

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

dpANS INCITS.xxx-200x
INCITS T11/Project 1466-D/Rev 6.0

FIBRE CHANNEL

**Backbone
(FC-BB-2)**

REV 6.0

INCITS working draft proposed
American National Standard
for Information Systems

February 4 2003

Secretariat: Information Technology Industry Council

NOTE: This is a working draft proposed American National Standard of the InterNational Committee for Information Technology Standards INCITS. As such, this is not a complete standard. The IINCITS T11.3 Technical Committee may modify this document as a result of comments received during public review and its approval as a standard.

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 rights are reserved. Any duplication of this document for commercial or for-profit use is strictly prohibited.

POINTS OF CONTACT:

Robert Snively (T11 Chairman)
Brocade Communications Systems, Inc.
1745 Technology Drive
San Jose, CA 95110
Phone: (408) 487-8135
Fax: (408) 392-6655
E-Mail: rsnively@brocade.com

Craig Carlson (T11.3 Chairman)
QLogic Corporation
6321 Bury Drive
Eden Prairie, MN 55346
Phone: (952) 932- 4064
Fax:(952) 932- 4037
E-Mail: craig.carlson@qlogic.com

Steve Wilson (Working Group Chairman)
Brocade Communications Systems, Inc.
1745 Technology Drive
San Jose, CA 95110
Phone: (408) 487-8128
Fax: (408) 392-6655
E-Mail: swilson@brocade.com

Murali Rajagopal (Technical Editor)
SV Systems
518 Valley Way
Milpitas, CA 92626
Phone: (949) 733-8412
Fax: (949) 733-8412
E-Mail: muralir@cox.net

Rev 6.0:

| - Fixed Bookmarks

draft proposed American National Standard
for Information Technology

Fibre Channel — Backbone
(FC-BB-2)

Secretariat

Information Technology Industry Council

Approved (not yet approved)

American National Standards Institute, Inc.

Abstract

This standard defines the functions and mappings necessary to bridge or tunnel between physically separate instances of the same network definition.

American National Standard

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

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

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give interpretation on any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION: 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.

PATENT STATEMENT

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.

Published by

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

Copyright © 2002 by American National Standards Institute
All rights reserved

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of the publisher.

Printed in the United States of America

Foreword (This Foreword is not part of American National Standards INCITS xxx-200x)

This Fibre Channel Backbone standard (FC-BB-2) defines the functions and mappings necessary to bridge (tunnel) between physically separate instances of the same network definition.

This standard was developed by Task Group T11.3 of the InterNational Committee for Information Technology Standards (INCITS) during 2000-2002. The standards approval process started in 2002. This standard includes annexes that are informative and are not considered part of the standard.

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

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

(to be filled in by INCITS)

Technical Committee T11 on Lower Level Interfaces, which reviewed this standard, had the following members:

Robert Snively, Chair
Edward L. Grivna, Vice Chair
Neil T. Wanamaker, Secretary

(to be filled in by INCITS)

Company	Name
----------------	-------------

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

Craig W. Carlson, Chair
George Penokie, Vice Chair
Bill Martin, Secretary

(to be filled in by INCITS)

Company	Name
----------------	-------------

Introduction

FC-BB-2 is one of the Fibre Channel family of standards. This family includes T11/Project 1331D, FC-FS, that specifies the Framing and Signalling Interface and ANSI NCITS 352, FC-PI, that specifies the Physical Interface. ANSI X3.289, FC-FG, and ANSI NCITS 355, FC-SW-2, are documents related to Fabric requirements. ANSI NCITS 348, FC-GS-3, specifies generic services. ANSI NCITS 332. FC-AL-2, specifies the arbitrated loop topology.

FC-BB-2 specifies mechanisms that allow extension of Fibre Channel switched networks across distances. FC-BB-2 defines three distinct Fibre Channel backbone mappings: FC over ATM, FC over SONET, and FC over TCP/IP.

Acknowledgements

The technical editor would like to thank the following individuals for their special contributions to this report:

Raj Bhagwat
Craig Carlson
Claudio DeSanti
Andy Helland
Don Fraser
Larry Lamers
Bret Ketchum
Scott Kipp
Bill Krieg
Charles Monia
Mike O'Donnell
Dave Peterson
Anil Rijhsinghani
Elizabeth Rodriguez
Bob Snively
Neil Wanamaker
Ralph Weber
Steve Wilson

Table of Contents

1	Scope	1
2	Normative References	3
2.1	Overview	3
2.2	Approved references	3
2.3	References under development	3
2.4	Other references	3
3	Definitions and Conventions	5
3.1	Common definitions	5
3.2	FC-BB-2_ATM Definitions	6
3.3	FC-BB-2_SONET Definitions	8
3.4	FCIP and TCP/IP Definitions	9
3.5	Editorial conventions	10
3.6	List of commonly used acronyms and abbreviations	11
3.6.1	General	11
3.6.2	FC-BB-2_ATM	12
3.6.3	FC-BB-2_SONET	12
3.6.4	FC-BB-2_IP	12
3.7	Symbols	12
3.8	Keywords	13
4	FC-BB-2 Structure and Concepts	15
4.1	FC-BB-2 Backbone Mappings	15
4.2	FC-BB-2 Reference Models	15
4.3	FC-BB-2 Specifications Overview	16
4.3.1	FC-BB-2_ATM	16
4.3.2	FC-BB-2_SONET	17
4.3.3	FC-BB-2_IP	17
4.4	FC-BB-2 Requirements	17
4.4.1	Fibre Channel Class support	17
4.4.2	Payload transparency	18
4.4.3	Latency delay and timeout value	18
4.4.4	QoS and Bandwidth	18
4.4.5	In-order delivery	18
4.4.6	Flow control	18
4.5	FC-BB-2 Link Service Codes	19

5 FC-BB-2_ATM and FC-BB-2_SONET messages and formats	20
5.1 Applicability	20
5.2 Message Formats	20
5.2.1 LLC/SNAP Header format	20
5.2.2 BBW_Header format	21
5.2.3 BBW message payload format for SFC	22
5.2.4 BBW message payload format for SR	22
5.2.4.1 SR_Header Formats	22
5.2.4.1.1 Field Formats	22
5.2.4.1.2 Information transfer I-format	23
5.2.4.1.3 Supervisory S-format	24
5.2.4.1.4 Unnumbered U-format	24
5.2.4.2 SR_BBW messages	24
5.2.4.3 Format Field Parameters	25
5.2.4.3.1 Modulus of SR	25
5.2.4.3.2 Send state variable V(S)	25
5.2.4.3.3 Send Sequence Number N(S)	25
5.2.4.3.4 Receive State Variable V(R)	25
5.2.4.3.5 Receive Sequence Number N(R)	25
5.2.4.3.6 Functions of the Poll/Final Bit (P/F)	25
5.2.4.4 SR Commands and Responses	25
5.2.4.4.1 Information (SR_I) command	25
5.2.4.4.2 Receive Ready (SR_RR) Command and Response	26
5.2.4.4.3 Receive Not Ready (SR_RNR) Command and Response	26
5.2.4.4.4 Selective Reject (SR_SREJ) Response	26
5.2.4.4.5 Set Mode (SR_SM) Command	27
5.2.4.4.6 Disconnect (SR_DISC) Command	27
5.2.4.4.7 Unnumbered Acknowledgement (SR_UA) Response	28
5.2.4.4.8 Disconnected Mode (SR_DM) Response	28
5.2.4.4.9 Message Reject (SR_FRMR) Response	28
5.2.4.5 Exception condition reporting and recovery	28
5.2.4.5.1 Exception Conditions	28
5.2.4.5.2 Busy Condition	29
5.2.4.5.3 N(S) Sequence Error Condition	29
5.2.4.5.3.1 SR_SREJ recovery	30

5.2.4.5.3.2	Time-out Recovery	30
5.2.4.5.4	Invalid message condition	30
5.2.4.5.5	Message rejection condition	30
6	The SR and SFC Protocol Procedures	32
6.1	Applicability	32
6.2	SR Protocol Overview	32
6.3	Description of the SR procedure	33
6.3.1	SR mode of operation	33
6.3.2	SR procedure for addressing	33
6.3.3	SR procedure for the use of the P/F bit	33
6.3.4	SR procedure for data link set-up and disconnection	33
6.3.4.1	Data link set-up	33
6.3.4.2	Information transfer Phase	34
6.3.4.3	Data link disconnection	34
6.3.4.4	Disconnected Phase	35
6.3.4.4.1		35
6.3.4.4.2		35
6.3.4.5	Collision of unnumbered commands	35
6.3.4.5.1		35
6.3.4.5.2		36
6.3.4.6	Collision of SR_DM response with SR_SM or SR_DISC command	36
6.3.4.7	Collision of SR_DM responses	36
6.3.5	Procedures for information transfer using multi-selective reject	36
6.3.5.1	Procedures for SR_I messages	36
6.3.5.2	Sending new SR_I messages	36
6.3.5.3	Receiving an in-sequence SR_I message	36
6.3.5.4	Reception of invalid messages	37
6.3.5.5	Reception of out-of-sequence SR_I messages	37
6.3.5.6	Receiving acknowledgement	38
6.3.5.7	Receiving a SR_SREJ response message	38
6.3.5.7.1	Receiving a SR_SREJ response message with the F bit set to 0	38
6.3.5.7.2	Receiving a SR_SREJ response message with the F bit set to 1	38

6.3.5.8	Receiving a SR_RNR message	39
6.3.5.9	BBW busy condition	39
6.3.5.10	Awaiting acknowledgement	40
6.3.5.11	Receiving a SR_RR response messages with the F bit set to 1	40
6.3.5.12	Responding to command messages with the P bit set to 1	41
6.3.6	SR conditions for data link resetting or data link re-initialization (data link set-up)	41
6.3.6.1	41
6.3.6.2	41
6.3.6.3	41
6.3.6.4	41
6.3.7	SR procedure for data link resetting	41
6.3.7.1	41
6.3.7.2	42
6.3.7.3	42
6.3.8	List of SR system parameters	43
6.3.8.1	Timer T1	43
6.3.8.2	Parameter T2	43
6.3.8.3	Maximum number of attempts to complete a transmission N2	43
6.3.8.4	Maximum number of outstanding SR_I messages <i>k</i>	43
6.4	Simple Flow Control (SFC)	44
7	FC-BB-2_ATM Structure and Concepts	45
7.1	Applicability	45
7.2	FC-BB-2_ATM Overview	45
7.3	FC-BB-2_ATM Functional Model	46
7.3.1	B_Port Network Interface	46
7.3.2	ATM Network Interface	46
7.3.3	Mapping and Encapsulation	46
7.3.4	FC-BB-2_ATM Forwarding	47
7.3.5	Call Handling and ATM Service	47
7.3.6	Frame Handling	48
8	Mapping and Message Encapsulation using AAL5	49
8.1	Applicability	49
8.2	Overview	49
8.3	Mapping BBW messages to AAL5	49

9 FC-BB-2_ATM Service Considerations	52
9.1 Applicability	52
9.2 ATM Service Type	52
9.3 Latency Delay and Timeout Value	52
9.4 Bandwidth Sharing and Allocation	52
9.5 Quality of Service (QoS)	53
9.6 Delivery Order	53
9.7 Loss and Flow Control	54
10 FC-BB-2_SONET Structure and Concepts	55
10.1 Applicability and Related Clauses	55
10.2 FC-BB-2_SONET Overview	55
10.3 FC-BB-2_SONET Functional Model	56
10.3.1 Fibre Channel Network Interface	56
10.3.2 SONET Network Interface	57
10.3.3 Mapping and Encapsulation	58
10.3.4 FC-BB-2_SONET Forwarding	58
10.3.5 Call Handling	58
10.3.6 Frame Handling	58
11 Mapping and Message Encapsulation using HDLC-like Framing	59
11.1 Applicability	59
11.2 Overview	59
11.3 Mapping of BBW messages to HDLC format	59
11.4 Mapping HDLC frames to SONET/SDH	61
12 FC-BB-2_SONET Service Considerations	65
12.1 Applicability	65
12.2 Latency Delay and Timeout Value	65
12.3 Delivery Order	65
12.4 Loss and Flow Control	65
13 FC-BB-2_IP Structure and Concepts	66
13.1 Applicability and Related Clauses	66
13.2 FC-BB-2_IP Overview	66
13.3 The VE_Port Functional Model	67
13.3.1 FC-BB-2_IP Interface Protocol Layers	67
13.3.2 E_Port/F_Port FC Interface	67
13.3.3 FC-BB-2_IP Protocol Interface	67

13.3.3.1	Major Components	67
13.3.3.2	FC Switching Element (SE) with FC Routing	68
13.3.3.3	FC and FCIP Entities	69
13.3.3.3.1	Function	69
13.3.3.3.2	FC Entity	69
13.3.3.3.3	FCIP Entity	70
13.3.3.3.4	VE_Port Virtual ISL and FCIP Link	70
13.3.3.4	VE_Port Virtual ISL Exchanges	71
13.3.3.4.1	SW_ILS Exchanges	71
13.3.3.4.2	Link Keep Alive (LKA) ELS Exchanges	71
13.3.3.5	Control and Service Module (CSM)	72
13.3.3.6	Platform Management Module (PMM)	73
13.3.3.6.1	Function	73
13.3.3.6.2	Time Synchronization	73
13.3.3.6.2.1	FCIP Transit Time (FTT)	73
13.3.3.6.2.2	Building outgoing FC frame encapsulation headers	73
13.3.3.6.2.3	Checking IP network transit times in incoming FC frame encapsulation headers.	73
13.3.3.6.3	Discovery	74
13.3.3.6.4	Security	74
13.3.4	IP Network Interface	75
13.4	The B_Access Functional Model	75
13.4.1	FC-BB-2_IP Interface Protocol Layers	75
13.4.2	B_Port FC Interface	75
13.4.3	FC-BB-2_IP Protocol Interface	76
13.4.3.1	Major Components	76
13.4.3.2	FC and FCIP Entities	76
13.4.3.2.1	Function	76
13.4.3.2.2	FC Entity	76
13.4.3.2.3	FCIP Entity	76
13.4.3.2.4	B_Access Virtual ISL and FCIP Links	76
13.4.3.3	B_Access Virtual ISL Exchanges	78
13.4.3.3.1	Exchange B_Access Parameters (EBP) SW_ILS Exchanges	78

13.4.3.3.2	B_Access Link Keep Alive (LKA) ELS Exchanges . . .	80
13.4.3.3.2.1	B_Access Initialization State Machine	80
13.4.3.4	B_Port Control and Service Module (CSM)	81
13.4.3.5	B_Port Platform Management Module (PMM)	81
13.4.4	IP Network Interface	81
13.5	FC-BB-2_IP Network Topologies	82
14	Mapping and Message Encapsulation using TCP/IP	84
14.1	Applicability	84
14.2	Encapsulated Frame Structures	84
14.2.1	FC frame Encapsulation Structure	84
14.2.2	Encapsulated FCIP Special Frame (FSF) structure	85
14.3	TCP/IP Encapsulation	86
15	The FC-BB-2_IP Protocol Procedures	87
15.1	Applicability	87
15.2	FC-BB-2_IP Protocol Procedures	87
15.3	Procedures for Platform Management	87
15.3.1	Function	87
15.3.2	Procedures for Discovery	87
15.3.3	Procedures for Extending FC-SP Security	87
15.3.3.1	Authentication Mechanisms	87
15.3.3.2	Authenticate Special Frame (ASF)	88
15.4	Procedures for Connection Management	89
15.4.1	Function	89
15.4.2	Procedures for Link Setup	89
15.4.3	Procedures for Data Transfer	90
15.4.4	Procedures for FCIP Link Disconnection	90
15.4.5	Procedures for Multiple Connection Management	91
15.5	Procedures for Error Detection Recovery	91
15.5.1	Procedures for Handling Invalid FC frames	91
15.5.2	Procedures for Error Recovery	91
15.6	List of FC-BB-2_IP System Parameters	92
15.6.1	FC Timers	92
15.6.2	TCP Timers	92
15.6.3	Maximum number of attempts to complete an Encapsulated FC Frame transmission	92

15.6.4	Maximum number of outstanding Encapsulated FC Frames . . .	92
16	FC-BB-2_IP Service Considerations	93
16.1	Applicability	93
16.2	Latency Delay	93
16.3	Throughput	93
16.3.1	How Timeouts affect Throughput	93
16.3.2	How loss affects Throughput	93
16.3.3	Other factors that affect Throughput	93
16.4	Reliability	94
16.4.1	Loss of Connectivity	94
16.4.2	Loss of Synchronization	94
16.4.3	Loss or Corruption of TCP Segments	94
16.4.4	Loss or Corruption of FC frames	94
16.4.5	FCIP Error Reporting	94
16.5	Quality of Service (QoS)	95
16.6	Delivery Order	95
16.7	IP Multicast and Broadcast	95
16.8	Security and Authentication	95
Annex A:	Encoded SOF and EOF Ordered Sets	96
A. 1	Ordered Sets	96
A.2	SOF and EOF OS-Codes	96
A.3	FC-BB-2 SOF OS codes	97
A.4	FC-BB-2 EOF OS codes	97
Annex B:	ATM Traffic Management and Signaling 98	
B.1	ATM Traffic Management	98
B.2	ATM QoS Parameters	98
B 2.1	Maximum Cell Transfer Delay (maxCTD) - Negotiated Parameter .	99
B 2.2	Peak-to-peak Cell Delay Variation (CDV)- Negotiated Parameter. .	99
B 2.3	Cell Loss Ratio (CLR) - Negotiated Parameter	100
B 2.4	Cell Error Ratio (CER) - Non-negotiated parameter	100
B 2.5	Severely Errored Cell Block Rate (SECBR) - Non-negotiated parameter	100
B 2.6	Cell Misinsertion Ratio (CMR) - Non-negotiated parameter	100
B.3	ATM Traffic Parameters	101

B.3.1 CDVT	101
B.3.2 Conformance Definition	101
B.3.3 Source Traffic Descriptor	102
B.4 ATM Service Categories	102
B.4.1 Real-time Services	102
B.4.2 Non-real-time Services	103
B.5 ATM Adaptation Layer (AAL) Types	104
B.6 ATM Multiplexing	104
B.7 Summary of ATM Services and Guarantees	105
B.8 ATM Signaling UNI Standard	105
B.8.1 Switched Virtual Connections	106
B.8.2 Address resolution	106
B.8.3 SVC UNI messages	106
B8.3.1 Message types	106
B8.3.2 Message information elements	106
B.9 References	107
Annex C: SR Protocol Parameter Guidelines and State Diagram	108
C.1 Assumptions	108
C.2 Calculated Ack Delay (T2) Timer variable	108
C.3 Calculated Window size	108
C.4 Calculated Ack (T1) Timer variable	108
C.5 SR Protocol State Diagram	109

List of Tables

Table 1 FC-BB-2 Organization	2
Table 2 ISO and American Conventions	11
Table 3 Specification and FC Port Types Supported	15
Table 4 FC-BB-2 SW_ILS Codes	19
Table 5 FC-BB-2 ELS Codes	19
Table 6 BBW message structure.	20
Table 7 LLC/SNAP Header	20
Table 8 SNAP PID	21
Table 9 BBW_Header	21
Table 10 Flow Control Protocol Type Encodings	21
Table 11 BBW message payload structure for SFC	22
Table 12 BBW message payload structure for SR	22
Table 13 SR_Header format	23
Table 14 SS bits encoding	23
Table 15 MMMMM bit encoding	23
Table 16 SR_BBW messages	24
Table 17 SR_I message format	26
Table 18 SR_SREJ payload format	27
Table 19 SR_FRMR payload format	29
Table 20 Mapping of BBW messages to AAL5 CPCS	49
Table 21 ATM VBR-NRT Service Specification	53
Table 22 SONET/SDH Data Rates	57
Table 23 Mapping of BBW messages to HDLC format	60
Table 24 FC-BB-2_SONET Protocol Stack	62
Table 25 LKA payload	72

Table 26 LKA Accept payload	72
Table 27 EBP Request payload	78
Table 28 EBP Accept payload.	79
Table 29 EBP Reject Reason Code Explanation	79
Table 30 TCP/IP Segment structure carrying Encapsulated FC Frame	84
Table 31 Encapsulated FC Frame structure	84
Table 32 TCP/IP Segment structure carrying Encapsulated FSF	85
Table 33 Encapsulated FSF structure	85
Table 34 ASF Request Payload	88
Table 35 ASF Accept Response Payload	89
Table A.1 Byte-encoded frame delimiter format	96
Table A.2 OS-Code Definition	96
Table A.3 FC-BB-2 SOF Codes	97
Table A.4 FC-BB-2 EOF Codes	97
Table B.1 I.356 defined QoS Parameters for different Traffic Classes	101
Table B.2 Service Categories and its Traffic and QoS Attributes	103
Table B.3 ATM Service Categories and Guarantees	105

List of Figures

Fig. 1 Scope and components of FC-BB-2 specification	1
Fig. 2 FC-BB-2_ATM Reference Model	15
Fig. 3 FC-BB-2-SONET Reference Model	16
Fig. 4 FC-BB-2_IP Reference Model	16
Fig. 5 SR Flow Control Protocol Between two BBWs	32
Fig. 6 FC-BB-2_ATM Network Configuration	45
Fig. 7 FC-BB-2_ATM Functional Block Diagram	47
Fig. 8 AAL5 Mapping of a BBW message with SFC	50
Fig. 9 AAL5 Mapping of a BBW message with SR	51
Fig. 10 Recommended ATM Bandwidth Allocation for multiple VCs	52
Fig. 11 FC-BB-2_SONET Network Configuration	55
Fig. 12 FC-BB-2_SONET Functional Block Diagram	57
Fig. 13 SONET SPE HDLC Mapping Example	62
Fig. 14 Path Signal label: C2	62
Fig. 15 Encapsulation of BBW message into HDLC frame using SFC	63
Fig. 16 Encapsulation of BBW message into HDLC frame using SR	64
Fig. 17 FC-BB-2_IP Network Configuration	66
Fig. 18 FC-BB-2_IP VE_Port Functional Model	68
Fig. 19 FC-BB-2_IP Protocol Layers	69
Fig. 20 Scope of VE_Port Virtual ISL	71
Fig. 21 Security Layers	74
Fig. 22 FC-BB-2_IP B_Access Functional Model	77
Fig. 23 Scope of B_Access Virtual ISL	78
Fig. 24 B_Access Initialization State Machine	81
Fig. 25 FC-BB-2_IP Network Topologies	83
Fig. 26 TCP/IP Encapsulation of an Encapsulated FC Frame	86

Fig. B.1 Cell Transfer Delay Distribution	100
Fig. B.2 SVC Signaling at the UNI and Switched payload.	107
Fig. C.1 SR Protocol State Diagram	109

draft proposed American National Standard
for Information Technology –

Fibre Channel – Backbone (FC-BB-2)

1 Scope

This American National Standard for Fibre Channel Backbone (FC-BB-2) consists of three distinct Fibre Channel mappings resulting in the following three specifications:

- FC-BB-2_ATM (FC over ATM backbone network)
- FC-BB-2_SONET (FC over SONET backbone network)
- FC-BB-2_IP (FC over TCP/IP backbone network)

Figure 1 illustrates the scope and the major components of the FC-BB-2 specification and its relationship to the FCIP [10] specification and the ATM Forum/ITU-T standards. Note 1 shows the organization of the document. FC-BB-2_IP, FC-BB-2_ATM and FC-BB-2_SONET do not interoperate in any way and are independent specifications

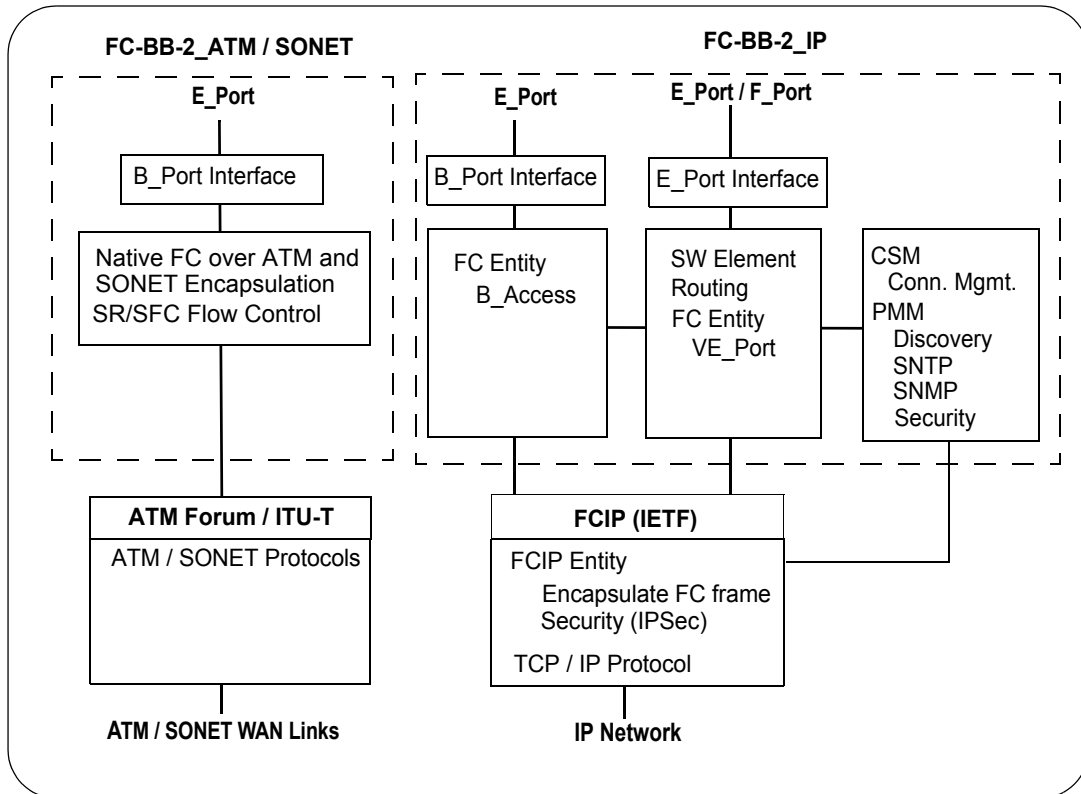


Figure 1 – Scope and components of FC-BB-2 specification

Table 1 – FC-BB-2 Organization

Specification Type	Applicable Clauses and Annexes
FC-BB-2_ATM, FC-BB-2_SONET, FC-BB-2_IP	1-4
FC-BB-2_ATM, FC-BB-2_SONET	5, 6
FC-BB-2_ATM	7, 8, 9 Annexes A, B, C
FC-BB-2_SONET	10, 11, 12 Annexes A, C
FC-BB-2_IP	13, 14, 15, 16 Annex A

2 Normative References

2.1 Overview

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

For electronic copies of ANSI standards listed here, visit ANSI's Electronic Standards Store (ESS) at:
<http://www.ansi.org>

For printed versions of ANSI standards listed here, contact:
Global Engineering Documents,
15 Inverness Way East, Englewood, CO 80112-5704
(800) 854-7179

Additional availability contact information is provided below as needed.

2.2 Approved references

- [1] ANSI T1.101-1999, *Telecommunications - Synchronization Interface Standard for Digital Networks*
- [2] ANSI T1.105-1995, *Telecommunications - Synchronous Optical Network (SONET) - Basic Description Including Multiplex Structures, Rates, and Formats*
- [3] ANSI T1.105.02-1995, *Telecommunications - Synchronous Optical Network (SONET) - Payload Mappings.*
- [4] ANSI X3.289-1996, *Information Technology - Fibre Channel - Fabric Generic Requirements (FC-FG)*
- [5] ANSI INCITS 342-2000, *Information Technology - Fibre Channel - Switch Fabric (FC-BB)*

2.3 References under development

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

- [6] T11/Project 1508D/Rev. 6.1, *Information Technology - Fibre Channel - Switch Fabric (FC-SW-3)*
- [7] T11/Project 1505D/Rev. 7.4, *Information Technology - Fibre Channel - Generic Services-4 (FC-GS-4)*
- [8] T11/Project 1331D/Rev. 1.7, *Information Technology - Fibre Channel - Framing and Signaling (FC-FS)*
- [9] T11/Project 1570D/Rev. x.x, *Information Technology - Fibre Channel - Security Protocol (FC-SP)*
- [10] IETF draft-ietf-ips-fcovertcpip-12.txt, *Fibre Channel Over TCP/IP (FCIP)*, August 2002
- [11] IETF draft-ietf-ips-fcip-slp-04.txt, *Finding FCIP Entities Using SLP*, September 2002
- [12] IETF draft-ietf-ips-fcencapsulation-08.txt, *FC frame Encapsulation*, May 2002

2.4 Other references

- [13] IETF draft-ietf-ips-security-16.txt, *Securing iSCSI, iFCP and FCIP*, September 2002
- [14] ATM Forum (1994), *ATM User Network Interface (UNI) Version 3.1. AF-UNI-0010.002.*

- [15] ATM Forum (1996), *ATM User Network Interface (UNI) Version 4.0*. AF-SIG-0061.000.
- [16] ATM Forum (1996), *Native ATM Services: Semantic Description Version 1.0*. AF-SAA-0048.000.
- [17] CCITT Recommendation G.707, *Synchronous Digital Hierarchy Bit Rates* (1990).
- [18] ISO/IEC 3309-1993, *Information Technology - Telecommunications and Information Exchange Between Systems - High-level Data Link Control (HDLC) Procedures - Frame Structure*.
- [19] ITU-T X.25-1997, *Interface between Data terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuit, X.25-1997*
- [20] ISO/IEC 13239:1997(E), *Information technology - Telecommunications and information exchange between systems - High-level data link control (HDLC) procedures*
- [21] ITU-T Q.2931 (1995), *Broadband Integrated Services Digital Network (B-ISDN) Digital Subscriber Signaling System No. 2 (DSS2); User-Network Interface (UNI) Layer 3 Specification for Basic Call/Connection Control*.
- [22] ITU-T Q.2971 (1995), *Broadband Integrated Services Digital Network (B-ISDN) Digital Subscriber Signaling System No. 2 (DSS2); User-Network Interface (UNI) Layer 3 Specification for Point-to-Multipoint Call/Connection Control*.
- [23] ITU-T Recommendation G.708, *Network Node Interface for the Synchronous Digital Hierarchy*, 1993.
- [24] ITU-T Recommendation G.709, *Synchronous Multiplexing Structure*, 1993.
- [25] IETF RFC 1619, *PPP over SONET/SDH*, May 1994.
- [26] IETF RFC 1662, *PPP in HDLC-like Framing*, July 1994.
- [27] IETF RFC 2030, *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*, October 1996
- [28] ITU-T Rec. G.7041/Y.1303, (2001), *Generic Framing Procedure (GFP)*

3 Definitions and Conventions

3.1 Common definitions

3.1.1 BBW: Refers to either FC-BB-2_ATM or FC-BB-2_SONET device

3.1.2 B_Port: A Bridge port on a device that implements FC-BB-2_ATM, FC-BB-2_SONET, and FC-BB-2_IP and connects to a FC switch on an E_Port

3.1.3 B_Port_Name: A Name_Identifier that identifies a B_Port for identification purposes. The format of the name is specified in FC-FS. Each B_Port shall provide a unique B_Port_Name within the Fabric.

3.1.4 BSW: Defined in FC-SW-3 and used as a generic term for a backbone Switch

3.1.5 E_Port: As defined in FC-SW-3, a Fabric expansion port that attaches to another E_Port to create an Inter-Switch Link.

3.1.6 E_Port_Name: A Name_Identifier that identifies an E_Port for identification purposes. The format of the name is specified in FC-FS. Each E_Port shall provide a unique E_Port_Name within the Fabric.

3.1.7 Fabric_Name: A Name_Identifier associated with a Fabric.

3.1.8 F_Port: As defined in FC-SW-3, a F_Port is a port to which non-loop N_Ports are attached to a Fabric, and does not include FL_Ports.

3.1.9 F_Port_Name: A Name_Identifier that identifies a F_Port for identification purposes. The format of the name is specified in FC-FS. Each F_Port shall provide a unique F_Port_Name within the Fabric.

3.1.10 Fabric Initialization: A process for configuring and building a Fabric, as defined in FC-SW-3.

3.1.11 FC-BB-2_ATM: A physical ATM WAN interface specification that interfaces with Fibre Channel switched network on one side and ATM on the other.

3.1.12 FC-BB-2_SONET: A SONET/SDH WAN interface specification that interfaces with Fibre Channel switched network on one side and SONET/SDH on the other.

3.1.13 FC-BB-2_IP: A logical specification that interfaces with Fibre Channel switched network on one side and the IP network on the other.

3.1.14 FC_Port: A port transmitting or receiving FC frames. FC_Ports include N_Ports, F_Ports, E_Ports, B_Ports, VE_Ports, and B_Access.

3.1.15 Fibre Channel BackBone Link: Defined to be any Fibre Channel Inter-switch Link over a non Fibre Channel transport. In FC-BB-2, this encompasses FC-BBW_ATM, FC-BBW_SONET and FC-BB-2_IP links. Note that a Fibre Channel BackBone Link may be comprised of more than one physical or logical connection.

