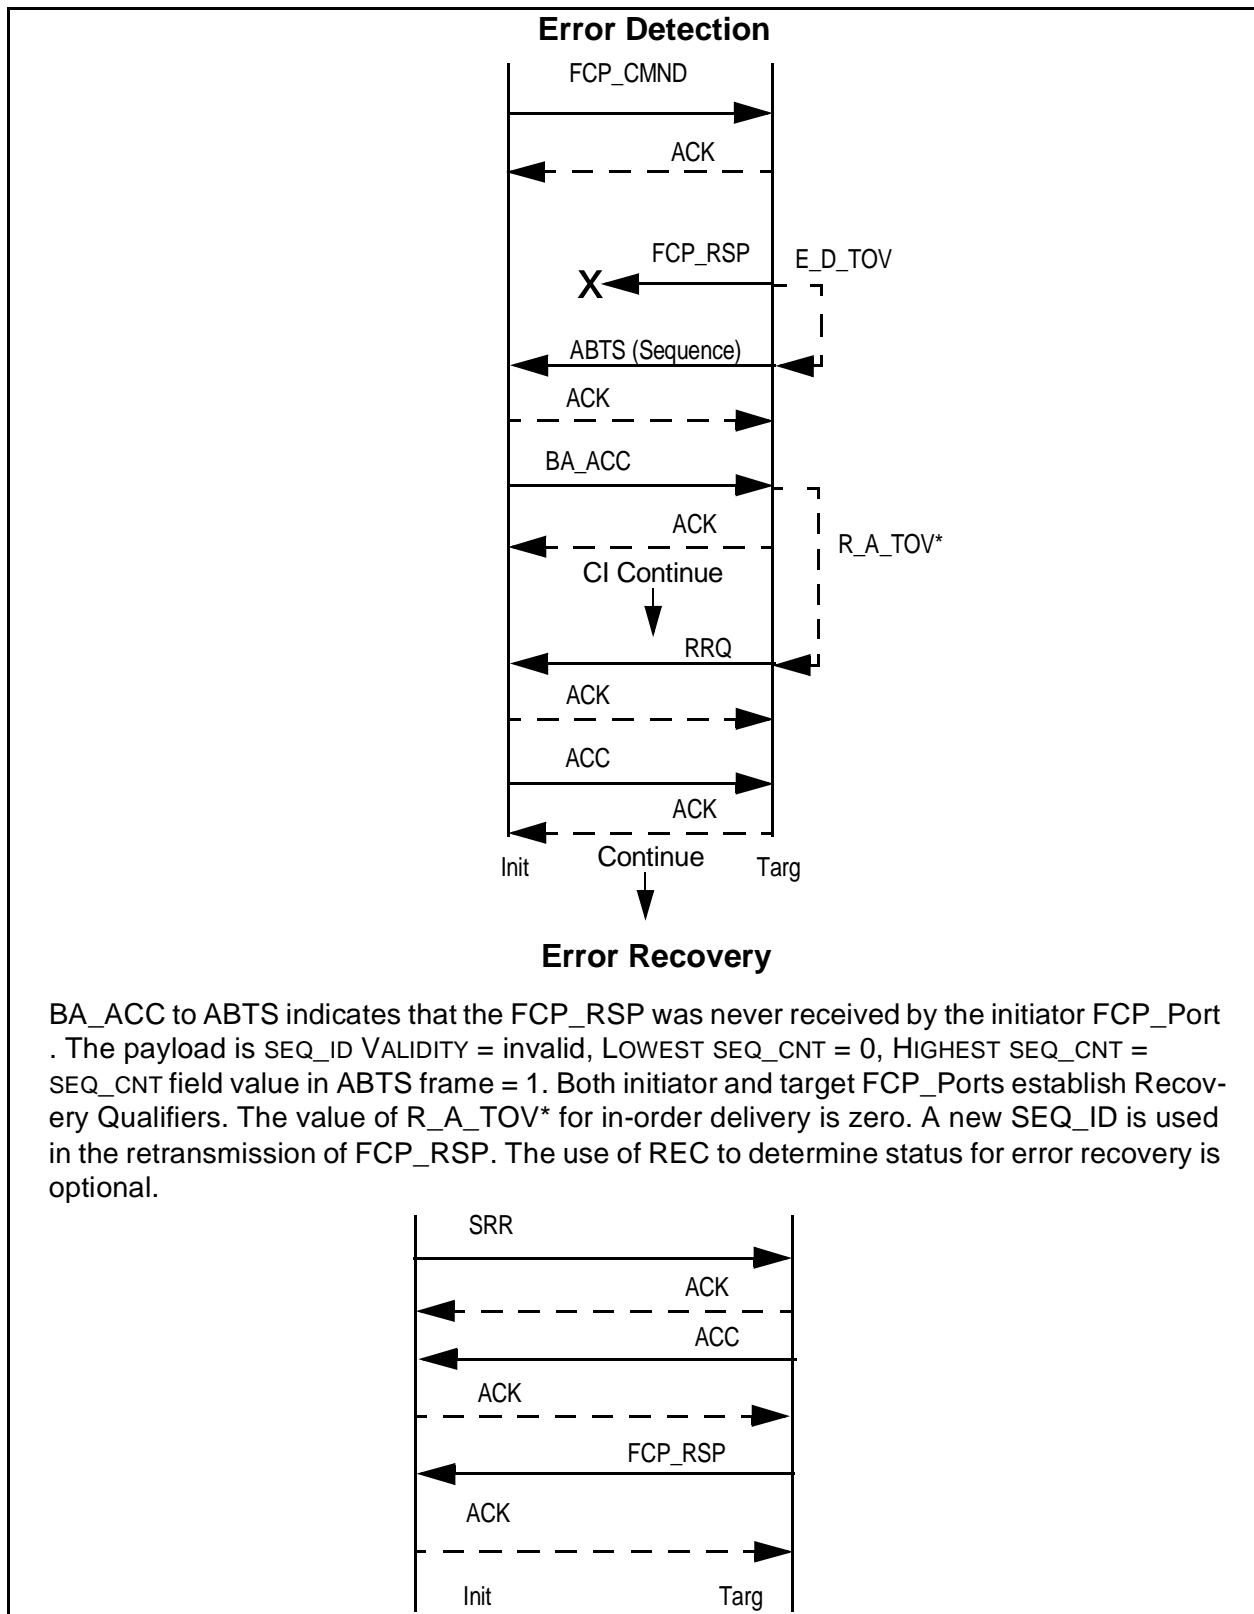
Figure C.9 - FCP_RSP Lost, FCP_CONF not requested, Acknowledged Classes

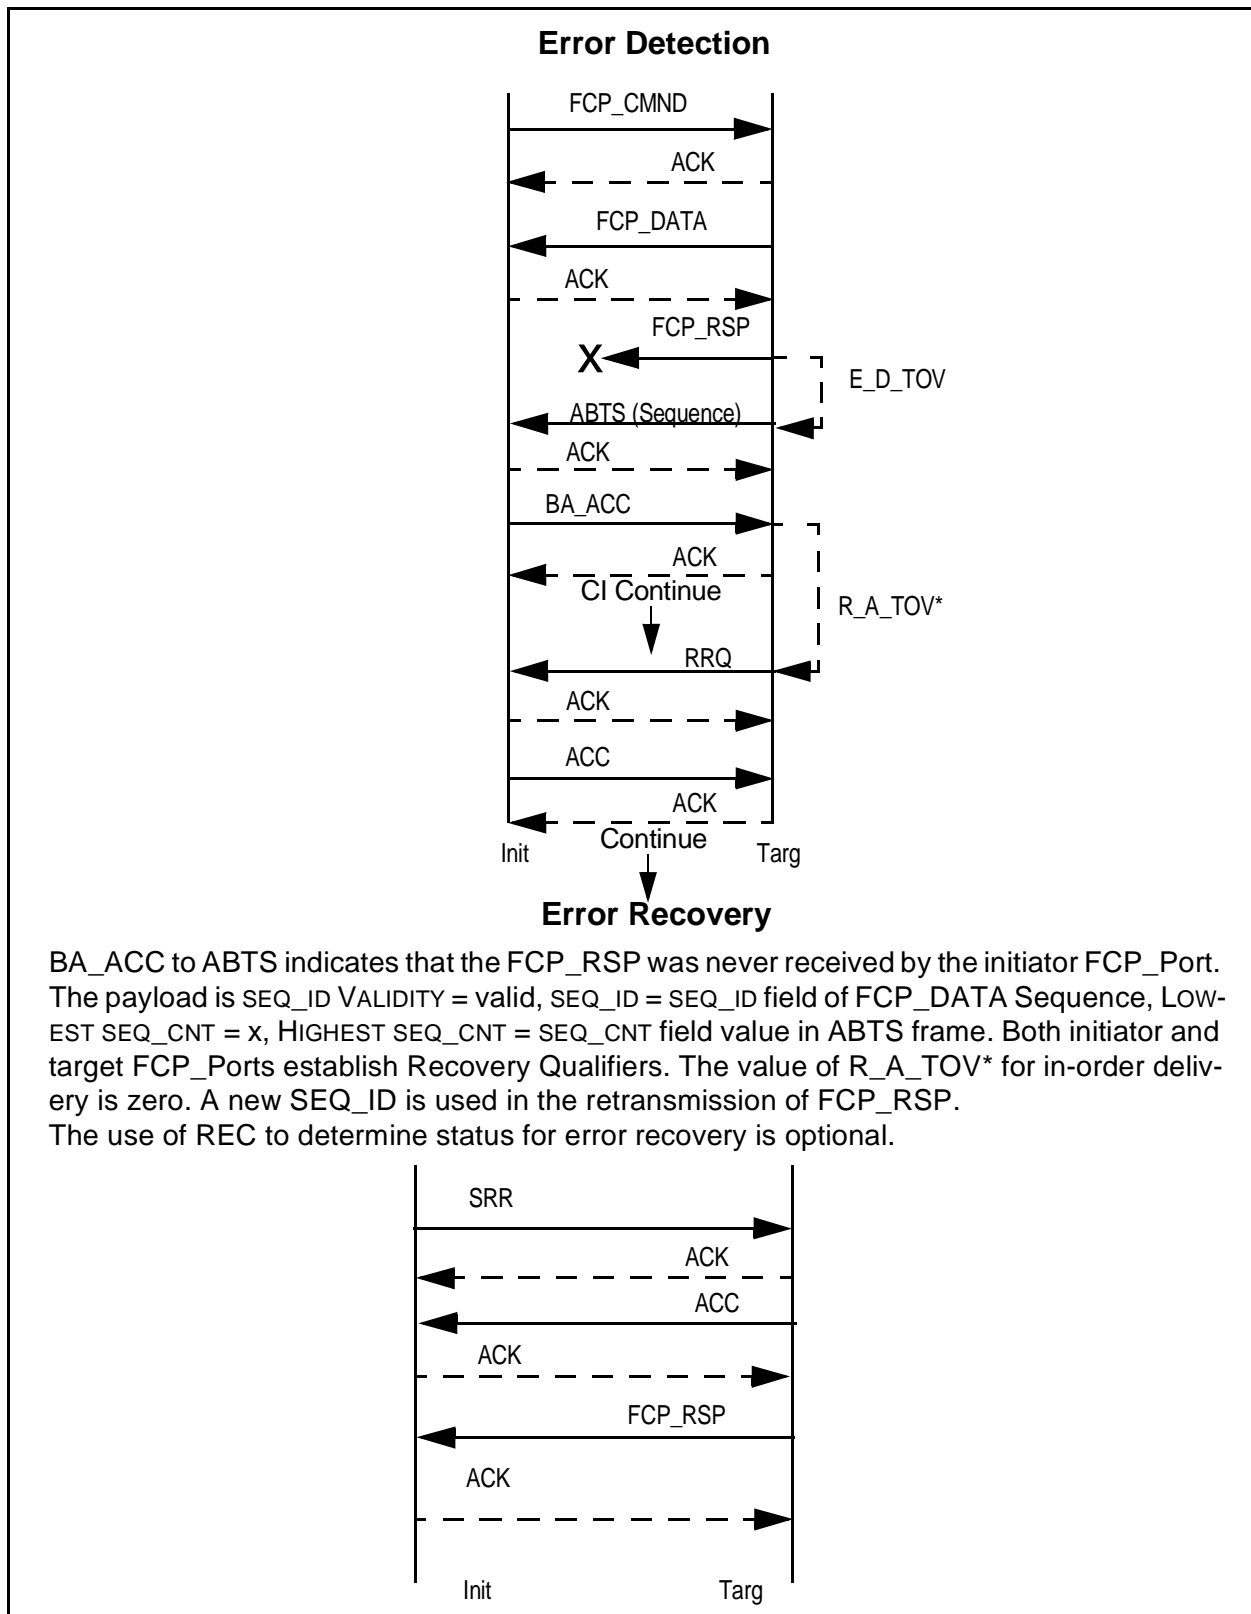
Figure C.10 - FCP_RSP Lost Read Command, no FCP_CONF, Acknowledged Classes