3.1.16 Keep Alive Timeout Value (K_A_TOV): The Keep Alive Timeout value is a timer defined in this document that is used by the Link Keep Alive (LKA) ELS as a trigger for issuing LKA. The LKA should be sent at least every K_A_TOV if no traffic has been sent and/or received on the connection. The default value for K_A_TOV is 1/2 E_D_TOV.

3.1.17 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 (e.g. N_Port, node, F_Port, or Fabric.).

3.1.18 Node Name: A Name_Identifier associated with a node.

3.1.19 Port_Name: An 8-byte identifier that identifies a port and used for such purposes as diagnostics that may be independent and unrelated to network addressing. Each FC_Port shall provide an unique Port_Name within the address domain of the Fabric and associated N_Ports.

3.1.20 Simple Flow Control (SFC): SFC is a flow control protocol applied between two FC-BB-2_ATM or FC-BB-2_SONET devices over the ATM/SONET WAN. The SFC protocol mechanism temporarily pauses the transmission of frames from the remote device.

3.1.21 Selective Retransmission (SR) Flow Control: SR Flow Control is a sliding window flow control protocol applied between two FC-BB-2_ATM or FC-BB-2_SONET devices over the ATM/SONET WAN and is used for both flow control and error recovery.

3.1.22 Switch_Name: A Name_Identifier that identifies a Switch or a Bridge device for identification purposes. The format of the name is specified in FC-FS. Each Switch and Bridge device shall provide an unique Switch_Name within the Fabric.

3.1.23 WAN Interface: An interface that connects to a Wide Area Network; specifically, a port that connects to ATM or SONET/SDH.

3.2 FC-BB-2_ATM Definitions

3.2.1 AAL: ATM Adaptation Layer. A collection of standardized protocols that adapt user traffic to 48-octet payloads that may be placed in a cell-formatted stream. The AAL is subdivided into the Convergence Sublayer (CS) and the Segmentation and Reassembly (SAR) sublayer. There are currently four types of AALs: AAL1, AAL2, AAL3/4, and AAL5 to support the various service categories.

3.2.2 AAL5: AAL Type 5 and specified for use in FC-BB-2_ATM. Protocol Standard originally intended for variable bit rate traffic, that does not require source-destination timing relation, is now also used with applications that have constant bit rate traffic where source-destination timing relation is important.

3.2.3 AAL Service Categories: The ATM Forum has defined 5 Traffic Service Categories supported by the AALs: Constant Bit Rate (CBR), Variable Bit Rate-Real Time (VBR-RT), Variable Bit Rate-Non Real Time (VBR-NRT), Available Bit Rate (ABR), and Unspecified Bit Rate (UBR). VBR_NRT is the recommended service for FC-BB.

3.2.4 ATM: Asynchronous Transfer Mode. Broadband-ISDN standards defined by ITU-T and the ATM Forum.

3.2.5 ATM QoS Parameters: QoS is a term used to refer to the set of performance characteristics of the contracted ATM connection. Six QoS parameters are defined: Peak-to-peak Cell Delay Variation (CDV), Maximum Cell Transfer Delay (maxCTD), Cell Loss Ratio (CLR), Cell Error Ratio (CER), Severely Errored Cell Block Ratio (SECBR), Cell Misinsertion Rate (CMR).

3.2.6 ATM Traffic Descriptor: A term used to describe the traffic characteristics of an ATM connection. A Connection Traffic Descriptor includes a Source Traffic Descriptor, CDV Tolerance (CDVT), and a Conformance definition. A Source Traffic Descriptor is described by three parameters: Peak Cell Rate (PCR), Sustainable Cell Rate (SCR), Maximum Burst Size (MBS), and a Minimum Cell Rate (MCR).

3.2.7 Cell Loss Priority (CLP): A one-bit field in the ATM cell header specifying whether the cell is more (CLP=1) or less (CLP=0) likely to be discarded by an ATM network experiencing congestion.

3.2.8 Cell Loss Ratio (CLR): A QoS parameter that gives the ratio of lost cells to the total number of transmitted cells.

3.2.9 Connection Admission Control (CAC): Actions taken by the ATM network to accept or reject a connection request based on its QoS and traffic parameter requirements and then to route this connection across the network.

3.2.10 Convergence Sublayer Protocol Data Unit: The PDU used at the CS for passing information between the higher layers and the SAR (that is located below the CS, where the cell conversion takes place).

3.2.11 GCRA: Generic Cell Rate Algorithm is a method applied at the network side of the UNI to test the conformance of an ATM cell to its traffic contract.

3.2.12 Operations, Administration, and Maintenance (OAM): Management framework defined by the ITU. OAM cells are special purpose ATM cells exchanged between an ATM end-system and an ATM switch and between ATM switches. OAM cells are used for network fault and performance management and analysis.

3.2.13 Permanent Virtual Circuit (PVC): A preconfigured logical connection between two ATM systems.

3.2.14 Permanent Virtual Connection (PVC): The ATM term for a Permanent Virtual Circuit between ATM switches. The terms may be used interchangeably.

3.2.15 Quality of Service: A term that refers to the set of ATM performance parameters that characterize the transmission quality over a given virtual connection (VC). These parameters include the CTD, CDV, CER, CLR, CMR, and SECBR.

3.2.16 Switched Virtual Call: A generic term that refers to Switched Virtual Circuits and connections.

3.2.17 Switched Virtual Circuit: A logical ATM connection established via signaling. End systems transmit UNI 3.1/4.0 signaling request via the Q.2931 Signaling Protocol.

3.2.18 Switched Virtual Connection (SVC): The ATM term for switched virtual circuit.

3.2.19 Usage Parameter Control (UPC): A set of policing mechanisms implemented by the network at the UNI to monitor and control traffic submitted by each end user.

3.2.20 User-Network Interface (UNI): The interface, defined as a set of protocols and traffic characteristics, between the customer premises equipment and the ATM networks. FC-BB requires ATM Forum UNI 3.1 or above.

3.2.21 Virtual Channel: A term used to describe one of several logical connections defined within one virtual path (VP) between two ATM devices.

3.2.22 Virtual Channel Connection (VCC): Defined as a concatenation of Virtual Channel Links (VCLs). Switching cells within an ATM switch for a given VCC is based on the VPI/VCI value indicated on the cell header.

3.2.23 Virtual Circuit (VC): A connection that is set up across the network between a source and a destination where a fixed route is chosen for the entire session and bandwidth and ID dynamically allocated to the user.

3.2.24 Virtual Path (VP): A logical connection between two ATM devices (CPEs, switches). A virtual path consists of a set of virtual channels.

3.2.25 Virtual Path Connection (VPC): Defined as a concatenation of virtual path links (VPLs). Switching cells within an ATM switch for a given VPC is based on the VPI value indicated on the cell header.

3.2.26 Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI): The combination of two numbers, namely the VPI and the VCI, used to identify a virtual connection (VC) and switch cells in an ATM network.

3.3 FC-BB-2_SONET Definitions

3.3.1 Administrative Unit (AU): A SDH-specific information structure, consisting of a STS SPE and its associated set of STS pointer/pointer action bytes.

3.3.2 Concatenated Synchronous Transport Signal Level N (STS- N c): A STS- N Line layer signal in which the STS Envelope Capacities from the N STS-1s have been combined to carry a STS- N c Synchronous Payload Envelope (SPE) that shall be transported not as several separate signals but as a single entity. The equivalent SDH term for a STS-3c SPE is a VC-4.

3.3.3 Container: A SDH term that is equivalent to the payload capacity of a synchronous payload envelope.

3.3.4 Generic Framing Procedure (GFP): A simple robust encapsulation method for frame-organized data traffic specified by ITU-T Study Group 15 [28].

3.3.5 Operations, Administration, and Maintenance (OAM): Management framework defined by the ITU. OAM cells are special purpose ATM cells exchanged between an ATM end-system and an ATM switch and between ATM switches. OAM cells are used for network fault and performance management and analysis.

3.3.6 Optical Carrier Level N (OC- N): The optical signal that results from an optical conversion of a STS- N signal. SDH does not make the distinction between a logical signal (e.g. STS-1 in SONET) and a physical signal (e.g. OC-1 in SONET). The equivalent SDH term for both logical and physical signals is synchronous transport module level M (STM- M), where $M=(N/3)$. There are equivalent STM- M signals only for values of $N=3,12,48$, and 192.

3.3.7 Path: A logical connection between the point at which a standard frame format for the signal at the given rate is assembled, and the point at which the standard frame format for the signal is disassembled. The equivalent SDH term is also Path.

3.3.8 Payload Pointer: The pointer that indicates the location of the beginning of the Synchronous Payload Envelope. The equivalent SDH term is pointer.

3.3.9 SONET: An acronym for Synchronous Optical NETWORK. SONET is a term in general usage, that refers to the rates and formats specified in ANSI T1.105

3.3.10 STS Path Terminating Equipment (STS PTE): Network Elements that multiplex/demultiplex the STS payload. STS PTEs may originate, access, modify, or terminate the STS Path Overhead necessary to transport the STS payload, or may perform any combination of these actions.

3.3.11 Super-rate Signals: A signal that has to be carried by a Concatenated Synchronous Transport Signal level N c (STS- N c). There is no equivalent SDH term.

3.3.12 Synchronous Digital Hierarchy (SDH): A family of ITU-T standards whose technical contents closely resemble that found for the SONET family of ANSI standards.

3.3.13 STS Synchronous Payload Envelope (STS SPE): A 125-microsecond frame structure composed of STS Path Overhead and bandwidth for payload. The term generically refers to STS-1 SPEs and STS- N c SPEs. The equivalent SDH term for STS-1 SPE is virtual container level 3 (VC-3). The equivalent SDH term for STS-3c SPE is virtual container level 4 (VC-4). The equivalent SDH term for STS- N c SPE ($N>3$) is virtual container level 4- X c (VC-4- X c), where $X=(N/3)$.

3.3.14 Synchronous Transport Module Level M (STM-M): These are the defined transport signals for the Synchronous Digital Hierarchy (SDH). Defined signals exist at rates of M times 155.52 Mbps, where M=1, 4, 16, or 64. These are equivalent to SONET OC-N signals, where N=3M.

3.3.15 Synchronous Transport Signal Level N (STS-N): This signal is obtained by byte interleaving N STS-1 signals together. The rate of the STS-N is N times 51.840 Mbps. SDH does not make the distinction between a logical signal (e.g. STS-N in SONET) and a physical signal (e.g. OC-N in SONET). The equivalent SDH term for both logical and physical signals is synchronous transport module level M (STM-M), where $M=(N/3)$. There are equivalent STM-M signals only for values of N=3, 12, 48, and 192.

3.3.16 Tributary Unit (TU): The SDH term for SONET virtual tributaries.

3.3.17 Virtual Container (VC): A SDH term for either a STS or VT SPE.

3.3.18 Virtual Tributary (VT): A structure designed for transport and switching of sub-STN-1 payloads. There are currently four sizes of VT. The equivalent SDH term is tributary unit (TU).

3.4 FCIP and TCP/IP Definitions

3.4.1 B_Access: A component of the FC Entity that interfaces with the FCIP_LEP component of the FCIP Entity on one side and the B_Port on the other.

3.4.2 B_Access_Name: Port_Name of the B_Access portal.

3.4.3 B_Access Virtual ISL: A Virtual ISL that connects two B_Access portals.

3.4.4 Control and Service Module (CSM): A control component of the FC-BB-2_IP interface that mainly handles Connection Management. CSM interfaces with the PMM.

3.4.5 Encapsulated FC Frame: A FC-BB-2_IP term to mean a SOF/EOF delimited FC frame prefixed with a 28-byte FC frame Encapsulation Header.

3.4.6 Encapsulated Frame Receiver Portal: The TCP access point through which an Encapsulated FC Frame is received from the IP network by a FCIP_DE.

3.4.7 Encapsulated Frame Transmitter Portal: The TCP access point through which an Encapsulated FC Frame is transmitted to the IP network by the FCIP_DE.

3.4.8 FC Entity: The FC Entity is the principal interface point to the FC switched network on one side and in combination with the FCIP Entity to the IP network on the other side. It is the data forwarding component of the FC-BB-2_IP interface consisting of VE_Port(s) and/or B_Access portals

3.4.9 FC Entity Protocol Layer: The protocol layer that lies between the Fibre Channel level FC-2 and the FCIP Entity protocol layer. Its primary function is supporting one or more Virtual E_Ports or B_Access portals and communicating with the FCIP Entity.

3.4.10 FC Receiver Portal: The access point through which a FC frame and time stamp enters a FCIP_DE from the VE_Port/B_Access.

3.4.11 FC Transmitter Portal: The access point through which a FC frame and time stamp leaves a FCIP_DE to the VE_Port/B_Access.

3.4.12 FC-BB-2_IP Device: A device that supports the FC-BB-2_IP, FC network, and the IP network interfaces.

3.4.13 FC-BB-2_IP Interface: The point that has interfaces to the FC switched network on one side and the IP network on the other. It consists of a Switching Element, FC/FCIP Entity pair(s), the CSM, and the PMM.

3.4.14 FCIP Data Engine (FCIP_DE): The data forwarding component of the FCIP Entity's FCIP_LEP that handles FC frame encapsulation, de-encapsulation, and transmission of encapsulated frames through a single TCP connection.

3.4.15 FCIP Entity: The data forwarding component of the FC-BB-2_IP interface consisting of the FCIP_LEP and is the principal interface point to the IP network on one side and in combination with the FC Entity to the FC switched network on the other. Its primary function is formatting, encapsulating, and forwarding Encapsulated FC Frames across the IP network interface.

3.4.16 FCIP Entity Protocol Layer: The protocol layer that lies between the FC Entity layer and the TCP layer.

3.4.17 FCIP frame: The FCIP [10] term for an Encapsulated FC Frame.

3.4.18 FCIP Link: A virtual link that connects a FCIP_LEP in one FC-BB-2_IP device with another. It consists of one or more TCP connections.

3.4.19 FCIP Link Originator and Acceptor: The FC-BB-2_IP FCIP_LEP that *originates* a FCIP Link is defined as the **FCIP Link Originator**. The corresponding FCIP_LEP that *accepts* this link is defined as the **FCIP Link Acceptor**.

3.4.20 FCIP Link Endpoint (FCIP_LEP): The component of a FCIP Entity that contains one or more FCIP_DEs.

3.4.21 FCIP Transit Time (FTT): The total transit time of an Encapsulated Fibre Channel frame in the IP network. The FCIP Transit Time is calculated by subtracting the time stamp value in the arriving Encapsulated FC Frame from the synchronized time in the FCIP Entity.

3.4.22 Platform Management Module (PMM): A management component of the FC-BB-2_IP interface that handles Time Synchronization, Discovery, and Security. It interfaces with the CSM.

3.4.23 Request For Comment (RFC): Documents put out by the Internet-Related organizations.

3.4.24 Virtual E_Port (VE_Port): The data forwarding component of the FC Entity that emulates an E_Port. The term *Virtual* indicates the use of a non Fibre Channel link connecting the VE_Ports. In the case of the FC-BB-2_IP specification, a VE_Port interfaces with the FCIP_LEP component of the FCIP Entity on one side and a Fibre Channel Switching Element on the other side.

3.4.25 VE_Port_Name: Port_Name of the VE_Port.

3.4.26 VE_Port Virtual ISL: A Virtual ISL that connects two VE_Ports.

3.4.27 Virtual ISL: An ISL that connects two VE_Ports or two B_Access portals across a non-Fibre Channel Link

3.5 Editorial conventions

In FC-BB-2, a number of conditions, mechanisms, sequences, parameters, events, states, or similar terms are printed with the first letter of each word in uppercase and the rest lowercase (e.g., Exchange, Class). Any lowercase uses of these words have the normal technical English meanings.

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

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

tions are shown in Tables 2.

Table 2 – ISO and American Conventions

ISO	American
0,6	0.6
1000	1,000
1 323 462,9	1,323,462.9

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

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

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

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

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

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

3.6 List of commonly used acronyms and abbreviations

3.6.1 General

BB	Backbone
BB-2	Backbone -2
BBW	Backbone (ATM or SONET) WAN
BSW	Border Switch
EBP	Exchange B_Access Parameters
ELP	Exchange Link Parameters
EOF	End of Frame
ESC	Exchange Switch Capabilities
FCIP	Fibre Channel Over TCP/IP [10]
FCS	Frame Check Sequence
FC-SP	Fibre Channel - Security Protocol [9]
FC-SW-3	Fibre Channel - Switched Fabric [6]
ISL	Inter-switch Link
ITU-T	Internat'l Telecomm. Union - Telecommunication Standardization Section
K_A_TOV	Keep Alive Timeout value
LKA	Link Keep Alive
MTU	Maximum Transfer Unit
PDU	Protocol Data Unit
SFC	Simple Flow Control
SOF	Start of Frame
SR	Selective Retransmission
SW_ACC	Switch Fabric Internal Link Service Accept
SW_CS	Switch Fabric Common Services
SW_ILS	Switch Fabric Internal Link Services
SW_RJT	Switch Fabric Internal Link Service Reject
WAN	Wide Area Network

3.6.2 FC-BB-2_ATM

AAL5	ATM Adaptation Layer 5
ATM	Asynchronous Transfer Mode
BER	Bit Error Rate
CLR	Cell Loss Ratio (ATM)
CPCS	Common Part Convergence Sublayer
PVC	Permanent Virtual Circuit, Permanent Virtual Connection
QoS	Quality of Service
SAAL	Signaling ATM Adaptation Layer
SVC	Switched Virtual Circuit, Switched Virtual Connection
UBR	Unspecified Bit Rate (ATM)
UNI	User Network Interface (ATM)
VBR-NRT	Variable Bit Rate - Non Real Time (ATM)
VC	Virtual Circuit

3.6.3 FC-BB-2_SONET

HDLC	High-level Data Link Control
nm	Nanometer
OC-N	Optical Carrier Level <i>N</i>
ppm	Parts per Million
PPP	Point-to-Point Protocol
PTE	Path Terminating Equipment
RFC	Request for Comment
SDH	Synchronous Digital Hierarchy
SMT	Station Management (FDDI)
SONET	Synchronous Optical Network
SPE	Synchronous Payload Envelope
STM-M	Synchronous Transport Module level <i>M</i>
STS	Synchronous Transport Signal
STS-N	Synchronous Transport Module level <i>N</i>
STS-Nc	Synchronous Transport Module level <i>Nc</i>
TU	Tributary Unit
ULA	48-bit Universal LAN MAC Address
ULP	Upper Level Protocol
ULP_TOV	Upper_Level_Protocol_Timeout value
VC	Virtual Container
VP	Virtual Path
VT	Virtual Tributary

3.6.4 FC-BB-2_IP

B_Access	B_Access Portals
CSM	Control and Service Module
FCIP	FC over TCP/IP [10]
FCIP_DE	FCIP Data Engine
FCIP_LEP	FCIP Link Endpoint
IETF	IETF Internet Engineering Task Force (www.ietf.org)
PMM	Platform Management Module
RFC	Request For Comment
VE_Port	Virtual E_Port

3.7 Symbols

Unless indicated otherwise, the following symbol has the listed meaning.

|| concatenation

3.8 Keywords

3.8.1 expected

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

3.8.2 ignored

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

3.8.3 invalid

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

3.8.4 mandatory

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

3.8.5 may

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

3.8.6 may not

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

3.8.7 obsolete

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

3.8.8 optional

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

3.8.9 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. Recipients are not required to check reserved bits, bytes, words or fields for zero values. In defined fields, receipt of reserved code values shall be reported as an error.

3.8.10 restricted

A keyword referring to bits, bytes, words, and fields that are set aside for use in other 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 technical report.

3.8.11 shall

A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure inter operability with other products that conform to this technical report.

3.8.12 should

A keyword indicating flexibility of choice with a strongly preferred alternative; equivalent to the phrase "it is strongly recommended".

3.8.13 x or xx

The value of the bit or field is not relevant.

4 FC-BB-2 Structure and Concepts

4.1 FC-BB-2 Backbone Mappings

The three distinct Fibre Channel mappings FC over ATM, FC over SONET, and FC over TCP/IP pertain to the extension of Fibre Channel switched networks across distances. An important distinction between the above mappings is the emphasis placed on the *backbone-type*. The FC over ATM and SONET network mappings layer on the ATM and SONET backbone technologies, resulting in the FC-BB-2_ATM and FC-BB-2_SONET specifications that map directly to physical connections. The FC over TCP/IP network mapping layers on the IP network, resulting in the FC-BB_IP specification that maps to a logical connection.

4.2 FC-BB-2 Reference Models

FC-BB-2 defines three reference models corresponding to the FC-BB-2_ATM, FC-BB-2_SONET, and FC-BB-2_IP specifications. These reference models are shown in Figures 2, 3, and 4. In the figures a B (E or F) port is the point at which a frame destined to a remote FC network enters the port and is forwarded on the backbone network to its destination. All frames arriving on the backbone network exit the B (E or F) port towards its ultimate destination.

The FC-BB-2_ATM and the FC-BB-2_SONET specifications supports the attachment of FC switches via one or more B_Ports. The FC-BB-2_IP specification supports the attachment of FC switches via one or more B or E ports and the attachment of Fibre channel Host Bus Adapters via one or more F ports. Tables 3 summarizes the FC port types support for the three types of specifications.

Table 3 – Specification and FC Port Types Supported

Port(s) Support	Specification Type		
	FC-BB-2_ATM	FC-BB-2_SONET	FC-BB-2_IP
B_Port	Specified	Specified	Specified
E_Port	Not Specified	Not Specified	Specified
F_Port	Not Specified	Not Specified	Specified

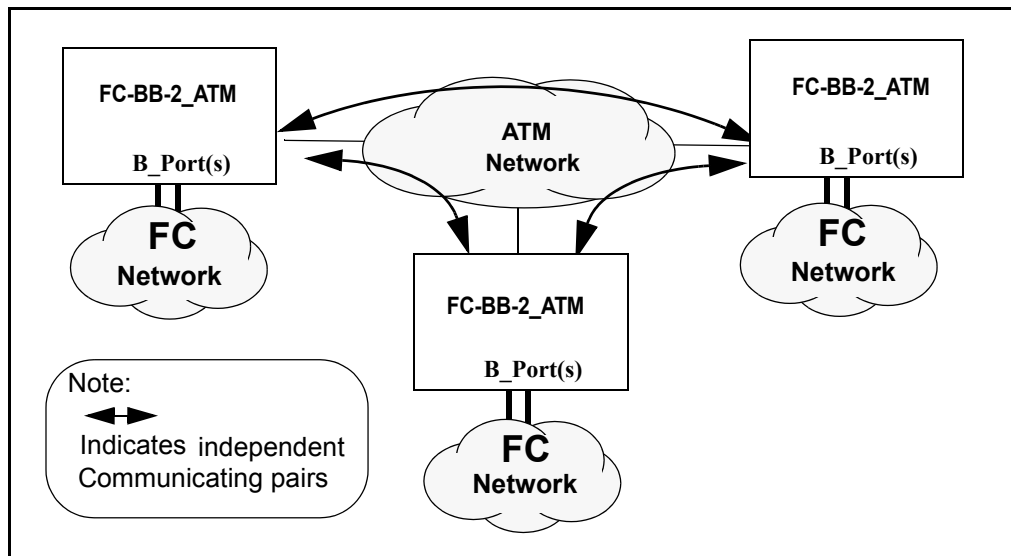


Figure 2 – FC-BB-2_ATM Reference Model

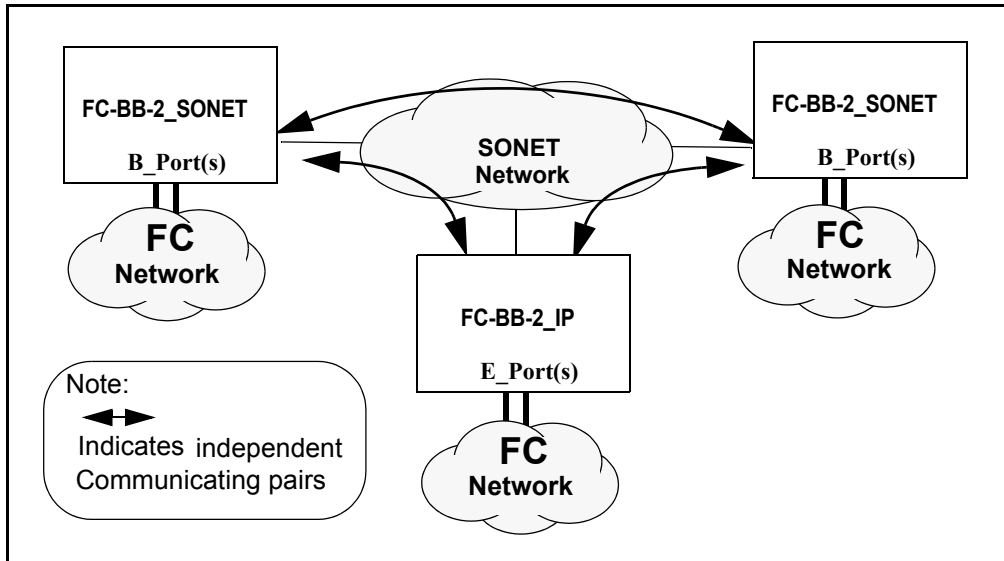


Figure 3 – FC-BB-2-SONET Reference Model.

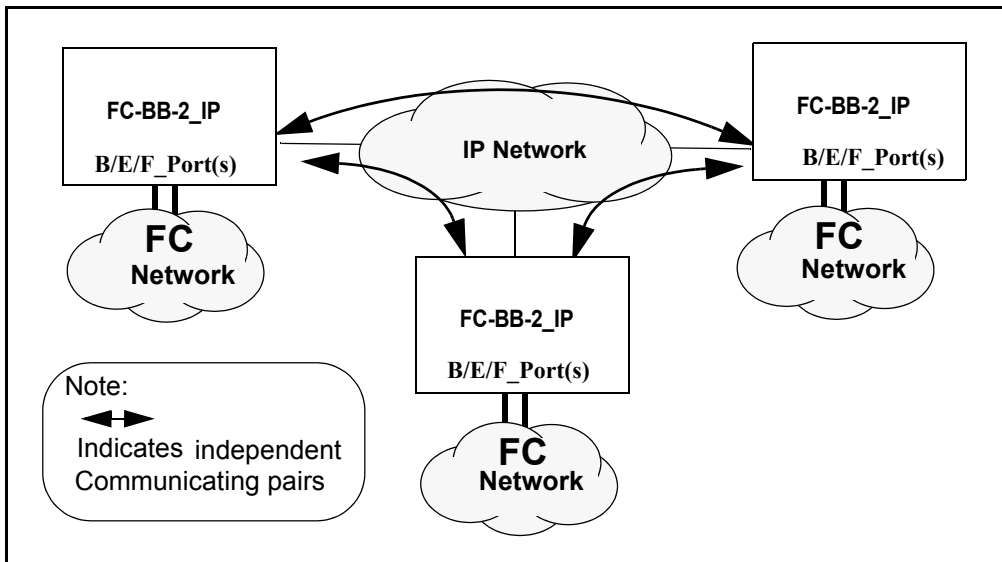


Figure 4 – FC-BB-2_IP Reference Model.

4.3 FC-BB-2 Specifications Overview

4.3.1 FC-BB-2_ATM

FC-BB-2_ATM (specification) is the means by which Fibre Channel networks interface with and connect across a wide-area ATM network. FC-BB-2_ATM defines the frame mapping, encapsulation, and any signaling required by the ATM protocols. FC-BB-2_ATM also defines the frame handling, call handling, addressing, flow control protocol and error recovery required to support the Fibre Channel mapping over ATM. FC-BB-2_ATM makes use of the ATM Adaptation Layer 5 for payload transport

FC-BB-2_ATM messages are formed by encapsulating byte-encoded Class 2, 3, 4 or F Fibre Channel frames into a suitable format for carriage over the WAN. Clause 5 describes the FC-BB-2_ATM message in detail.

The SR and SFC are two flow control protocols that may be used over the ATM networks. The SR protocol provides a reliable transport of frame between two FC-BB-2_ATM devices. Use of the SR protocol is optional. The SR protocol is an efficient sliding window link-layer full-duplex protocol that supports data transport with flow control and error recovery functions. SR has been adopted from ITU's Link Access Protocol B (LAPB), that is derived from ISO/IEC's High-level Data Link Control (HDLC) (Balanced Classes). Use of LAPB in SR is limited to a subset of the synchronous modulo 32768 super sequence numbering service option. Clause 6 describes the SR protocol in detail. The SFC Protocol (see 6.4) provides a mechanism to temporarily pause the transmission of frames from a remote FC-BB-2_ATM device. Use of the SFC protocol is optional.

4.3.2 FC-BB-2_SONET

FC-BB-2_SONET (specification) is the means by which Fibre Channel networks interface with and connect across a wide-area SONET/SDH network. FC-BB-2_SONET defines the frame mapping, encapsulation, and any signaling required by the SONET protocols. FC-BB-2_SONET also defines the frame handling, call handling, addressing, flow control protocol and error recovery required to support the Fibre Channel mapping over SONET/SDH. FC-BB-2_SONET makes use of the High-level data link control (HDLC) for payload transport

NOTE 1– The FC over SONET mapping defined in FC-BB-2_SONET is different from the native FC mapping over the Transparent Mapped Generic Framing Procedure (GFP-T) defined in [28]. These native mappings are essentially FC “wire-extenders” between two FC ports using other technologies and their protocol operations are completely transparent to the connecting FC ports. As such, these native mappings provide a transparent point-to-point Fibre Channel extension between any of the FC port types (e.g., N_Port-to-N_Port, N_Port-to-E_Port, E_Port-to-E_Port, etc.), and will not be discussed further in the FC-BB-2_IP specification.

The SR and SFC flow control protocols may be used over the SONET networks.

4.3.3 FC-BB-2_IP

FC-BB-2_IP specification is the means by which Fibre Channel networks interface with and connect across an IP network. FC-BB-2_IP makes use of the FCIP specification [10] to define the mapping and control required by the TCP/IP protocol and the FC frame Encapsulation specification [12] to define the encapsulation. FC-BB-2_IP also defines the connection management, addressing, time synchronization, discovery, security, switching, routing, and error recovery required to support Fibre Channel over TCP/IP. FC-BB-2_IP is agnostic about the underlying physical technology that exists beneath the IP layer. In this sense, the IP network could use ATM, SONET, Gigabit Ethernet or any other link level technology below it.

FC-BB-2_IP encapsulates byte-encoded Class 2, 3, 4 or F Fibre Channel frames into a suitable format (Encapsulated FC Frames) for carriage over the IP network. Clause 14 describes Encapsulated FC Frames in detail. The TCP/IP protocol suite provides a reliable transport of frames over the IP network. TCP provides flow control and error recovery.

The FC-BB-2_IP protocol provides mechanisms to create Virtual E_Port or B_Access connectivity as the case may be, over the IP network (see Clause 13).

4.4 FC-BB-2 Requirements

4.4.1 Fibre Channel Class support

- i) Class F shall be supported between FC-BB-2 devices. Class 2, 3 or 4 may be supported between FC-BB-2 devices.

NOTE 2– FC-BB-2 does not support Class 1

4.4.2 Payload transparency

- i) Arriving Class 2, 3, 4, and F Fibre Channel frames from a FC network and destined to a remote FC network shall be encapsulated using the FC-BB-2 defined mechanisms and transmitted to the appropriate FC-BB-2 device.
- ii) Arriving Encapsulated frames received from remote FC-BB-2 device shall be deencapsulated and sent to a FC network.

4.4.3 Latency delay and timeout value

- i) FC-BB-2_IP shall ensure that the incoming Encapsulated FC Frames whose FTT exceeds $1/2 E_D_TOV$ shall be discarded and not admitted into the FC network. Fibre Channel Timeout values shall be administratively set to accommodate the FTT.
- ii) FC-BB-2_IP shall allow Class F Encapsulated FC Frames to be transmitted with a zero timestamp value.

4.4.4 QoS and Bandwidth

- i) FC-BB-2_ATM shall use the VBR-NRT ATM Service. VBR-NRT ATM Service provides cell loss and bandwidth guarantees. It is recommended that FC-BB-2_ATM make use of a single Virtual Circuit (VC). Use of additional VCs to address special traffic QoS requirements is allowed but not recommended. FC-BB-2_ATM recommends allocating a minimum bandwidth for each FC-BB-2_ATM VC that is used in order to avoid starvation; however, the Service discipline (prioritization) for the VCs is implementation specific and beyond the scope of this standard.
- ii) FC-BB-2_SONET has no specific SONET service required
- iii) FC-BB-2_IP recommends that some form of preferential QoS be used for the FCIP traffic in the IP network to minimize latency and packet drops although no particular form of QoS is recommended. See [10].

4.4.5 In-order delivery

- i) FC-BB-2_ATM shall guarantee in-order delivery of frames within each ATM VC. No other ordering relationship across ATM VCs need be preserved.
- ii) FC-BB-2_SONET shall guarantee in-order delivery of frames within each SONET/SDH provisioned path. No other ordering relationship across SONET/PDH provisioned paths need be preserved.
- iii) FC-BB-2_IP shall guarantee in-order delivery of frames within the scope of any TCP connection.