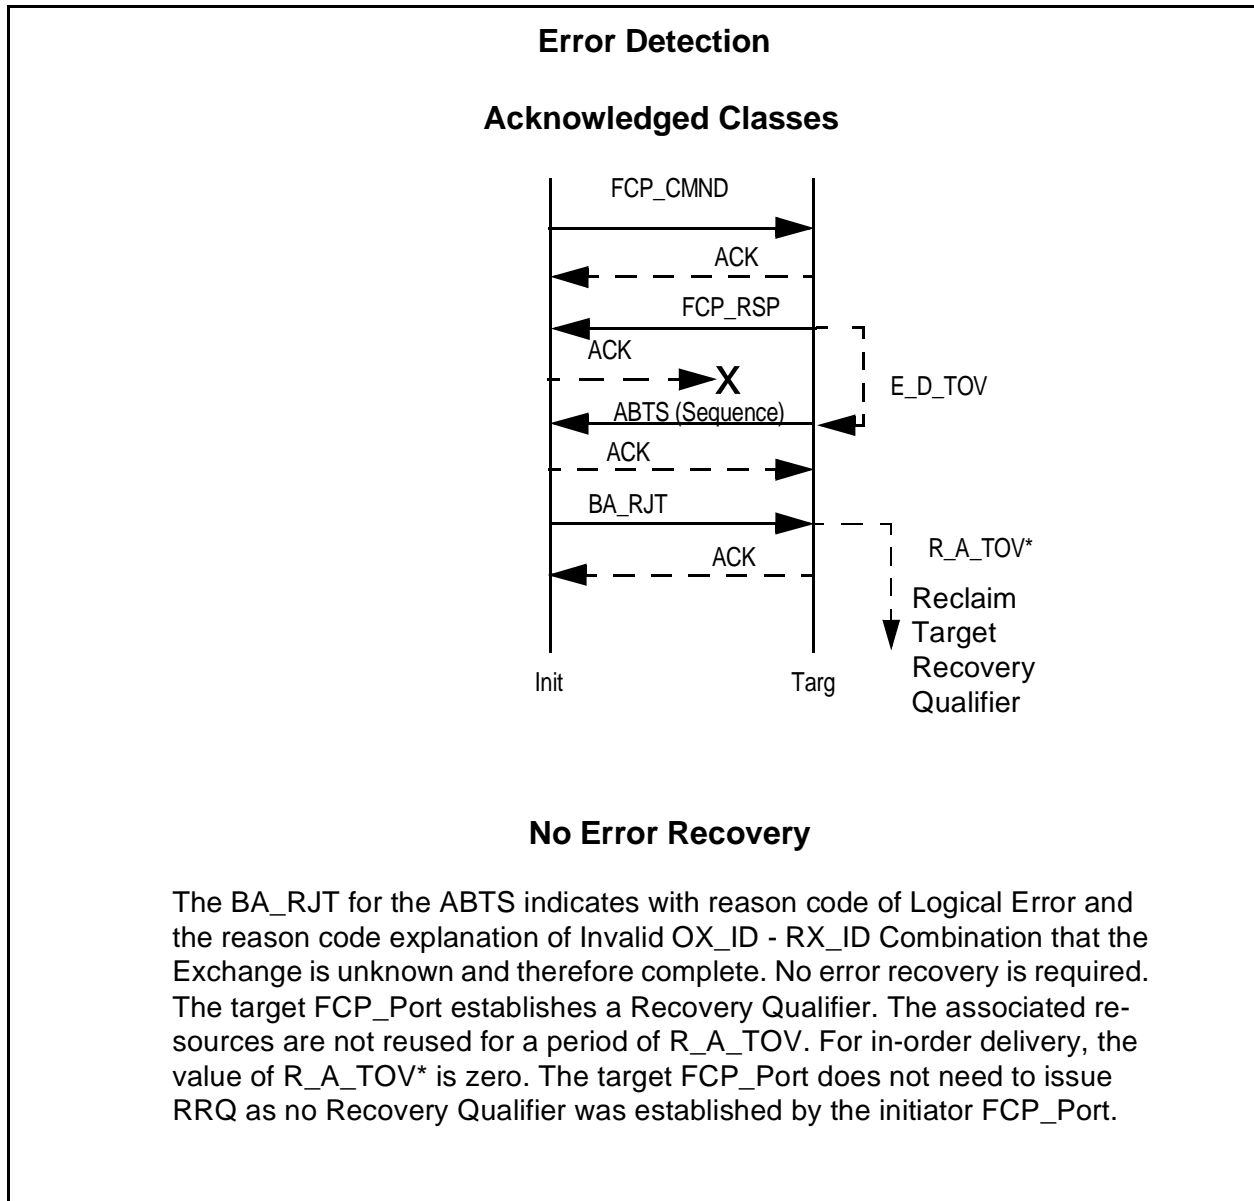
Figure C.11 - FCP_RSP Received, ACK Lost, Acknowledged Classes, Example 1


Figure C.12 - FCP_RSP Received, ACK Lost, Acknowledged Classes, Example 2

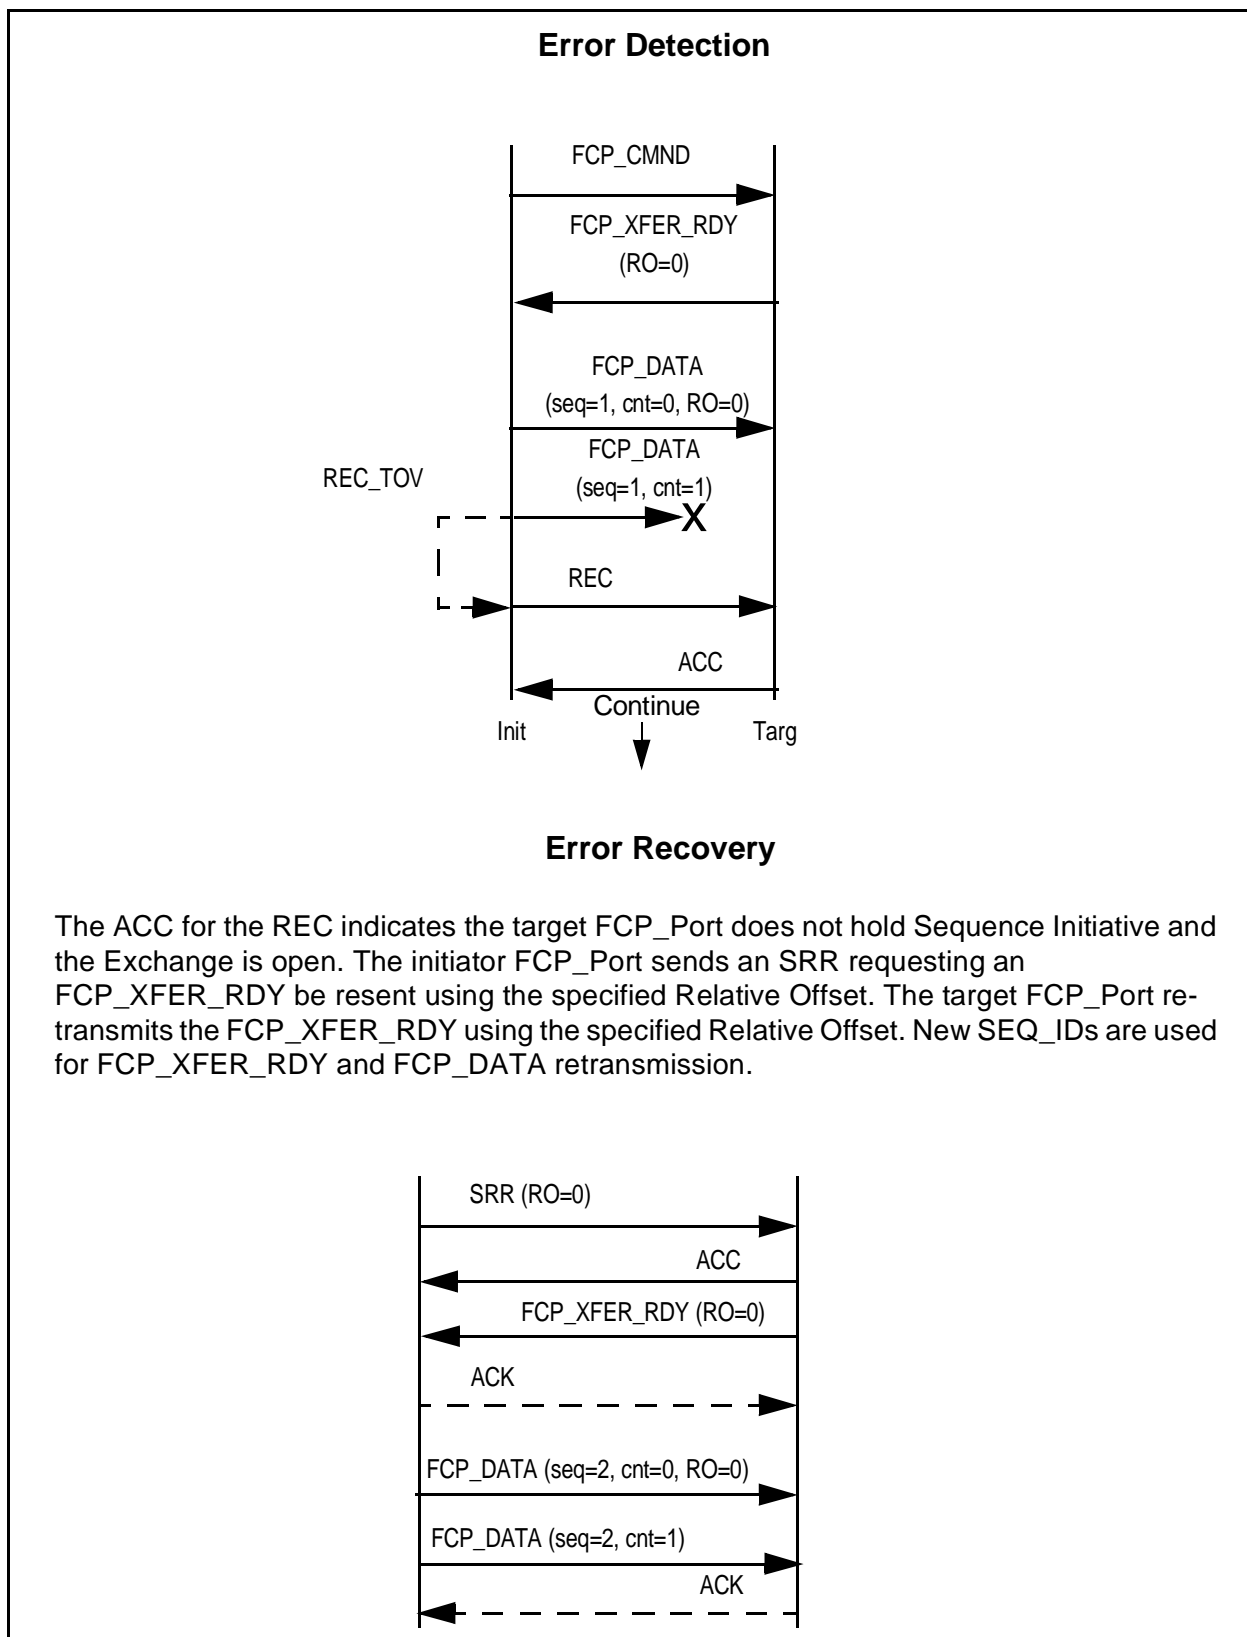
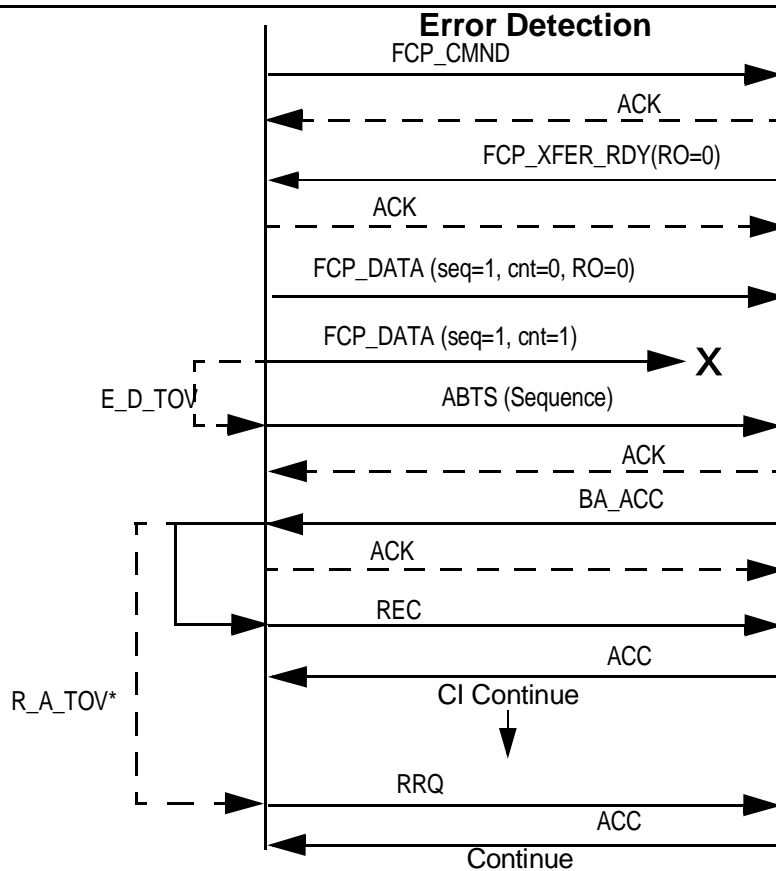
Figure C.13 - Lost Write Data, Last Frame of Sequence, Unacknowledged Classes

Figure C.14 - Lost Write Data, Last Frame of Sequence, Acknowledged Classes**Error Recovery**

Using the information contained in the REC, the recovery may be performed as in figure C.13. The BA_ACC may also be used. The payload is SEQ_ID VALIDITY = valid, SEQ_ID = SEQ_ID field of FCP_CMND, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = 2.

New SEQ_IDs are used for retransmitting FCP_XFER_RDY and FCP_DATA. ACKs for REC/ACC and RRQ/ACC are not shown. The value of R_A_TOV* for in-order delivery is zero.

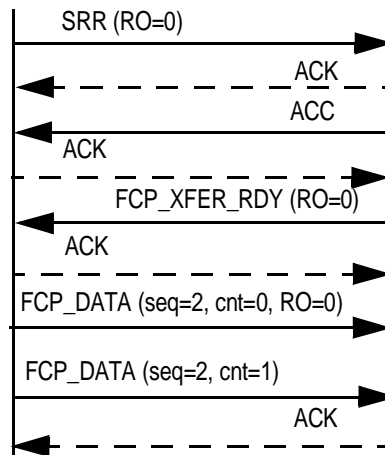


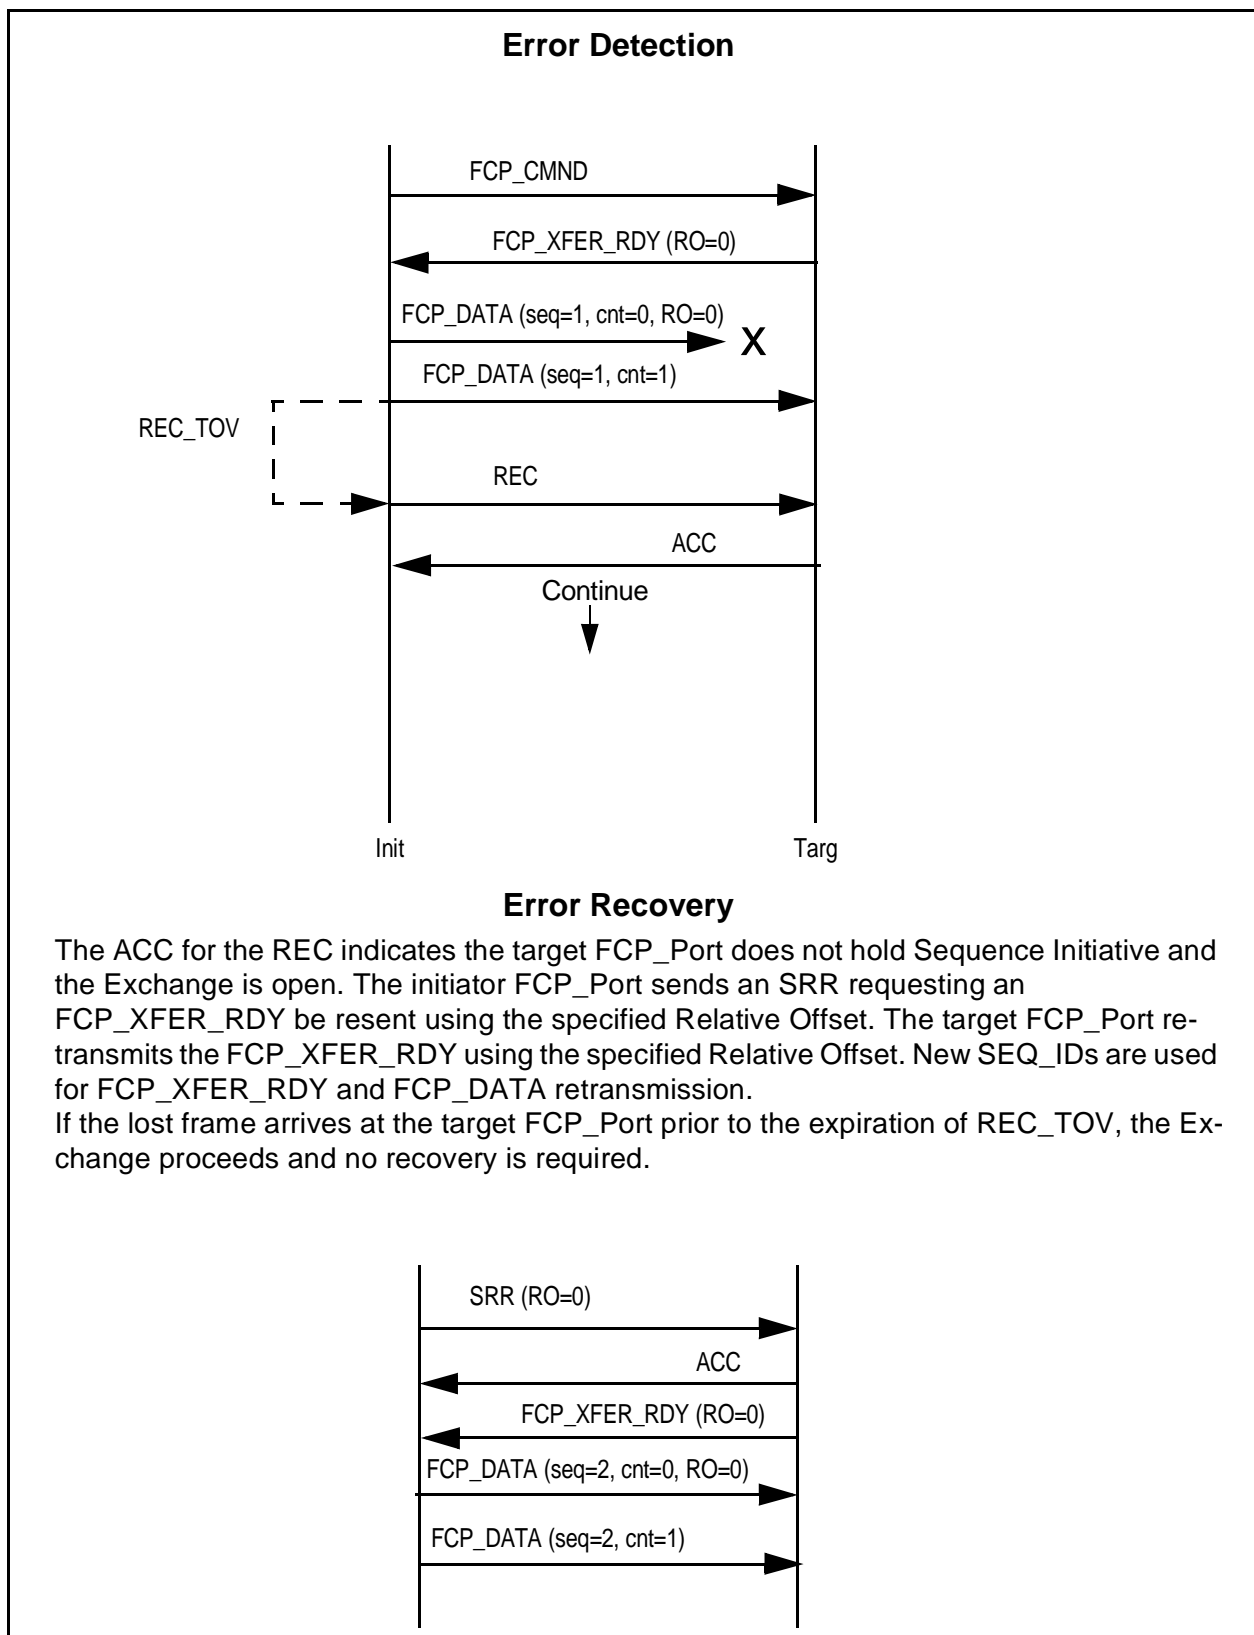
Figure C.15 - Lost Write Data, Not Last Frame of Sequence, Unacknowledged Classes

Figure C.16 - Lost Write Data, Not Last Frame of Sequence, Acknowledged Classes

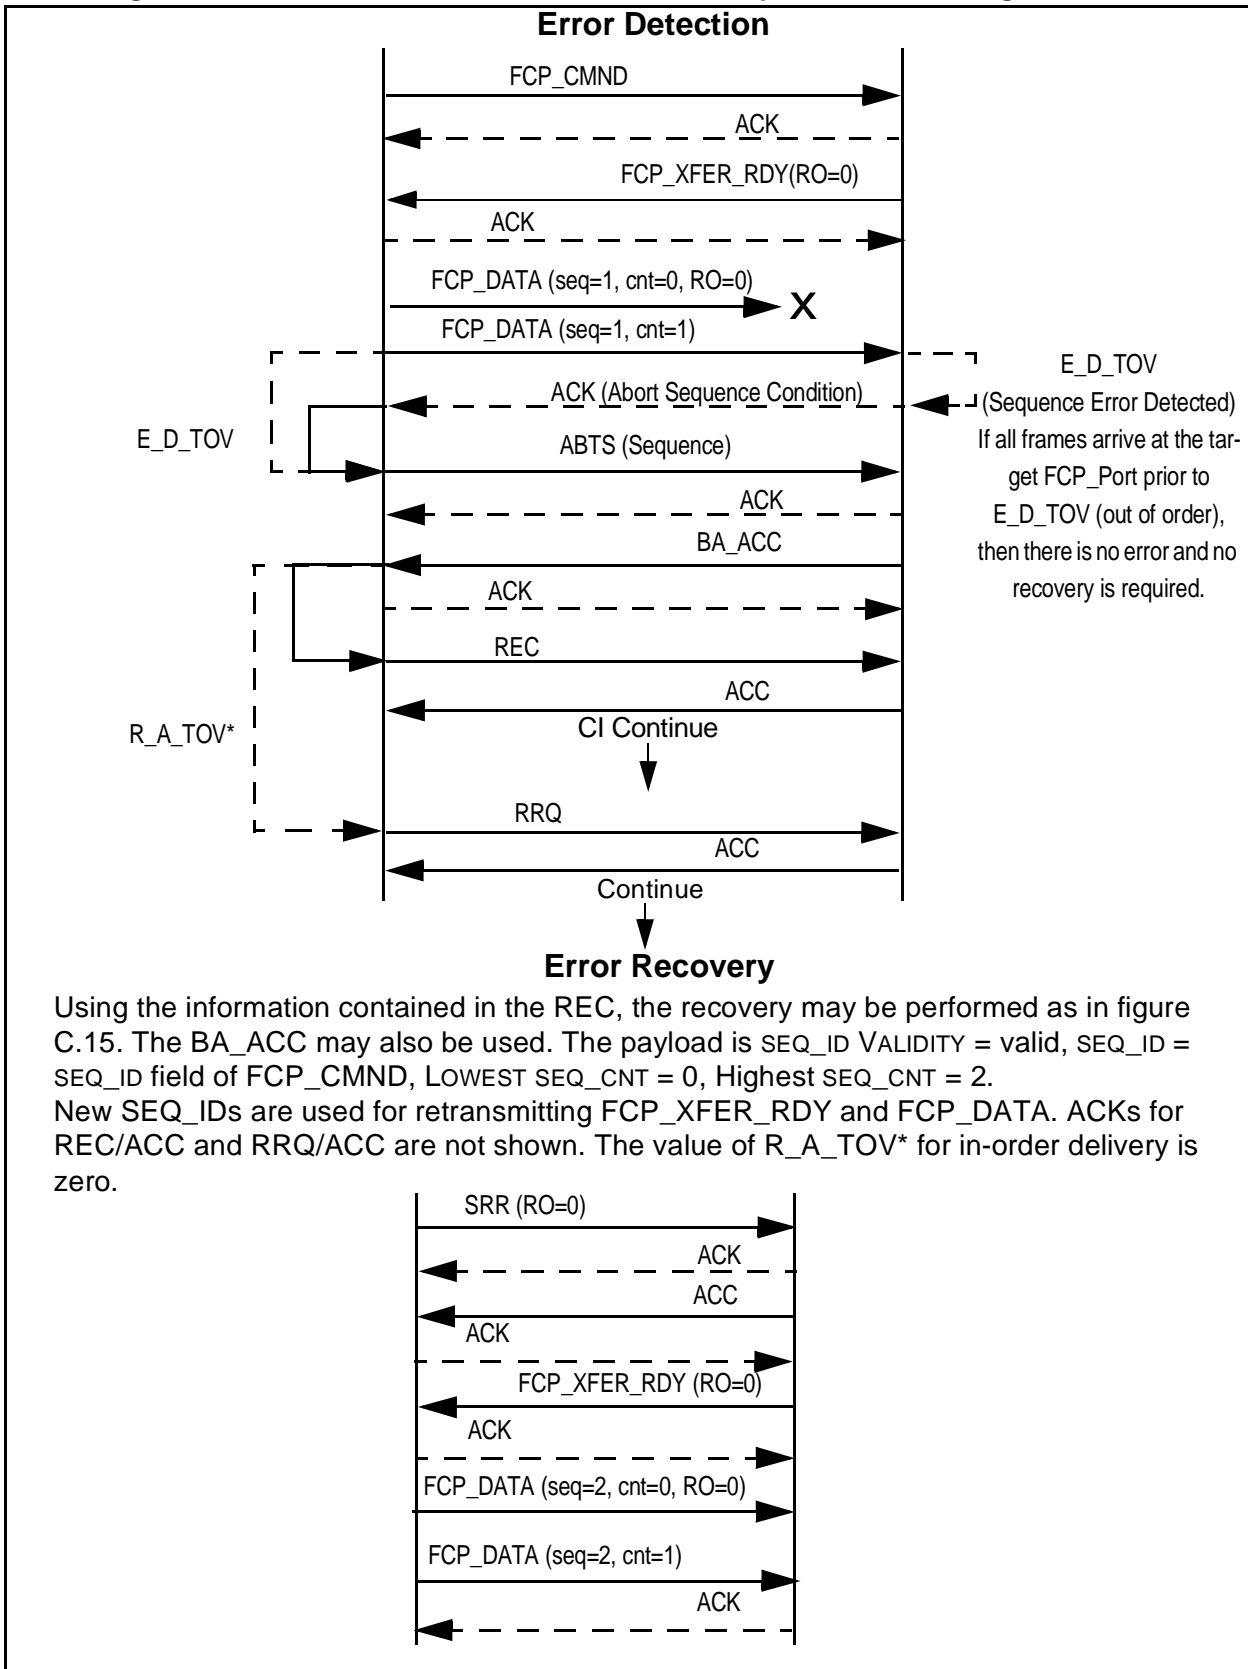


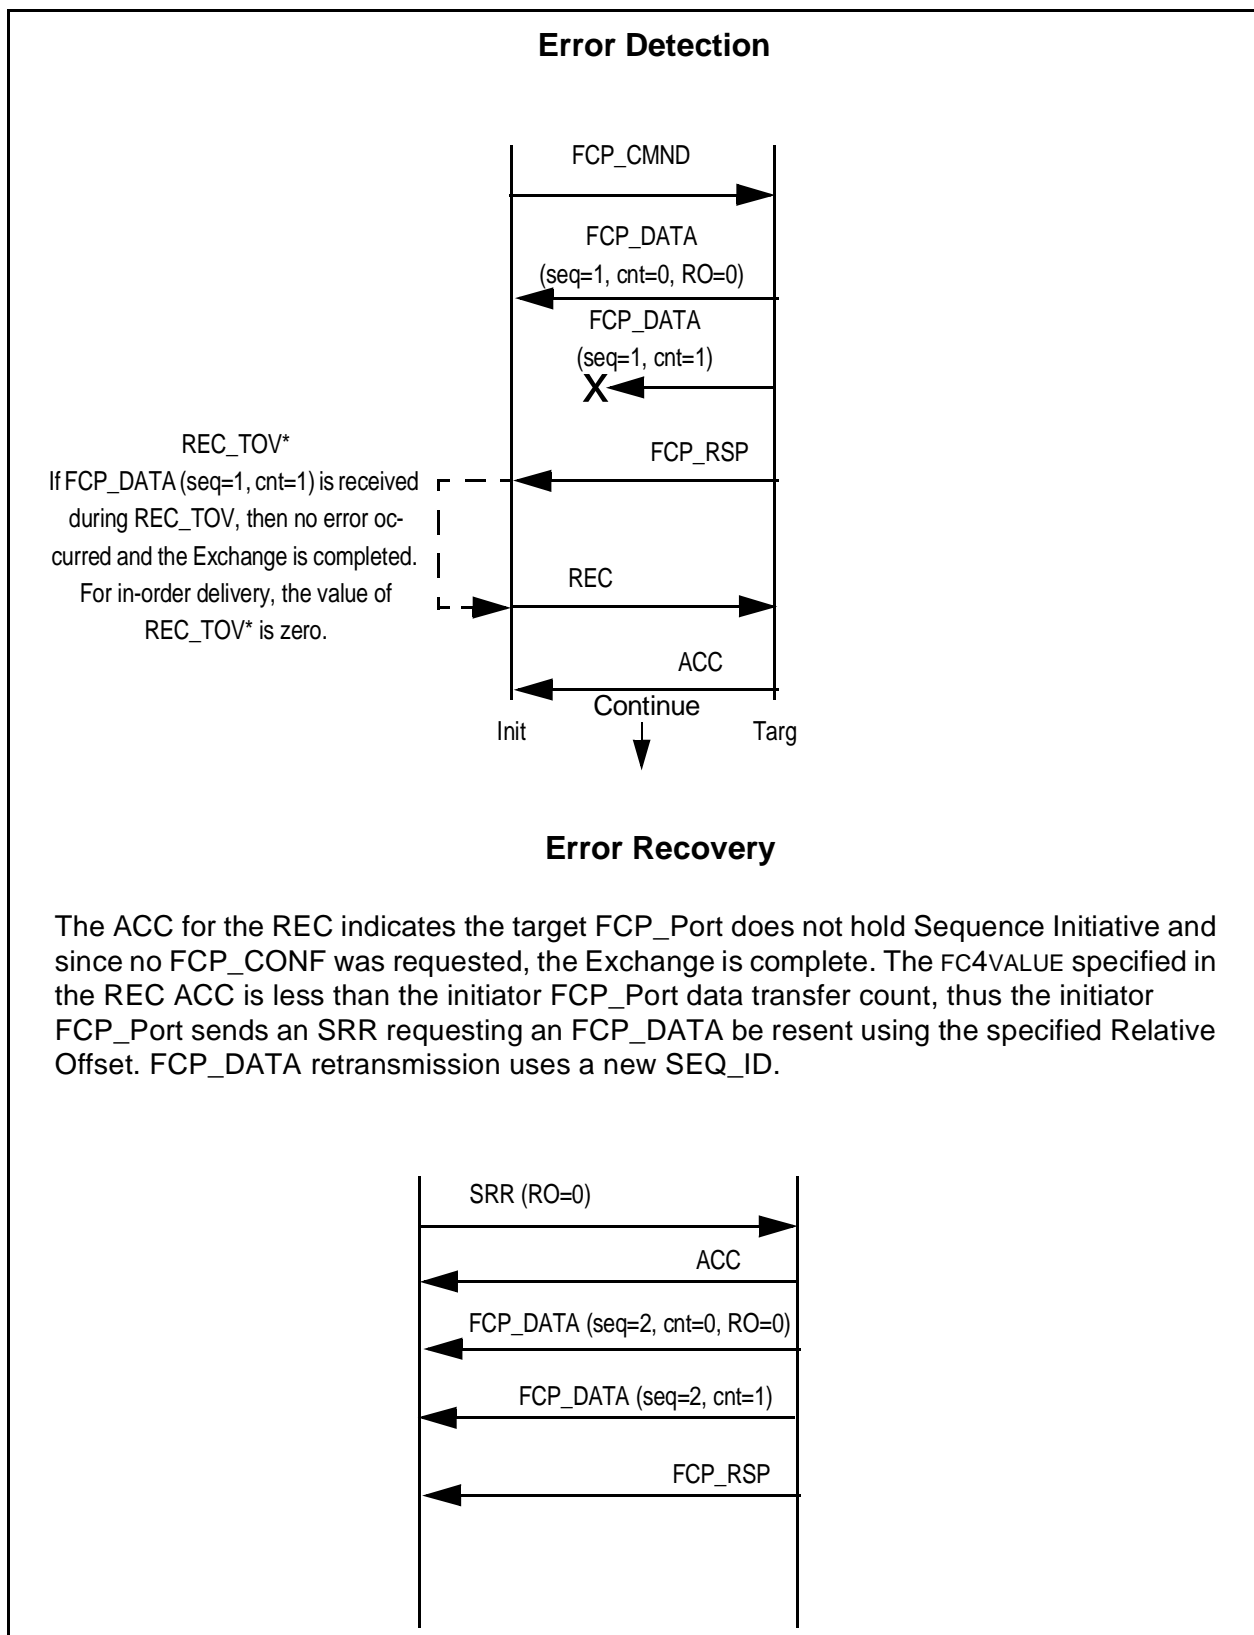
Figure C.17 - Lost Read Data, Last Frame of Sequence, Unacknowledged Classes