4.4.6 Flow control

- i) FC-BB-2_ATM or FC-BB-2_SONET devices may use the Selective Retransmission (SR) protocol to provide a reliable delivery of frames over the WAN between two devices. In the case of FC-BB-2_ATM, if SR protocol is used, then the flow control is separately applied to each ATM VC.
- ii) FC-BB-2_ATM or FC-BB-2_SONET devices may use Simple Flow Control (SFC) protocol to temporarily pause the transmission of frames from a remote device. In the case of FC-BB-2_ATM, if SFC protocol is used, then the flow control is separately applied to each ATM VC.
- iii) FC-BB-2_IP devices using the TCP Flow control and error recovery should act in concert with the Fibre Channel Buffer-to-Buffer Credit mechanism.

- iv) Flow Control at the E_Ports, F_Ports, VE_Ports, and B_Ports will operate as defined in FC-SW-3.

4.5 FC-BB-2 Link Service Codes

Table 4 shows all the SW_ILS codes allocated for FC-BB-2 use.

Table 4 – FC-BB-2 SW_ILS Codes

Encoded Value (hex)	Description	Abbr.
28 03 00 00	Authentication Special Frame Request	ASF
28 01 00 00	Exchange B_Access Parameter	EPB

Table 5 shows all the ELS codes used in FC-BB-2.

Table 5 – FC-BB-2 ELS Codes

Encoded Value (hex)	Description	Abbr.
80 00 00 00	Link Keep Alive Request	LKA

5 FC-BB-2_ATM and FC-BB-2_SONET messages and formats

5.1 Applicability

This Clause only applies to FC-BB-2_ATM and FC-BB-2_SONET. See Clause 14 for a description of the FC-BB-2_IP messages and formats.

5.2 Message Formats

In all text to follow, the term BBW applies to both FC-BB-2_ATM and FC-BB-2_SONET. The structure of a BBW message is given in Tables 6. It consists of 3 fields: LLC/SNAP Header, BBW_Header, and the BBW message payload. The structures of the fields are given in Tables 7, 9, 11, and 12.

Table 6 – BBW message structure

	Field	Size (Bytes)
BBW message	LLC/SNAP Header	8
	BBW_Header	4
	BBW message payload	Max: 2148

5.2.1 LLC/SNAP Header format

LLC/SNAP Header (8 bytes): The Logical Link Control (LLC)/Sub Network Access Protocol (SNAP) Header consists of a 3-byte LLC field and a 5-byte SNAP sub-field.

Table 7 – LLC/SNAP Header

Field		Word	Byte	Encoded Value (hex)
LLC	DSAP	0	0	AA
	SSAP		1	AA
	CTRL		2	03
SNAP	OUI	0	3	00
		1	0	00
		1	1	00
	PID	1	2	xx
		1	3	xx

LLC (3 bytes): The LLC field consists of 3 1-byte sub-fields: Destination Service Access Point (DSAP), Source Service Access Point (SSAP), and Control (CTRL). The encoding for LLC given in Tables 7 indicates that an IEEE 802.2 SNAP follows.

SNAP (5 bytes): The SNAP field consists of 2 sub fields: A 3-byte Organizationally Unique Identifier (OUI) sub-field and a 2-byte Protocol Identification (PID) sub-field. The encoding for OUI given in Tables 7 indicates the presence of an IEEE 802.2 Routed protocol in the payload. The structure of the SNAP/PID sub-field is given in Tables 8.

Table 8 – SNAP PID

Encoded Value <Byte 2:3> (hex)	Payload Protocol Type
0800	IP
888D	Fibre Channel
Others	<i>Reserved</i>

SNAP PID (2 bytes): The SNAP PID sub-field indicates the payload protocol type. Tables shows encodings for IP and the Fibre Channel payload protocol types. All BBW message make use of only the Fibre Channel payload protocol type.

5.2.2 BBW_Header format

The 4-byte BBW_Header (Tables 9) consists of 3 fields: A 1-byte Flow Control Type field, a 2-byte PAUSE field and a 1-bit Address Bit field. The structure of the Flow Control Field is given in Tables 10.

Table 9 – BBW_Header

Word	Byte	Field	Size (Bytes)	Remarks
2	0	Flow Control Type	1	
	1-2	PAUSE	2	Applicable only when SFC protocol is specified as the Flow Control Type
	3	<bit 0>: Address bit 1= Command; 0 = Response <bits 1-7>: <i>Reserved</i>	1	Applicable only when SR protocol is specified as the Flow Control Type

Flow Control Type (1 byte): This field defines encodings for Simple Flow Control (SFC) and SR Flow Control.

Table 10 – Flow Control Protocol Type Encodings

Encodings (hex)	Flow Control Type
00	Simple Flow Control
01	SR Flow Control
Others	<i>Reserved</i>

PAUSE (2 bytes): This 2-byte PAUSE field is applicable only when the Flow Control Type is SFC. The PAUSE field defines the number of 512-bit time units to pause transmission. A value of zero indicates zero pause transmission time units. This field is also set to zero value when no flow control is desired with the Flow Control Type specified as SFC.

Address Bit (bit 0: Byte 3): This field is applicable only when the Flow Control Type is SR Flow Control. This bit identifies the SR message as either a Command or a Response. Messages containing Commands shall set this bit to 1; messages containing Responses shall set this bit to 0. This field is used in conjunction with the Poll Bit of the SR protocol.

5.2.3 BBW message payload format for SFC

The general structure of the BBW message payload when SFC is specified as the Flow Control Type is given in Tables 11. It consists of the following fields: 4-byte SOF, 24-byte FC-Header, FC frame-payload, 4-byte CRC, and 4-byte EOF. The SOF and EOF byte encodings are defined in Annex A. The FC-Header, FC frame-payload, and the CRC are contents of standard Fibre Channel frame fields that arrive at the B_Port or the E_Port interface.

NOTE 3– The format also applies when the Flow Control Type is specified as SFC and the PAUSE field carries a zero value, that is when no flow control is desired.

Table 11 – BBW message payload structure for SFC

Field	Size (Bytes)
SOF	4
FC-Header	24
FC frame-payload (includes optional header)	Min: 0 Max:2112
CRC	4
EOF	4

5.2.4 BBW message payload format for SR

The general structure of the BBW message payload when SR is specified as the Flow Control Type is given in Tables 12. It consists of a 4-byte SR_Header and the SR_BBW message payload. The

Table 12 – BBW message payload structure for SR

Field	Field	Size (Bytes)
SR_BBW message	SR_Header	4
	SR_BBW message payload	Max: 2148

The SR_Header encodings specify the SR_BBW message Type. SR_Header format is described in 5.2.4.1.

The SR_BBW message payload format depends on the type of the SR_BBW message and is described in 5.2.4.4.1, 5.2.4.4.4, and 5.2.4.4.9.

5.2.4.1 SR_Header Formats

5.2.4.1.1 Field Formats

The SR_Header defines three types of field formats that are used to perform numbered information transfer (I-format), numbered supervisory functions (S-format), and unnumbered control functions (U-format). See Tables 13. SR makes use of 9 different types of messages:

- a) I-format (1): SR_I

- b) S-format (3): SR_RR, SR_RNR, SR_REJ
- c) U-format (5): SR_SM, SR_DISC, SR_FRMR, SR_UA, SR_DM

Bits 0 and 1 of the SR_Header provide the encodings for the three field formats. Bits 2 and 3 provide the encodings for the different S-format messages and are shown in Tables 14. Bits 23, 5,6, and 7 provide the encodings for the different U-format messages and are shown in Tables 15.

Table 13 – SR_Header format

For- mat	3 1	3 0	2 9	2 8	2 7	2 6	2 5	2 4	2 3	2 2	2 1	1 0	1 9	1 8	1 7	1 6	1 5	1 4	1 3	1 2	1 1	1 0	9	8	7	6	5	4	3	2	1	0
I	N(R)												P	N(S)															0			
S	N(R)												P / F	Reserved												S	S	0	1			
U	Reserved													Reserved						M	M	M	P / F	M	M	1	1					

N(S): Transmitter send sequence number; N(R): Transmitter receive sequence number
 SS: Supervisory function bits defined in table below
 M: Modifier function bits defined in table below
 P/F: Poll/Final bit; P Poll bit (1 = Poll)

Table 14 – SS bits encoding

SS Bits		Supervisory message
3	2	
0	0	SR_RR
0	1	Reserved
1	0	SR_RNR
1	1	SR_SREJ

Table 15 – MMMMM bit encoding

MMMMM Bits					Unnumbered message
8	7	6	4	3	
1	1	0	0	0	SR_SM
0	1	0	0	0	SR_DISC
1	0	0	0	1	SR_FRMR
0	1	1	0	0	SR_UA
0	0	0	1	1	SR_DM

5.2.4.1.2 Information transfer I-format

The I-format is used to perform an information transfer. The functions of the N(S), N(R), and P fields are independent; i.e. each SR_I message has a N(S), a N(R), that may or may not acknowledge additional SR_I messages received by the BBWs, and a P-bit that may be set to a 0 or 1.

5.2.4.1.3 Supervisory S-format

The S-format is used to perform data link supervisory control functions such as acknowledge SR_I messages, request retransmission of SR_I messages, and to request a temporary suspension of transmission of SR_I messages. The functions of the N(R) and P/F fields are independent; i.e. each supervisory message has a N(R) that may or may not acknowledge additional SR_I messages received by the BBW, and a P/F bit that may be set to a 0 or 1.

5.2.4.1.4 Unnumbered U-format

The U format is used to provide additional data link control functions. This format contains no sequence numbers, but does include a P/F bit that may be set to a 0 or a 1

5.2.4.2 SR_BBW messages

A description of the 9 different SR_BBW messages appears in Tables 16. Only the SR_I, SR_SREJ, and SR_FRMR messages carry a payload; all other messages carry a null payload.

Table 16 – SR_BBW messages

Purpose	Message	Command/Response	Description
Information transfer	SR_I	Command	Carries encapsulated Class 2, 3, 4 or F frames as payload
Control (Supervisory messages)	SR_RR	Command or Response	Indicates Ready to Receive SR_I messages (negates busy condition) and Acknowledges previous SR_I messages; carries no payload
	SR_RNR	Command or Response	Indicates Receiver Not Ready to accept more SR_I messages (busy condition) and acknowledges previous SR_I messages; carries no payload
	SR_SREJ	Command or Response	Indicates Selective Retransmission of errored SR_I messages; carries a payload
Control (Unnumbered messages)	SR_SM	Command	Mode setting command to set up link and resets all messages counters to 0; carries no payload
	SR_UA	Response	Unnumbered response to the SR_SM command and indicates an acceptance and information transfer phase; carries no payload
	SR_DM	Response	Unnumbered response to the SR_SM command and indicates a disconnected phase; carries no payload
	SR_FRMR	Response	Unnumbered response to the SR_SM command and indicates message reject for the SR_SM message; carries a payload
	SR_DISC	Command	Command indicates the sender is suspending operation and enters the disconnected mode after receiving a SR_UA response; carries no payload
* Command/Response indicated by the Address Bit in Byte 3 of the BBW_Header			

The following describes the different format fields and other related aspects of the SR protocol.

5.2.4.3 Format Field Parameters

5.2.4.3.1 Modulus of SR

Each SR_I message is sequentially numbered and may have the value 0 through modulus minus 1, where “modulus” is equal to 32768 the modulus of the sequence numbers. The sequence numbers cycle through the entire range.

5.2.4.3.2 Send state variable V(S)

The send state variable V(S) denotes the sequence number of the next in-sequence SR_I message to be transmitted. V(S) may take on the values 0 through modulus minus 1. The value of V(S) is incremented by 1 with each successive SR_I message transmission, but cannot exceed the N(R) of the last received SR_I or supervisory message by more than the maximum number of outstanding SR_I messages k . The value of k is defined in 6.3.8.4, Maximum number of outstanding SR_I messages k .

5.2.4.3.3 Send Sequence Number N(S)

Only SR_I messages contain N(S), the send sequence number of the transmitted SR_I message. At the time that an in-sequence SR_I message is designated for transmission, the value of N(S) is set equal to the value of the send state variable V(S).

5.2.4.3.4 Receive State Variable V(R)

The receive state variable V(R) denotes the sequence number of the next in-sequence SR_I message expected to be received. V(R) may take on the values 0 through modulus minus 1. The value of V(R) is incremented by 1 by the receipt of an error-free, in-sequence SR_I message whose send sequence number N(S) equals the receive state variable V(R).

5.2.4.3.5 Receive Sequence Number N(R)

All SR_I messages and supervisory messages, except SR_SREJ messages with the F bit set to 0, shall contain N(R), the expected send sequence number of the next received SR_I message. At the time that a message of the above types is designated for transmission, the value of N(R) is set equal to the current value of the receive state variable V(R). N(R) indicates that the BBW transmitting the N(R) has received correctly all SR_I messages numbered up to and including N(R)-1.

5.2.4.3.6 Functions of the Poll/Final Bit (P/F)

All messages contain P/F, the Poll/Final bit. In command messages, the P/F bit is referred to as the P bit. In response messages it is referred to as the F bit.

The Poll bit set to 1 is used by the BBW to solicit (poll) a response from the remote BBW.

The Final Bit set to 1 is used by the BBW to indicate the response message transmitted by the remote BBW, as a result of the soliciting (poll) command.

The use of the P/F bit is described 6.3.3.

5.2.4.4 SR Commands and Responses

5.2.4.4.1 Information (SR_I) command

The function of the information (SR_I) command is to transfer across a data link a sequentially numbered message containing an information field.

The SR_I message command carries the mapped byte-encoded Class 2, 3, 4 or F frames. The following steps are involved in generating the SR_I message:

- a) Constructing the SR_I message payload by prefixing the proper 32-bit SOF delimiter to the incoming FC-Header, FC frame-payload, and the CRC, and appending the corresponding 32-bit EOF delimiter to the CRC.

NOTE 4– The original Fibre Channel frame CRC field remains and the sender does not have to send a valid CRC and the receiver does not have to validate the CRC

- b) Constructing the 4-byte SR_Header and prefixing it to the SR_I message payload

Tables 17 illustrates the format of the SR_I message information field (payload). The maximum size of the BBW message is 2152 bytes corresponding to a maximum size FC frame-payload of 2112 bytes. The FC frame payload uses the SOF and EOF codes defined in Annex A.

Table 17 – SR_I message format

Field	Description	Size (Bytes)
SR_Header		4
SR_I message payload	SOF	4
	FC-Header	24
	FC frame-payload (includes optional header)	Min: 0 Max:2112
	CRC	4
	EOF	4

NOTE 5– SR protocol generated control messages do not carry the SOF and the EOF fields nor the FC headers in the payload

5.2.4.4.2 Receive Ready (SR_RR) Command and Response

The Receive Ready (SR_RR) supervisory message is used by the BBW to:

- a) indicate it is ready to receive a SR_I message; and
- b) acknowledge previously received SR_I messages numbered up to and including N(R)-1.

A SR_RR message may be used to indicate the clearance of a busy condition that was reported by the earlier transmission of a SR_RNR message by the same device. In addition to indicating the BBW status, the SR_RR message with the P-bit set to 1 may be used to ask for the status of the remote BBW.

5.2.4.4.3 Receive Not Ready (SR_RNR) Command and Response

The Receive Not Ready (SR_RNR) supervisory message is used to indicate a busy condition; i.e. temporary inability to accept additional incoming SR_I messages. SR_I messages numbered up to and including N(R)-1 are acknowledged. SR_I message N(R) and any subsequent SR_I messages received, if any, are not acknowledged; the acceptance status of these SR_I messages shall be indicated in subsequent exchanges.

In addition to indicating the status, the SR_RNR command with the P-bit set to 1 may be used by a BBW to ask for the status of the remote BBW.

5.2.4.4.4 Selective Reject (SR_SREJ) Response

The SR_REJ supervisory message shall be used by a BBW to request retransmission of one or more (not necessarily contiguous), SR_I messages. The N(R) field shall contain the sequence number of the earliest SR_I message to be retransmitted and the information field (payload) shall contain, in as-

ending order (32767 is higher than 32766 and 0 is higher than 32767 for modulo 32768), the sequence numbers of additional SR_I message(s), if any, in need of retransmission.

The payload field shall be encoded such that there is a 2-octet field for each standalone SR_I message in need of retransmission, and a 4-octet span list for each sequence of two or more contiguous numbered SR_I messages in need of retransmission, as depicted in Tables 18. In the case of the standalone SR_I messages, their identity in the payload field consists of the appropriate N(R) value preceded by a 0 bit in the 2-octet field used. In the case of span lists, their identity in the payload field consists of the N(R) value of the first SR_I message in the span list preceded by a 1 bit in the 2-octet field used, followed by the N(R) value of the last message in the span list preceded by a 1 bit in the 2-octet field used.

Table 18 – SR_SREJ payload format

	Field	Size (Bytes)
<Bit 1> = 0	<Bits 2-16> = 1-N(R) of standalone SR_I message	2
<Bit 1> = 1	<Bits 2-16> = N(R) of first SR_I message in span list	2
<Bit 1> = 1	<Bits 2-16> = N(R) of last SR_I message in span list	2
<Bit 1> = 0	<Bits 2-16> = N(R) of standalone SR_I message	2
<Bit 1> = 1	<Bits 2-16> = N(R) of first SR_I message in span list	2
<Bit 1> = 1	<Bits 2-16> = N(R) of last SR_I message in span list	2
	...	

NOTE 6– The maximum size of the BBW message payload carrying the SR_SREJ message is 2148 bytes corresponding to a maximum possible encoding of 1074 standalone SR_I messages or a maximum possible encoding of 537 span list sets.

If the P/F bit in a SREJ message is set to 1, then SR_I messages numbered up to N(R)-1 inclusive (N(R) being the value in the SR_Header field), are considered as acknowledged. If the P/F bit in a SREJ message is set to 0, then the N(R) in the SR_Header field of the SREJ message does not indicate acknowledgement of SR_I messages.

The procedures to be followed on receipt of Set Mode (SR_SM) command are specified in 6.3.5.7, Receiving a SR_SREJ response message.

5.2.4.4.5 Set Mode (SR_SM) Command

The SR_SM unnumbered command is used to initialize the BBW device.

No information field is permitted with the SR_SM command. The transmission of a SR_SM command indicates the clearance of a busy condition that was reported by the earlier transmission of a SR_RNR message by the same BBW device. The BBW device confirms the acceptance of the SR_SM command by the transmission, at the first opportunity, of a SR_UA response. Upon acceptance of this command, the BBW device send state variable V(S) and receive state variable V(R) are set to 0.

Previously transmitted SR_I messages that are unacknowledged when this message is actioned remain unacknowledged. It is the responsibility of a higher layer to recover from the possible loss of the contents of such SR_I messages.

5.2.4.4.6 Disconnect (SR_DISC) Command

The SR_DISC unnumbered command is used to terminate the link that had been previously set. It is used to inform the BBW receiving the SR_DISC command that the remote BBW is suspending operation. No information field is permitted with the SR_DISC command. Prior to actioning the SR_DISC command, the BBW receiving the SR_DISC command confirms the acceptance of the SR_DISC

command by the transmission of a SR-UA response. The BBW sending the SR_DISC command enters the disconnected phase when it receives the acknowledging SR-UA response.

Previously transmitted SR_I messages that are unacknowledged when this command is actioned remain unacknowledged. It is the responsibility of a higher layer to recover from the possible loss of the contents of such SR_I messages.

5.2.4.4.7 Unnumbered Acknowledgement (SR-UA) Response

The SR-UA unnumbered response is used by the BBW device to acknowledge the receipt and acceptance of the SR_SM mode setting command. Received mode-setting command is not actioned until the SR-UA response is transmitted. The transmission of a SR-UA response indicates the clearance of a busy condition that was reported by the earlier transmission of a SR_RNR message by that same BBW device. No information field is permitted with the SR-UA response.

5.2.4.4.8 Disconnected Mode (SR_DM) Response

The SR_DM unnumbered response is used to report a status where a BBW is logically disconnected from the data link, and is in the disconnected phase. The SR_DM response may be sent to indicate that the BBW has entered the disconnected phase without benefit of having received a SR_DISC command, or, if sent in response to the SR_SM mode setting command, is sent to inform the remote BBW that the BBW is still in the disconnected phase and cannot execute the SR_SM set mode command. No information field is permitted with the SR_DM response.

A BBW in a disconnected phase shall monitor received commands and shall react to a SR_SM command as outlined in 6.3.4, SR procedure for data link set-up and disconnection, and shall respond with a SR_DM response with the F bit set to 1 to any other command received with the P bit set to 1.

5.2.4.4.9 Message Reject (SR_FRMR) Response

The SR_FRMR unnumbered response is used by the BBW device to report an error condition not recoverable by retransmission of the identical message, i.e., at least one of the following conditions, that results from the receipt of a valid message (see Tables 19):

- a) the receipt of a command or response SR_Header sub-field that is undefined
- b) the receipt of an invalid N(R) (defined below); or
- c) the receipt of a message with an information field that is not permitted or the receipt of a supervisory or unnumbered message with incorrect length.

A valid N(R) shall be within the range from the lowest send sequence number N(S) of the still unacknowledged message(s) to the current BBW send state variable inclusive (or to the current internal variable x if the BBW is in the timer recovery condition as described in 6.3.5.10, Awaiting acknowledgement).

An information field that immediately follows the SR_Header consists of 9 octets, is returned with this response and provides the reason for the SR_FRMR response. Tables 19 shows the payload format.

5.2.4.5 Exception condition reporting and recovery

5.2.4.5.1 Exception Conditions

The error recovery procedures that are available to effect recovery following the detection/occurrence of an exception condition are described below. Exception conditions described are those situations that may occur as the result of transmission errors, BBW device malfunction, or operational situations.

5.2.4.5.2 Busy Condition

The busy condition results when the BBW is temporarily unable to continue to receive SR_I messages due to internal constraints, e.g. receive buffering limitations. In this case a SR_RNR message is transmitted from the busy BBW. SR_I messages pending transmission may be transmitted from the busy BBW prior to or following the SR_RNR message.

An indication that the busy condition has cleared is communicated by the transmission of a SR-UA (only in response to a SR_SM command), SR_RR, SR_SREJ, or SR_SM message.

Table 19 – SR_FRMR payload format

Word	Bit Number	Field		Size (bits)
0	1 -32		Rejected SR_Header Field	32
1	33		Set to 0	1
	34-48	V(S)	V(S) is the current send state variable value at the BBW reporting the rejection condition (bit 34= low-order bit)	15
	49	C/R	C/R set to 1 indicates the rejected message was a response; C/R set to 0 indicates the rejected message was a command	1
	50-64	V(R)	V(R) is the current receive state variable value at the BBW reporting the rejection condition (bit 50= low-order bit)	15
2	65	W	W set to 1 indicates that the SR_Header field received and returned in bits 1 through 32 was undefined	1
	66	X	X set to 1 indicates that the SR_Header field received and returned in bits 1 through 32 was considered invalid because the message contained a payload that was not permitted with this type of message or is a supervisory or unnumbered message with incorrect length. Bit W shall be set to 1 in conjunction with this bit.	1
	67	Y	Reserved	1
	68	Z	Z set to 1 indicates that the SR_Header field received and returned in bits 1 through 32 contained an invalid N(R)	1
	69-72		Set to 0	4
	73-96		Reserved	24

5.2.4.5.3 N(S) Sequence Error Condition

The information field of all SR_I messages received whose N(S) is not in the range V(R) and V(R)+k-1 inclusive, shall be discarded. The information field of all SR_I messages received by the BBW whose N(S) is in the range V(R) and V(R) + k -1 inclusive, shall be saved in the receive buffer.

A N(S) sequence error exception condition occurs in the receiver when a SR_I message received contains a N(S) that is not equal to the receive state variable V(R) at the receiver. The receiver does not acknowledge (increment its receive state variable), the SR_I message causing the sequence error, or any SR_I message that may follow, until a SR_I message with the correct N(S) is received.

A BBW device that receives one or more valid SR_I messages having sequence errors or subsequent supervisory messages (SR_RR, SR_RNR, and SR_SREJ) shall accept the N(R) field and the P or F bit to perform data link control functions, e.g. to receive acknowledgement of previously transmitted SR_I messages and to cause the BBW to respond (P bit set to 1).

The means specified in 5.2.4.5.3.1, SR_SREJ recovery and 5.2.4.5.3.2, Time-out recovery shall be available for initiating the retransmission of lost or errored SR_I messages following the occurrence of a N(S) sequence error condition.

5.2.4.5.3.1 SR_SREJ recovery

The SR_SREJ message shall be used to initiate more efficient error recovery by selectively requesting the retransmission of one or more (not necessarily contiguous), lost or errored SR_I message(s) following the detection of sequence errors, rather than requesting the retransmission of all SR_I messages. When a BBW receives an out-of-sequence message, the SR_I message shall be saved in a receive buffer. The SR_I message shall be delivered to the upper layer only when all SR_I messages numbered below N(S) are correctly received. If message number N(S) -1 has not been received previously, then a SR_SREJ response message with the F bit set to 0 shall be transmitted, that contains the sequence numbers of the block of consecutive missing SR_I messages ending at N(S)-1. The BBW on receiving such a SR_SREJ message shall retransmit all requested SR_I messages. After having retransmitted these SR_I messages, the BBW may transmit new SR_I messages, if they become available.

When a BBW receives a command message with the P bit set to a 1, if there are out-of-sequence SR_I messages saved in the receive buffer, it shall transmit a SR_SREJ message with the F bit set to 1, that contains a complete list of missing sequence numbers. The BBW on receiving such a SR_SREJ message shall retransmit all requested SR_I messages, except those that were transmitted subsequent to the last command message with the P bit set to 1.

5.2.4.5.3.2 Time-out Recovery

If a BBW, due to a transmission error, does not receive (or receives and discards) a single SR_I message or the last SR_I message in a sequence of SR_I messages, it shall not detect a N(S) sequence error condition and, therefore, shall not transmit a SR_SREJ message.

The BBW that transmitted the unacknowledged SR_I message(s) shall, following the completion of a system specified time-out period (see 6.3.5.2, Sending new SR_I messages and 6.3.5.10, Awaiting acknowledgements), send a supervisory command message (SR_RR or SR_RNR) with the P bit set to 1. SR_I messages shall be retransmitted on the receipt of a SR_RR response message with the F bit set to 1 or a SR_SREJ message.

5.2.4.5.4 Invalid message condition

Any message that is invalid shall be discarded, and no action is taken as the result of that message. An invalid message is defined as one that contains:

- a) the BBW_Header defined in Tables 9 with an invalid encoding
- b) the SR_Header defined in Tables 13 with an invalid encoding.

5.2.4.5.5 Message rejection condition

A message rejection condition is established upon the receipt of an error-free message with one of the conditions listed in 5.2.4.4.9, Message Reject (SR_FRMR) response. At the BBW, this message rejection exception condition is reported by a SR_FRMR response for an appropriate BBW action.

Once a BBW has established such an exception condition, no additional SR_I messages are accepted until the condition is reset by the remote BBW, except for examination of the P bit. The SR_FRMR response may be repeated at each opportunity as described in 6.3.7.3, until recovery is effected by the remote BBW, or until the BBW initiates its own recovery in case the remote BBW does not respond.

6 The SR and SFC Protocol Procedures

6.1 Applicability

This clause only applies to FC-BB-2_ATM and FC-BB-2_SONET. SR Protocol.

The SR protocol is described in 6.2 and 6.3, and the SFC protocol in 6.4.

6.2 SR Protocol Overview

The Selective Retransmission (SR) protocol is an efficient sliding window link-layer full-duplex protocol that supports both the flow control and error recovery functions. SR has been adopted from ITU's Link Access Protocol B (LAPB), that was derived from ISO/IEC's High-level Data Link Control (HDLC) (Balanced Classes). Use of LAPB in SR is limited to a subset of the synchronous modulo 32768 super sequence numbering service option.

SR works between two BBW devices. See Figure 5. SR flow control works by streaming multiple messages within an allowed window (bounded by the system parameter k), and awaits acknowledgements before sending more messages. Acknowledgements indicate which messages were correctly received and there is a provision for requesting retransmission of "selected" messages in the current window. Fibre Channel Sequences and Exchanges are not visible to the SR Flow Control protocol and it only sees the BBW messages constructed from the FC frames.

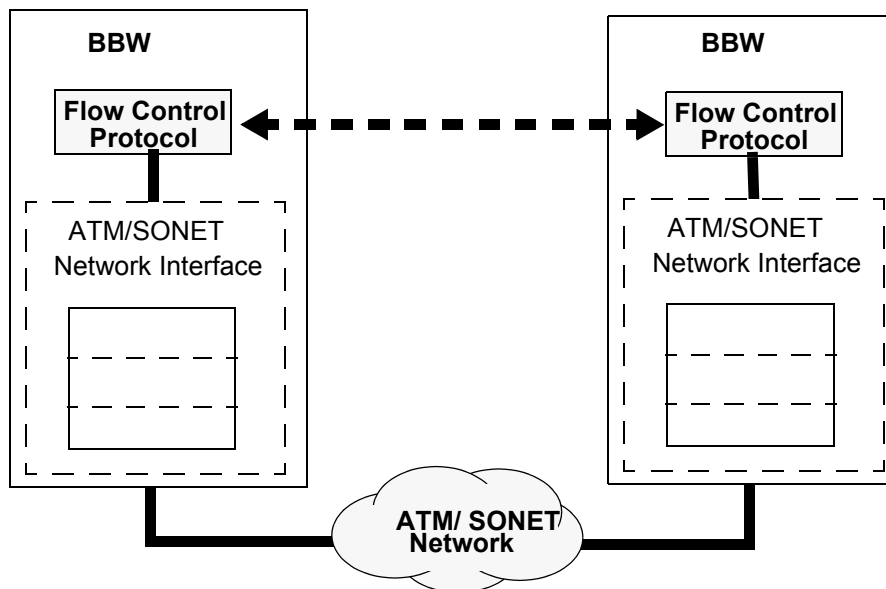


Figure 5 – SR Flow Control Protocol Between two BBWs

Some benefits of the SR protocol are summarized below:

- a) It is used for reliable transport of all Class 2, 3, 4 and F frames between two BBW devices
- b) It synchronizes the BBW Sender and the BBW Receiver at the BBW message level
- c) It optimizes buffer management at the BBW devices
- d) It acts as a congestion avoidance technique to match the capacity of the sender to the capacity of the network that carries the payload
- e) It ensures correct delivery of messages (an error control and recovery function) and

- f) It provides a continuous stream of traffic across the WAN thus leading to a higher throughput, i.e., optimizes bandwidth utilization at each BBW device.

The 9 different SR messages listed in Tables 16 have a correspondence to the LAPB frame-types. Note that only the information transfer SR_I message is flow-controlled while all other messages are control messages of the protocol.

The SR protocol specifies the maximum number (k) of outstanding messages at any given time. k is a system parameter that is not negotiated and fixed in a given implementation. The value of this system parameter depends on the WAN delay characteristics and the number of buffers available. Typically, the value of k is expected to be far below the maximum number of 32767.

The following sub clauses describe the SR protocol procedures and reference the different message fields discussed in Clause 5.

6.3 Description of the SR procedure

6.3.1 SR mode of operation

The SR protocol shall be limited to a subset of the synchronous modulo 32768 super sequence numbering service option operation of the LAPB protocol. See 5.2.4 for a description of the SR BBW message formats. The mode-setting command employed to initialize (set-up) or reset the protocol is the SR_SM command.

6.3.2 SR procedure for addressing

An Address bit-field identifies a message as either a command or a response.

This field is used in conjunction with the Poll/Final bit.

6.3.3 SR procedure for the use of the P/F bit

The BBW receiving a SR_SM, SR_DISC, supervisory command (SR_RR, SR_RNR, SR_SREJ), or SR_I message with the P bit set to 1 shall set the F bit to 1 in the next response message it transmits.

The response message returned by the BBW to a SR_SM or SR_DISC command with the P bit set to 1 shall be a SR_UA or SR_DM response with the F bit set to 1.

The response message returned by the BBW to a SR_I message with the P bit set to 1, received during the information transfer phase, shall be a SR_RR, SR_SREJ, SR_RNR, or SR_FRMR response with the F bit set to 1.

The response message returned by the BBW to a supervisory command with the P bit set to 1, received during the information transfer phase, shall be a SR_RR, SR_RNR, SR_SREJ or SR_FRMR response with the F bit set to 1.

The response message returned by the BBW to a SR_I message or supervisory message with the P bit set to 1, received during the disconnected phase, shall be a SR_DM response with the F bit set to 1.

The P bit may be used by the BBW in conjunction with the timer recovery condition (see 6.3.5.10, Awaiting Acknowledgement).

6.3.4 SR procedure for data link set-up and disconnection

6.3.4.1 Data link set-up

The BBW shall indicate to the SR protocol layer that it is able to set up the data link after it has provisioned an underlying ATM VC or a SONET Path.

Either BBW may initiate data link set-up. Prior to initiation of data link set-up, either BBW may initiate data link disconnection (see 6.3.4.3, Data link disconnection) for the purpose of ensuring that both

BBW devices are in the same phase. A BBW may also transmit an unsolicited SR_DM response to request the remote BBW to initiate data link set-up.

The BBW shall initiate data link set-up by transmitting a SR_SM command. If, upon correct receipt of the SR_SM command, the BBW device determines that it may enter the information transfer phase, it shall return a SR_UA response to the sender, reset its send and receive state variables V(S) and V(R) to zero and shall consider that the link is set up.

If, upon receipt of the SR_SM command, the BBW device determines that it cannot enter the information transfer phase, it shall return a SR_DM response as a denial to the link set up initialization and shall consider that the data link is not set up. In order to avoid misinterpretation of the SR_DM response received, it is suggested that the BBW always send its SR_SM command with the P bit set to 1. Otherwise, it is not possible to differentiate a SR_DM response intended as a denial to data link set up from a SR_DM response that is issued in a separate unsolicited sense as a request for a mode-setting command (as described in 6.3.4.4.2).

The BBW device shall initiate link set up by transmitting a SR_SM command and starting its Timer T1 in order to determine when too much time has elapsed waiting for a reply (see 6.3.8.1, Timer T1). Upon reception of a SR_UA response, the BBW shall reset its send and receive state variables V(S) and V(R) to zero, shall stop its Timer T1, and shall consider that the link is set up. Upon reception of a SR_DM response as a denial to the link set-up initialization, the BBW shall stop its Timer T1 and shall consider that the link is not set up.

The BBW having sent the SR_SM command, shall ignore and discard any messages except a SR_SM, or SR_DISC command, or a SR_UA or SR_DM response received from the remote BBW. The receipt of a SR_SM or SR_DISC command results in a collision situation that is resolved per 6.3.4.5, Collision of unnumbered commands below. Messages other than the SR_UA and the SR_DM responses sent in response to a received SR_SM or SM_DISC command shall be sent only after the link is set up and if no outstanding SR_SM command exists.

After the BBW sends the SR_SM command, if a SR_UA or SR_DM response is not received correctly, Timer T1 shall run out in the BBW. The BBW shall then resend the SR_SM command and shall restart Timer T1. After transmission of the SR_SM command N2 times by the BBW, appropriate higher layer recovery action shall be initiated. The value of N2 is defined in , Maximum number of attempts to complete a transmission N2.

6.3.4.2 Information transfer Phase

After having transmitted the SR_UA response to the SR_SM command or having received the SR_UA response to a transmitted SR_SM command, the BBW shall accept and transmit SR_I messages and supervisory messages (SR_RR, SR_RNR, and SR_SREJ) according to the procedures defined in 6.3.5, Procedures for information transfer when using multi-selective reject.

When receiving the SR_SM command while in the information transfer phase, the BBW shall conform to the data link resetting procedure described in 6.3.7, SR procedure for data link resetting.

6.3.4.3 Data link disconnection

The BBW shall initiate a disconnect of the link by transmitting a SR_DISC command. On correctly receiving a SR_DISC command in the information transfer phase, the BBW shall send a SR_UA response and enter the disconnected phase. On correctly receiving a SR_DISC command in the disconnected phase, the remote BBW shall send a SR_DM response and remain in the disconnected phase. In order to avoid misinterpretation of the SR_DM response received, it is suggested that the BBW always sends its SR_DISC command with the P-bit set to 1. Otherwise, it is not possible to differentiate a SR_DM response intended as an indication that the device is already in the disconnected phase from a SR_DM response that is issued in a separate unsolicited sense as a request for a mode setting command (SR_SM) as described in 6.3.4.4.2.

The BBW shall initiate a disconnect of the data link by transmitting a SR_DISC command and starting its Timer T1 (see 6.3.8.1, Timer T1 below). Upon reception of a SR_UA response from the remote BBW, the BBW shall stop its Timer T1 and shall enter the disconnected phase. Upon reception of a SR_DM response from the remote BBW as an indication that the remote BBW was already in the disconnected phase, the BBW shall stop its Timer T1 and shall enter the disconnected phase.

The BBW having sent the SR_DISC command shall ignore and discard any messages except a SR_SM or SR_DISC command, or a SR_UA or SR_DM response received from the remote BBW. The receipt of a SR_SM or SR_DISC command from the remote BBW shall result in a collision situation that is resolved per 6.3.4.5, Collision of unnumbered commands below.

After the BBW sends the SR_DISC command, if a SR_UA or SR_DM response is not received correctly, Timer T1 shall run out in the BBW. The BBW shall then resend the SR_DISC command and shall restart Timer T1. After transmission of the SR_DISC command N2 times by the BBW, appropriate higher layer recovery action shall be initiated. The value of N2 is defined in , Maximum number of attempts to complete a transmission N2.

6.3.4.4 Disconnected Phase

6.3.4.4.1

After having received a SR_DISC command from the remote BBW and returned a SR_UA response to the remote BBW, or having received the SR_UA response to a transmitted SR_DISC command, the BBW shall enter the disconnected phase.

In the disconnected phase, the BBW may initiate data link set-up. In the disconnected phase, the BBW shall react to the receipt of a SR_SM command as described in 6.3.4.1, Data link set-up above and shall transmit a SR_DM response in answer to a received SR_DISC command. When receiving any other command (defined or undefined) with the P-bit set to 1, the BBW shall transmit a SR_DM response with the F-bit set to 1. Other messages received in the disconnected phase shall be ignored by the BBW.

6.3.4.4.2

When the BBW enters the disconnected phase after detecting error conditions as listed in 6.3.6, SR conditions for data link resetting or data link re initialization (data link set-up) below, or after an internal malfunction, it may indicate this by sending a SR_DM response rather than a SR_DISC command. In these cases, the BBW shall transmit a SR_DM response and start its Timer T1 (see 6.3.8.1, Timer T1 below).

If Timer T1 runs out before the reception of a SR_SM or SR_DISC command from the remote BBW, the BBW shall retransmit the SR_DM response and restart Timer T1. After retransmission of the SR_DM response N2 times, the BBW shall remain in the disconnected phase and appropriate recovery actions shall be initiated. The value of N2 is defined in , Maximum number of attempts to complete a transmission N2.

Alternatively, after an internal malfunction, the BBW may either initiate a data link resetting procedure (see 6.3.7, SR procedure for data link resetting) or disconnect the data link (see 6.3.4.3, Data link disconnection) prior to initiating a data link set-up procedure (see 6.3.4.1, Data link set-up).

6.3.4.5 Collision of unnumbered commands

6.3.4.5.1

If the sent and received unnumbered commands are the same, the BBWs shall each send the SR_UA response at the earliest possible opportunity. The BBW shall enter the indicated phase either:

- 1) after receiving the SR_UA response;
- 2) after sending the SR_UA response; or

- 3) after timing out waiting for the SR-UA response having sent a SR-UA response.

In the case of 2) above, the BBW shall accept a subsequent SR-UA response to the mode-setting command it issued without causing an exception condition if received within the time-out interval.

6.3.4.5.2

If the sent and received unnumbered commands are different, the BBWs shall each enter the disconnected phase and issue a SR_DM response at the earliest possible opportunity.

6.3.4.6 Collision of SR_DM response with SR_SM or SR_DISC command

When a SR_DM response is issued by the BBW as an unsolicited response to request the remote BBW to issue a mode-setting command as described in 6.3.4.4, Disconnected Phase, a collision between a SR_SM or SR_DISC command and the unsolicited SR_DM response may occur. In order to avoid misinterpretation of the SR_DM response received, the remote BBW always sends its SR_SM or SR_DISC command with the P-bit set to 1.

6.3.4.7 Collision of SR_DM responses

A contention situation may occur when both the BBWs issue a SR_DM response. In this case, either BBW may issue a SR_SM command to resolve the contention situation.

6.3.5 Procedures for information transfer using multi-selective reject

6.3.5.1 Procedures for SR_I messages

The procedures that apply to the transmission of SR_I messages in each direction during the information transfer phase using multi-selective reject are described below.

In the following, “number one higher” is in reference to a continuously repeated sequence series, i.e. 32767 is one higher than 32766 and 0 is one higher than 32767 for modulo 32768 series.

The term “outstanding poll condition” is used to indicate the condition when the BBW has sent a command message with the P bit set to 1 and has not yet received a response message with the F bit set to 1.

6.3.5.2 Sending new SR_I messages

When the BBW has a new SR_I message to transmit (i.e., a SR_I message not already transmitted), it shall transmit it with a N(S) equal to its current send state variable V(S), and a N(R) equal to its current receive state variable V(R). At the end of the transmission of the SR_I message, it shall increment its send state variable V(S) by 1.

If the BBW Timer T1 is not running at the time of transmission of the SR_I message, it shall be started.

If the BBW send state variable V(S) is equal to the last value N(R) received plus k (where k is the maximum number of outstanding SR_I frames; see 6.3.8.4, Maximum number of outstanding SR_I messages k below), the BBW shall not transmit any new SR_I frames.

If the remote BBW is busy, the BBW shall not transmit any new SR_I messages.

When the BBW is in the busy condition, it may still transmit SR_I messages, provided that the remote BBW is not busy.

6.3.5.3 Receiving an in-sequence SR_I message

When the BBW is not in a busy condition and receives a valid SR_I message whose send sequence number N(S) is equal to its receive state variable V(R), the BBW shall accept the information field of this message and increment by one the receive state variable V(R). If the SR_I message, whose N(S) is equal to (the incremented value of) V(R), is present in the receive buffer, then the BBW shall remove it from the receive buffer, deliver it to the upper layer increment V(R) by one; the BBW shall

repeat this procedure until $V(R)$ reaches a value such that the SR_I message whose $N(S)$ is equal to $V(R)$ is not present in the receive buffer. The BBW shall then take one of the following actions:

- a) If the BBW is still not in a busy condition:
 - i) If the P-bit is set to 1, then the BBW shall transmit a response message with the F bit set to 1, as specified in 6.3.5.12, Responding to command messages with the P bit set to 1.
 - ii) Otherwise, if a SR_I message is available for transmission (as specified in 6.3.8.4, Maximum number of outstanding SR_I messages k), the BBW shall act as described in 6.3.5.2, Sending new SR_I messages and acknowledge the received SR_I message by setting $N(R)$ in the SR_Header field of the next transmitted SR_I message to the value of the BBW receive state variable $V(R)$, or the BBW shall acknowledge the received SR_I message by transmitting a SR_RR message with the $N(R)$ equal to the value of the BBW receive state variable $V(R)$.
 - iii) Otherwise, the BBW shall transmit a SR_RR message with $N(R)$ equal to the value of the BBW receive state variable $V(R)$
- b) If the BBW is now in the busy condition, it shall transmit a SR_RNR message with $N(R)$ equal to the value of the BBW receive variable $V(R)$ (see 6.3.5.9, BBW busy condition)

When the BBW is in a busy condition, it may ignore the information field contained in any received SR_I message.

6.3.5.4 Reception of invalid messages

When the BBW receives an invalid message (see 5.2.4.5.4), it shall discard the message.

6.3.5.5 Reception of out-of-sequence SR_I messages

When the BBW is not in a busy condition and it receives a valid SR_I message whose send sequence number $N(S)$ is out-of-sequence, i.e. not equal to the receive state variable $V(R)$, then it shall perform one of the following actions:

- a) If $N(S)$ is less than $V(R)$ or greater than or equal to $V(R) + k$, then it shall discard the information field of the SR_I message. If the P bit of the SR_I message is set to 1, then the BBW shall transmit a response message with the F bit set to 1, as specified in 6.3.5.12, Responding to command messages with the P bit set to 1.
- b) If $N(S)$ is greater than $V(R)$ and less than $V(R) + k$, then it shall save the SR_I message in the receive buffer. It shall then perform one of the following actions:
 - 1) If the P bit of the SR_I message is set to 1, then the BBW shall transmit a response message with the F bit set to 1, as specified in 6.3.5.12, Responding to command messages with the P bit set to 1.
 - 2) Otherwise, if the BBW is now in a busy condition, it shall transmit an SR_RNR message with $N(R)$ equal to the value of the receive variable $V(R)$, as specified in 6.3.5.9, BBW busy condition.
 - 3) Otherwise, if the SR_I message numbered $N(S)-1$ has not yet been received, then the BBW shall transmit a SR_SREJ response message with the F bit set to 0. The BBW shall create a list of contiguous sequence numbers $N(X)$, $N(X)+1$, $N(X)+2$, ..., $N(S)-1$, where $N(X)$ is greater than or equal to $V(R)$ and none of the SR_I messages $N(X)$ to $N(S)-1$ have been received. The $N(R)$ field of the SR_SREJ message shall be set to $N(X)$ and the information field set to the list $N(X)+1$, ..., $N(S)-1$. If the list of sequence numbers is too large to fit into the information field of the SR_SREJ message, then the list shall be truncated to fit in one SR_SREJ message, by including only the earliest sequence numbers.

When the BBW is in busy condition, it may ignore the information field contained in any received SR_I message.

6.3.5.6 Receiving acknowledgement

When correctly receiving a SR_I message or a supervisory message (SR_RR, SR_RNR, or SR_SREJ with the F bit set to 1), even in the busy condition, the BBW shall consider the N(R) contained in this message as an acknowledgement for all the SR_I messages it has transmitted with a N(S) up to and including the received N(R)-1. The BBW shall stop the Timer T1 if the received supervisory message has the F bit set to 1 or if there is no outstanding poll condition and the N(R) is higher than the last received N(R) (actually acknowledging some SR_I messages).

If Timer T1 has been stopped by the receipt of a SR_I message, a SR_RR command message, a SR_RR response message with the F bit set to 0 or a SR_RNR message, and if there are outstanding SR_I messages still unacknowledged, the BBW shall restart Timer T1. If Timer T1 has been stopped by the receipt of a SR_SREJ message with the F bit set to 1, the BBW shall follow the retransmission procedure in 6.3.5.7.2, Receiving a SR_SREJ response message with the F bit set to 1. If Timer T1 has been stopped by the receipt of a SR_RR message with the F bit set to 1, the BBW shall follow the retransmission procedure in 6.3.5.11, Receiving a SR_RR response message with the F bit set to 1.

6.3.5.7 Receiving a SR_SREJ response message

6.3.5.7.1 Receiving a SR_SREJ response message with the F bit set to 0

When receiving a SR_SREJ response message with the F bit set to 0, the BBW shall retransmit all SR_I messages, whose sequence numbers are indicated in the N(R) field and the information field of the SR_SREJ message, in the order specified in the SR_SREJ message. Retransmission shall conform to the following:

- a) If the BBW is transmitting a supervisory or SR_I message when it receives the SR_SREJ message, it shall complete that transmission before commencing transmission of the requested SR_I messages.
- b) If the BBW is transmitting an unnumbered command or response message when it receives the SR_SREJ message, it shall ignore the request for retransmission.
- c) If the BBW is not transmitting any message when it receives the SR_SREJ message, it shall commence transmission of the requested SR_I messages immediately.

If there is no outstanding poll condition, then a poll shall be sent, either by transmitting a SR_RR command (or SR_RNR command if the BBW is in the busy condition) with the P bit set to 1 or by setting the P bit in the last retransmitted SR_I message and Timer T1 shall be restarted.

If there is an outstanding poll condition, then Timer T1 shall not be restarted.

6.3.5.7.2 Receiving a SR_SREJ response message with the F bit set to 1

When receiving a SR_SREJ response message with the F bit set to 1, the BBW shall retransmit all SR_I messages, whose sequence numbers are indicated in the N(R) field and the information field of the SR_SREJ message, in the order specified in the SR_SREJ message, except those messages that were sent after the message with the P bit set to 1 was sent. Retransmission shall conform to the following:

- a) If the BBW is transmitting a supervisory message or SR_I message when it receives the SR_SREJ message, it shall complete that transmission before commencing transmission of the requested SR_I messages.
- b) If the BBW is transmitting an unnumbered command or response when it receives the SR_SREJ message, it shall ignore the request for retransmission.

- c) If the BBW is not transmitting any message when it receives the SR_SREJ message, it shall commence transmission of the requested SR_I messages immediately.

If any messages are retransmitted, then a poll shall be sent, either by transmitting a SR_RR command (or SR_RNR command if the BBW is in the busy condition) with the P bit set to 1 or by setting the P bit in the last retransmitted SR_I message.

Timer T1 shall be restarted.

6.3.5.8 Receiving a SR_RNR message

After receiving a SR_RNR message, the BBW shall stop transmission of SR_I messages until a SR_RR or SR_SREJ message is received.

The BBW shall start Timer T1, if necessary, as specified in 6.3.8.1.

When Timer T1 runs out before receipt of a busy clearance indication, the BBW shall transmit a supervisory message (SR_RR, SR_RNR), with the P bit set to 1 and shall restart Timer T1, in order to determine if there is any change in the receive status of the remote BBW. The remote BBW shall respond to the P bit set to 1 with a supervisory response message (SR_RR, SR_RNR, SR_SREJ) with the F bit set to 1 indicating continuation of the busy condition (SR_RNR message) or clearance of the busy condition (SR_RR, SR_SREJ). Upon receipt of the remote BBW response, Timer T1 shall be stopped.

- a) If the response is a SR_RR message, the busy condition shall be assumed to be cleared and the BBW may retransmit messages as specified in 6.3.5.11, Receiving a SR_RR response message with the F bit set to 1. New SR_I messages may be transmitted as specified in 6.3.5.2, Sending new SR_I messages.
- b) If the response is a SR_SREJ message, the busy condition shall be assumed to be cleared and the BBW may retransmit messages as specified in 6.3.5.7.2, Receiving a SR_SREJ response message with the F bit set to 1. New SR_I messages may be transmitted as specified in 6.3.5.2, Sending new SR_I messages.
- c) If the response is a SR_RNR message, the busy condition shall be assumed to still exist and the BBW, after a period of time (for example the duration of Timer T1), shall repeat the enquiry of the remote BBW receive status.

If Timer T1 runs out before a status response is received, the enquiry process above shall be repeated. If N2 attempts to get a status response fail, the BBW shall initiate link resetting procedure as described in 6.3.7, SR procedures for data link resetting.

If, at any time during the enquiry process, an unsolicited SR_RR or SR_SREJ message is received from the remote BBW, it shall be considered to be an indication of clearance of the busy condition. Should the unsolicited SR_RR message be a command message with the P bit set to 1, the appropriate response message with the F bit set to 1 shall be transmitted (see 6.3.5.12, Responding to command messages with the P bit set to 1) before the BBW may resume transmission of SR_I messages. The BBW shall not clear the poll outstanding condition. The BBW shall not stop Timer T1. If an unsolicited SR_SREJ message is received, then the BBW shall perform retransmissions as specified in 6.3.5.7.1, Receiving a SR_SREJ response message with the F bit set to 0.

6.3.5.9 BBW busy condition

When the BBW enters a busy condition, it shall transmit a SR_RNR message at the earliest opportunity. The SR_RNR message shall be a command frame with the P bit set to 1 if an acknowledged transfer of the busy condition indication is required; otherwise the SR_RNR message may be a command or response message. While in the busy condition, the BBW shall accept and process supervisory messages, accept and process the N(R) field of SR_I, SR_RR and SR_SREJ messages with the F bit set to 1, and return a SR_RNR response with the F bit set to 1 if it receives a supervisory

command or SR_I command message with the P bit set to 1. Received SR_I messages may be discarded or saved as specified in 6.3.5.3, Receiving an in-sequence SR_I message, and 6.3.5.5, Reception of out-of-sequence SR_I messages; however, SR_RR or SR_SREJ messages shall not be transmitted. To clear the busy condition, the BBW shall transmit a SR_RR message, with the N(R) field set to the current receive state variable V(R). The SR_RR message shall be a command message with the P bit set to 1 if an acknowledged transfer of the busy-to-non-busy transition is required; otherwise the SR_RR message may be either a command or response message.

6.3.5.10 Awaiting acknowledgement

If the Timer T1 runs out while waiting for the acknowledgement of a SR_I message from the remote BBW, the BBW shall restart Timer T1 and transmit an appropriate supervisory command message (SR_RR, SR_RNR) with the P bit set to 1. The BBW may transmit new SR_I messages after sending this enquiry message.

If the BBW receives a SR_SREJ response message with the F bit set to 1, the BBW shall restart Timer T1 and retransmit SR_I messages as specified in 6.3.5.7.2, Receiving a SR_SREJ response message with the F bit set to 1

If the BBW receives a SR_SREJ response message with the F bit set to 0, the BBW shall retransmit SR_I messages as specified in 6.3.5.7.2, Receiving a SR_SREJ response message with the F bit set to 1

If the BBW receives a SR_RR response message with the F bit set to 1, the BBW shall restart Timer T1 and retransmit SR_I messages as specified in 6.3.5.11, Receiving a SR_RR response message with the F bit set to 1.

If the BBW receives a SR_RR response message with the F bit set to 0, or a SR_RR command message or SR_I message with the P bit set to 0 or 1, the BBW shall not restart Timer T1, but use the received N(R) as an indication of acknowledgement of transmitted SR_I messages up to and including SR_I message numbered N(R)-1.

If Timer T1 runs out before a supervisory response message with the F bit set to 1 is received, the BBW shall retransmit an appropriate supervisory command message (SR_RR, SR_RNR) with the P bit set to 1. After N2 such attempts, the BBW shall initiate a link resetting procedure as described in 6.3.7, SR procedure for data link resetting.

6.3.5.11 Receiving a SR_RR response messages with the F bit set to 1

When receiving a SR_RR response message with the F bit set to 1, the BBW shall process the N(R) field as specified in 6.3.5.6, Receiving acknowledgements. If there are outstanding SR_I messages that are unacknowledged and no new SR_I messages have been transmitted subsequent to the last message with the P bit set to 1, then the BBW shall retransmit all outstanding SR_I messages except those that were sent after the message with the P bit set to 1 was sent. Retransmission shall conform to the following:

- a) If the BBW is transmitting a supervisory or SR_I message when it receives the SR_RR message, it shall complete that transmission before commencing transmission of the requested SR_I messages.
- b) If the BBW is transmitting an unnumbered command or response when it receives the SR_RR message, it shall ignore the request for retransmission.
- c) If the BBW is not transmitting any message when it receives the SR_RR message, it shall commence transmission of the requested SR_I messages immediately.

If any messages are retransmitted, then a poll shall be sent, either by transmitting a SR_RR command (or SR_RNR command if the BBW is in the busy condition) with the P bit set to 1 or by setting the P bit in the last retransmitted SR_I message.

The Timer T1 shall be stopped. If any SR_I messages are outstanding, then Timer T1 shall be started.

6.3.5.12 Responding to command messages with the P bit set to 1

When receiving a SR_RR or SR_RNR or SR_I command message with the P bit set to 1, the BBW shall generate an appropriate response message as follows:

- a) If the BBW is in the busy condition, it shall transmit a SR_RNR response message with the F bit set to 1.
- b) If there are some out-of-sequence messages in the receive buffer, then it shall transmit a SR_SREJ message with the F bit set to 1; N(R) shall be set to the receive state variable V(R) and the information field set to the sequence numbers of all missing SR_I messages, except V(R). If the list of sequence numbers is too large to fit in the information field of the SR_SREJ message, then the list shall be truncated by including only the earliest sequence numbers.
- c) If there are no out-of-sequence messages in the receive buffer, then a SR_RR response message with the F bit set to 1 shall be sent.

6.3.6 SR conditions for data link resetting or data link re-initialization (data link set-up)

6.3.6.1

When a BBW receives, during the information transfer phase, a message that is not valid (see 5.2.4.5.4) with one of the conditions listed in 5.2.4.4.9, the BBW shall request the remote BBW to initiate a data link resetting procedure by transmitting a SR_FRMR response to the remote BBW as described in 6.3.7.3.

6.3.6.2

When the BBW receives, during the information transfer phase, a SR_FRMR response from the remote BBW, the BBW shall either initiate the data link resetting procedures itself as described in 6.3.7.2 or return a SR_DM response to ask the remote BBW to initiate the data link set-up (initialization) procedure as described in 6.3.4.1, Data link set-up. After transmitting a SR_DM response, the BBW shall enter the disconnected phase as described in 6.3.4.4.2.

6.3.6.3

When the BBW receives, during the information transfer phase, a SR_UA response, or an unsolicited response with the F bit set to 1, the BBW may either initiate the data link resetting procedures itself as described in 6.3.7.2, or return a SR_DM response to ask the remote BBW to initiate the data link set-up (initialization) procedure as described in 6.3.4.1, Data link set-up. After transmitting a SR_DM response, the BBW shall enter the disconnected phase as described in 6.3.4.4.2.

6.3.6.4

When the BBW receives, during the information transfer phase, a SR_DM response from the remote BBW, the BBW shall either initiate the data link set-up (initialization) procedure as described in 6.3.4.1, Data link set-up, or return a SR_DM response to ask the remote BBW to initiate the data link set-up (initialization) procedures as described in 6.3.4.1, Data link set-up. After transmitting a SR_DM response, the BBW shall enter the disconnected phase as described in 6.3.4.4.2.

6.3.7 SR procedure for data link resetting

6.3.7.1

The data link resetting procedure is used to initialize both directions of information transfer according to the procedure below. The data link resetting procedure only applies during the information transfer phase.

6.3.7.2

Either BBW may initiate a data link reset procedure. The data link reset procedure indicates a clearance of a BBW and/or remote BBW busy condition, if present.

The remote BBW shall initiate a data link resetting by transmitting a SR_SM command to the BBW. If, upon correct receipt of the SR_SM command, the BBW determines that it is able to continue in the information transfer phase, it shall return a SR_UA response to the remote BBW, shall reset its send and receive state variables V(S) and V(R) to zero, and shall remain in the information transfer phase. If, upon the receipt of the SR_SM command, the BBW determines that it cannot remain in the information transfer phase, it shall return a SR_DM response as a denial to the resetting request and shall enter the disconnected phase.

The BBW shall initiate a data link resetting by transmitting a SR_SM command to the remote BBW and starting its Timer T1 (see 6.3.8.1, Timer T1). Upon reception of a SR_UA response from the remote BBW, the BBW shall reset its send and receive state variables V(S) and V(R) to zero, shall stop its Timer T1, and shall remain in the information transfer phase. Upon reception of a SR_DM response from the remote BBW as a denial to the data link resetting request, the BBW shall stop its Timer T1 and shall enter the disconnected phase.

The BBW, having sent a SR_SM command shall ignore and discard any messages received from the remote BBW except a SR_SM or SR_DISC command, or a SR_UA or SR_DM response. The receipt of a SR_SM or SR_DISC command from the remote BBW shall result in a collision situation that is resolved per 6.3.4.5, Collision of unnumbered commands above. Messages other than the SR_UA or SR_DM response sent in response to a received SR_SM or SR_DISC command shall be sent only after the data link is reset and if no outstanding SR_SM command exists.

After the BBW sends the SR_SM command, if a SR_UA or SR_DM response is not received correctly, Timer T1 shall run out in the BBW. The BBW shall then resend the SR_SM command and shall restart Timer T1. After N2 attempts to reset the data link, the BBW shall initiate appropriate higher layer recovery action and shall enter the disconnected phase. The value of N2 is defined in , Maximum number of attempts to complete a transmission N2 below.

6.3.7.3

The BBW may ask the remote BBW to reset the data link by transmitting a SR_FRMR response (see 6.3.6.1). After transmitting a SR_FRMR response, the BBW shall enter the message rejection condition.

The message rejection condition is cleared when the BBW receives a SR_SM command, a SR_DISC command, a SR_FRMR response, or a SR_DM response; or if the BBW transmits a SR_SM command, a SR_DISC command, or a SR_DM response. Other commands received while in the message rejection condition shall cause the BBW to retransmit the SR_FRMR response with the same information field as originally transmitted.

The BBW may start Timer T1 on transmission of the SR_FRMR response. If Timer T1 runs out before the message rejection condition is cleared, the BBW may retransmit the SR_FRMR response, and restart T1. After N2 attempts (time outs) to get the remote BBW to reset the data link, the BBW may reset the data link itself as described in 6.3.7.2. The value of N2 is defined in , Maximum number of attempts to complete a transmission N2 below.

In the message rejection condition, SR_I messages and supervisory messages shall not be transmitted by the BBW. Also, the BBW shall ignore and discard the N(S) and information fields of any received SR_I messages and the N(R) fields of any received SR_I messages and supervisory messages. When an additional SR_FRMR response shall be transmitted by the BBW as a result of the receipt of a command message while Timer T1 is running, Timer T1 shall continue to run. Upon reception of a SR_FRMR response (even during a message rejection condition), the BBW shall initiate a resetting procedure by transmitting a SR_SM command as described in 6.3.7.2, or shall trans-

mit a SR_DM response to ask the remote BBW to initiate the data link set-up procedure as described in 6.3.4.1, Data link set-up, and enter the disconnected phase.

6.3.8 List of SR system parameters

6.3.8.1 Timer T1

The same value of the Timer T1 shall be made known and agreed to by all BBWs.

The period of Timer T1, at the end of which retransmission of a message may be initiated (see 6.3.4, SR procedures for data link set-up and disconnection and 6.3.5, Procedures for information transfer using multi-selective reject above), shall take into account whether T1 is started at the beginning or the end of the transmission of a message.

The proper operation of the procedure requires that the transmitter's Timer T1 be greater than the maximum time between transmission of a message (SR_SM, SR_DISC, SR_I, or supervisory command, or SR_DM or SR_FRMR response) and the reception of the corresponding message returned as an answer to that message (SR_UA, SR_DM, or acknowledging message). Therefore, the receiver should not delay the response or acknowledging message returned to one of the above messages by more than a value T2, where T2 is a system parameter (6.3.8.2, Parameter T2).

The BBW shall not delay the response or acknowledging message returned to one of the above remote BBW messages by more than a period T2.

6.3.8.2 Parameter T2

The same value of the Parameter T2 shall be made known and agreed to by all BBWs.

The period of parameter T2 shall indicate the amount of time available at the BBW before the acknowledging message shall be initiated in order to ensure its receipt by the remote BBW, prior to Timer T1 running out at the BBWs (parameter T2 < Timer T1).

NOTE 7— The period of parameter T2 shall take into account the following timing factors: the transmission time of the acknowledging message, the propagation time over the access link, the stated processing times at the BBWs, and the time to complete the transmission of the message(s) in the BBW transmit queue that are neither displaceable nor modifiable in an orderly manner.

Given a value for Timer T1 for the BBWs, the value of parameter T2 shall be no larger than T1 minus 2 times the propagation time over the access data link, minus the message processing time at the BBW, minus the message processing time at the remote BBW, and minus the transmission time of the acknowledging message by the BBW.

Annex C provides guidelines for tuning the SR protocol parameters

6.3.8.3 Maximum number of attempts to complete a transmission N2

The same value of the N2 system parameter shall be made known and agreed to by the BBWs.

The value of N2 shall indicate the maximum number of attempts made by the BBW to complete the successful transmission of a message to the remote BBW.

6.3.8.4 Maximum number of outstanding SR_I messages k

The same value of the *k* system parameter shall be made known and agreed to by the BBWs.

The value of *k* shall indicate the maximum number of sequentially numbered SR_I messages that the BBWs may have outstanding (i.e., unacknowledged) at any given time. The value of *k* shall never exceed 32767 for modulo 32768 operation.

NOTE 8— Annex C provides guidelines for selecting appropriate values of *k* and message size to maximize the efficiency of links with long propagation delays.

6.4 Simple Flow Control (SFC)

The Simple Flow Control (SFC) is a mechanism that requests the remote BBW from pausing transmission for a time period defined by the number of time units in the PAUSE bytes of the BBW_Header. Each time unit corresponds to a 512-bit (64 bytes) transmission time. A zero value in the PAUSE bytes indicates that the remote BBW does not need pause transmission. (The effect is the same as non use of flow control). If a subsequent message is received with the PAUSE field set to a value, then the pause time is reset to this new value.

Use of SFC is optional and may result in the remote BBW in simply ignoring the PAUSE bytes. In this case, pausing is not accomplished.

7 FC-BB-2_ATM Structure and Concepts

7.1 Applicability

This clause only applies to FC-BB-2_ATM.

Clause 4 discussed the FC-BB-2_ATM Reference Model. This clause discusses the FC-BB-2_ATM Functional Model. Other FC-BB-2_ATM applicable clauses include Clause 5 (Messages and Formats), Clause 6 (SR Protocol Procedures), Clause 8 (Mapping and Encapsulation), and Clause 9 (Service Considerations).

7.2 FC-BB-2_ATM Overview

FC-BB-2_ATM is a Fibre Channel backbone transport protocol that tunnels AAL5 encapsulated FC frames across the ATM network. Figure 6 shows a network configuration consisting of three FC-BB-2_ATM devices. A FC-BB-2_ATM device has interfaces to both the ATM and the FC Network. The FC network interface supports multiple B_Ports. The model applies equally to both private and public ATM networks.

FC-BB-2_ATM devices that support B_Port do not require FC Switching. The FC-BB-2_ATM protocol communication occurs between pairs of FC-BB-2_ATM devices. Although, the communication occurs between pairs of FC-BB-2_ATM devices, a single FC-BB-2_ATM device may communicate with more than one device at the same time.

NOTE 9– The current scheme allows a FC-BB-2_ATM device to independently connect to more than one FC-BB-2_ATM device, but does not specify a point-to-multipoint connection.