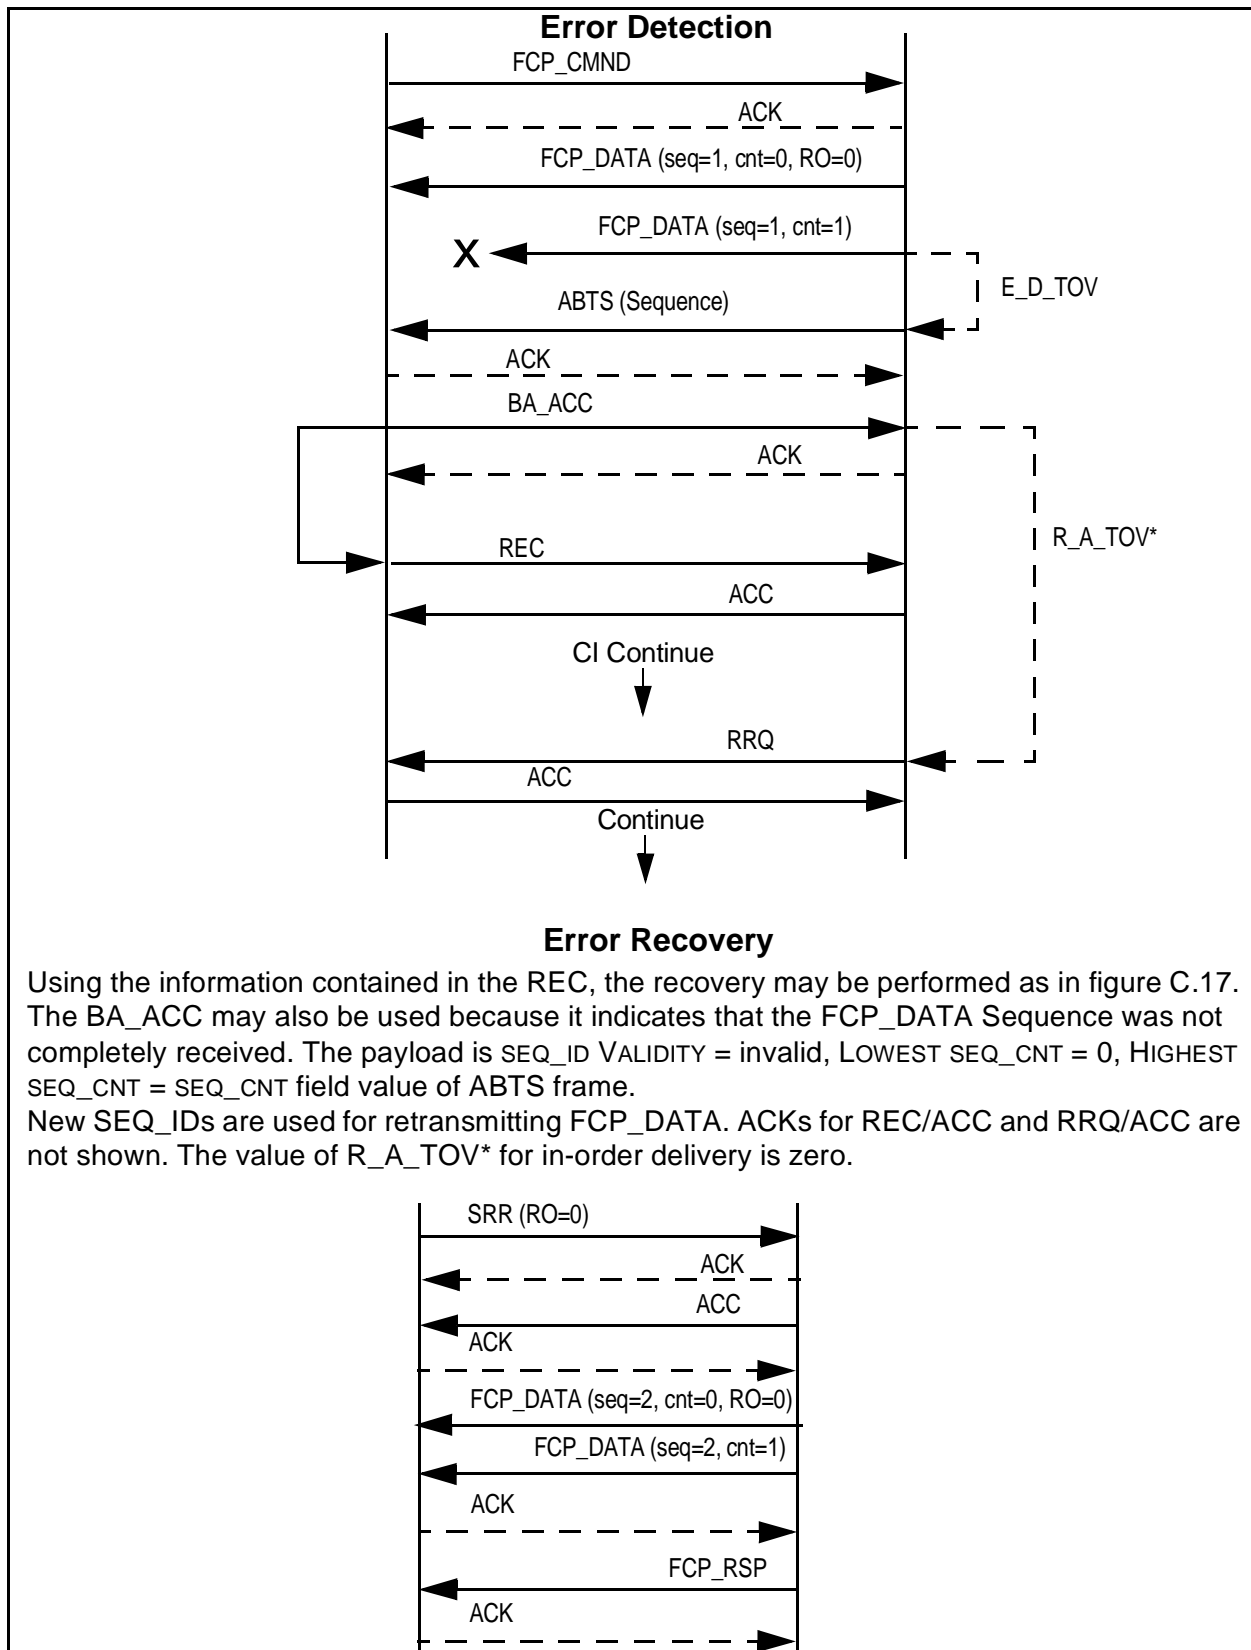
Figure C.18 - Lost Read Data, Last Frame of Sequence, Acknowledged Classes

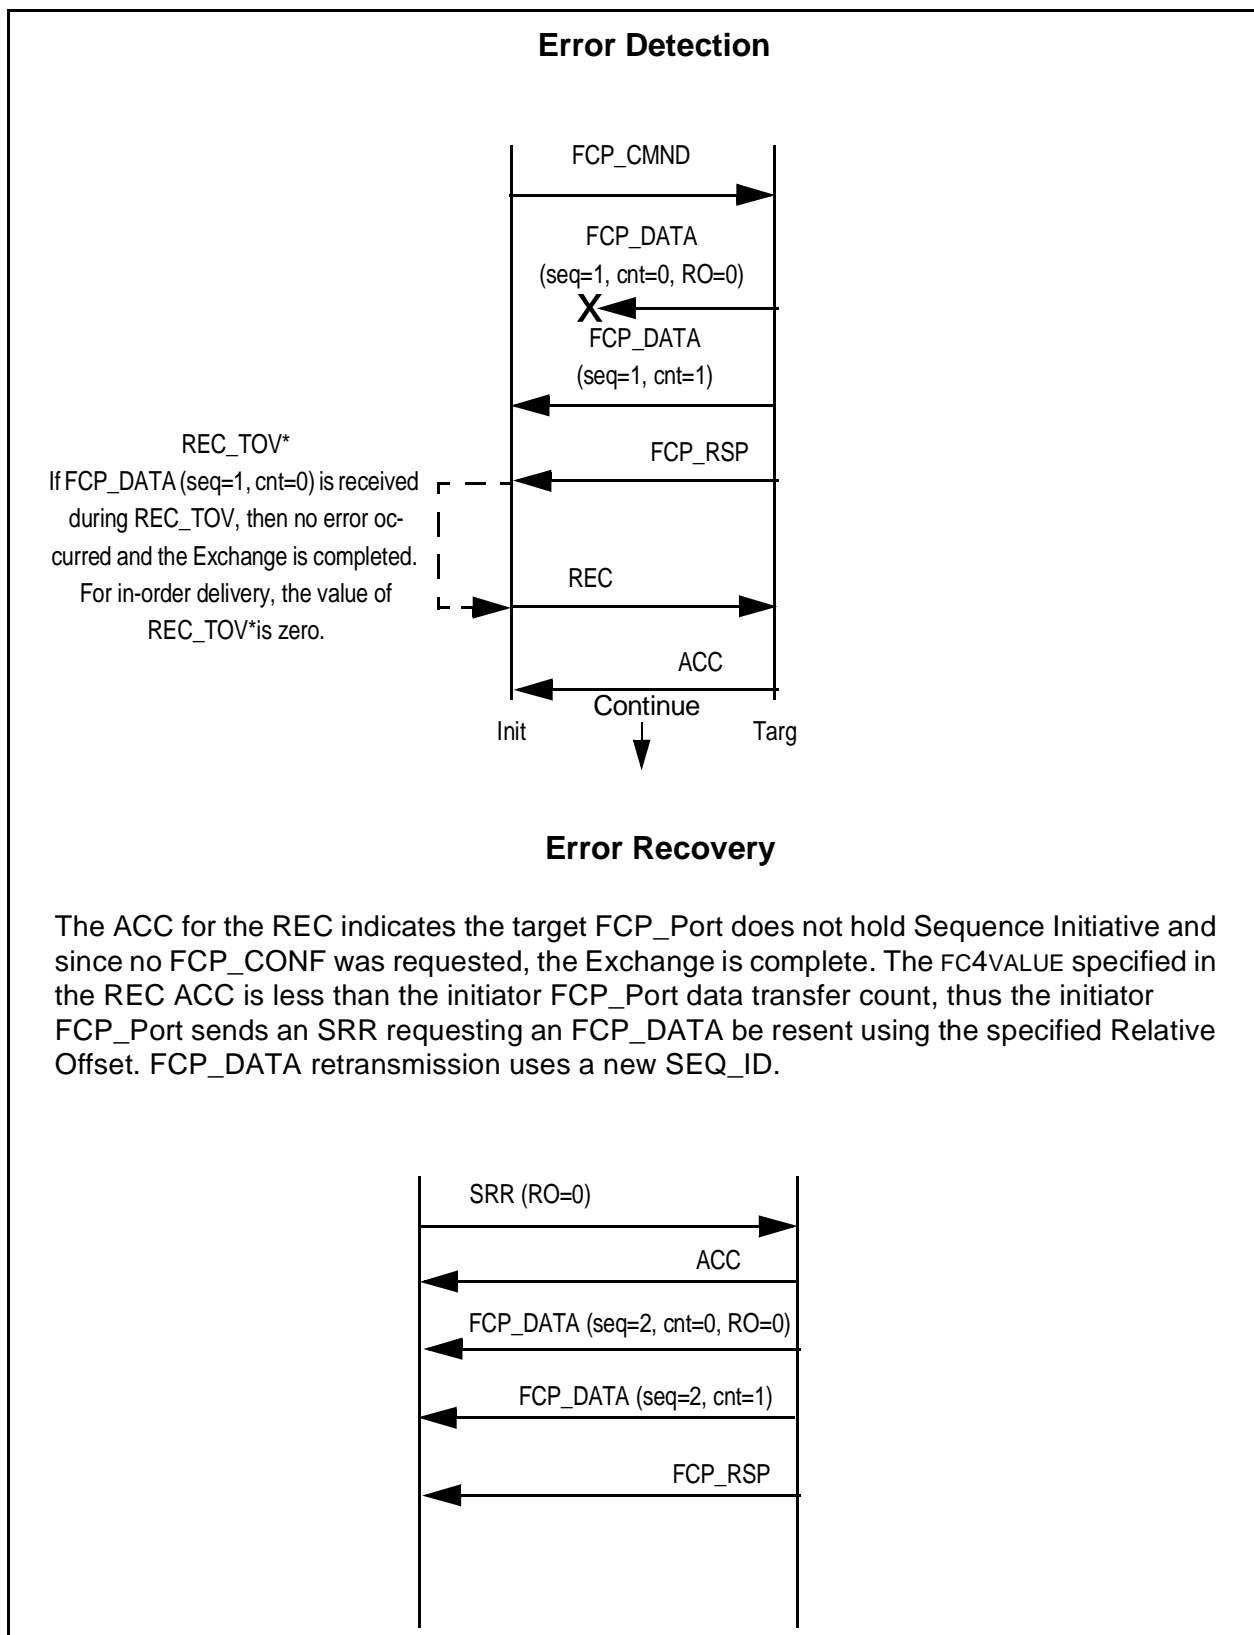
Figure C.19 - Lost Read Data, Not Last Frame of Sequence, Unacknowledged Classes

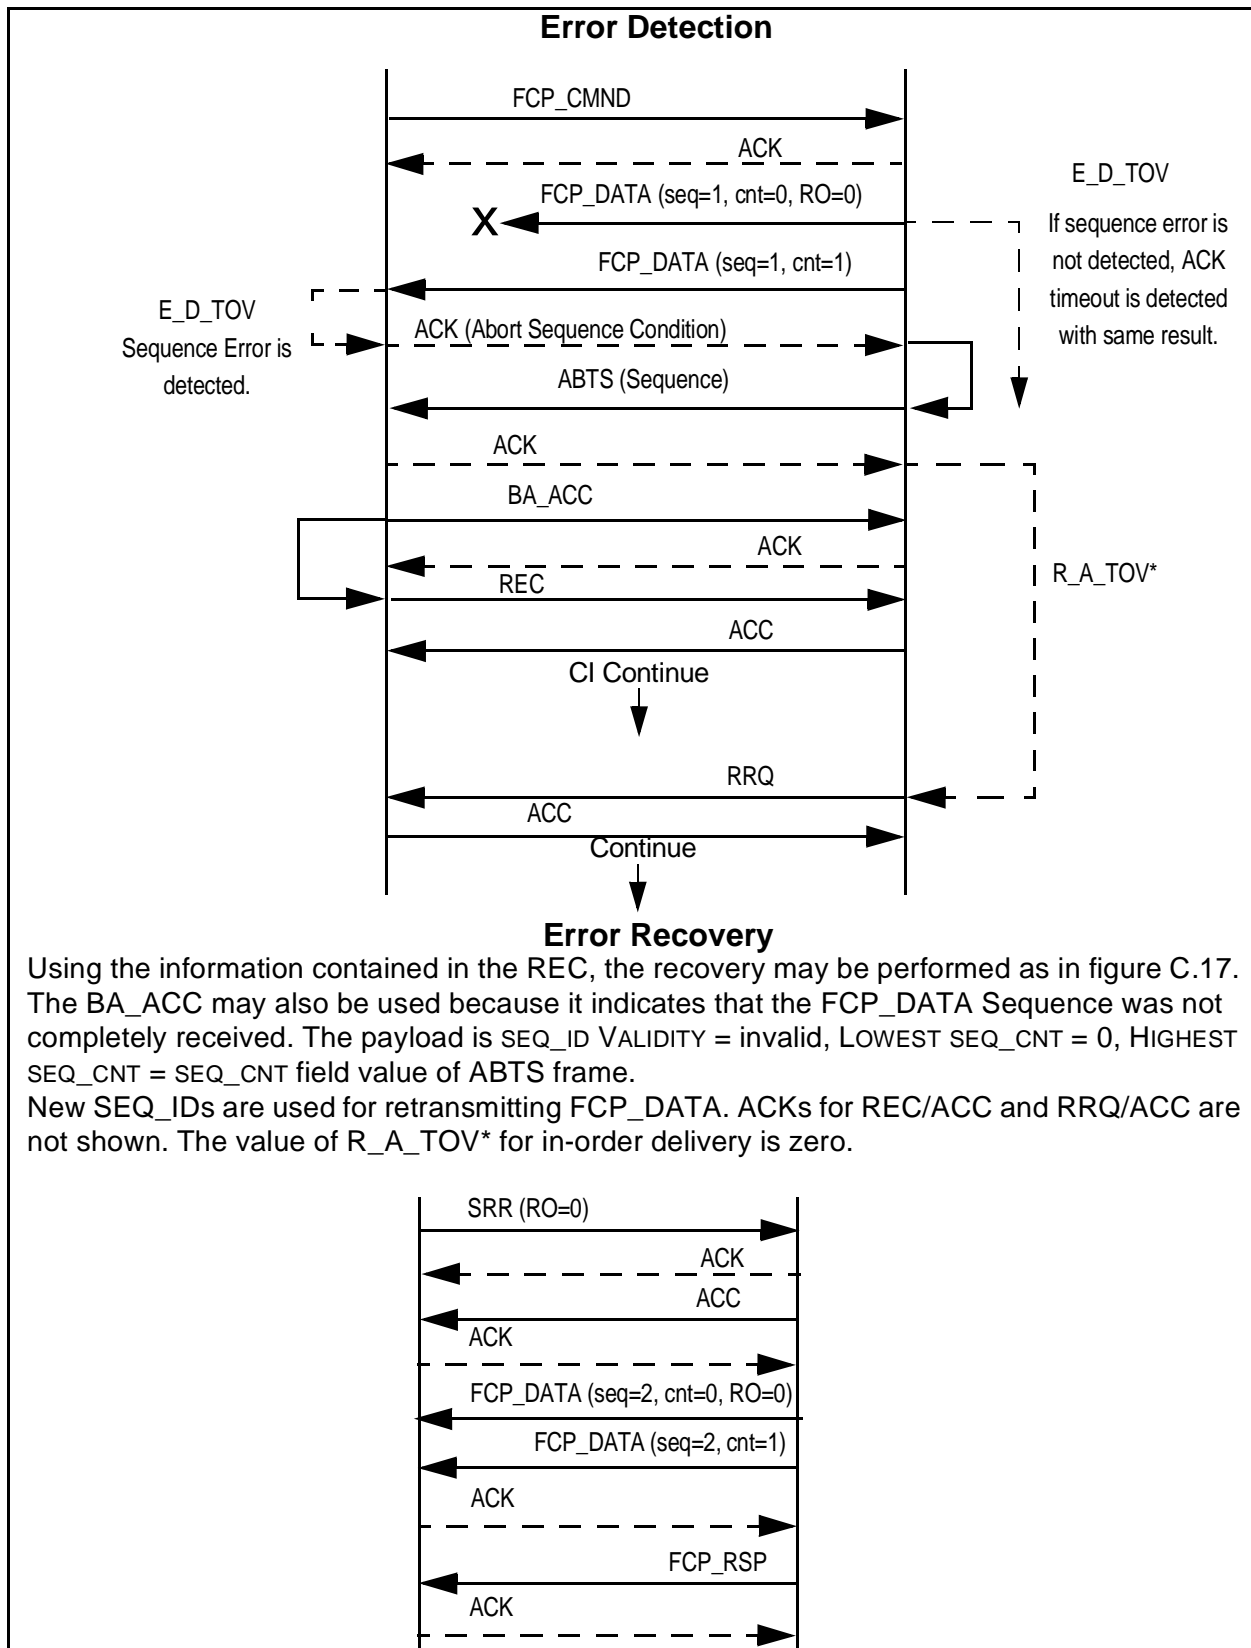
Figure C.20 - Lost Read Data, Not Last Frame of Sequence, Acknowledged Classes

Figure C.21 - ACK Lost on Read (Acknowledged Classes)