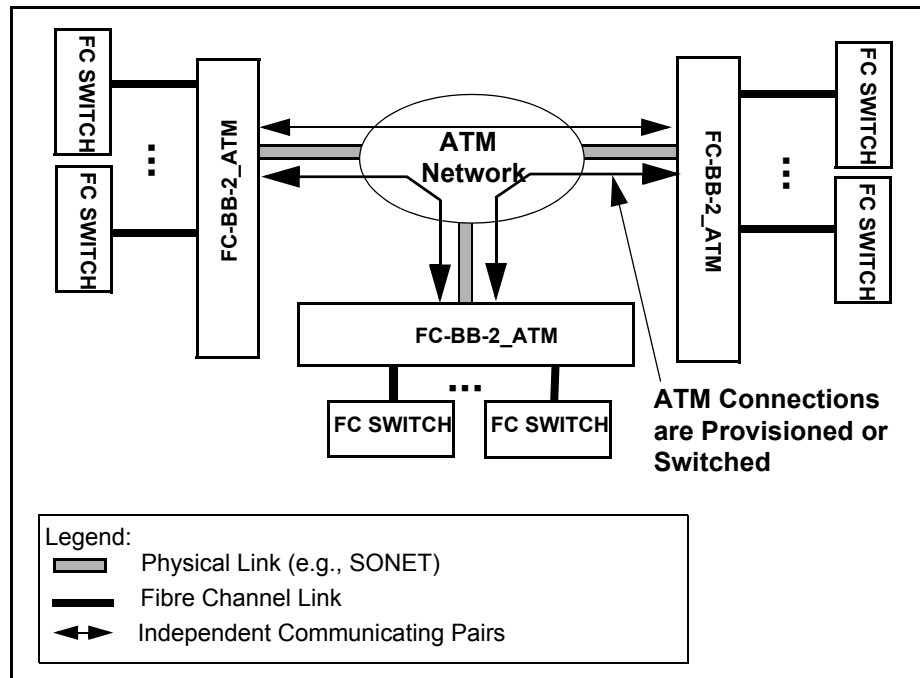


Figure 6 – FC-BB-2_ATM Network Configuration

The FC-BB-2_ATM protocol creates BBW messages that consist of an 8-byte LLC/SNAP Header and a 4-byte BBW_Header followed by the BBW message payload. The specific format and content of the BBW message payload depends on the type of flow control protocol used. The BBW message payloads carry byte-encoded SOF/EOF delimited Class 2, 3, 4 or F FC frames.

The BBW messages are encapsulated in ATM Adaptation Layer 5 (AAL5) format for carriage over the ATM Network. The AAL5 encapsulated BBW messages are segmented into ATM cells and for-

warded to the proper destination ATM address. FC-BB-2_ATM does not interpret the data content of the FC frames other than capturing and retaining their SOF/EOF identities in the encapsulated FC frame. As such, FC Sequences and Exchanges are not visible to the FC-BB-2_ATM Protocol. All AAL5 encapsulated FC frames are transparently transported over the ATM network.

The LLC/SNAP Header indicates the payload type as Fibre Channel. (See 5.2.1.) The BBW_Header indicates the type of flow control used - Selective Retransmission (SR), Simple Flow Control (SFC) or none. The SR Protocol makes the transport of FC frames between two FC-BB-2_ATMs reliable. The SR Protocol supports both flow control and error recovery functions. Use of the SR protocol is optional. When SR Flow Control is used, the 4 byte BBW_Header is followed by a 4-byte SR_Header which is prefixed at the begin of the BBW message payload. (See 5.2.4). The SFC Protocol provides a mechanism to temporarily pause the transmission of frames from a remote BBW device. Use of the SFC protocol is optional. When SFC is used, the 4 byte BBW_Header is directly followed by the BBW message payload. No SFC header is prefixed or used. (See 5.2.3.)

In-order delivery is guaranteed within the scope of an ATM Virtual Connection (VC) and frames are transmitted from the FC-BB-2_ATM in the same order as they are received.

7.3 FC-BB-2_ATM Functional Model

7.3.1 B_Port Network Interface

Figure 7 shows the Functional Model of the FC-BB-2_ATM. The Fibre Channel interface nominal port rate is assumed to be full-rate, unless otherwise specified.

The FC-BB-2_ATM FC Interface supports one or more B-Ports thus requiring the support of the FC-0, FC-1, and FC-2 Levels. The B_Ports in general connect to different external FC switches, but connectivity to the same FC switch is also allowed. B_Ports are uniquely identified by an 8-byte B_Port_Name.

The FC-BB-2 initialization occurs across the B_Port interface facing the FC network. The initialization of any generic B_Port is described in FC-SW-3. A B_Port indicates its support for the ELP/ESC Parameters using the ELP/ESC exchange protocol that is capable of parameter negotiation. Since FC-BB-2 does not support Class 1, the Class 1 Port Parameter VAL bit in the ELP shall be set to 0 (invalid). An ELP received at a B_Port may be rejected (SW_RJT) due to many reasons, including Port-mismatch.

NOTE 10– Initialization across the ATM WAN interface may use mechanisms similar to the one described in 13.4.3.3.2.1.

7.3.2 ATM Network Interface

The ATM Network Interface includes the PHY, ATM, and Adaptation Layers. The basic FC-BB-2_ATM Reference Model supports one ATM port using different media types and/or different rates. The ATM Adaptation Layer-5 (AAL5) is used for BBW message transport while the SAAL Adaptation is used for ATM signaling. FC-BB-2_ATM may use either provisioned Permanent Virtual Circuit (PVC) or Switched Virtual Connection (SVC) to transport messages. SVC requires the use of the User Network Interface (UNI) Signaling Protocol specifying the desired Service Category, QoS and Traffic Parameters. Both Public UNI and Private UNI shall be supported.

7.3.3 Mapping and Encapsulation

The FC-BB-2_ATM creates the 8-byte LLC/SNAP Header and the 4-byte BBW_Header that are prefixed to the BBW message payload. The specific format and content of the BBW message payload depends on the type of flow control protocol used. The BBW message payload carries the byte-encoded SOF/EOF delimited Class 2, 3, 4 or F FC frames.

When flow control is not used, the FC-BB-2_ATM sets the PAUSE field in the BBW_Header to a zero value and the Flow Control Type to SFC.

NOTE 11– This setting of Flow Control Type in combination with a zero value in the PAUSE field amounts to non use of any flow control protocol and avoids specifying another flow control type encoding

When SFC is used, the FC-BB-2_ATM sets the PAUSE field to an appropriate value indicating the number of 512-time units to pause transmission. See 6.4.

When SR Protocol is used, the FC-BB-2_ATM prefixes a 4 byte SR Header at the begin of an encapsulated frame that is mapped into the payload of the SR_I message. The SR Header indicates the type of SR message type along with other control information. See 6.2 and 6.3.

See Clause 8 for details on encapsulation using AAL5.

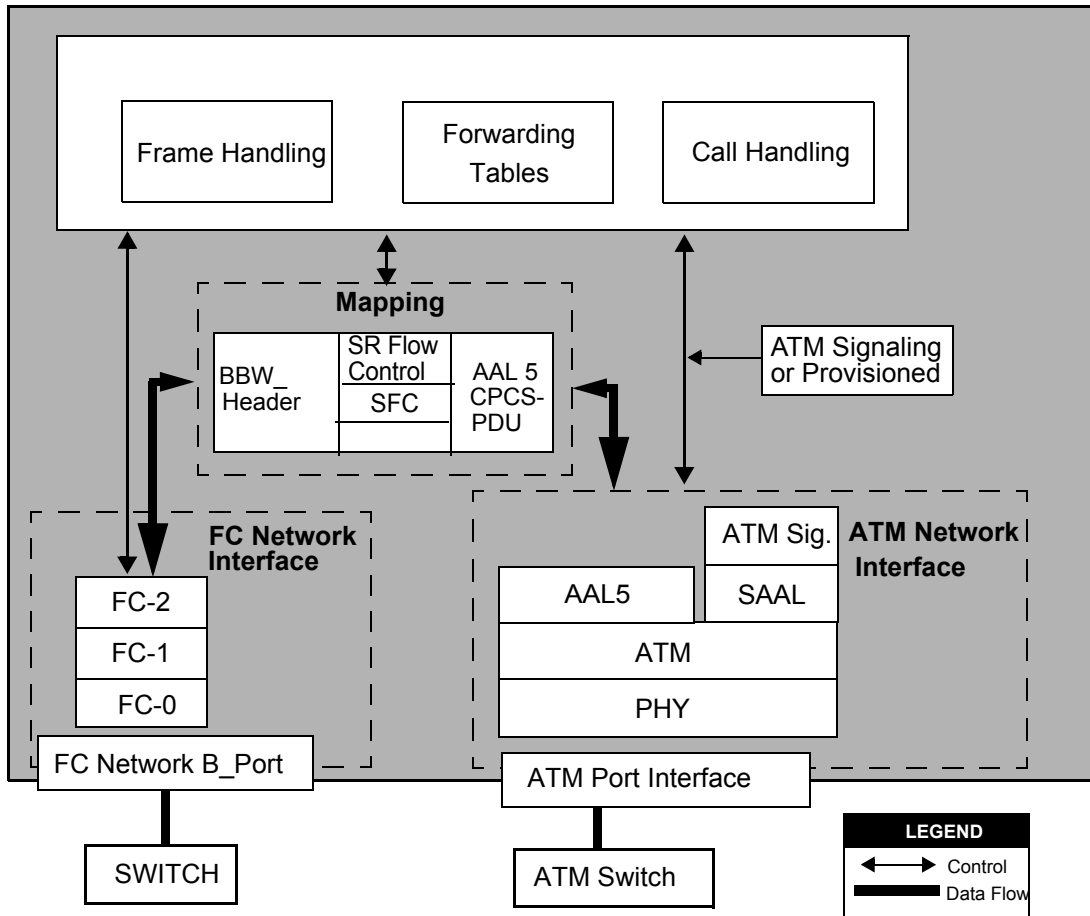


Figure 7 – FC-BB-2_ATM Functional Block Diagram

7.3.4 FC-BB-2_ATM Forwarding

FC-BB-2_ATM forwards FC frames that enter its B_Ports to the remote FC-BB-2_ATM using a mapping table that contains a list of FC-BB-2_ATM ATM addresses corresponding to a list of D_ID addresses.

7.3.5 Call Handling and ATM Service

FC-BB-2_ATMs supports the use of both Provisioned Permanent Virtual Connections (PVCs) and Switched Virtual Connections (SVCs) to transport messages. If PVCs are used then no ATM signaling is required and connections are provisioned (preconfigured).

If SVC is used then Call Handling initiates the ATM User Network Interface (UNI) Signaling protocol to set up a Virtual Connection (VC) [21]; the VC is torn down after its use. The FC-BB-2_ATM shall use the ATM UNI signaling connection request messages to establish a connection and a traffic contract.

The traffic contract establishes the FC-BB-2_ATM defined QoS and traffic parameters. If the requested connection is acceptable to the network, then a connection is set up between the FC-BB-2_ATMs. FC-BB-2_ATM shall support UNI 3.1 and higher. A dedicated channel (Virtual Path Identifier (VPI) = 0 and Virtual Channel Identifier (VCI) =5) is reserved for signaling between the end user and the interfacing ATM device (switch). ATM connections allow traffic to flow in one or both directions (unidirectional or bi-directional) with the bandwidth the same or different in each direction. FC-BB-2_ATM requires bidirectional connectivity.

FC-BB-2_ATM uses Variable Bit Rate Non Real Time traffic (VBR-NRT) service (see 9.5). AAL5 is particularly well suited for carriage of VBR-NRT. VBR-NRT ATM Service provides cell loss and bandwidth guarantees.

FC-BB-2_ATM recommends use of a single Virtual Circuit (VC) (see 9.4). Use of additional VCs to address special traffic QoS requirements is allowed but not recommended. If SR or SFC Flow Control is used, then flow control is separately applied to each VC.

FC-BB-2_ATM recommends allocating a minimum bandwidth for each VC that is used in order to avoid starvation (see 9.4). However, the Service discipline (prioritization) for the VCs is implementation specific and beyond the scope of this standard.

In-order delivery is guaranteed within the scope of the ATM Virtual Connection (VC). Frames shall be shipped from the FC-BB-2_ATM in the same order as they are received.

7.3.6 Frame Handling

Frame Handling is mainly concerned with the following two tasks:

- Processing the incoming FC frames from the external switch that emerges from the FC-2 Level that has to be transported across the WAN. Processing includes tasks such as BBW header and message generation, and mapping to AAL5 CPCS.
- Processing the FC-BB-2_ATM message that has successfully made it across the WAN and that is to be sent to the external switch. Processing includes decoding the AAL5 CPCS, decoding the BBW message and removing the message headers.

8 Mapping and Message Encapsulation using AAL5

8.1 Applicability

This Clause only applies to FC-BB-2_ATM.

8.2 Overview

BBW messages are transparently transported over the ATM WAN. However, before it may be transported, it first has to be adapted. This adaptation is done using the ATM Adaptation Layer (AAL5). The AAL5 encapsulated BBW message is then segmented into ATM cells and routed to the proper destination ATM address.

8.3 Mapping BBW messages to AAL5

The BBW message is first mapped to a null AAL5 Service Specific Convergence Sublayer (SSCS) and then to a Common Part Convergence Sublayer (CPCS) to form the AAL5 CPCS-PDU (max size 2160/2164 bytes). (See Note 12 below.) The AAL5 CPCS-PDU is padded (if necessary up to 47 bytes) and then appended with an 8-byte CPCS Trailer. The CPCS-PDU, Pad, and CPCS-Trailer is then segmented into 48 bytes to form the Segmentation and Reassembly PDU (SAR -PDU). A 5-byte ATM Cell Header is attached to each SAR PDU to form an ATM cell.

CPCS-PDU: The BBW message maps into this field that consists of the LLC/SNAP Header, BBW_Header, and the BBW message payload.

CPCS-Pad: A CPCS-Pad ensures an exact mapping of the CPCS-PDU into SAR 48-byte payloads. A CPCS-Pad may range from 0-47 bytes. The maximum Pad value of 47 bytes never occurs when the CPCS-PDU carries the BBW message payload because the payload is always a multiple of 4 bytes and aligned on a 4-byte boundary

CPCS-Trailer: A CPCS-Trailer is 8 bytes long and consists of a 1-byte User-to-User (UU) field, a 1-byte Common Part Indicator (CPI) field, a 2-byte length field, and a 4-byte CRC check sum field.

The UU and CPI fields are currently not used. The CPCS-PDU length field indicates the length in bytes of the CPCS-PDU payload. The length indicates the useful payload size. Therefore, the CPCS-PDU size may vary with byte increments. The AAL5 CRC field is set as defined in ITU Recommendation.

Table 20 – Mapping of BBW messages to AAL5 CPCS

Field	Item	Size Bytes
CPCS-PDU	LLC/SNAP Header	8
	BBW_Header	4
	BBW message payload (See Note 12)	Max: 2148/2152 (See Note 13)
CPCS-Pad		0-47
CPCS-Trailer	Reserved (CPCS-UU, not used)	1
	Reserved (CPI, not used)	1
	CPCS-PDU Length (in bytes)	2
	CPCS-PDU CRC	4

NOTE 12– If SR is used, then only the SR_I, SR_SREJ, and SR_FRMR carry a non-zero payload.

NOTE 13– The maximum CPCS-PDU value indicated in the table is based on the maximum Fibre Channel frame size. CPCS-PDU for other non Fibre Channel transport may be much larger and up to 65535 bytes. The maximum of 2148 bytes of BBW message payload is due to a maximum of 2112 bytes of FC frame payload, 4 bytes of SOF, 24 bytes of FC Header, 4 bytes of EOF, 4 bytes of CRC. If SR is used then 4 bytes of SR_Head-er yields a total maximum of 2152 bytes.

Figure 8 illustrates the AAL5 Mapping for a FC frame when SFC is used.

Figure 9 illustrates the AAL5 mapping for a FC frame when SR is used.

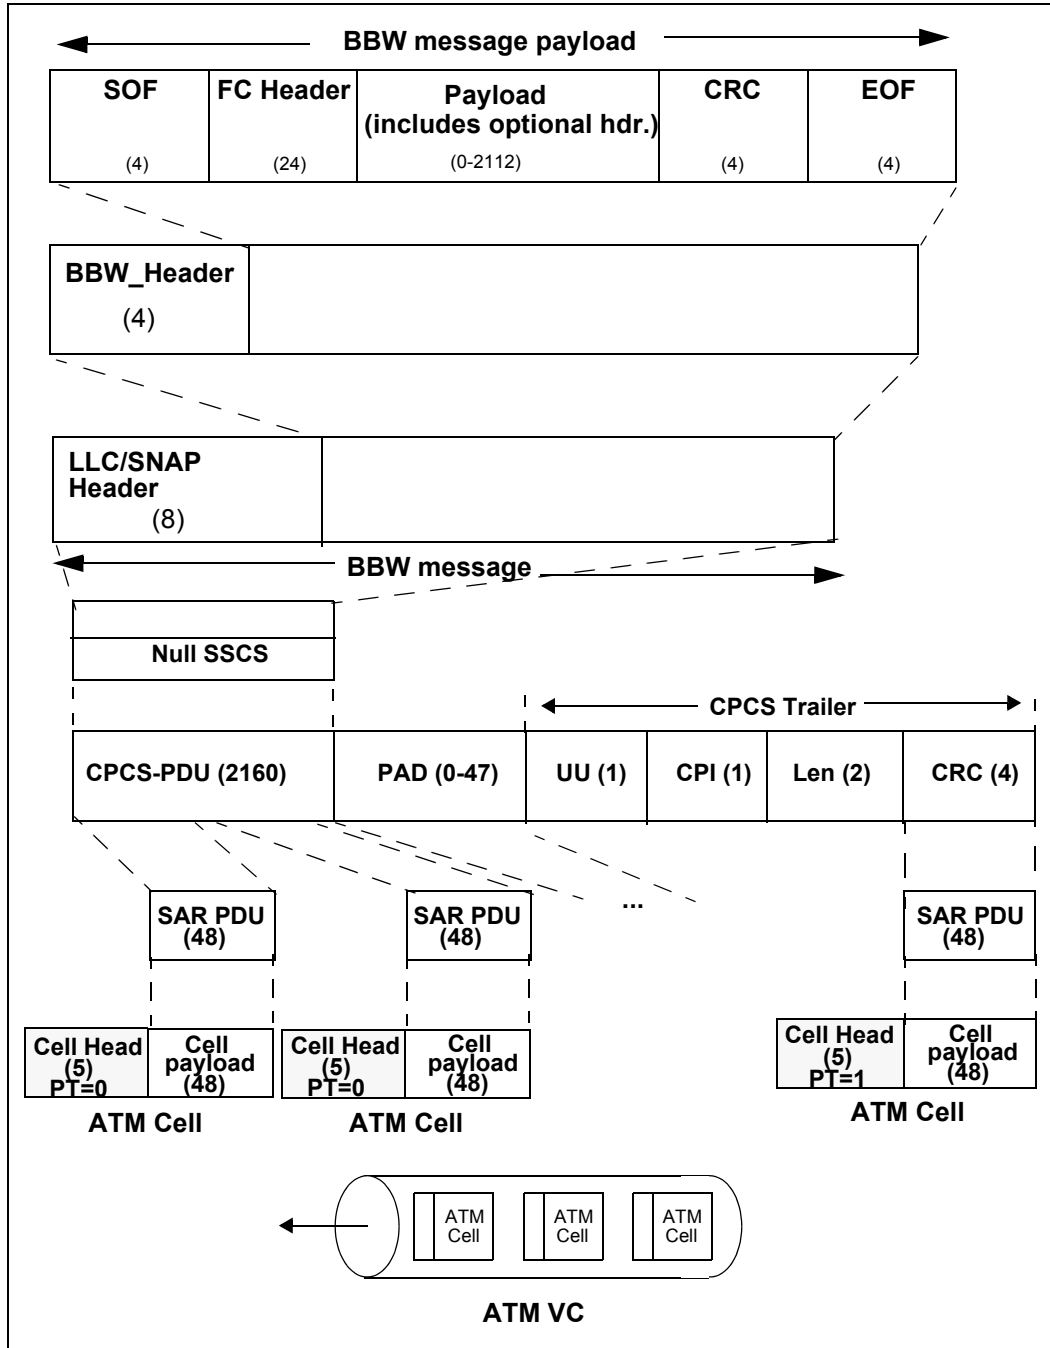


Figure 8 – AAL5 Mapping of a BBW message with SFC

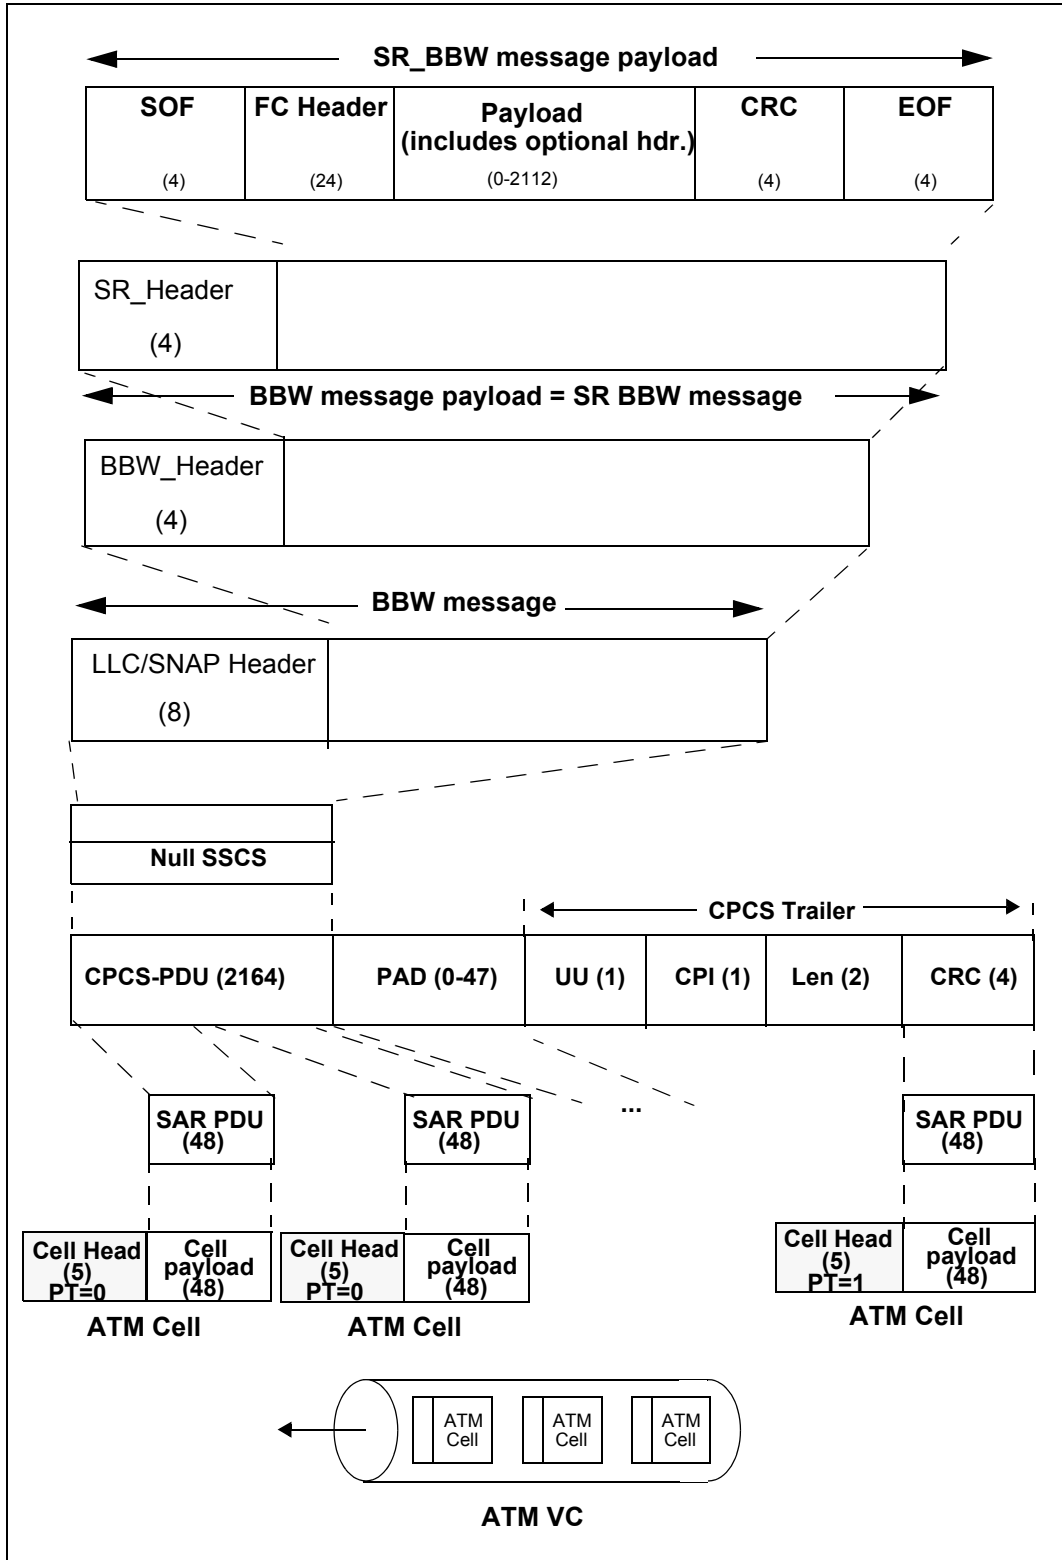


Figure 9 – AAL5 Mapping of a BBW message with SR

9 FC-BB-2_ATM Service Considerations

9.1 Applicability

This Clause only applies to FC-BB-2_ATM.

9.2 ATM Service Type

Different types of ATM service provide different levels of service features. FC-BB-2_ATM recommends use of the VBR-NRT ATM Service or better. (See AnnexNote B.)

9.3 Latency Delay and Timeout Value

FC-BB-2_ATM and the ATM network introduce latency delays that warrant special considerations with respect to Fibre Channel E_D_TOV and R_A_TOV values. The total path delay between a source FC-BB-2_ATM and a destination FC-BB-2_ATM consists of the latency delay components due to the queuing time at the FC-BB-2_ATM devices and all intermediate ATM switches, cell transmission time, propagation time and SVC setup time if applicable. It is recommended that this total path delay be less than 1/2 E_D_TOV to conform to normal Fibre Channel time out values.

NOTE 14– VBR-NRT does not provide delay guarantees; delay guarantees are practically realized by Service Level Agreements (SLAs) with the ATM Service Provider.

9.4 Bandwidth Sharing and Allocation

In ATM, bandwidth sharing is accomplished by multiplexing different upper layer traffic (e.g., Fibre Channel, IP). Multiplexing different upper layer protocol traffic may occur in two ways: multiplexing within a single VC, multiplexing using different VCs. The latter method is not recommended.

Multiplexing within a single VC, also referred to as VC Multiplexing, is applicable for all traffic intended for the same destination and when using the same ATM service category. Upper layer multiplexed protocol data is distinguished based on the BBW_Header. The biggest reason to use VC multiplexing is to minimize the number of VCCs established especially in a PVC environment. FC-BB-2_ATM recommends using a single VC to multiplex all traffic.

Use of more than one VC to the same destination to address special traffic QoS requirements is allowed but introduces an increased level of complexity and is therefore not recommended. The SR Flow Control protocol is separately applied to each VC in such a case. A minimum bandwidth allocation is recommended for each VC that is used in order to avoid starvation. FC-BB-2_ATM does not specify any particular service discipline when more than one VC is used, but recommends a minimum bandwidth for each VC, to protect it from starvation. This is illustrated in Figure 10. The Service discipline (prioritization) for the VCs is implementation specific and outside the scope of FC-BB-2.

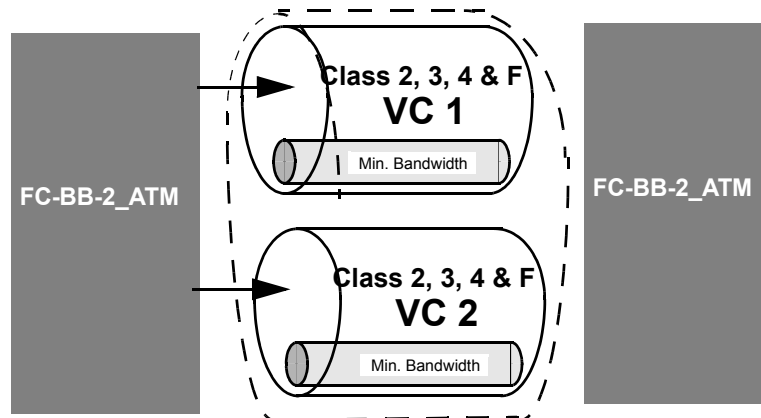


Figure 10 – Recommended ATM Bandwidth Allocation for multiple VCs

9.5 Quality of Service (QoS)

FC-BB-2_ATM specifies the use of VBR-NRT service for all VCs. Table 21 shows the QoS parameters and traffic descriptors specific to VBR-NRT and the guarantees provided by this service.

VBR-NRT service is best suited for non-time-based critical data, that require guarantees for loss and bandwidth but not delay. This service matches the requirements of FC Classes 2, 3, 4, and F.

QoS is a term used to refer to the set of performance characteristics of the contracted ATM connection. Although, a total of 6 QoS parameters are defined and available with other ATM Services, VBR-NRT only specifies a single QoS parameter - Cell Loss Ratio.

NOTE 15- Some QoS parameters specified with other ATM Services include: Peak-to-peak Cell Delay Variation (CDV), Maximum Cell Transfer Delay (maxCTD). maxCTD provides delay guarantees. See B for details.

ATM Traffic Descriptor is a term used to describe the traffic characteristics of an ATM connection. A Connection Traffic Descriptor includes a Source Traffic Descriptor, CDV Tolerance (CDVT), and a Conformance definition. A Source Traffic Descriptor is described by four parameters: Peak Cell Rate (PCR), Sustainable Cell Rate (SCR), Maximum Burst Size (MBS), and a Minimum Cell Rate (MCR). See Annex B for more details. Service guarantees are realized by these traffic descriptors.

Bandwidth guarantees are achieved by SCR, PCR, and MBR

Table 21 – ATM VBR-NRT Service Specification

ATM Traffic Descriptors	VBR-NRT Service Category	Remark
QoS Parameters	CLR*	Cell Loss Ratio; guarantees Loss
Source Traffic Descriptors: (SCR, PCR, MBS guarantee bandwidth)	PCR*, CDVT*, SCR, MBS MCR	Peak Cell Rate, CDV Tolerance, Sustainable Cell Rate, Maximum Burst Size, Minimum Cell Rate;
Conformance Definition	GCRA*	Generic Cell Rate Algorithm (Leaky Bucket Algorithm)
NOTE 16- * Items are supplied by telco and are negotiable		
NOTE 17- Cell Transfer Delay (CTD) (not in table) a QoS parameter associated with VBR-RT is also negotiable.		

9.6 Delivery Order

FC-BB-2_ATM shall guarantee in-order delivery of frames within a VC. No other ordering relationship between VCs is normally preserved or assumed. When the number of VCs is greater than 1, than the traffic management entity within the FC-BB-2_ATM device shall ensure that using separate VCs does not result in out-of-order delivery. In other words, once message transmission begins on a VC, then it shall continue using the same VC until completion of the message.

NOTE 18- The out-of-sequence delivery problem associated with datagram networks is not present here. This benefit is a consequence of the strict requirement in ATM that requires all cells to always follow the same route during the Call's duration. However, the possibility of missing or errored messages still remains and is addressed by the SR protocol.

9.7 Loss and Flow Control

ATM networks are lossy and they may drop cells, typically due to network congestion. When a cell loss occurs, the end applications are expected to recover from this loss. Recovery from such losses occurs at the FC-BB-2_ATM devices using the SR protocol that supports error recovery.

NOTE 19– The SFC protocol has no error recovery support

Use of a flow control protocol (SFC or SR) at the FC-BB device allows to cope up with the speed mismatches between the FC and the ATM interface.

10 FC-BB-2_SONET Structure and Concepts

10.1 Applicability and Related Clauses

This Clause only applies to FC-BB-2_SONET.

Clause 4 discussed the FC-BB-2_SONET Reference Model. This clause discusses the FC-BB-2_SONET Functional Model. Other FC-BB-2_SONET applicable clauses include Clause 5 (Messages and Formats), Clause 6 (SR Protocol Procedures), Clause 11 (Mapping and Encapsulation), and Clause 12 (Service Considerations).

10.2 FC-BB-2_SONET Overview

FC-BB-2_SONET is a Fibre Channel backbone transport protocol that tunnels HDLC encapsulated FC frames across the SONET/SDH network. Figure 11 shows a network configuration consisting of three FC-BB-2_SONET devices. A FC-BB-2_SONET device has interfaces to both the SONET and the FC Network. The FC network interface supports multiple B_Ports. The model applies equally to both private and public SONET/SDH networks.

FC-BB-2_SONET devices that support B_Port do not require FC Switching. The FC-BB-2_SONET protocol communication occurs between pairs of FC-BB-2_SONET devices. Although, the communication occurs between pairs of FC-BB-2_SONET devices, a single FC-BB-2_SONET device may communicate with more than one device at the same time.

NOTE 20– The current scheme allows a FC-BB-2_SONET device to independently connect to more than one FC-BB-2_SONET device, but does not specify a point-to-multipoint connection

No distinction is made in this document regarding the topology of the SONET/SDH network; be it point-to-point using PTE pairs, HUB networks or ring architectures. The current model supports a configuration of one or more point-to-point connections only.