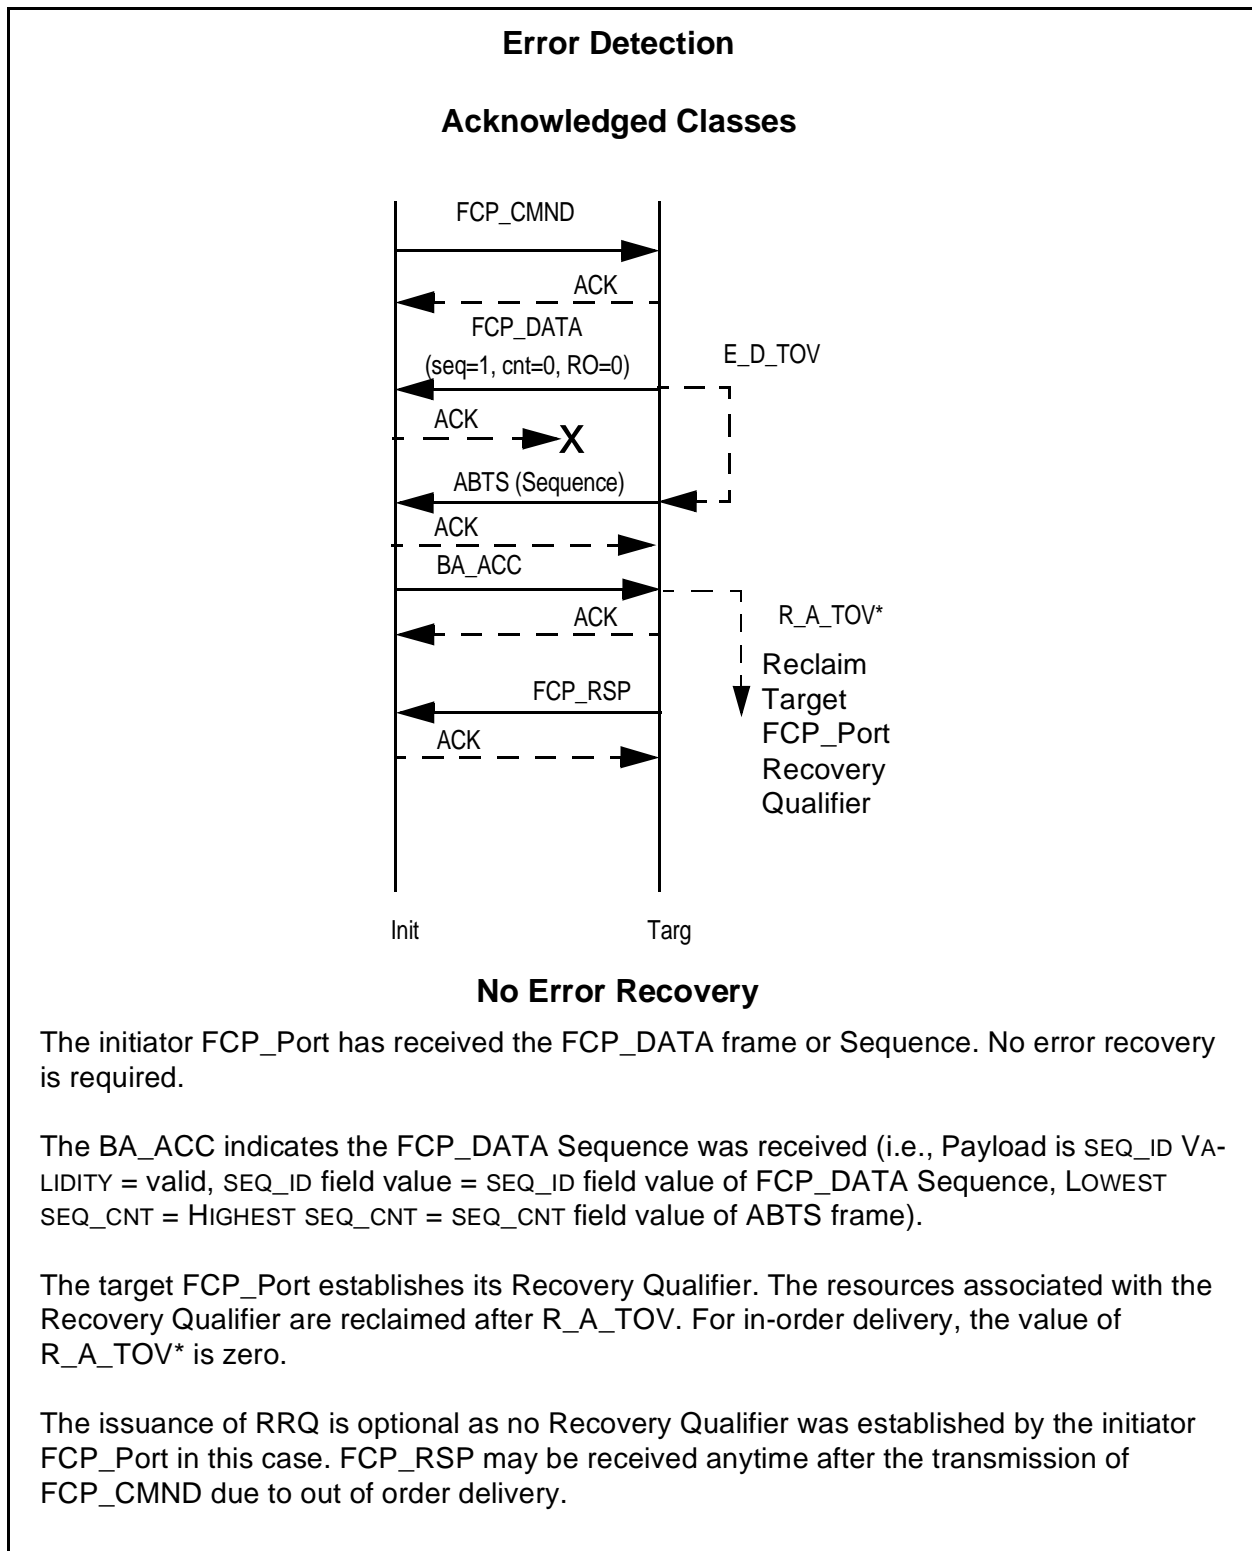


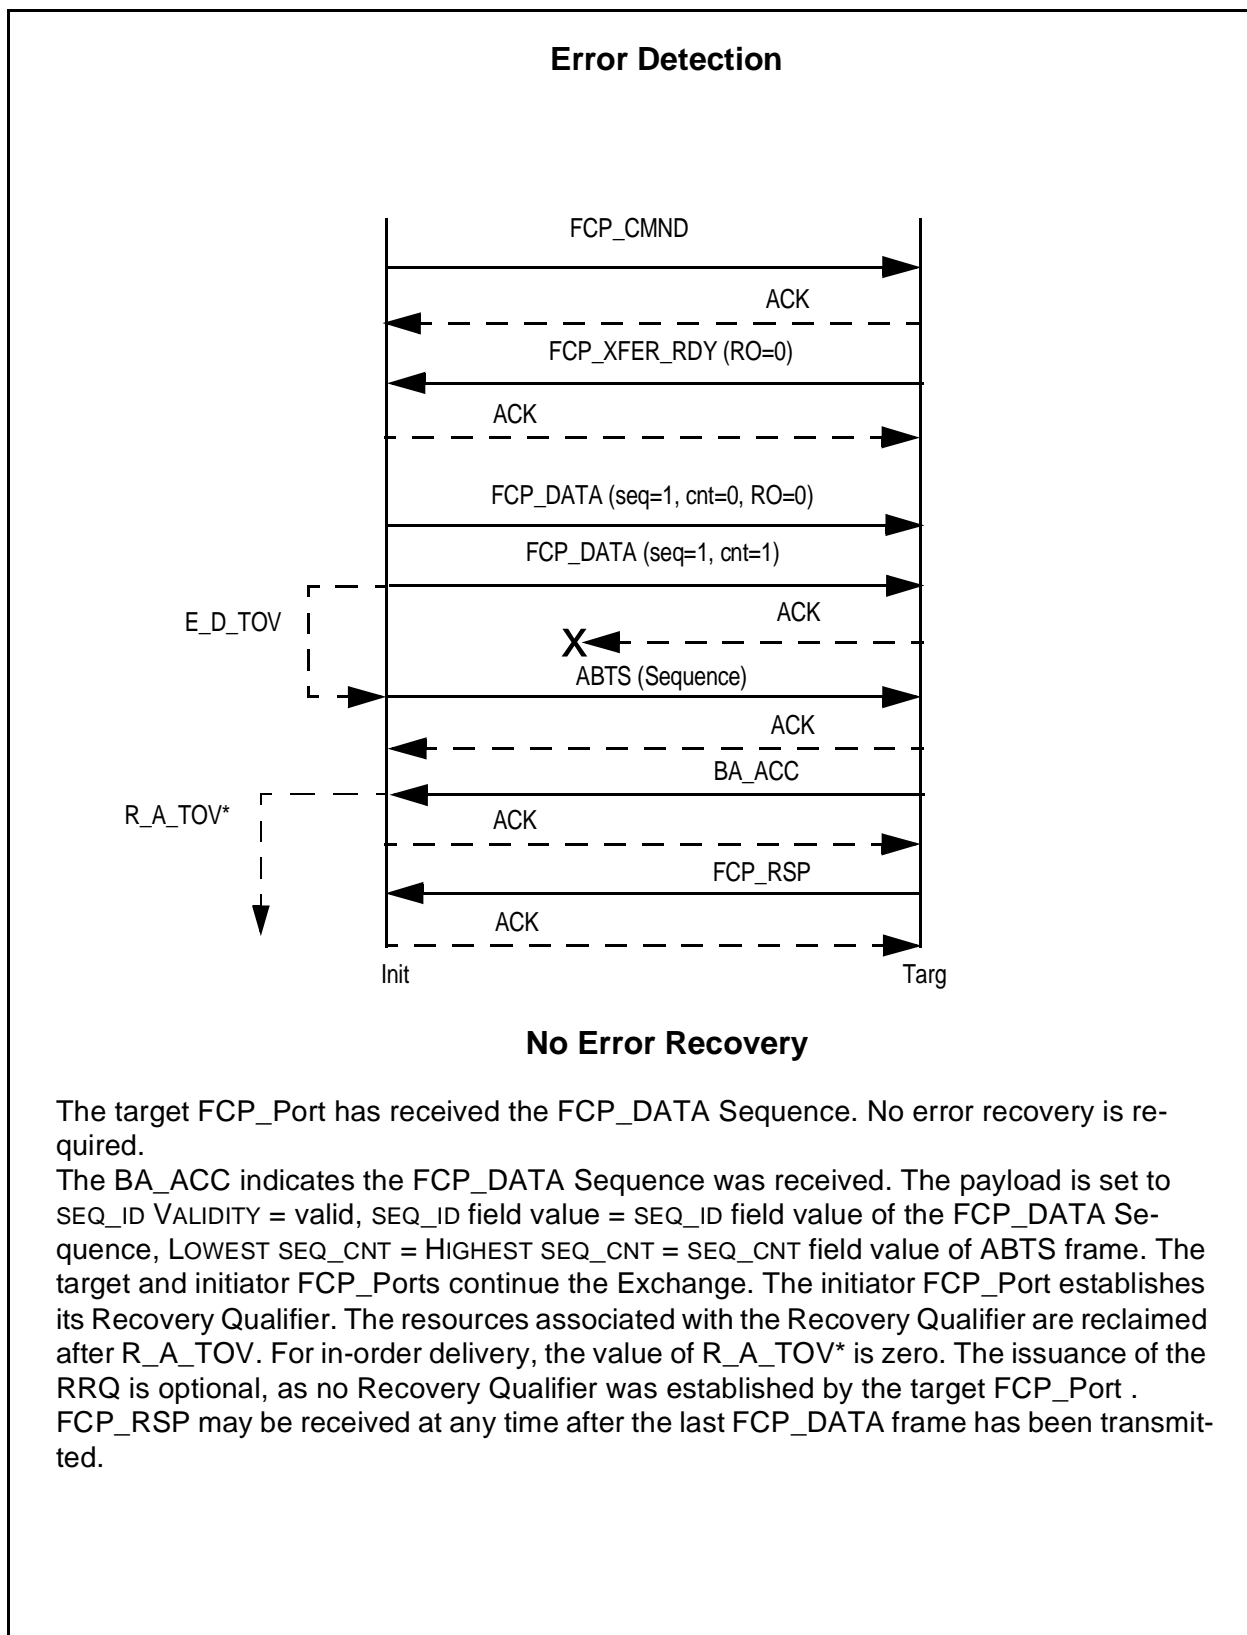
Figure C.22 - ACK Lost on Write (Acknowledged Classes)

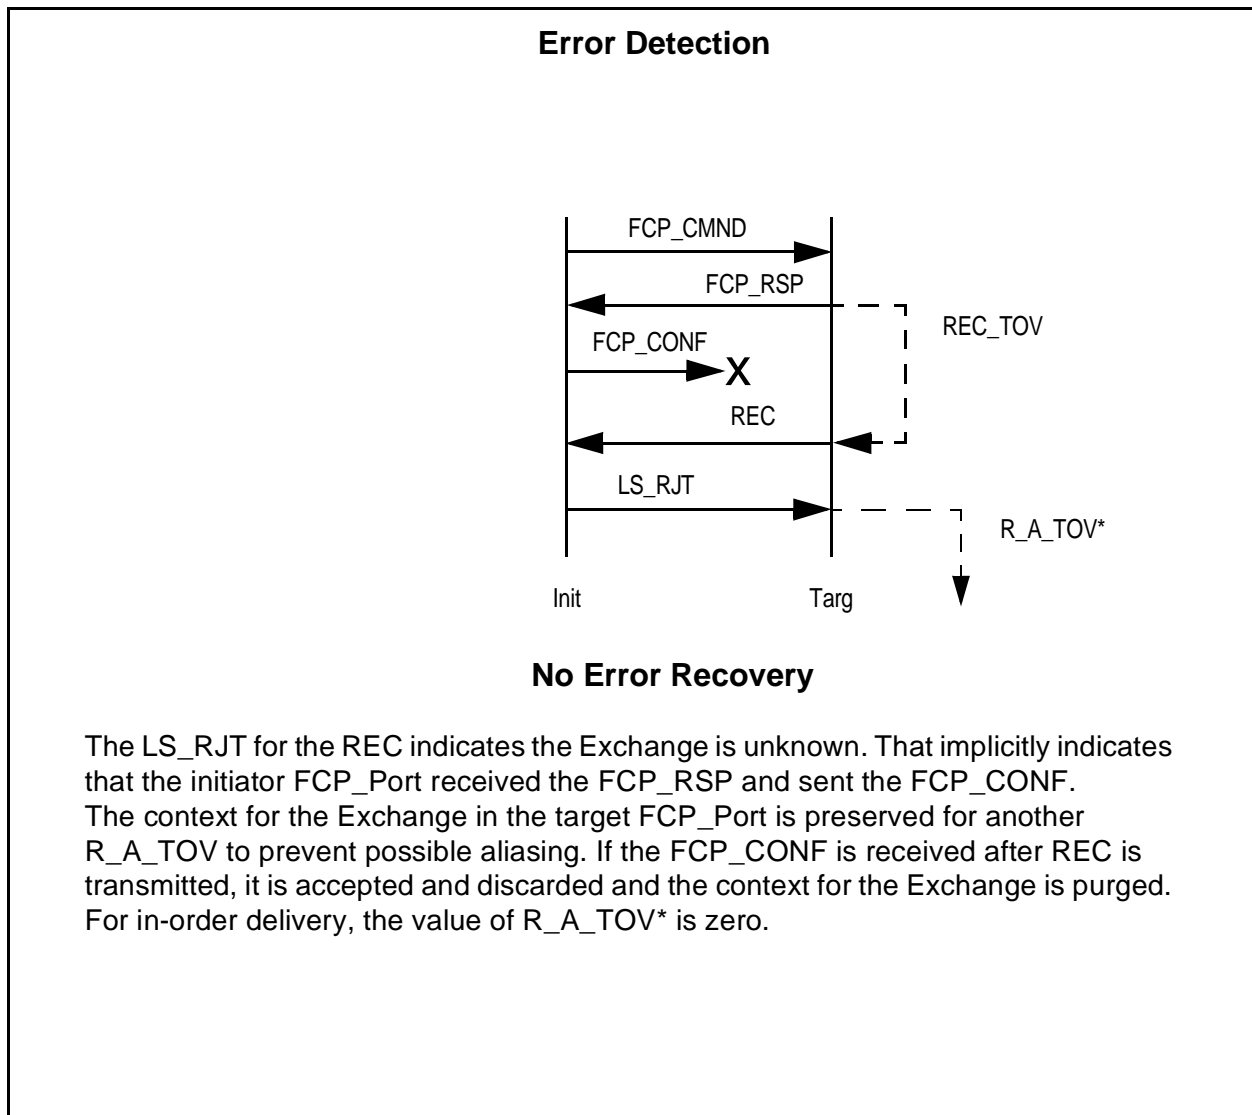
Figure C.23 - FCP_CONF Lost, Unacknowledged Classes

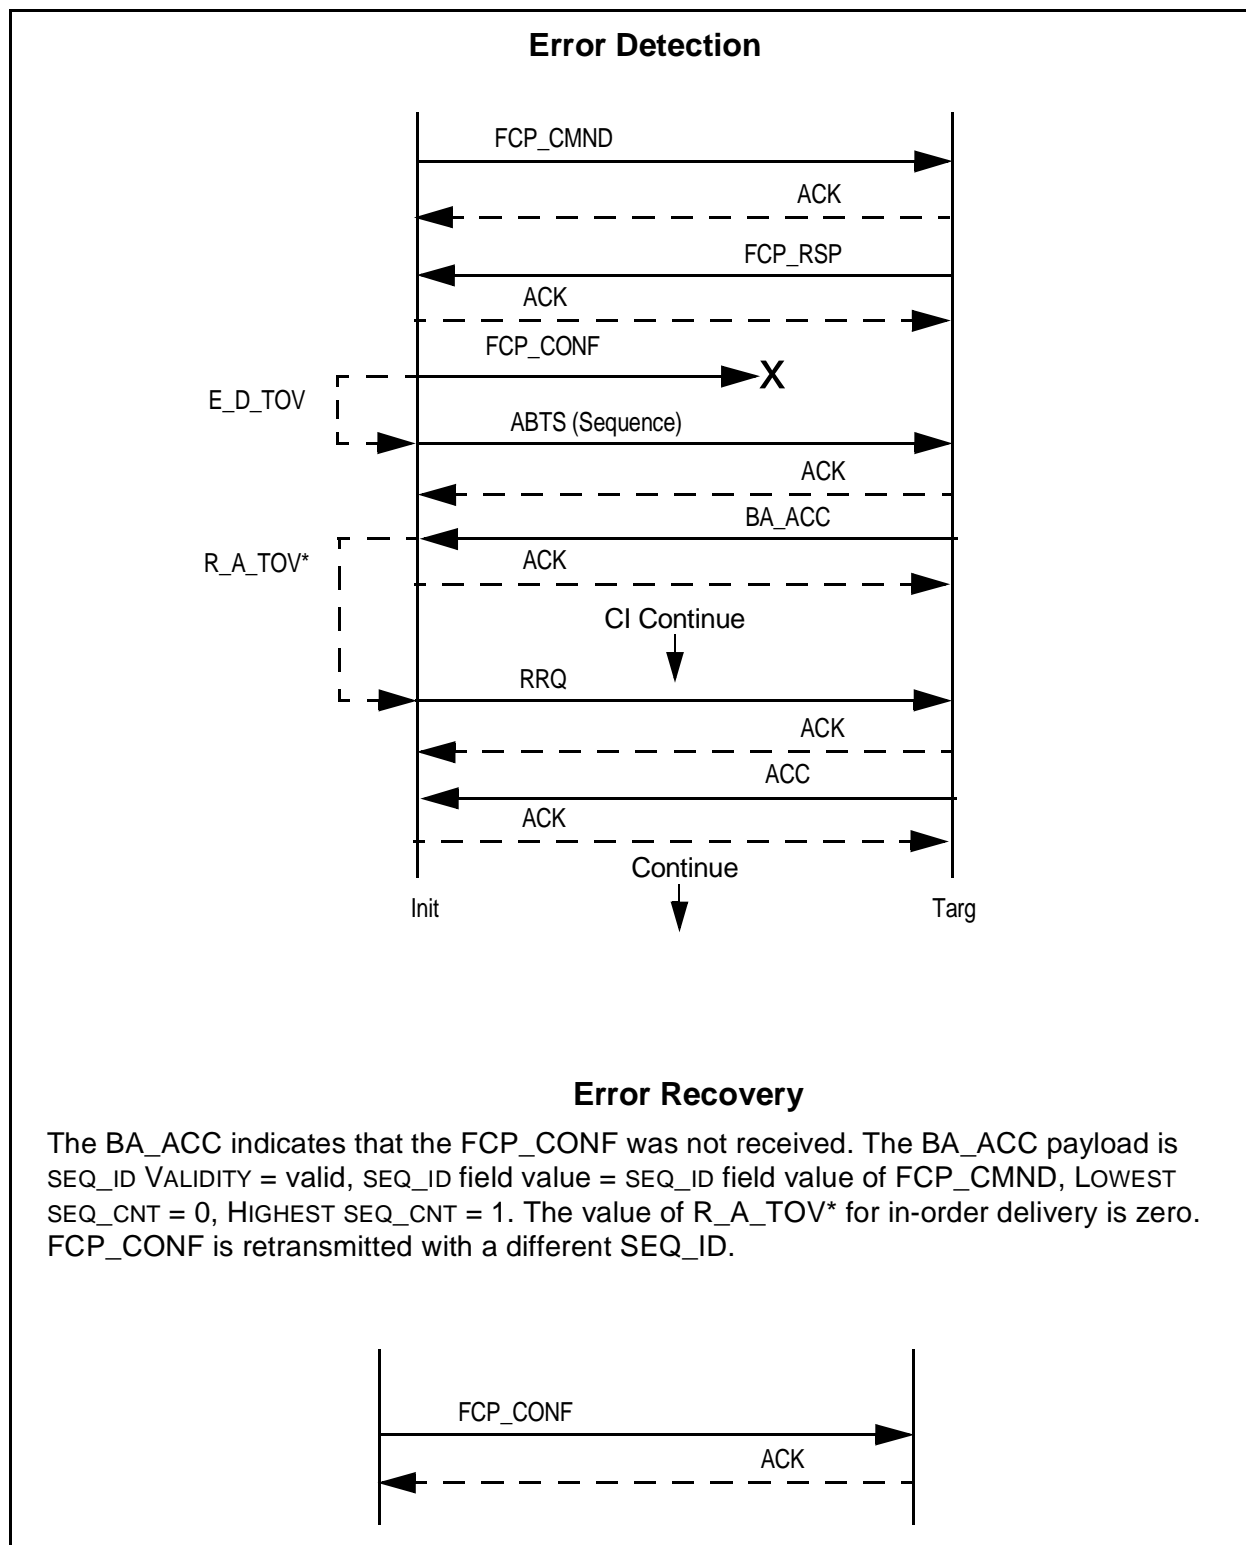
Figure C.24 - FCP_CONF Lost, Acknowledged Classes


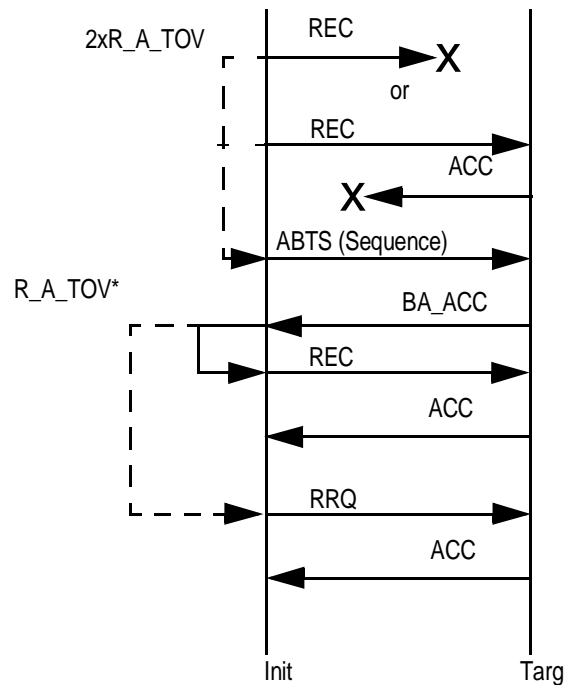
Figure C.25 - ACK lost on FCP_CONF, Acknowledged Classes

Figure C.26 - REC or REC Response Lost, Unacknowledged Classes**Error Recovery Explanation**

For the case of the REC never having been received, the BA_ACC payload is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = SEQ_CNT field value of ABTS = 1.

For the case of the ACC response to REC never having been received, the target FCP_Port views the ABTS as having been issued on a new Exchange. The BA_ACC payload is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = HIGHEST SEQ_CNT = SEQ_CNT field value of ABTS.

In both cases, a Recovery Qualifier is established. The second REC is issued in a new Exchange. For in-order delivery, the value of R_A_TOV* is zero.

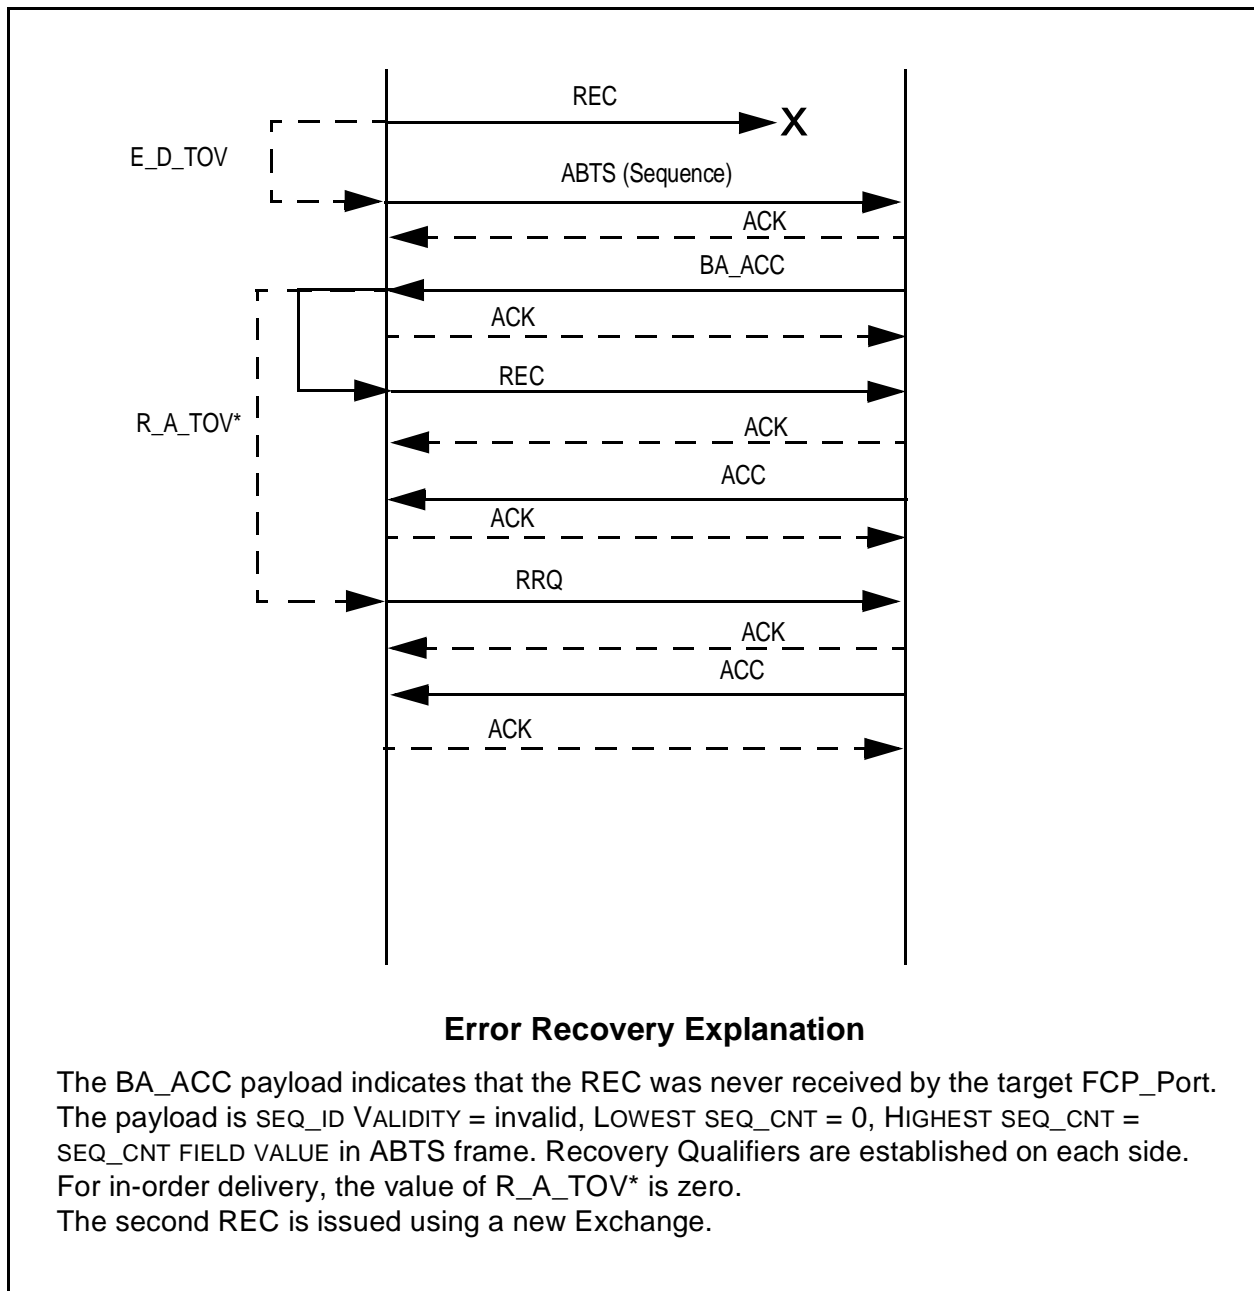
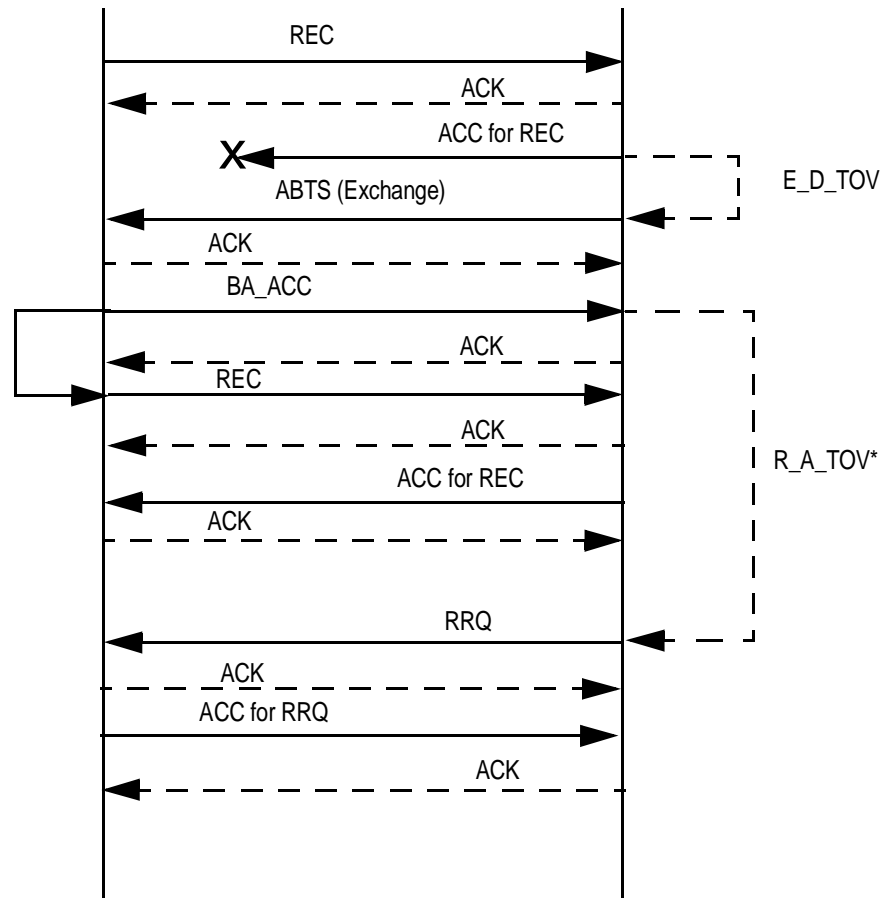
Figure C.27 - REC Lost, Acknowledged Classes

Figure C.28 - REC Response Lost, Acknowledged Classes**Error Recovery Explanation**

The BA_ACC payload indicates that the ACC 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. After responding to the ABTS, the initiator FCP_Port reissues the REC in a new Exchange.

Recovery Qualifiers are established on each side.


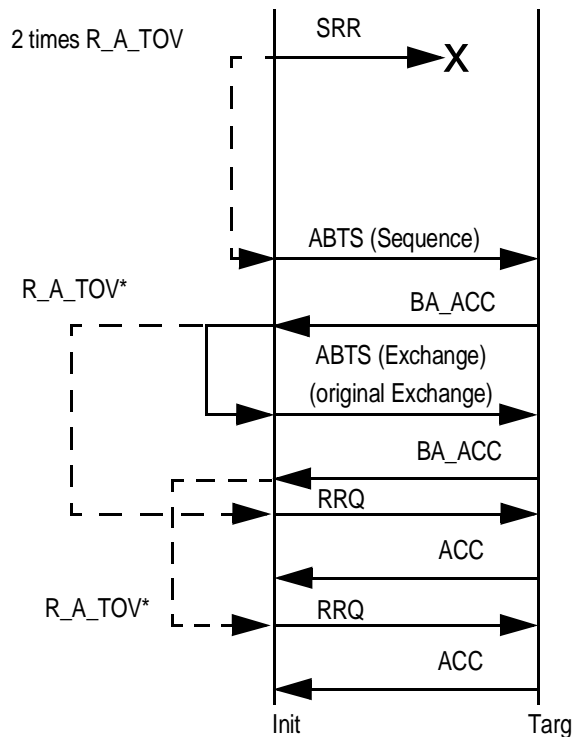
Figure C.29 - Two RECs Lost, Unacknowledged Classes, Abort the original Exchange

Figure C.30 - SRR Lost, Unacknowledged Classes, Abort original Exchange

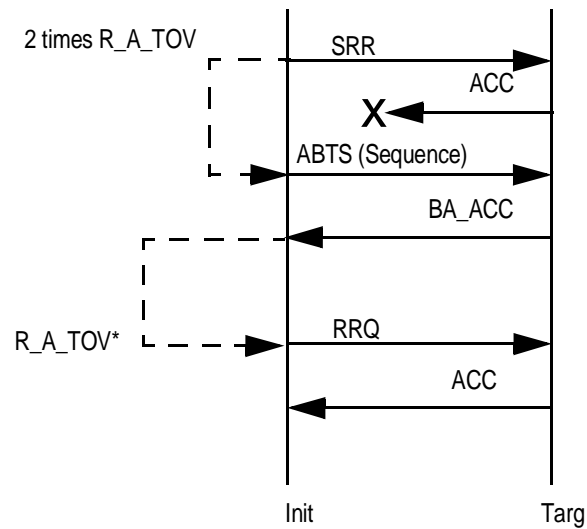
Error Recovery Explanation

The payload for the BA_ACC associated with the ABTS of the SRR is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = SEQ_CNT field value of the ABTS frame.

The ABTS for the original Exchange uses the previous SEQ_ID field value and a SEQ_CNT field value one greater than the count used in the previous Sequence and Bit 0 = 0 in the PARAMETER field. The payload for the BA_ACC associated with the ABTS for the original Exchange is SEQ_ID VALIDITY = valid, SEQ_ID = SEQ_ID of the last deliverable Sequence of the original Exchange received, LOWEST SEQ_CNT = 0, and HIGHEST SEQ_CNT = FFFFh.

Recovery Qualifiers are established on both sides for each Exchange.

For in-order delivery, the value of R_A_TOV* is zero.