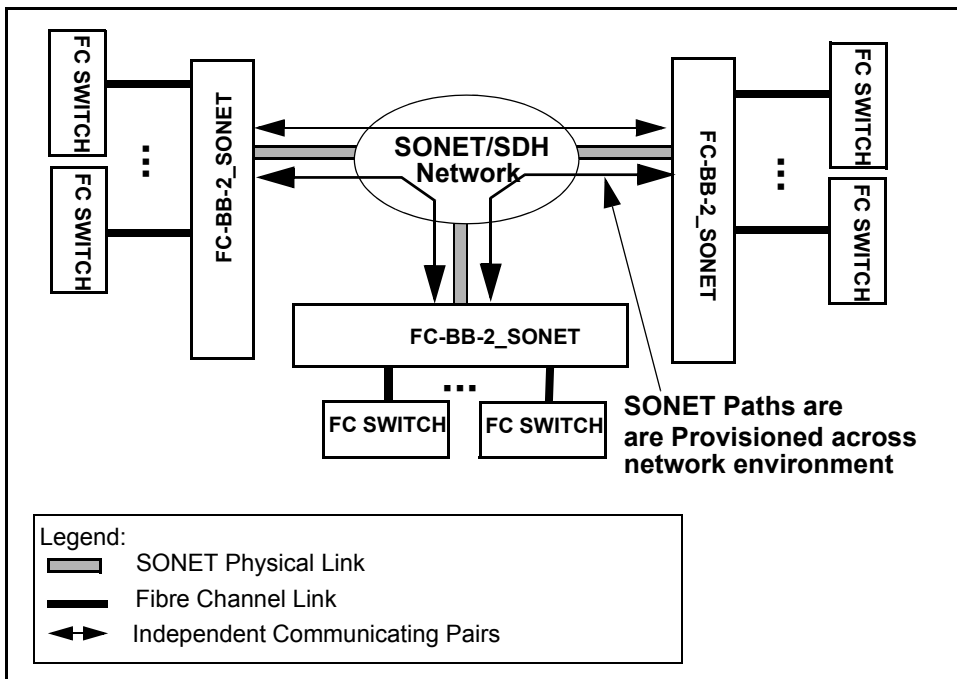


Figure 11 – FC-BB-2_SONET Network Configuration

The FC-BB-2_SONET protocol creates BBW messages that consist of an 8-byte LLC/SNAP Header and a 4-byte BBW_Header followed by the BBW message payload. The specific format and content

of the BBW message payload depends on the type of flow control protocol used. The BBW message payloads carry byte-encoded SOF/EOF delimited Class 2, 3, 4 or F FC frames.

The BBW messages are encapsulated in HDLC-like format for carriage over SONET/SDH network. The HDLC encapsulated BBW messages are mapped into SPE/Virtual Containers and finally transmitted to the destination. HDLC encapsulation is the typical method of preparing frames for transmission over SONET/SDH and is described in RFCs 1662 and 2615. FC-BB-2_SONET does not interpret the data content of the FC frames other than capturing and retaining their SOF/EOF identities in the encapsulated FC frame. As such, FC Sequences and Exchanges are not visible to the FC-BB-2_SONET Protocol. All HDLC encapsulated FC frames are transparently transported over the SONET/SDH network.

Prior to a FC-BB-2_SONET transmitting data to a remote FC-BB-2_SONET, the required provisioning of the SONET/SDH Path to the remote BBW needs to be completed. The details of this configuration are dependent upon the network topology and are beyond the scope of this document.

All FC_frames are encapsulated with the All-Stations address a binary sequence 11111111b (hexadecimal FFh) in the HDLC header therefore there is no requirement for a FC-BB-2_SONET to examine the destination address (D_ID) field in the Fibre Channel frame header. Frames are simply forwarded to the attached FC-BB-2_SONET egress device across the SONET network.

The LLC/SNAP Header indicates the payload type as Fibre Channel. (See 5.2.1.) The BBW_Header indicates the type of flow control used - Selective Retransmission (SR), Simple Flow Control (SFC) or none. The SR Protocol makes the transport of FC frames between two FC-BB-2_SONETs reliable. The SR Protocol supports both flow control and error recovery functions. Use of the SR protocol is optional. When SR Flow Control is used, the 4 byte BBW_Header is followed by a 4-byte SR_Header which is prefixed at the begin of the BBW message payload. (See 5.2.4.) The SFC Protocol provides a mechanism to temporarily pause the transmission of frames from a remote BBW device. Use of the SFC protocol is optional. When SFC is used, the 4 byte BBW_Header is directly followed by the BBW message payload. No SFC header is prefixed or used. (See 5.2.3.)

In-order delivery is guaranteed for each BBW message and frames shall be transmitted from the FC-BB-2_SONET in the same order as they are received.

10.3 FC-BB-2_SONET Functional Model

10.3.1 Fibre Channel Network Interface

Figure 12 shows a Functional Model of the FC-BB-2_SONET. The Fibre Channel interface nominal port rate is assumed to full-rate, unless otherwise specified.

The FC-BB-2_SONET FC Interface supports one or more B-Ports thus requiring the support of the FC-0, FC-1, and FC-2 Levels. The B_Ports in general connect to different external FC switches, but connectivity to the same FC switch is also allowed. B_Ports are uniquely identified by an 8-byte B_Port_Name.

The FC-BB-2 initialization occurs across the B_Port interface facing the FC network. The initialization of any generic B_Port is described in FC-SW-3. A B_Port indicates its support for the ELP/ESC Parameters using the ELP/ESC exchange protocol that is capable of parameter negotiation. Since FC-BB-2 does not support Class 1, the Class 1 Port Parameter VAL bit in the ELP shall be set to 0 (invalid). An ELP received at a B_Port may be rejected (SW_RJT) due to many reasons, including Port-mismatch.

NOTE 21– Initialization across the SONET WAN interface may use mechanisms similar to the one described in 13.4.3.3.2.1.

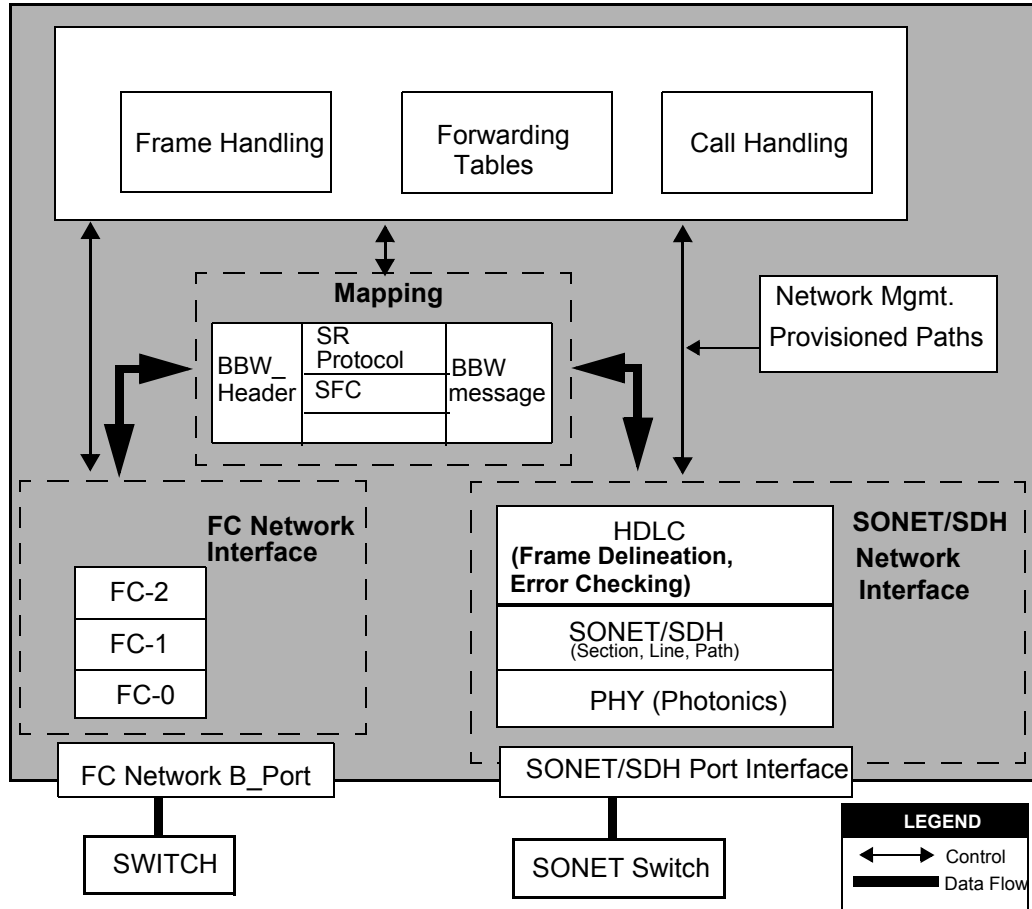


Figure 12 – FC-BB-2_SONET Functional Block Diagram

10.3.2 SONET Network Interface

The SONET rate is nominally assumed to be STS-3c/STM-1 at 155.52 Mbps and higher. In the case of STS-3c/STM-1 the available information bandwidth is 149.760 Mbps, which is the STS-3c/STM-1 SPE with Section, Line and Path overhead removed. This is the same super-rate mapping that is used for ATM and FDDI. While the STS-3c/STM-1 rate is specified as the basic rate, the mapping specified within this document is extended down to the STS-1 SONET rate (51.84 Mbps).

Higher signal rates shall conform to the SDH STM series, rather than the SONET STS series. The STM series progresses in powers of 4 (instead of 3), and employs fewer steps, which simplifies multiplexing and integration. For applications of Fibre Channel over SONET/SDH, it is envisioned that higher rates such as 622.08 Mbit/s and 2488.32 Gbit/s may be developed and deployed as indicated in Table 22.

Table 22 – SONET/SDH Data Rates

SONET	SDH Equivalent	Basic Rate
STS-3c-SPE	VC-4	155.52 Mbps
STS-12c-SPE	VC-4-4c	622.08 Mbps
STS-48c-SPE	VC-4-16c	2.4 Gbps
STS-192c-SPE	VC-4-64c	9.95 Gbps

Mappings for sub STS-1 rates and rates of STS-192c/STM-48 or greater requires further study and are beyond the scope of this document.

The SONET/SDH interface includes the Photonics, Section/Line/Path, and HDLC encapsulation layers. The basic FC-BB-2_SONET Reference Model supports one SONET port using different rates.

The HDLC layer is used to prepare FC frame payloads for transport in SONET/SDH payload envelopes.

10.3.3 Mapping and Encapsulation

The FC-BB-2_SONET creates the 8-byte LLC/SNAP Header and the 4-byte BBW_Header that are prefixed to the BBW message payload. The specific format and content of the BBW message payload depends on the type of flow control protocol used. The BBW message payload carries the byte-encoded SOF/EOF delimited Class 2, 3, 4 or F FC frames.

When flow control is not used, the FC-BB-2_SONET sets the PAUSE field in the BBW_Header to a zero value and the Flow Control Type to SFC.

NOTE 22– This setting of Flow Control Type in combination with a zero value in the PAUSE field amounts to non use of any flow control protocol and avoids specifying another flow control type encoding

When SFC is used, the FC-BB-2_SONET sets the PAUSE field to an appropriate value indicating the number of 512-time units to pause transmission. See 6.4.

When SR Protocol is used, the FC-BB-2_SONET prefixes a 4 byte SR Header at the begin of an encapsulated frame that is mapped into the payload of the SR_I message. The SR Header indicates the type of SR message type along with other control information. See 6.2 and 6.3.

See Clause 11 for details on encapsulation using HDLC-like framing.

10.3.4 FC-BB-2_SONET Forwarding

FC-BB-2_SONET forwards FC frames that enter its B_Ports to a remote FC-BB-2_SONET using a mapping table that contains a list of FC-BB-2_SONET HDLC corresponding to a list of D-ID addresses.

10.3.5 Call Handling

FC-BB-2_SONET provides a Point-to-point service for all classes of FC frames transmitted between two Switches.

10.3.6 Frame Handling

Frame Handling is mainly concerned with the following two tasks:

- a) Processing the incoming FC frames from the external switch that emerges from the FC-2 Level that has to be transported across the WAN. Processing includes tasks such as BBW header and message generation, and mapping to HDLC and SONET SPE.
- b) Processing the FC-BB-2_SONET message that has successfully made it across the WAN and that is to be sent to the external switch. Processing includes decoding the SONET SPE containing the HDLC frames, decoding the BBW message and removing the message headers.

11 Mapping and Message Encapsulation using HDLC-like Framing

11.1 Applicability

This clause only applies to FC-BB-2_SONET.

11.2 Overview

BBW messages are transparently transported over the SONET WAN. However, before it may be transported, it has to be first adapted. This adaptation is done using the HDLC layer. Similar to Packet over SONET and Frame Relay over SONET, the FC-BB-2_SONET specification is based on the HDLC-like framing used in PPP-over-SONET/SDH, and described in RFC-1662. The BBW messages form the payload of the HDLC frame that is mapped into SPE/Virtual Containers.

11.3 Mapping of BBW messages to HDLC format

Table 23 shows the mapping of the BBW message to HDLC format according to RFC 1662. The contents of the fields are transmitted from left to right. HDLC framing provides for the delineation of the SONET payloads using a technique called 'stuffing/unstuffing.' Each HDLC frame begins and ends with the flag sequence. During transmission, if the flag sequence occurs anywhere within the information field of the HDLC frame, it is changed to an escape sequence. At the receiver, the escape sequences are removed and replaced with the original fields. A 32-bit FCS is calculated across the HDLC frame for error checking purposes.

The HDLC frames are then mapped byte synchronously into the SONET SPE / SDH Virtual container including any necessary inter-frame byte stuffing. The STS-SPE/SDH Higher Order VC is then scrambled using the self-synchronizing $x^{43}+1$ scrambler. Since the FC-BB-2_SONET interface is comprised of Path Terminating Equipment, the SONET Section, Line and Path layers (Regenerator, Multiplex and Path layers for SDH) are required. Any of the many physical interfaces specified for SONET and SDH may be accommodated depending on the distances required and the WAN service offering being utilized.

Flag sequence:

The Flag sequence is used to encapsulate and delineate the HDLC frame (frame synchronization). Each frame begins and ends with the Flag sequence 0x7E. If a frame immediately follows another, one flag sequence may be treated as the end of the preceding frame and the beginning of the immediately following frame (e.g., there does not need to be two Flags separating the frames). When there are no HDLC frames to be transmitted, the Flag sequence is to be transmitted continuously in the SONET/SDH envelope/VC. Back-to-back Flags are considered Empty frame indications.

Address:

The Address field contains the destination HDLC address. The address 0xFF is an All Stations Address / Broadcast Address. Any station on the link connection shall accept this address. Frames with invalid addresses are silently ignored.

Control:

The Control field identifies the HDLC frame type (information, supervisory, unnumbered). The Control field of 0x03 is the Unnumbered Information (UI) command. Unnumbered frames are used for transferring data when the location of the data in a sequence of frames is not to be checked (no send or receive counts are utilized).

Protocol:

The protocol field is as defined in RFC 1661. It is one or two octets, and its value identifies the payload encapsulated in the Information field. The field is transmitted and received most significant octet first.

Table 23 – Mapping of BBW messages to HDLC format

Field		Encoding (hex)	Size (Bytes)	Remarks
Begin Flag		7Eh	1	
Address		FFh	1	Set to FFh for Broadcast
Control		03h	1	Only Information Type used
Protocol			2	
BBW message	LLC/SNA P Header		8	
	BBW Header		4	
	BBW message payload		Max: 2148/2152 (See Note 23)	Variable Length
FCS			4	
End Flag		7E	1	
Fill or Address			>=1	Inter-frame Fill or next Address

NOTE 23– The maximum of 2148 bytes is due to a maximum of 2112 bytes of FC frame payload, 4 bytes of SOF, 24 bytes of FC Header, 4 bytes of EOF, 4 bytes of CRC. If SR is used then 4 bytes of SR_Header yields a total maximum of 2152 bytes.

Information:

The information field contains the BBW message.

Frame Check Sequence (FCS):

By default, the 32-bit frame check sequence (FCS) field is required as described in RFC 1662. The FCS is calculated most-significant byte to least-significant byte and from least-significant bit to most-significant bit within each such byte over all bits of the address, control, and information fields prior to escape conversions. The least significant byte of the result is transmitted first as it contains the coefficient of the highest term. The FCS is calculated based upon the following polynomial:

$$x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

Inter-frame fill:

A sending FC-BB-2_SONET shall continuously transmit the Flag sequence as inter-frame fill after the FCS field. The inter-frame Flag sequences shall be silently discarded by the receiving station. When an under-run occurs during DMA in the sending FC-BB-2_SONET, it shall abort the frame transfer by continuously transmitting the flag sequence.

Framing and byte stuffing:

The framing and byte stuffing for octet-oriented synchronous links are described in RFC 1662, PPP in HDLC-like Framing. HDLC frames (octet streams) are mapped into the SONET STS-SPE/SDH Higher Order VC with octet boundaries aligned using $x^{43}+1$ scrambling.

Escape sequences are defined to minimally escape the Flag Sequence and Control Escape octet. Prior to sending the frame, but after the FCS computation, every occurrence of The Flag Sequence, Control Escape octet or Async-Control-Character-Map (ACCM) found within the octets of the payload are converted to a two-octet sequence that includes the Control Escape octet followed by the original octet exclusive-or'd with hexadecimal 20h. For example:

- a) 7Eh is encoded as 7Dh, 5Eh. (Flag Sequence)
- b) 7Dh is encoded as 7Dh, 5Dh. (Control Escape)
- c) 03h is encoded as 7Dh, 23h. (ETX)

Upon receiving a frame, this conversion shall be reversed prior to FCS computation.

Abort sequence:

A Flag sequence inserted into the octet stream between the initial frame Flag sequence and the FCS constitutes sequence abort. The receiver considers the frame invalid until a subsequent Flag Sequence is found in the octet stream.

For example, when an under-run condition occurs at the sending station (the sending station cannot complete the data transfer for one reason or another) the sending station transmits a Control Escape octet followed immediately by the Flag Sequence, the frame is ignored and not counted as a FCS error.

11.4 Mapping HDLC frames to SONET/SDH

The mapping of HDLC framed signals according to ISO/IEC 3309 is performed by aligning the byte structure of every HDLC frame with the byte structure of the SONET SPE / SDH Virtual Container. The HDLC frames are located by row within the SPE payload. Since the HDLC frames are of variable length (this mapping does not impose any restrictions on the maximum length) a frame may cross the SPE/Virtual Container frame boundary. See Figure 13.

HDLC Flag sequence shall be used for inter-frame fill to buffer out the asynchronous nature of the arrival of the HDLC framed SONET PDUs according to the effective payload of the SPE/Virtual Container used (this excludes any fixed stuff bytes).

The HDLC framed signal plus the inter-frame fill shall be scrambled before they are inserted as payload of the SPE/Virtual Container used. In the reverse operation, following termination of the SPE/Virtual Container signal, the payload shall be descrambled before it is passed on to the HDLC Mapping layer. A self-synchronizing scrambler with generator polynomial $x^{43} + 1$ [25] shall be used. Scrambling of the HDLC framed signal is required to provide security against emulation of the SONET/SDH set-reset scrambler pattern and replication of the STM-N frame alignment word.

The $x^{43} + 1$ scrambler shall operate continuously through the bytes of the SPE, bypassing bytes of SONET Path Overhead. The scrambling state at the beginning of a SPE shall be the state at the end of the previous SPE. Thus, the scrambler runs continuously and is not reset per frame. An initial seed of the scrambler is unspecified. Consequently, the first 43 transmitted bits following start-up or a SONET/SDH re-frame operation shall not be descrambled correctly.

The $x^{43} + 1$ scrambler operates on the input data stream with Most Significant Bit (MSB) first, consistent with the bit ordering and transmission ordering defined for SONET in T1.105.

The above mapping procedure shall be used for the mapping of HDLC framed signals in SONET STS-3c, STS-12c and STS-48c SPEs and for equivalent SDH Virtual Containers.

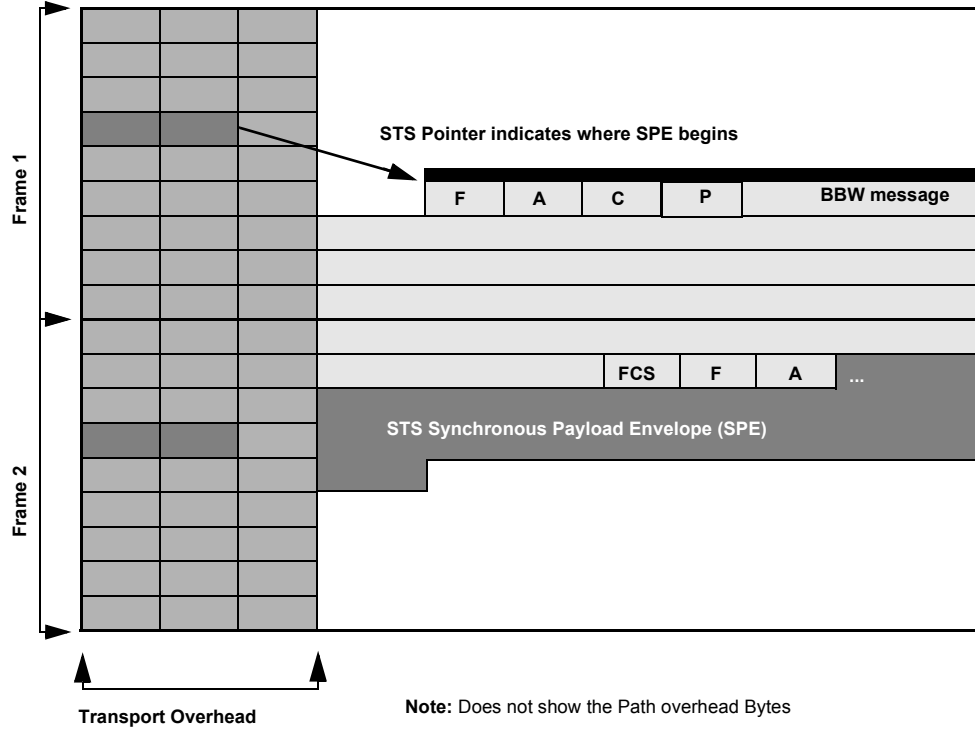


Figure 13 – SONET SPE HDLC Mapping Example

The Path Signal Label (C2) (see Figure 14) indicates the contents of the SPE/Virtual Container. The value of 22d (16h) shall be used to indicate a variable-length HDLC frame with $x^{43} + 1$ scrambling enabled. Implementations shall not use a Path Signal Label (C2) value of 207 (CFh) that indicates a variable-length packet or frame without scrambling. The Multi-frame Indicator (H4) is unused, and shall be zero. Table 24 shows the FC-BB-2_SONET protocol stack.

00010110b	16h	Mapping of HDLC framed signal
-----------	-----	-------------------------------

Figure 14 – Path Signal label: C2

Table 24 – FC-BB-2_SONET Protocol Stack

Interface Layer	Functionality
HDLC Mapping	<ul style="list-style-type: none"> – Frame Delineation – Link & Mapping Error Checking
SONET/SDH (Section, Line, path)	SONET/SDH <ul style="list-style-type: none"> – Section layer – Line layer – Path layer
Photonics	Optical layer

Figure 15 illustrates the encapsulation of BBW message into HDLC frame using SFC. Figure 16 illustrates the encapsulation of the BBW message into HDLC Frame using SR.

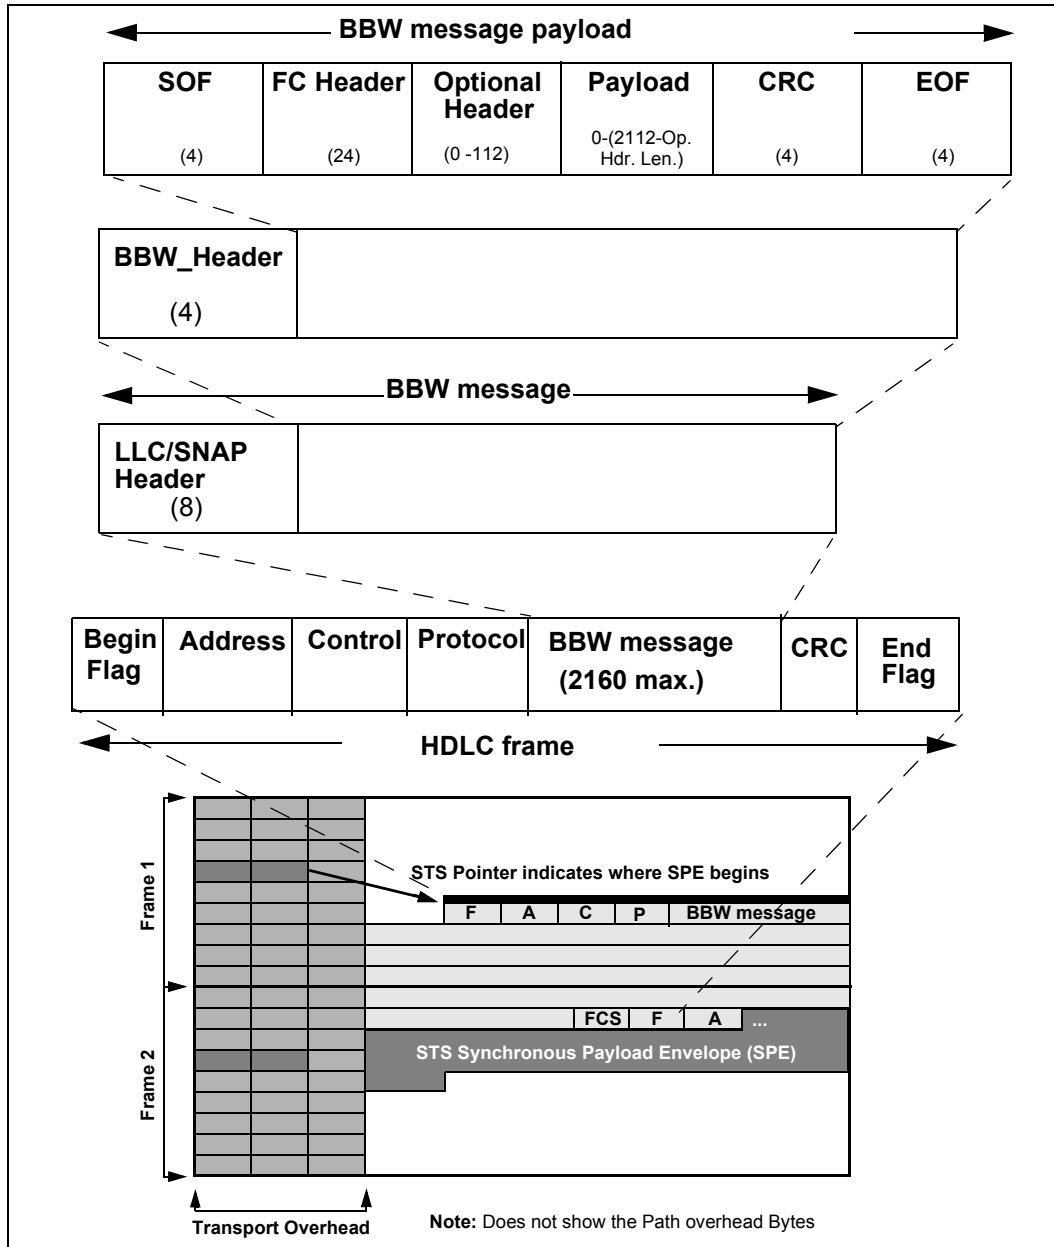


Figure 15 – Encapsulation of BBW message into HDLC frame using SFC

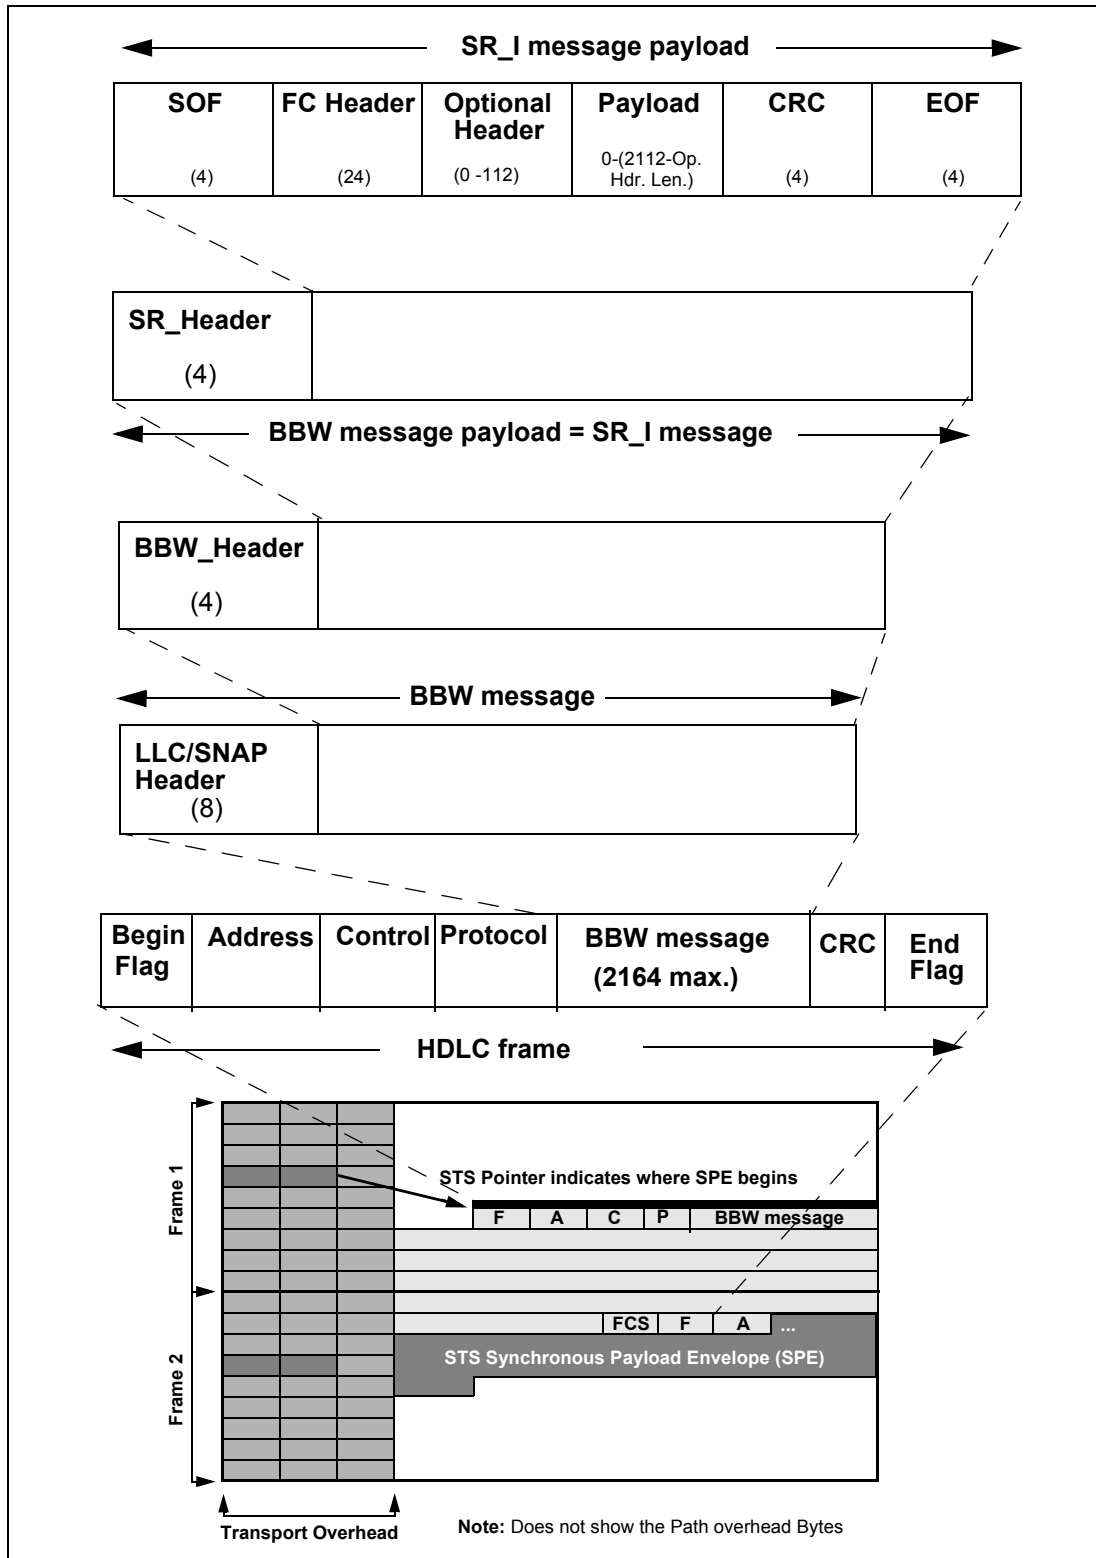


Figure 16 – Encapsulation of BBW message into HDLC frame using SR

12 FC-BB-2_SONET Service Considerations

12.1 Applicability

This clause only applies to FC-BB-2_SONET.

12.2 Latency Delay and Timeout Value

FC-BB-2_SONET and the SONET/SDH network introduce latency delays that warrant special considerations with respect to Fibre Channel E_D_TOV and R_A_TOV values. The total path delay between a source FC-BB-2_SONET and a destination FC-BB-2_SONET consists of the latency delay components due to queuing time at the two FC-BB-2_SONET devices and all intermediate SONET switches, transmission time, and propagation time. It is recommended that this total path delay be less than 1/2 E_D_TOV to conform to normal Fibre Channel time out values.

12.3 Delivery Order

FC-BB-2_SONET shall guarantee in-order delivery of frames.

NOTE 24– The out-of-sequence delivery problem associated with datagram networks is not present here. This benefit is a consequence of SONET Technology. However, the possibility of missing or errored messages still remains and is addressed by the SR protocol.

12.4 Loss and Flow Control

SONET/SDH networks are not lossy but may suffer from occasional loss of frame due to BER. When such a loss occurs, the end application is expected to recover from this loss. Recovery from such losses occurs at the FC-BB-2_SONET devices using the SR protocol that supports error recovery.

NOTE 25– The SFC protocol has no error recovery support

Use of a flow control protocol (SFC or SR) at the FC-BB device allows to cope up with the speed mismatches between the FC and the SONET/SDH interface.

Following is a typical list of reliability specifications for SONET/SDH networks.

- a) MTTF Mean time to frame (approximately 1.5 packets)
- b) MTTS Mean time to synchronization (same as MTTF)
- c) PFF Probability of false frame ($232.8E-12$)
- d) PFS Probability of false synchronization (same as PFF)
- e) PLF Probability of loss of frame (square of the BER multiplied by 500)

13 FC-BB-2_IP Structure and Concepts

13.1 Applicability and Related Clauses

This clause only applies to FC-BB-2_IP.

Clause 4 discussed the FC-BB-2_IP Reference Model. This clause discusses the FC-BB-2_IP Functional Model. Other FC-BB-2_IP applicable clauses include Clause 14 (Mapping and Encapsulation), Clause 15 (Protocol Procedures), and Clause 16 (Service Considerations).

13.2 FC-BB-2_IP Overview

Figure 17 shows a network configuration consisting of three FC-BB-2_IP devices. FC-BB-2_IP is a Fibre Channel backbone transport protocol that tunnels Encapsulated FC Frames across the IP network. A FC-BB-2_IP device has interfaces to both the IP and the FC network. The FC network interface supports multiple E_Ports/F_Ports (Figure 18) or multiple B_Ports (Figure 22).

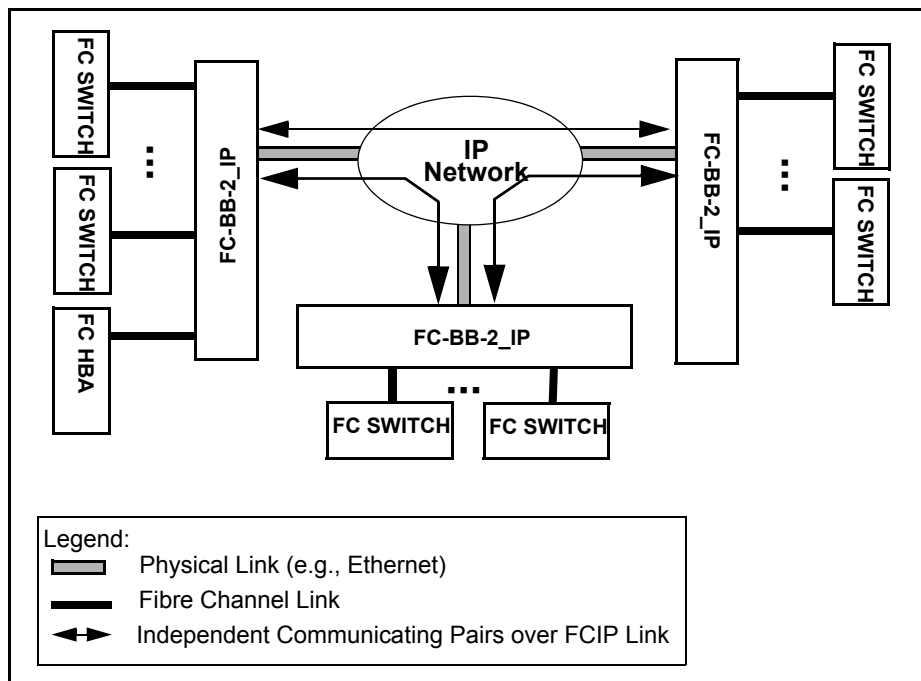


Figure 17 – FC-BB-2_IP Network Configuration

Only FC-BB-2_IP devices that support E_Ports or F_Ports require FC switching.

The FC-BB-2_IP protocol provides mechanisms to create Virtual E_Port or B_Access connectivity over the IP network. The FC-BB-2_IP protocol communication occurs between pairs of FC-BB-2_IP devices over virtual constructs (FCIP Links) that are described in 13.3.3.3.4. Although, the communication occurs between pairs of FC-BB-2_IP devices, a single FC-BB-2_IP device may communicate with more than one device at the same time (Figure 25).

NOTE 26– Although the current scheme allows a FC-BB-2_IP device to independently connect to more than one FC-BB-2_IP device, it does not specify a point-to-multipoint connection.

The FC-BB-2_IP protocol creates Encapsulated FC Frames by prefixing a (28-byte) FC Encapsulation Header to the incoming SOF/EOF delimited FC frame. See FC Frame Encapsulation [12]. FC-BB-2_IP does not interpret the data content of the FC frames other than capturing and retaining their SOF/EOF identities in the Encapsulated FC Frame. As such, FC Sequences and Exchanges are not visible to the FC-BB-2_IP protocol. All Encapsulated FC Frames are transparently transported over the IP network.

FC-BB-2_IP devices also exchange SW_ILS control information using Class F FC frames (see Figures 20 and 23). These FC frames are encapsulated and tunneled in the same way as the incoming FC frames.

Encapsulated FC Frames join the TCP byte stream in order (see Figure 25). TCP Segments are created from TCP byte streams without any visibility or regard to Encapsulated FC Frame boundaries.

The TCP flow control between two FC-BB-2_IP devices provides a reliable transport of Encapsulated FC Frames across the IP network. The only delivery order guarantee provided by TCP with respect to the FCIP protocol is the correctly ordered delivery of Encapsulated FC Frames within a single TCP connection. The FC Entity is expected to specify and handle all other FC frame delivery ordering requirements.

Functional Models describing the VE_Port and B_Access are separately discussed in 13.3 and 13.4.

13.3 The VE_Port Functional Model

13.3.1 FC-BB-2_IP Interface Protocol Layers

Figure 18 shows the VE_Port functional model of a FC-BB-2_IP device that consists of the E_Port/F_Port FC interface, the FC-BB-2_IP interface, and the IP network interface. The protocol layers at these interfaces is listed below:

- a) E_Port/F_Port FC interface: FC-0, FC-1, and FC-2 levels
- b) FC-BB-2_IP interface: FC Entity and FCIP Entity protocol layers
- c) IP network interface: TCP and IP layers

Figure 19 illustrates the protocol layers across these interfaces.

13.3.2 E_Port/F_Port FC Interface

The FC-BB-2_IP FC interface supports one or more E_Ports or F_Ports thus requiring the support of the FC-0, FC-1, and FC-2 Levels. The E_Ports in general connect to different external FC switches, but connectivity to the same FC switch is also allowed. The data emerging from the FC Levels is fed into a FC Switching Element.

The initialization of any generic E_Port or F_Port is described in FC-SW-3. An E_Port indicates its support for the ELP/ESC Parameters using the ELP/ESC exchange protocol that is capable of parameter negotiation. Since FC-BB-2 does not support Class 1, the Class 1 Port Parameter VAL bit in the ELP shall be set to 0 (invalid). An ELP received at an E_Port may be rejected (SW_RJT) due to many reasons, including Port-mismatch.

An E_Port/F_Port is uniquely identified by an 8-byte E_Port_Name/F_Port_Name.

13.3.3 FC-BB-2_IP Protocol Interface

13.3.3.1 Major Components

The FC-BB-2_IP protocol interface is a point that has interfaces to the FC network on one side and the IP network on the other. In addition to the two network interfaces, it consists of the following major components:

- a) FC Switching Element (SE) with FC Routing
- b) FC and FCIP Entities
- c) Control and Service Module (CSM)
- d) Platform Management Module (PMM)

13.3.3.2 FC Switching Element (SE) with FC Routing

The FC Switching Element (SE) switches and routes the incoming FC frames from the E_Port or F_Port to the proper Virtual E_Port (see [6]). Routing is accomplished with the support of the FSPF routing protocol. Conversely, the FC SE switches and routes the data arriving from a VE_Port to the proper E_Port or F_Port.

The switch is uniquely identified by an 8-byte Switch_Name.

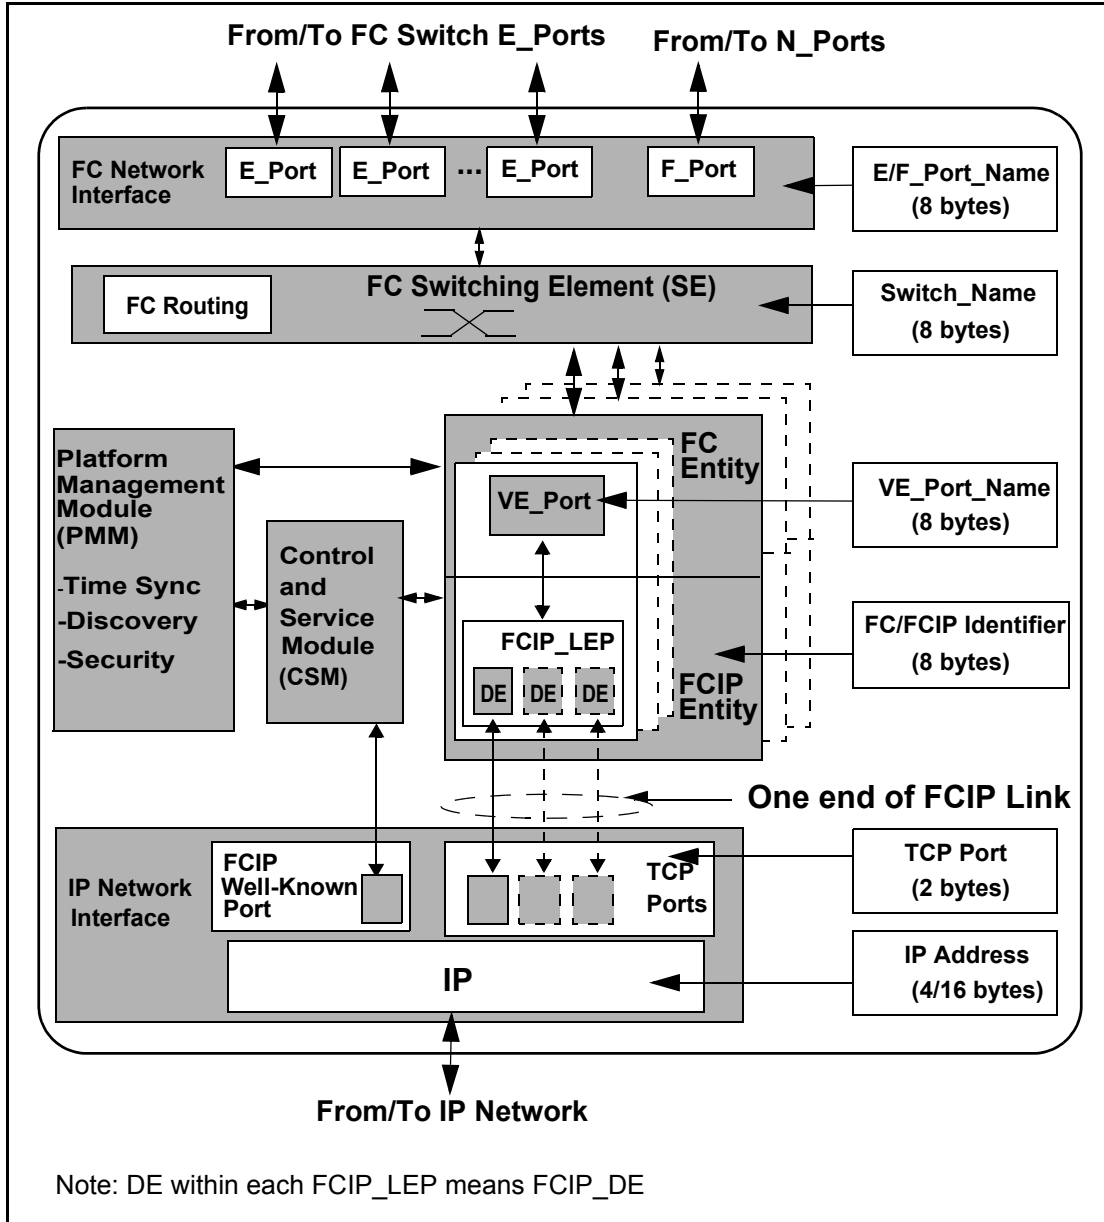


Figure 18 – FC-BB-2_IP VE_Port Functional Model

FC Routing occurs at a higher level than IP Routing. FC/FCIP Entities themselves do not actively participate in FC frame routing. FC Routing uses the FSPF protocol described in SW-3 [6]. FSPF Routes are mapped onto the FCIP Links connecting FC-BB-2_IP devices. A FC frame's FSPF route decides the selection of the VE_Port/FCIP_LEP pair within a selected FC/FCIP Entity pair (when multiple pairs are in use). When multiple DEs (within a FCIP_LEP) are in use the selection of which FCIP_DE to use is described in 15.4.5 (Procedures for Multiple Connection Management).

13.3.3.3 FC and FCIP Entities

13.3.3.3.1 Function

The FC Entity is the principal interface point to the FC network on one side and in combination with the FCIP Entity to the IP network on the other side. The primary function of the FC Entity is supporting one or more Virtual E_Ports and communicating with the FCIP Entity. The FC Entity layer lies between the FC-2 FC level and the FCIP Entity layer as shown in Figure 19.

The FCIP Entity is the principal interface point to the IP network on one side and in combination with the FC Entity to the FC network on the other. The primary function of the FCIP Entity is formatting, encapsulating, and forwarding Encapsulated FC Frames across the IP network interface.

The FC/FCIP Entity pair interfaces with the CSM and the PMM through an implementation defined interface.

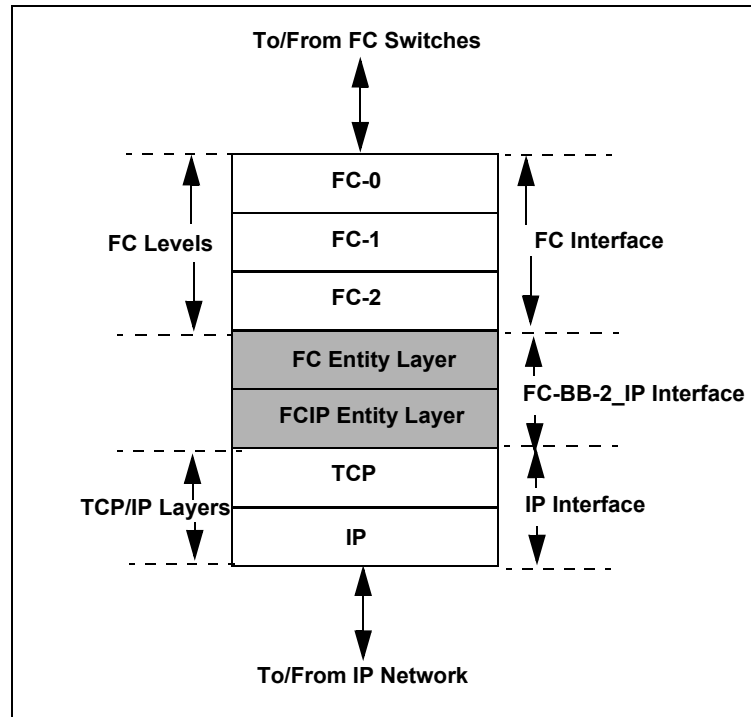


Figure 19 – FC-BB-2_IP Protocol Layers

13.3.3.3.2 FC Entity

The FC-BB-2_IP interface may support multiple instances of the FC/FCIP Entity pair. Each instance of the FC/FCIP Entity pair consists of one or more Virtual E_Port (VE_Port) and Link End Point (FCIP_LEP) pair. A VE_Port emulates an E_Port and interfaces with the Link End Point (FCIP_LEP) component of the FCIP Entity. The term “Virtual” in VE_Port indicates the use of a non Fibre Channel link connecting the VE_Ports.

The VE_Port receives FC frames from the FC side and sends them to the FCIP_LEP for encapsulation and transmission on the IP network. The VE_Port may also exchange Class F control frames with the remote VE_Port via the LEPs. There is a one-to-one relationship between a VE_Port and a FCIP_LEP. VE_Ports communicate via VE_Port Virtual ISLs (described in 13.3.3.3.4).

NOTE 27– The term Virtual ISL when used unqualified refers to both VE_Port Virtual ISL and B_Access Virtual ISL.

A VE_Port is uniquely identified by an 8-byte VE_Port_Name.

Within a FC-BB-2_IP device each FC/FCIP Entity pair instance is uniquely identified by an 8-byte Identifier called the FC/FCIP Identifier. The FC/FCIP Identifier uses the Name_Identifier format.

Initialization at the FC-BB-2 Protocol Interface occurs with the ELP, EFP, ESC, etc. SW_ILS exchanges between VE_Ports in a manner identical to standard E_Ports and is described in 13.3.3.4

13.3.3.3.3 FCIP Entity

The FCIP_LEP is a component of the FCIP Entity that formats, encapsulates, and forwards Encapsulated FC Frames. Encapsulated FC Frames are sent as TCP segments over the IP network.

The FCIP_LEP receives byte-encoded SOF/EOF delimited FC frames and a time stamp (see 13.3.3.6.2.2) from its VE_Port. The FCIP Data Engine (FCIP_DE) is the data forwarding component of the FCIP_LEP. The FCIP_DE handles all encapsulation (de-encapsulation), and transmission (reception) of the Encapsulated FC Frames on the FCIP Link. The FCIP_LEP contains one or more FCIP_DEs, each corresponding to a TCP connection.

The FCIP_DE has 4 interface points (see [10]):

- a) **FC Receiver Portal:** The access point through which a byte-encoded SOF/EOF delimited FC frame and time stamp enters a FCIP_DE from the VE_Port
- b) **FC Transmitter Portal:** The access point through which a reconstituted byte-encoded SOF/EOF delimited FC frame and time stamp leaves a FCIP_DE to the VE_Port.
- c) **Encapsulated Frame Receiver Portal:** The TCP access point through which an Encapsulated FC Frame is received from the IP network by the FCIP_DE.
- d) **Encapsulated Frame Transmitter Portal:** The TCP access point through which an Encapsulated FC Frame is transmitted to the IP network by the FCIP_DE.

13.3.3.3.4 VE_Port Virtual ISL and FCIP Link

The FC/FCIP Entity pair provides a data forwarding path between itself and a remote FC/FCIP Entity pair via virtual constructs. Two types of virtual constructs are defined:

- a) A VE_Port Virtual ISL (Inter Switch Link) is a logical construct that is created between two FC Entity VE_Ports for the explicit purpose of sending and receiving byte-encoded SOF/EOF delimited FC frames via the FCIP Entity. Conceptually, communication between two VE_Ports is similar to communication between E_Ports.
- b) A FCIP Link is a logical construct that is created between two FCIP Entity LEPs for the explicit purpose of sending and receiving Encapsulated FC Frames and Encapsulated FCIP control information. Conceptually, communication between two LEPs is similar to the communication between two instances of a TCP application.

There is a one-to-one mapping between a VE_Port Virtual ISL and a FCIP Link. Each FCIP Link consists of one or more TCP Connections (all between the same two FC-BB-2_IP devices). Although, more than one FCIP Link may be formed between a pair of FC-BB-2_IP devices, a typical configuration may only consist of a single FCIP Link. See Figure 25 for some examples of allowed network topologies.

The FCIP_LEP that originates a FCIP Link is defined as the FCIP Link Originator. The corresponding FCIP_LEP that accepts this link is defined as the FCIP Link Acceptor. A FCIP Link is fully characterized by its FCIP Link Originator and FCIP Link Acceptor identities. A FCIP Link Originator or FCIP Link Acceptor is fully identified by all of the following:

- a) An 8-byte Switch_Name
- b) An 8-byte VE_Port_Name

- c) An 8-byte FC/FCIP Entity Identifier

To uniquely identify a FCIP Link, all the following are required:

- a) The 8-byte Switch_Name of the FCIP Link Originator
- b) The 8-byte VE_Port_Name of the FCIP Link Originator
- c) The 8-byte FC/FCIP Entity Identifier of the FCIP Link Originator
- d) The 8-byte Switch_Name of the FCIP Link Acceptor

NOTE 28– The FCIP Link Acceptor’s 8-byte FC/FCIP Entity Identifier and the VE_Port_Name of the Acceptor provide additional information about a FCIP Link but are not required to uniquely identify it.

13.3.3.4 VE_Port Virtual ISL Exchanges

13.3.3.4.1 SW_ILS Exchanges

VE_Ports exchange SW_ILSs on the VE_Port Virtual ISL. The SW_ILSs that occur on the VE_Port Virtual ISL are the standard E_Port SW_ILSs (ELP, ESC, EFP, etc.), and in addition the LKA SW_ILS (see 13.3.3.4.2). Figure 20 shows the scope of the VE_Port Virtual ISLs.

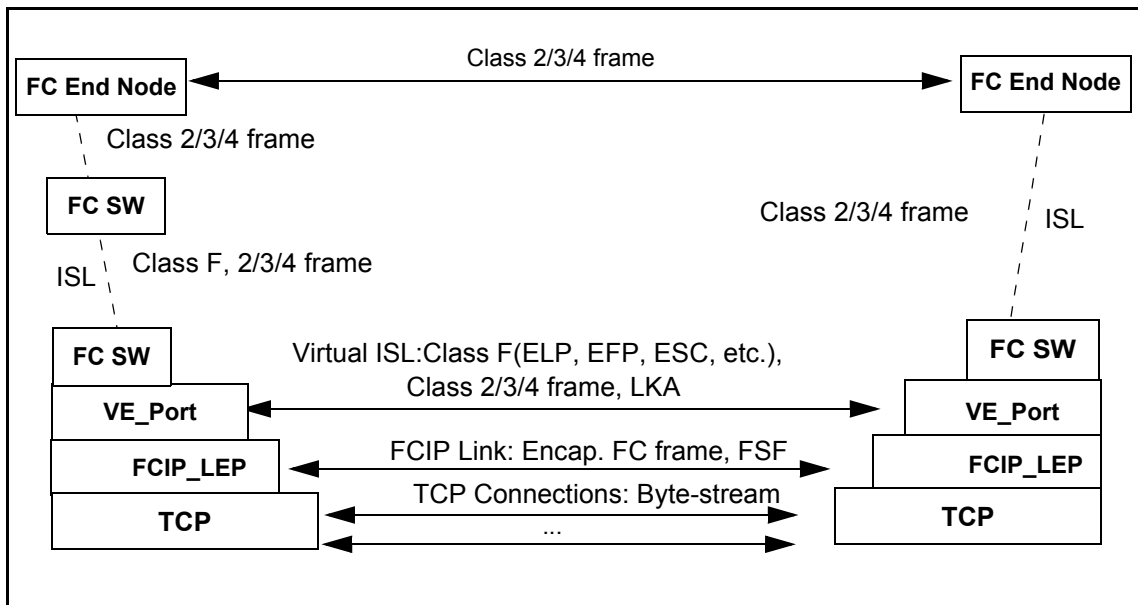


Figure 20 – Scope of VE_Port Virtual ISL

13.3.3.4.2 Link Keep Alive (LKA) ELS Exchanges

The LKA ELS is an ELS for traffic generation. It provides a means to generate traffic in order to confirm that the link is still intact and/or to ensure the link is not terminated due to lack of traffic. The LKA ELS was specifically designed to keep Fibre Channel BackBone Links alive, which are prone to being terminated due to lack of traffic.

The Link Keep Alive ELS is sent by a VE_Port or B_Access portal to a remote peer in order to determine the health of a link between them, or simply to generate traffic to keep a link from being terminated. Should a link be comprised of more than one physical or virtual connection, this LKA may be transmitted on each of the connections. If a connection is configured to handle only specific class(es) of traffic, the LKA must be sent on a class of service the connection is configured for.

The Link Keep Alive ELS request Sequence shall consist of a single frame requesting the Recipient to reply using the ACC reply Sequence consisting of a single frame. The Link Keep Alive ELS request frame shall indicate End_Sequence and Sequence Initiative transfer as well as other appropriate F_CTL bits as defined in FC-FS. The Link Keep Alive command shall be transmitted as a one frame Sequence and the ACC reply Sequence is also a one frame Sequence. The Link Keep Alive Protocol shall be transmitted as an Exchange that is separate from any other Exchange. The Link Keep Alive Protocol is applicable to Class F, 2, 3 and 4.

The LKA may be sent at any time. The LKA should be sent at least every K_A_TOV if no traffic has been sent and/or received on the connection. The default value for K_A_TOV shall be 1/2 E_D_TOV.

If Accept is not received within E_D_TOV, a new LKA may be transmitted on a new exchange. The exchange used for the previous LKA request may be aborted.

Upon discovering an error, e.g. due to Service Reject or failure to receive a timely Accept in response to one or more LKA requests, the initiator should initiate appropriate exception handling. The definition of appropriate exception handling is topology specific.

Protocol:

Link Keep Alive Request Sequence

LS_ACC or LS_RJT Reply Sequence

Format: FT_1

Addressing: The S_ID field shall be set to FFFFDh, indicating the Fabric Controller of the VE_Port or B_Access Portal originating the request. The D_ID field shall be set to FFFFDh, indicating the Fabric Controller of the remote peer.

Payload: The format of the payload is shown in Table 25.

Table 25 – LKA payload

Bits Word	31... 24	23... 16	15... 08	07... 00
0	80h	00h	00h	00h

Reply Sequence:

LS_RJT: LS_RJT signifies rejection of the LKA command.

LS_ACC: LS_ACC signifies that the connection is intact. The format of the LS_ACC payload is found in Table 26.

Table 26 – LKA Accept payload

Bits Word	31... 24	23... 16	15... 08	07... 00
0	02h	00h	00h	00h

13.3.3.5 Control and Service Module (CSM)

The CSM is a control component of the FC-BB-2_IP interface that mainly deals with Connection Management. The CSM creates the FC/FCIP Entity pair during the Virtual ISL/FCIP Link setup. The

CSM processes all requests for a link setup via the FCIP Registered TCP Port 3225 or optionally another TCP Port. CSM also processes requests to add additional TCP connections over the same FCIP Link. CSM is also responsible for tearing down existing FCIP Links and TCP connections and deleting the FC/FCIP Entity pair.

NOTE 29– Some aspects of the CSM functions are discussed only in [10].

13.3.3.6 Platform Management Module (PMM)

13.3.3.6.1 Function

The PMM is a management component of the FC-BB-2_IP interface that handles Time Synchronization, Discovery, and Security. The PMM is also the intended component for any miscellaneous housekeeping functions such as maintenance of event logs (see 16.4.5)

13.3.3.6.2 Time Synchronization

13.3.3.6.2.1 FCIP Transit Time (FTT)

FCIP Transit Time (FTT) is defined as the total transit time of an Encapsulated Fibre Channel frame in the IP network. The FCIP Transit Time is calculated by subtracting the time stamp value in the arriving Encapsulated FC Frame from the synchronized time in the FCIP Entity.

13.3.3.6.2.2 Building outgoing FC frame encapsulation headers

The FC Entity shall establish and maintain a synchronized time value in Simple Network Time Protocol (SNTP) Version 4 format [27] for use in computing the IP network transit times. The FC Entity shall use suitable internal clocks and one of the following mechanisms to establish and maintain the synchronized time value:

- a) Fibre Channel time services; or
- b) IP network SNTP server(s).

Each byte-encoded SOF/EOF delimited FC frame that the FC Entity delivers to the FCIP_DE through the FC Receiver Portal shall be accompanied by a time stamp value obtained from the synchronized time service. The FCIP_DE places the time stamp in the encapsulation header part of the Encapsulated FC Frame that carries FC frame (See FC frame Encapsulation [12]). If no synchronized time stamp value is available to accompany an entering Class 2, 3, or 4 FC frame, the frame should not be delivered to the FCIP_DE. However, FC-BB-2_IP shall allow any class F Encapsulated FC Frames to be transmitted with a zero timestamp value.

13.3.3.6.2.3 Checking IP network transit times in incoming FC frame encapsulation headers.

Each byte-encoded SOF/EOF delimited FC frame delivered to the FC Entity through the FCIP_DE FC Transmitter Portal is to be accompanied by the time stamp value taken from the Encapsulation Header of the Encapsulated FC Frame. As noted in 13.3.3.6.2.2, the time stamp may be zero indicating that no valid time stamp was supplied by the sending FC Entity. Any frame other than a Class F frame whose time stamp is zero shall be discarded. A Class F frame whose time stamp is zero shall be processed as if it meets all Fibre Channel timeout requirements.

When the time stamp is non-zero, the FTT of the arriving Encapsulated Fibre Channel frame shall be compared to $1/2$ ED_TOV. If the FTT exceeds $1/2$ ED_TOV, then the frame shall be discarded. Otherwise the frame shall be processed normally. Fibre Channel Timeout values shall be administratively set to accommodate the FTT.

13.3.3.6.3 Discovery

Discovery of FC-BB-2_IP devices is handled in accordance with the procedures outlined in 15.3.2 and IPS WG specifications FCIP [10], FCIP SLP [11].

13.3.3.6.4 Security

Security in FC-BB-2_IP is defined at two levels: Fibre Channel and FCIP. The Fibre Channel Level is secured through FC-SP [9] mechanisms that are extended by FC-BB-2_IP. The FCIP Level is secured through IPsec mechanisms (see FCIP [10]). Figure 21 illustrates the scope of the two security mechanisms.

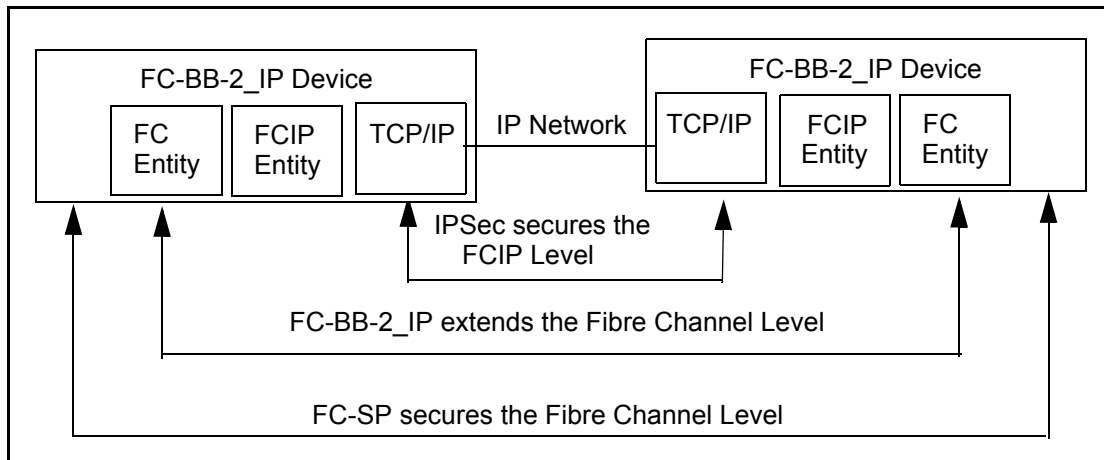


Figure 21 – Security Layers

In most cases, the security requirements of a FC/FCIP Entity pair are satisfied outside the scope of this standard as follows:

- Security for the Fibre Channel Fabric is provided by the FC-SP [9] capabilities (e.g., switch-to-switch authentication, frame authentication and confidentiality), and
- Security for the TCP connections used to transit the IP network is provided by the security features described in FCIP [10] (e.g., IPsec packet authentication and confidentiality).

Depending on the security requirements of a given configuration, any or all of the security capabilities described in other standards may be enabled or disabled. However, it is important to note that the Public IP network is subject to a large variety of security attacks, meaning that serious consideration should be given to enabling the full suite of security features described in FCIP [10] whenever the Public IP network is to be used to transit FCIP frames.

A FC/FCIP entity pair has a potential security vulnerability where interactions may not be fully secured by either the FC-SP or FCIP security features. This vulnerability occurs when two or more TCP connections are aggregated in a single FCIP Link. The first TCP connection in a FCIP Link and its associated Virtual ISL may be authenticated using the FC-SP mechanisms. However, no such authentication is defined for the second, third, etc. TCP connection, since to Fibre Channel they all appear to be part of an already authenticated Virtual ISL.