Figure C.31 - SRR Response Lost, Unacknowledged Classes**Error Recovery Explanation**

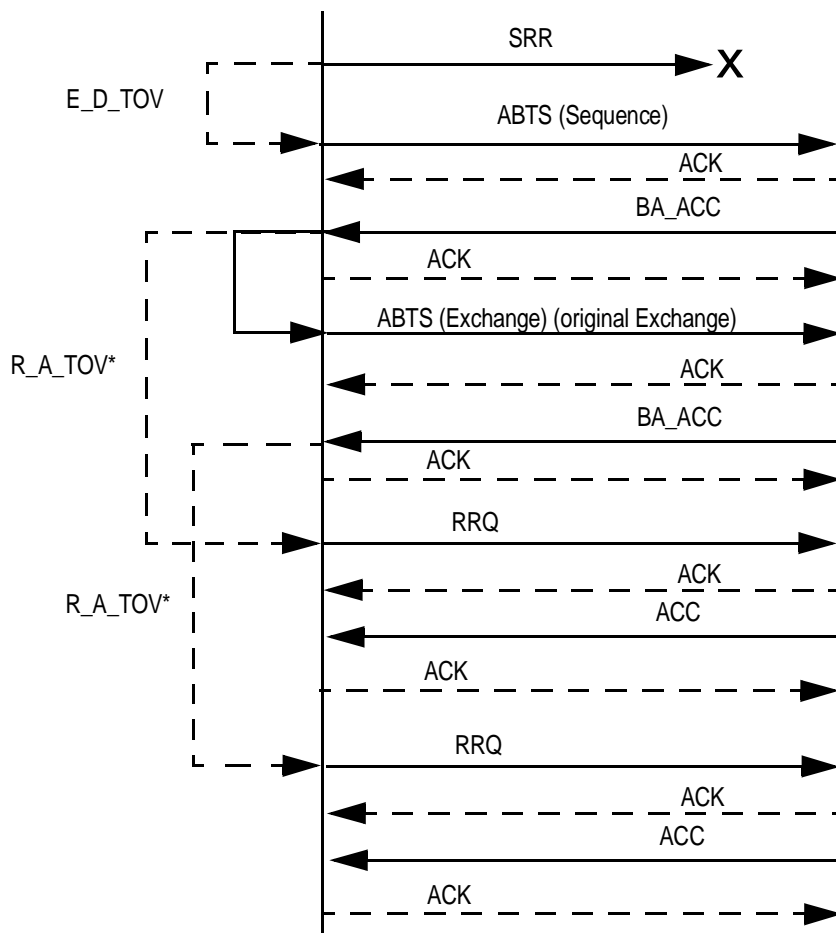
If the SRR Exchange is unknown to the Recipient, the Exchange was completed and the context purged. The payload for the BA_ACC is SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, HIGHEST SEQ_CNT = FFFFh. Recovery Qualifiers are established on both sides.

If the SRR Exchange is still known to the Recipient, the payload for the BA_ACC is SEQ_ID VALIDITY = valid, SEQ_ID = SEQ_ID field value of the SRR, LOWEST SEQ_CNT = HIGHEST SEQ_CNT = SEQ_ID field value of the ABTS frame. Since no Recovery Qualifier is established, RRQ need not be issued. The Recovery Qualifier is established on the initiator FCP_Port side and is timed out for R_A_TOV.

For in-order delivery, the value of R_A_TOV* is zero.

In either case, the original Exchange need not be aborted.

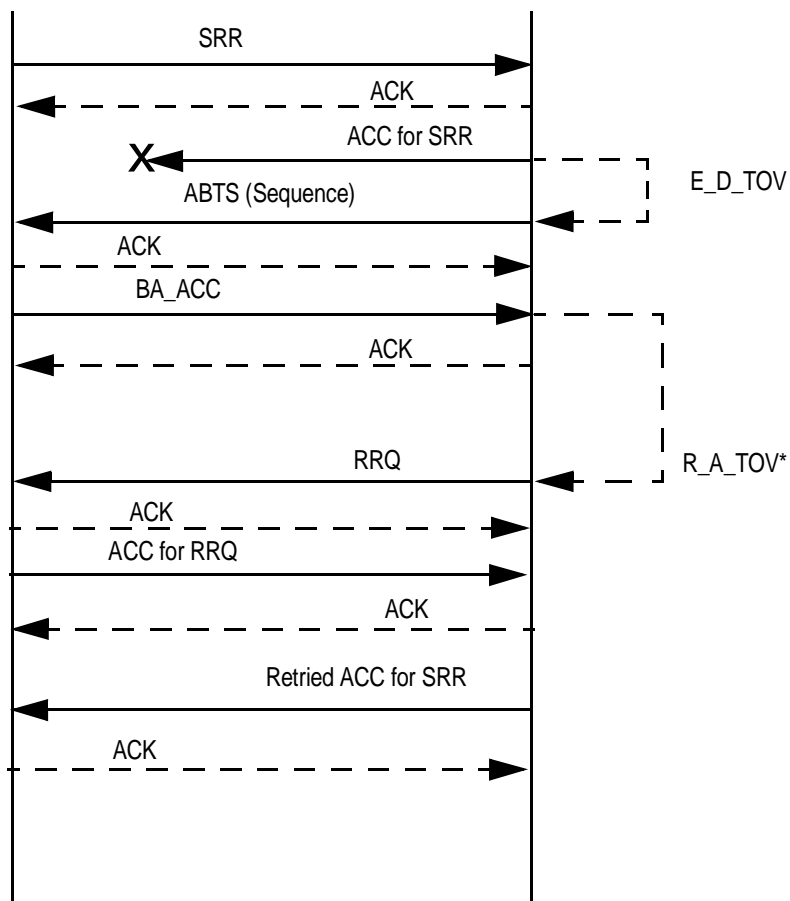
The RRQ references the Exchange of the SRR.

Figure C.32 - SRR Lost, Acknowledged Classes**Error Recovery Explanation**

Since the ABTS on the SRR is on a new Exchange, Recovery Qualifiers are established. The BA_ACC payload indicates SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, and HIGHEST SEQ_CNT = SEQ_CNT field value of the ABTS.

An error on an SRR is a second error and the original Exchange is also aborted. SRR is not retried.

For in-order delivery, the value of R_A_TOV* is zero.

Figure C.33 - SRR Response Lost, Acknowledged Classes**Error Recovery Explanation**

The BA_ACC of the ABTS associated with the SRR indicates that the ACC for the SRR was not received and is discarded if it is later received. The BA_ACC payload indicates SEQ_ID VALIDITY = invalid, LOWEST SEQ_CNT = 0, and HIGHEST SEQ_CNT = SEQ_CNT field value of the ABTS.

The retry of the ACC for SRR is issued with a new SEQ_ID.

Recovery Qualifiers are established on each side.

For in-order delivery, the value of R_A_TOV* is zero.

Annex D

(informative)

FCP Device Discovery Procedure

D.1 FCP Device Discovery Procedure

D.1.1 Initiator discovery of Fabric-attached target FCP_Ports

The following procedure may be used by initiator FCP_Ports for discovering and authenticating FCP devices in a Fabric switch environment. The FCP device discovery procedure also applies to NL_Ports attached to an FL_Port that provides the Name Server capability described in FC-GS-4. The Node_Name and Port_Name are Name_Identifiers that are Worldwide_Names, assuring that they are uniquely identifiable. Logical units are also assigned a Worldwide_Name that may be examined using the INQUIRY command with the EVPD bit set to one and the PAGE CODE set to 83h (i.e., the Device Identification VPD page) (see SPC-3).

Depending on the specific configuration and the management requirements, any step other than step 2 and step 3 may be omitted and may be performed using actions outside this standard or the referenced standards.

- 1) Perform FLOGI;
- 2) Login with the Name Server;
- 3) Register information with Name Server:
 - a) FC-4 TYPEs object (see 7.2); and
 - b) FC-4 Features object (see 7.3).
- 4) Register for State Change Notification with the Fabric Controller (see FC-LS);
- 5) Issue a GID_FF query to the Name Server with the Domain_ID Scope and Area_ID Scope fields set to zero, the FC-4 Feature bits field set to 01h, and the TYPE Code field set to 08h. This query obtains a list of the Port Identifiers of devices that support the FCP protocol and FCP target function (see FC-GS-4);
- 6) For each Port Identifier returned in the Accept CT_IU for the GID_FF, perform login and Process Login, and issue an INQUIRY command to LUN 0 to identify the type of target (see 6.2 and SPC-3);
- 7) If the INQUIRY succeeds, issue a REPORT LUNS command to LUN 0 to obtain a list of the logical units accessible through the target FCP_Port (see SPC-3);
- 8) Issue an INQUIRY command for each reported LUN to determine the type of peripheral device and supported command set for the logical unit (see SPC-3); and
- 9) Issue an INQUIRY command with the EVPD bit set to one and the PAGE CODE set to 83h for each reported LUN to obtain the logical unit's Worldwide_Name. This allows higher level applications to identify possible redundant paths to a logical unit (see SPC-3).

D.1.2 Initiator discovery of loop-attached target FCP_Ports

The following procedure may be used by initiator FCP_Ports for discovering and authenticating FCP devices in a loop environment where no fabric switch is attached to the loop.

The following steps are all optional. Depending on the specific configuration and the management requirements, any step may be omitted and any step may be performed using actions outside this standard or the referenced standards.

- 1) Obtain a map of the loop or poll all possible addresses if a loop map is not available to identify devices that are present on the loop (see FC-AL-2);

- 2) For each loop ID found in step 1, perform login and Process Login, and if the device is determined to be an FCP target, issue an INQUIRY command to LUN 0 (see 6.2 and SPC-3);
- 3) If the INQUIRY command succeeds, issue a REPORT LUNS command to LUN 0 to obtain a list of the logical units accessible through the target FCP_Port (see SPC-3); and
- 4) Issue an INQUIRY command with the EVPD bit set to one and the PAGE CODE set to the Device Identification VPD page for each reported LUN to obtain the logical unit's Worldwide_Name. This allows higher level programs to identify possible redundant paths to a logical unit (see SPC-3).

D.2 Fabric and Device Authentication

The following mechanisms are used by any Fibre Channel device to verify its relationship with other devices attached to the fabric. Such verification may be required after initialization or other temporary fabric disturbances. The following steps are all optional. Depending on the specific configuration and the management requirements, any step may be omitted by any of the attached devices and any step may be performed using actions outside this standard or the referenced standards.

- 1) N_Ports or NL_Ports retain the Fabric Port Name and Fabric Name of the Fabric from the information exchanged during FLOGI and associate that information with the Loop Fabric Address of that Fabric. This information is retained by the N_Port or NL_Port for as long as the login with the Fabric is active (see FC-FS-2);
- 2) All N_Ports and NL_Ports, including initiator FCP_Ports and target FCP_Ports, validate the current Fabric login following every Loop Initialization by comparing the Loop Fabric Address, Fabric Port Name, and Fabric Name received during FLOGI with those reported by the FAN performed during the initialization. If all three identifiers reported by the FAN do not match the values reported during FLOGI, a configuration change has occurred and an explicit LOGO is performed and all open Exchanges are terminated (see 4.10 and FC-DA);
- 3) N_Ports and NL_Ports retain the Node Name and Port Name of the other port from each PLOGI and associate that information with the Address Identifier of that port. This information is retained for as long as the PLOGI with the other port is active (see FC-FS-2); and
- 4) Initiator FCP_Ports and target FCP_Ports validate N_Port and NL_Port logins following every Loop Initialization by comparing the Port Name, Node Name, and Address Identifier received during the PLOGI with those reported by the Name Server (see FC-GS-4) or the ADISC/ACC (see FC-LS) that follows loop initialization. If all three identifiers reported by the Name Server or ADISC/ACC do not match the values reported during PLOGI, a configuration change has occurred and a LOGO is performed and all open Exchanges with that initiator FCP_Port or target FCP_Port are terminated (see 4.10 and FC-DA).

D.3 Logical unit authentication

A logical unit's identity is optionally verified and monitored by performing an INQUIRY command with the EVPD bit set to one and the PAGE CODE set to the Device Identification VPD page to obtain the logical unit's Worldwide_Name. The same Worldwide_Name is presented by a logical unit regardless of the port by which it is accessed and the value of the LUN field used to access it.

Annex E

(informative)

FCP-3 examples of link service usage

E.1 Formats for recovery link services

Examples of the formats for recovery ELSs are described in this annex.

E.2 Abort Sequence (ABTS) Request

E.2.1 Abort Sequence (ABTS) Request fields

The initiator FCP_Port or target FCP_Port may transmit an ABTS Frame. When it does so, the specified fields should be set as shown in table E.1.

Table E.1 - ABTS Frame

	Field	Sub-field	Content
Frame Header	F_CTL	Sequence Context bit	Sequence Initiator (even though the ABTS Originator may not have Sequence Initiative for the Sequence being aborted).
		Sequence Initiative bit	Transferred, even if the ABTS Originator did not hold Sequence Initiative prior to the ABTS.
	SEQ_ID		If ABTS is sent by the Sequence Initiator and the Sequence is still open, the SEQ_ID field value of the open Sequence in the Exchange being aborted is used. Otherwise, the SEQ_ID field value is any SEQ_ID not currently open (for any Exchange) between that pair of FCP_Ports.
	SEQ_CNT		SEQ_CNT field value of last Frame transmitted in an Open Sequence + 1. If no Sequence is open, then SEQ_CNT = zero or SEQ_CNT field value of last Frame transmitted + 1.
	OX_ID		OX_ID field value = same as that assigned by the Exchange Originator (i.e., initiator FCP_Port) for the Exchange being aborted.
	RX_ID		Set to FFFFh or the same as that assigned by the Exchange Responder (i.e., target FCP_Port) for the task being aborted.
	PARAMETER	Bit 0 = 0	Abort Exchange
		Bit 0 = 1	Abort Sequence

E.2.2 Basic Accept (BA_ACC) Frame to ABTS

An initiator FCP_Port or target FCP_Port may accept ABTS with BA_ACC. When it does so, the BA_ACC should be as shown in table E.2.

Table E.2 - BA_ACC Frame to ABTS

	Field	Sub-field	Content
Frame Header	OX_ID		OX_ID field value from ABTS Frame
	RX_ID		RX_ID field value from ABTS Frame
	F_CTL	Last_Sequence bit	Set to one for Abort Exchange or Set to zero for Abort Sequence
		Sequence Context bit	Recipient
Payload	SEQ_ID VALIDITY		Set to 00h for Abort Exchange or Set to 80h for Abort Sequence
	SEQ_ID		Invalid (don't care) for Abort Exchange or Set to SEQ_ID field value of last deliver- able Sequence received from ABTS Originator for Abort Sequence
	OX_ID		OX_ID field value from ABTS Frame
	RX_ID		RX_ID field value from ABTS Frame
	LOWEST SEQ_CNT		Set to 0000h for Abort Exchange or Refer to FC-FS-2 for Abort Sequence
	HIGHEST SEQ_CNT		Set to FFFFh for Abort Exchange or Set to SEQ_CNT field value of the ABTS Frame for Abort Sequence

E.2.3 Basic Reject (BA_RJT) Frame to ABTS

An initiator FCP_Port or target FCP_Port may reject ABTS with BA_RJT. When it does, the BA_RJT should be as shown in table E.3.

Table E.3 - BA_RJT Frame to ABTS

	Field	Sub-field	Content
Frame Header	OX_ID		OX_ID field value from ABTS Frame
	RX_ID		RX_ID field value from ABTS Frame
	F_CTL	Last_Sequence bit	1
		Sequence Context bit	Recipient
Payload	Byte 0	Reserved	00h
	Byte 1	Reason Code	refer to FC-LS
	Byte 2	Reason Code Explanation	refer to FC-LS
	Byte 3	Vendor Unique	00h

E.3 Reinstate Recovery Qualifier (RRQ)

E.3.1 RRQ request format

The format of the RRQ is shown in table E.4.

Table E.4 - Reinstate Recovery Qualifier

	Field	Content
Frame Header	OX_ID	Identifier of a new Exchange
	RX_ID	FFFFh
Payload	ORIGINATOR S_ID	Source_ID of the initiator FCP_Port
	OX_ID	OX_ID field value of Exchange that was previously aborted with ABTS
	RX_ID	RX_ID field value of Exchange that was previously aborted with ABTS

Following successful completion of the RRQ, the target FCP_Port responds with ACC.