To prevent attacking entities in the IP network from forging additional invalid TCP Connections, the FC-BB-2_IP mechanism described in 15.3.3 extends the protection of FC-SP authentication to subsequently added TCP connections. The extension to FC-SP authentication described in 15.3.3.2 is based on the exchange of Class F requests and responses between FC Entities. This mechanism

works in concert with the FC-SP Virtual ISL authentication mechanism, transacting the Class F requests and responses over a previously authenticated TCP connection. In some configurations, this overhead may be unnecessary. However, in cases where fabric entities are capable of being authenticated without having their behavior fully trusted, the extension to FC-SP authentication should be used in combination with other FC-SP and FCIP security mechanisms to assure trustworthy formation of FCIP Links and Virtual ISLs.

13.3.4 IP Network Interface

The FC-BB-2_IP VE_Port Reference Model supports one logical IP interface and allows sharing a 4-byte IPv4 or 16-byte IPv6 address in the following ways:

- a) A single IP address per FC-BB-2_IP device
 - a single IP address shared by all FC/FCIP Entity pairs
- b) Multiple IP addresses per FC-BB-2_IP device
 - A single IP address per FC/FCIP Entity pair
- c) Multiple IP addresses per FC/FCIP Entity pair
 - A single IP address per VE_Port/FCIP_LEP pair
- d) Multiple IP Addresses per FCIP Link
 - A single IP address per TCP Port

Use of different IP address schemes at the two ends of a FCIP Link is not expected to cause inter operability problems.

As shown in Figure 19, the IP network interface consists of the TCP and IP layers. The Encapsulated FC Frame emerging out of the FCIP_DE, interfaces with the TCP layer. The IP layer interfaces with the TCP layer above it and the IP network below it. The TCP layer supports multiple TCP connections each corresponding to a FCIP_DE. Each client side TCP connection within a FCIP Link is assigned an unique TCP Port Number. Either the FCIP Well-known TCP Port 3225 or optionally another TCP Port is used for accepting connection requests. These ports interface with the CSM through an implementation defined interface.

IP Routing occurs inside the IP network. Within the IP network, the route taken by an Encapsulated FC Frame follows the normal routing procedures of the IP network.

13.4 The B_Access Functional Model

13.4.1 FC-BB-2_IP Interface Protocol Layers

Figure 22 shows the Functional Model of a FC-BB-2_IP device that consists of the B_Port FC interface, the FC-BB-2_IP protocol interface, and the IP network interface. Figure 19 shows the details of the protocol layers across these interfaces. The following sub clauses describes each of the above interfaces.

NOTE 30– Because of the similarity between the E_Port and B_Port Functional Models this sub clause only describes any unique definitions for the B_Access. Other definitions and descriptions from 13.3 apply equally well and remain unchanged.

13.4.2 B_Port FC Interface

The FC-BB-2_IP FC network interface supports one or more B_Ports thus requiring the support of the FC-0, FC-1, and FC-2 Levels. These ports in general connect to different external FC switches, but connectivity to the same external FC switch is also allowed.

B_Ports are uniquely identified by an 8-byte B_Port_Name.

13.4.3 FC-BB-2_IP Protocol Interface

13.4.3.1 Major Components

The B_Port FC-BB-2_IP interface consists of all the components of the VE_Port Functional Model (see 13.3.3.1) except FC Switching Element with FC Routing.

13.4.3.2 FC and FCIP Entities

13.4.3.2.1 Function

The primary function of the FC Entity is supporting one or more B_Access portals and communicating with the FCIP Entity.

The function of the FCIP Entity is identical to its function in the VE_Port Functional Model described in 13.3.3.3.

The FC/FCIP Entity pair interfaces with the CSM and the PMM through an implementation defined interface.

13.4.3.2.2 FC Entity

The FC-BB-2_IP interface may support multiple instances of the FC/FCIP Entity pairs. Each instance of the FC/FCIP Entity pair consists of one or more B_Access/FCIP_LEP pairs. A B_Access portal is a component of the FC Entity that interfaces with the FCIP_LEP component of the FCIP Entity. The B_Access portal receives FC frames from the B_Port and sends them to the FCIP_LEP for encapsulation and transmission on the IP network. The B_Access portal may also exchange Class F control frames with the remote B_Access portal via the LEPs. There is a one-to-one relationship between a B_Access portal and a FCIP_LEP. B_Access portals communicate via B_Access Virtual ISLs (described in 13.4.3.2.4).

There is no switching and routing required in the case of the B_Port Functional Model. However, the forwarding of FC frames across the B_Access/FCIP_LEP pair is still required. When multiple DEs (within a FCIP_LEP) are in use the selection of which FCIP_DE to use is described in 15.4.5 (Procedures for Multiple Connection Management).

Initialization at the FC-BB-2 Protocol Interface occurs with the EBP SW_ILS exchanges between B_Access portals in a manner identical to standard E_Ports and is described in 13.4.3.2.4. The B_Access Initialization State Machine is described in 13.4.3.3.2.1.

13.4.3.2.3 FCIP Entity

The FCIP_LEP receives byte-encoded SOF/EOF delimited FC frames and a time stamp from its B_Access portals. All other functions are identical to the functions of the FCIP Entity in the VE_Port Functional Model (see 13.3.3.3.3).

13.4.3.2.4 B_Access Virtual ISL and FCIP Links

A **B_Access Virtual ISL** is a logical construct that is created between two FC Entity B_Access portals for the explicit purpose of sending and receiving byte-encoded SOF/EOF delimited FC frames via the FCIP Entity. Conceptually, communication between two B_Access portals is similar to communication between two VE_Ports.

There is a one-to-one mapping between a B_Access Virtual ISL and a FCIP Link.

A FCIP Link Originator or FCIP Link Acceptor is fully identified by all of the following:

- a) An 8-byte Fabric_Name
- b) An 8-byte B_Access_Name
- c) An 8-byte FC/FCIP Entity Identifier

To uniquely identify a FCIP Link, the following items are required:

- a) The 8-byte Fabric_Name of the FCIP Link Originator
- b) The 8-byte B_Access_Name of the FCIP Link Originator
- c) The 8-byte FC/FCIP Entity Identifier of the FCIP Link Originator
- d) The 8-byte Fabric_Name of the FCIP Link Acceptor.

NOTE 31– The FCIP Link Acceptor’s 8-byte FC/FCIP Entity Identifier and the B_Access_Name of the Acceptor provide additional information about a FCIP Link but are not required to uniquely identify it.

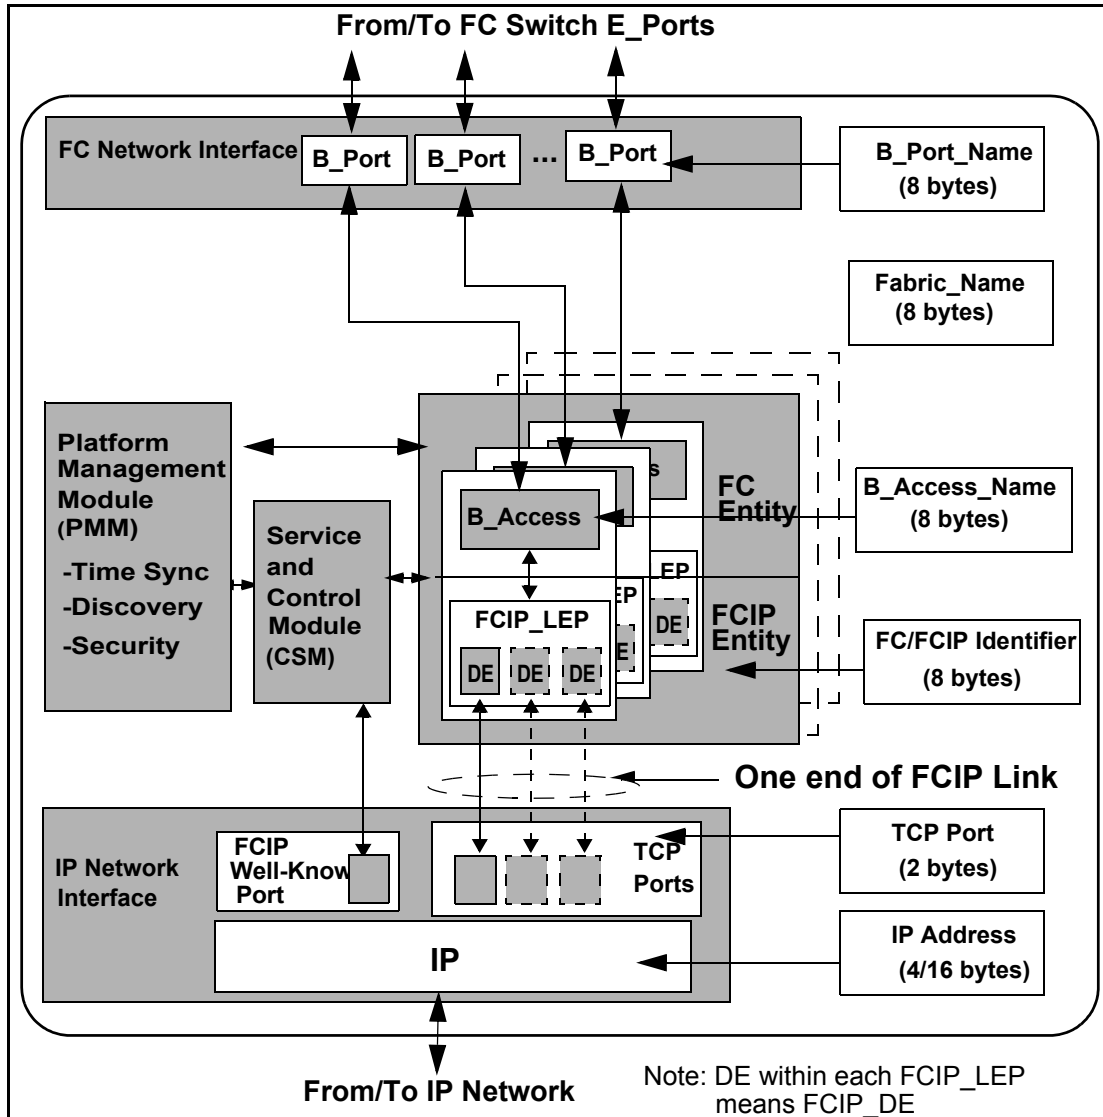


Figure 22 – FC-BB-2_IP B_Access Functional Model

13.4.3.3 B_Access Virtual ISL Exchanges

13.4.3.3.1 Exchange B_Access Parameters (EBP) SW_ILS Exchanges

B_Access portals exchange SW_ILSs on the B_Access Virtual ISL. The SW_ILSs that occur on the B_Access Virtual ISL are the EBP and LKA. Figure 23 shows the scope of the B_Access Virtual ISLs.

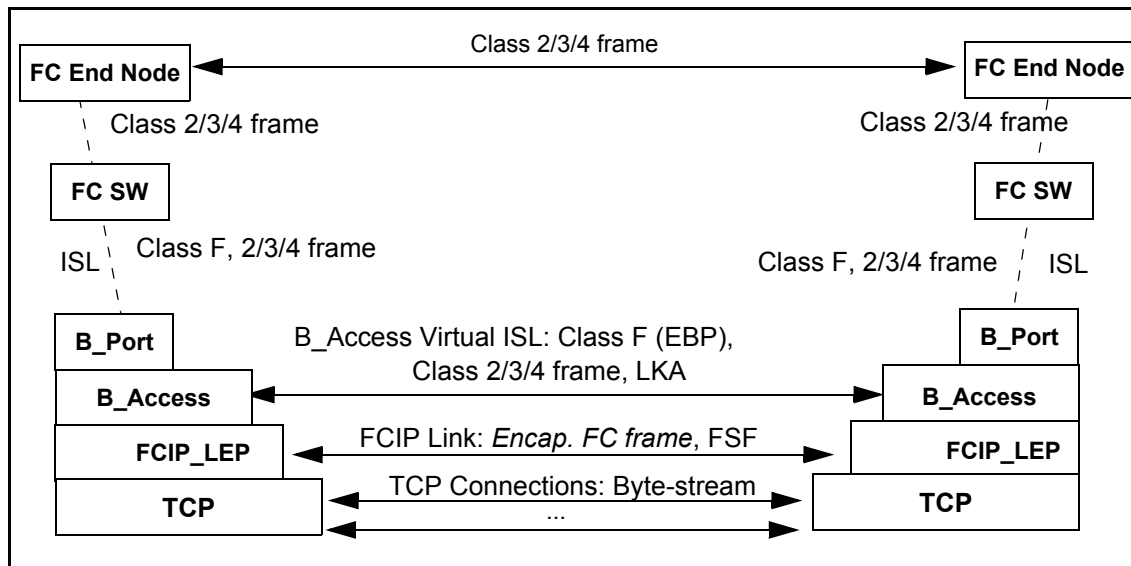


Figure 23 – Scope of B_Access Virtual ISL

The Exchange B_Access Parameters (EBP) Switch Fabric Internal Link Service (SW_ILS) is sent by a B_Access portal to a remote B_Access portal in order to establish operating Link Parameters and port capabilities for the B_Access Virtual ISL formed by the two B_Access portal peers. Successful acceptance of EBP SW_ILS shall be completed before the B_Ports begin Switch Port Mode Initialization.

Protocol: Exchange B_Access Parameters (EBP) Request Sequence

Reply Switch Fabric Internal Link Service Sequence

Format: FT_1

Addressing: For use in Switch Port Configuration, the S_ID field shall be set to FFFFFFFDh, indicating the Fabric Controller of the originating B_Access; the D_ID field shall be set to FFFFFFFDh, indicating the Fabric Controller of the destination B_Access.

Payload: The format of the EBP request payload is shown in Table 27.

Table 27 – EBP Request payload

Item	Size Bytes	Remarks
28 01 00 00h	4	
R_A_TOV	4	Value in milliseconds
E_D_TOV	4	Value in milliseconds
K_A_TOV	4	Value in milliseconds
Requester B_Access_Name	8	
Class F Service Parameters	16	

Requester B_Access_Name: This field shall contain the B_Access_Name of the device that originated the EBP request.

R_A_TOV: This field shall be set to the value (in milliseconds) of R_A_TOV required by the FC-BB-2_IP device.

E_D_TOV: This field shall be set to the value (in milliseconds) of E_D_TOV required by the FC-BB-2_IP device.

K_A_TOV: This field shall be set to the value (in milliseconds) of K_A_TOV required by the FC-BB-2_IP device.

Class F Service Parameters: This field shall contain the B_Access Class F Service Parameters and its format is identical with its use in the ELP SW_ILS [6].

Reply Switch Fabric Internal Link Service Sequence:

Service Reject (SW_RJT)

Signifies the rejection of the EBP command

Accept (SW_ACC)

Signifies acceptance of the EBP command

-Accept payload

Payload: The format of the EBP Accept payload is shown in Table 28.

Table 28 – EBP Accept payload

Item	Size Bytes	Remarks
02 00 00 00h	4	
R_A_TOV	4	Value in milliseconds
E_D_TOV	4	Value in milliseconds
K_A_TOV	4	Value in milliseconds
Responder B_Access_Name	8	
Class F Service Parameters	16	

The fields in Table 28 are the same as defined for Table 27 except for the Responder B_Access_Name field.

Responder B_Access_Name: This field shall contain the B_Access_Name of the remote device that responds to the EBP request.

The SW_RJT Reply payload format is given in [6]. The EBP Reject Reason Code Explanation is shown in Table 29.

Table 29 – EBP Reject Reason Code Explanation

Encoded Value (Bits 23-16)	Description
0000 0000	No additional explanation
0000 0001	Class F Service Parameter error
0000 0010	Invalid B_Access_Name

Table 29 – EBP Reject Reason Code Explanation

Encoded Value (Bits 23-16)	Description
0000 0011	K_A_TOV mismatch
0000 0100	E_D_TOV mismatch
0000 0101	R_A_TOV mismatch
others	Reserved

13.4.3.3.2 B_Access Link Keep Alive (LKA) ELS Exchanges

See 13.3.3.4.2.

13.4.3.3.2.1 B_Access Initialization State Machine

The B_Access initialization state machine is shown in Figure 24.

State P0: Exchange B_Access Parameters. This state marks the beginning of the B_Access initialization. Activity other than that described within the state machine is suspended until initialization is complete.

Transition P0:P1. The B_Access resets the RX_EBP flag.

State P1: Wait for ACK. In this state the B_Access waits until an ACK for the B_Access's transmitted EBP is received.

Transition P1:P0. This transition occurs when the B_Access has not received an ACK within E_D_TOV after the transmission of an EBP.

Transition P1:P2. This transition occurs when the B_Access receives an ACK before E_D_TOV expires.

Transition P1:P4. This transition occurs when the B_Access receives an EBP while waiting for an ACK.

State P2: Wait for Response. In this state the B_Access has received an ACK for its EBP and is waiting for a response.

Transition P2:P0. This transition occurs when the B_Access has not received a response within E_D_TOV after the transmission of an EBP or receives a SW_RJT.

Transition P2:P3. This transition occurs when the B_Access receives a SW_ACC and has not received an EBP.

Transition P2:P4. This transition occurs when the B_Access receives an EBP while waiting for a response.

Transition P2:P5. This transition occurs when the B_Access receives a SW_ACC and has received an EBP.

State P3: Wait for EBP. In this state the B_Access has received an ACK for its EBP and is waiting for an EBP.

Transition P3:P0. This transition occurs when the B_Access has not received an EBP within E_D_TOV of the transmission of an EBP.

Transition P3:P4. This transition occurs when a B_Access receives an EBP while waiting for a response.

State P4: Receive EBP. In this state the B_Access has received an EBP. The B_Access responds with an ACK and transmits a SW_ACC or SW_RJT depending upon whether or not the received configuration parameters contained within the EBP are acceptable. The B_Access sets RX_EBP to indicate an EBP has been received and is accepted.

Transition P4:P1. This transition occurs when a B_Access receives an EBP from its peer yet hasn't received an ACK for a previously transmitted EBP

Transition P4:P2. This transition occurs when a B_Access receives an EBP from its peer yet hasn't received a response for a previously transmitted EBP.

Transition P4:P3. This transition should be removed from the diagram as it is termination point of the machine.

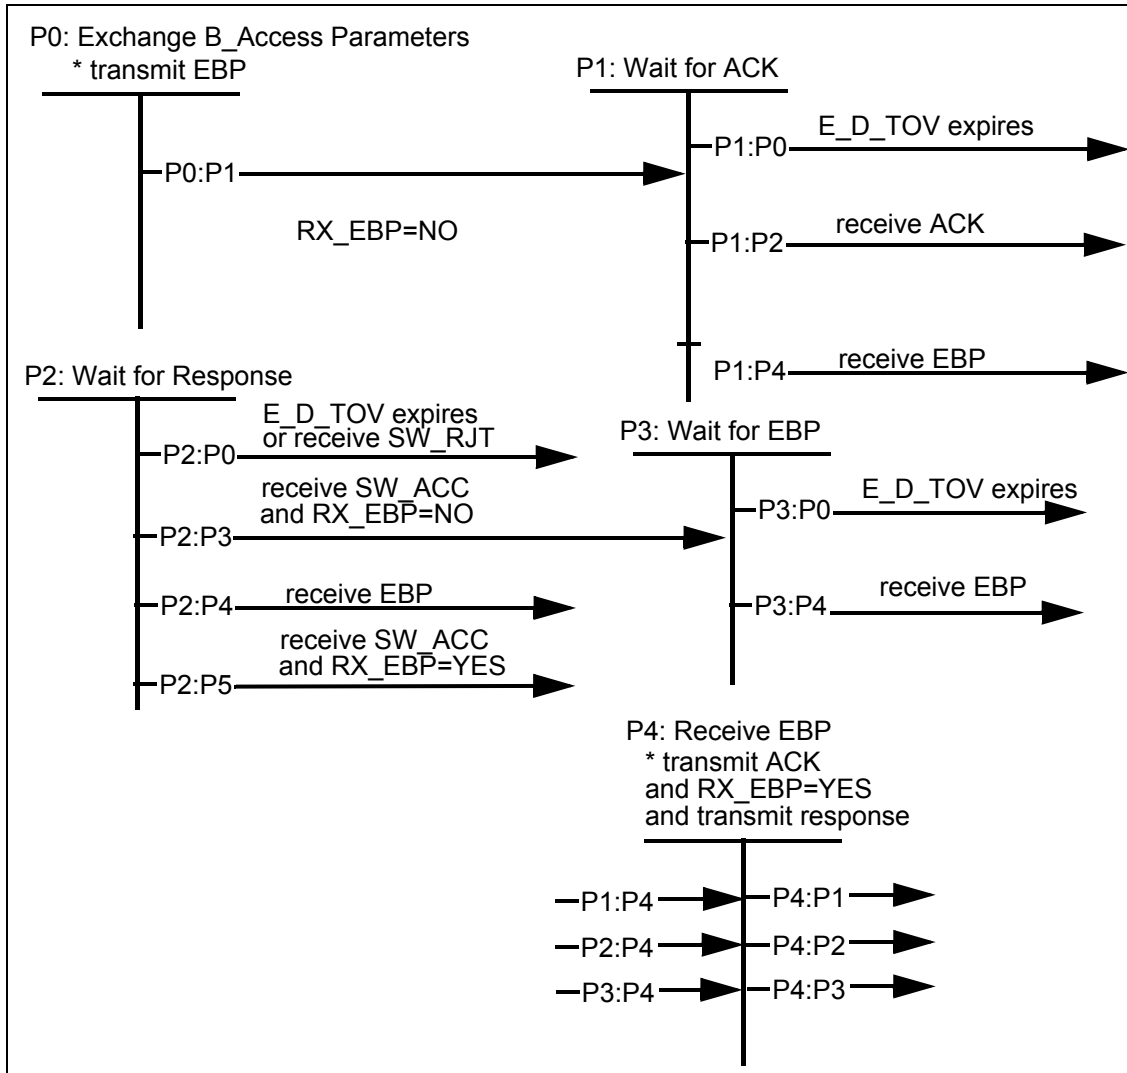


Figure 24 – B_Access Initialization State Machine

13.4.3.4 B_Port Control and Service Module (CSM)

The B_Port CSM is identical to the E_Port CSM described in 13.3.3.5.

13.4.3.5 B_Port Platform Management Module (PMM)

The B_Port PMM is identical to the E_Port PMM described in 13.3.3.6.

13.4.4 IP Network Interface

The B_Port IP network interface is identical to the E_Port IP network interface described in 13.3.4 with a change in Item c), where a single IP address is per B_Access/FCIP_LEP pair.

13.5 FC-BB-2_IP Network Topologies

Figure 25 shows some example FC-BB-2_IP network topologies that exists between 3 FC-BB-2_IP sites:

- a) FCIP Link 1 connects Sites 1 and 2 and consists of 3 TCP connections
- b) FCIP Link 2 connects Sites 1 and 2 and consists of 2 TCP connections. FCIP Link 2 however is distinct from Link 1 although it exists between the same two FC/FCIP Entity pairs (FC/FCIP_Entity_1 and FC/FCIP_Entity_2)
- c) FCIP Link 3 connects Sites 1 and 3 and consists of 2 TCP connections. FCIP Link 3 exists between FC/FCIP_Entity_3 and FC/FCIP_Entity_5.
- d) FCIP Link 4 connects Sites 2 and 3 and consists of 1 TCP connection. FCIP Link 4 exists between FC/FCIP_Entity_4 and FC/FCIP_Entity_6.

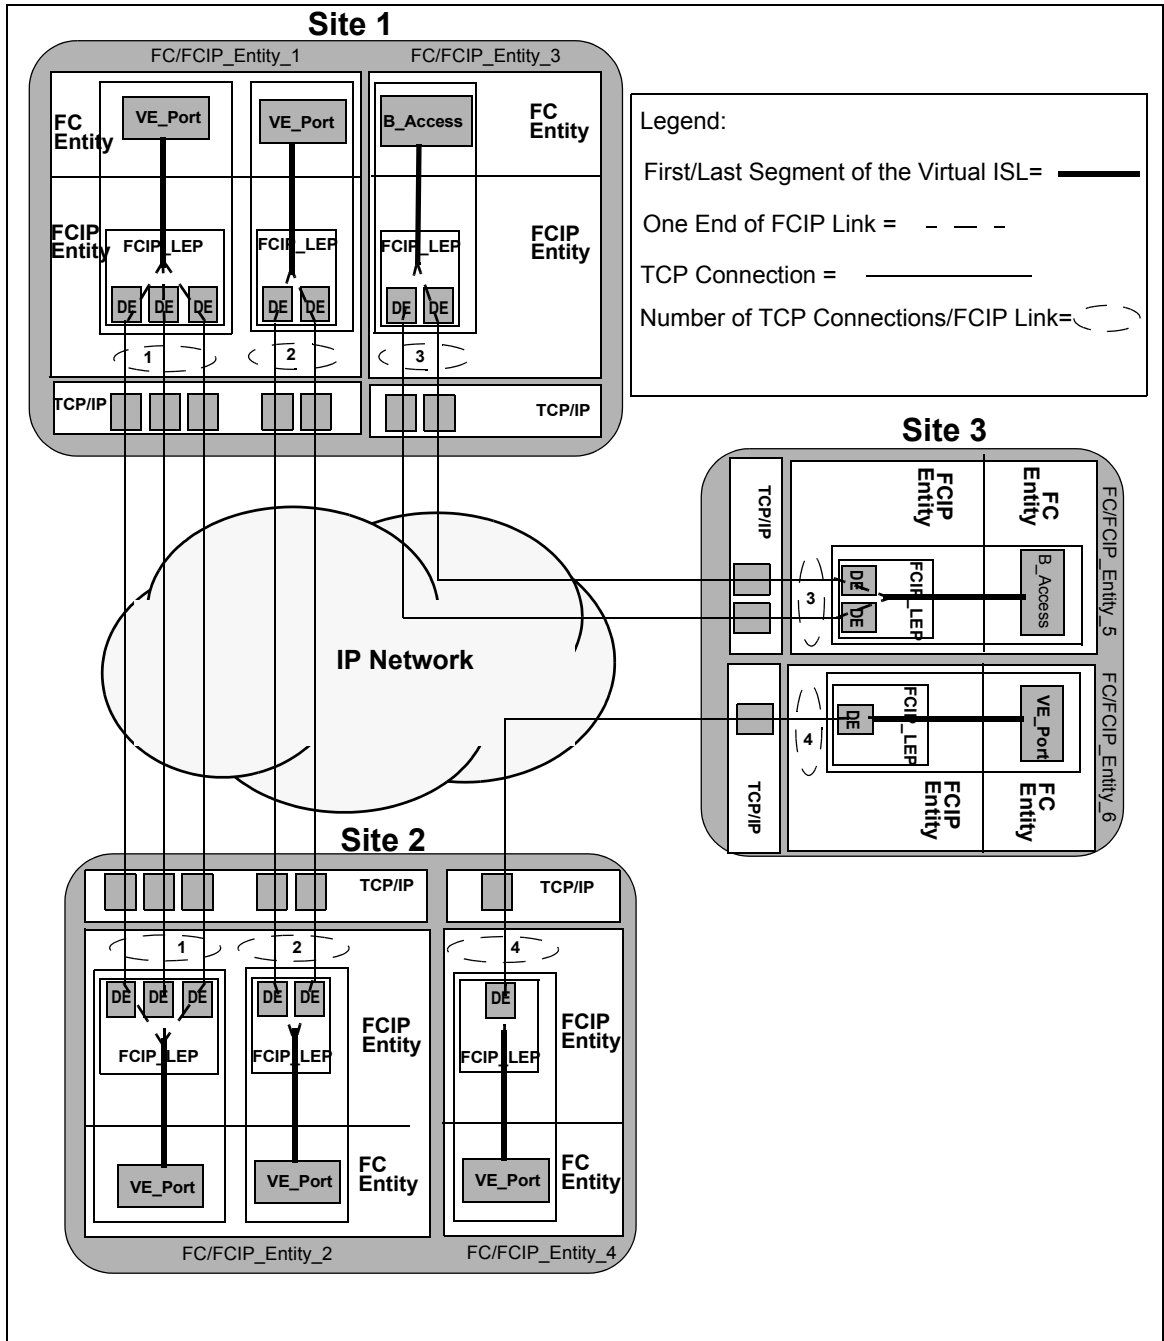


Figure 25 – FC-BB-2_IP Network Topologies

14 Mapping and Message Encapsulation using TCP/IP

14.1 Applicability

This Clause only applies to FC-BB-2_IP.

14.2 Encapsulated Frame Structures

14.2.1 FC frame Encapsulation Structure

An Encapsulated FC Frame is carried as a TCP segment as shown in Table 30. The structure of an Encapsulated FC Frame is shown in Table 31 and consists of a FC Encapsulation Header and a byte-encoded SOF/EOF delimited Class 2, 3, 4 or F FC frame.

Table 30 – TCP/IP Segment structure carrying Encapsulated FC Frame

Field	Sub-field	Size (Bytes)
IP Header		Min:20 Max:40
TCP Header		Min:20 Max:40
TCP payload	Encapsulated FC Frame	Min:64 Max:2176

Table 31 – Encapsulated FC Frame structure

Field	Size (Bytes)
FC Encapsulation Header	28
SOF (see Note below)	4
FC-Header	24
FC frame-payload (includes optional header)	Min: 0 Max:2112
CRC	4
EOF (see Note below)	4

FC frame Encapsulation [12] describes the structures of the 4-byte SOF/EOF values fields and the FC Encapsulation Header. The FC Encapsulation Header consists of several fields: Protocol#, Version, pFlags, Flags, Frame Length, Time Stamp, and CRC. Following is a brief description of these fields. See FC Frame Encapsulation [12] for details.

- a) The Protocol# and Version fields indicates the FCIP protocol and its version number.
- b) The pFlags field defines flag bits FSF and Ch that distinguish Encapsulated FC Frames from FCIP originated or echoed control frames.
- c) The Flag CRCV bit value indicates if the contents of the CRC field are valid or invalid. For FC-BB-2_IP protocol the CRCV bit shall be zero (invalid).
- d) The Frame Length field contains the length of the entire FC Encapsulated Frame including the FC Encapsulation Header and the FC frame (including SOF and EOF words).

e) The (two) Time Stamp fields contain time at which the FC Encapsulated Frame was sent as known to the sender. The format of integer and fraction Time Stamp word values is specified in Simple Network Time Protocol (SNTP) Version 4 [27]. The contents of the Time Stamp [integer] and Time Stamp [fraction] words shall be set as described in 13.3.3.6.2.

f) For FC-BB-2_IP protocol the CRC shall be zero.

14.2.2 Encapsulated FCIP Special Frame (FSF) structure

An Encapsulated FCIP Special Frame (FSF) is carried as a TCP segment as shown in Table 32. The structure of an Encapsulated FSF is shown in Table 33 and consists of a FC Encapsulation Header and a FCIP Special Frame (FSF).

Table 32 – TCP/IP Segment structure carrying Encapsulated FSF

Field	Sub-field	Size (Bytes)
IP Header		Min:20 Max:40
TCP Header		Min:20 Max:40
TCP payload	Encapsulated FSF	76

Table 33 – Encapsulated FSF structure

Field	Size (Bytes)
FC Encapsulation Header	28
FCIP Special Frame (FSF)	48

See 14.2.1 for a description of the FC Encapsulation Header structure and format.

The FSF structure is defined in FCIP [10] and consists of several fields: Source FC Fabric _Name, Source FC/FCIP Entity Identifier, Connection Nonce, Connection Usage Flags, Connection Usage Code, Destination FC Fabric_Name, and K_A_TOV. Following is a brief description of these fields. See FCIP [10] for details.

- a) The Source FC Fabric_Name is the identifier for the FC Fabric associated with the FC/FCIP Entity pair that generates the FCIP Special Frame. If the FC Fabric is a FC Switch, then the field contains the Switch_Name.
- b) The Source FC/FCIP Entity Identifier is an unique identifier for the FC/FCIP Entity pair that generates the FSF. The value is assigned by the FC Fabric whose name appears in the Source FC Fabric_Name field.
- c) The Connection Nonce field contains a 64-bit random number generated to uniquely identify a single TCP connect request. In order to provide sufficient security for the nonce, the randomness recommendations described in FCIP [10] should be followed.
- d) Connection Usage Flag field identifies the types of SOF values to be carried on the connection. All or none of the bits corresponding to Class F, 2, 3, or 4 may be set to one. If all of the bits are zero, then the types of FC frames intended to be carried on the connection has no specific relationship to SOF code.

- e) The Connection Usage Code field is to contain Fibre Channel defined information regarding the intended usage of the connection. The FCIP Entity uses the contents of the Connection Usage Flags and the Connection Usage Code fields to locate appropriate QoS settings in the shared database of TCP connection information and apply those settings to a newly formed connection. No values have been defined for this field at this time and shall carry a 0 value.
- f) The Destination FC Fabric_Name field may contain the Fibre Channel identifier for the FC Fabric associated with the FC/FCIP Entity pair that echoes (as opposed to generates) the FSF.
- g) The K_A_TOV field contains the FC Keep Alive Timeout value to be applied to the new TCP Connection.

14.3 TCP/IP Encapsulation

Figure 26 illustrates the TCP/IP encapsulation of an Encapsulated FC Frame. The TCP/IP encapsulation of an Encapsulated FSF is similar.

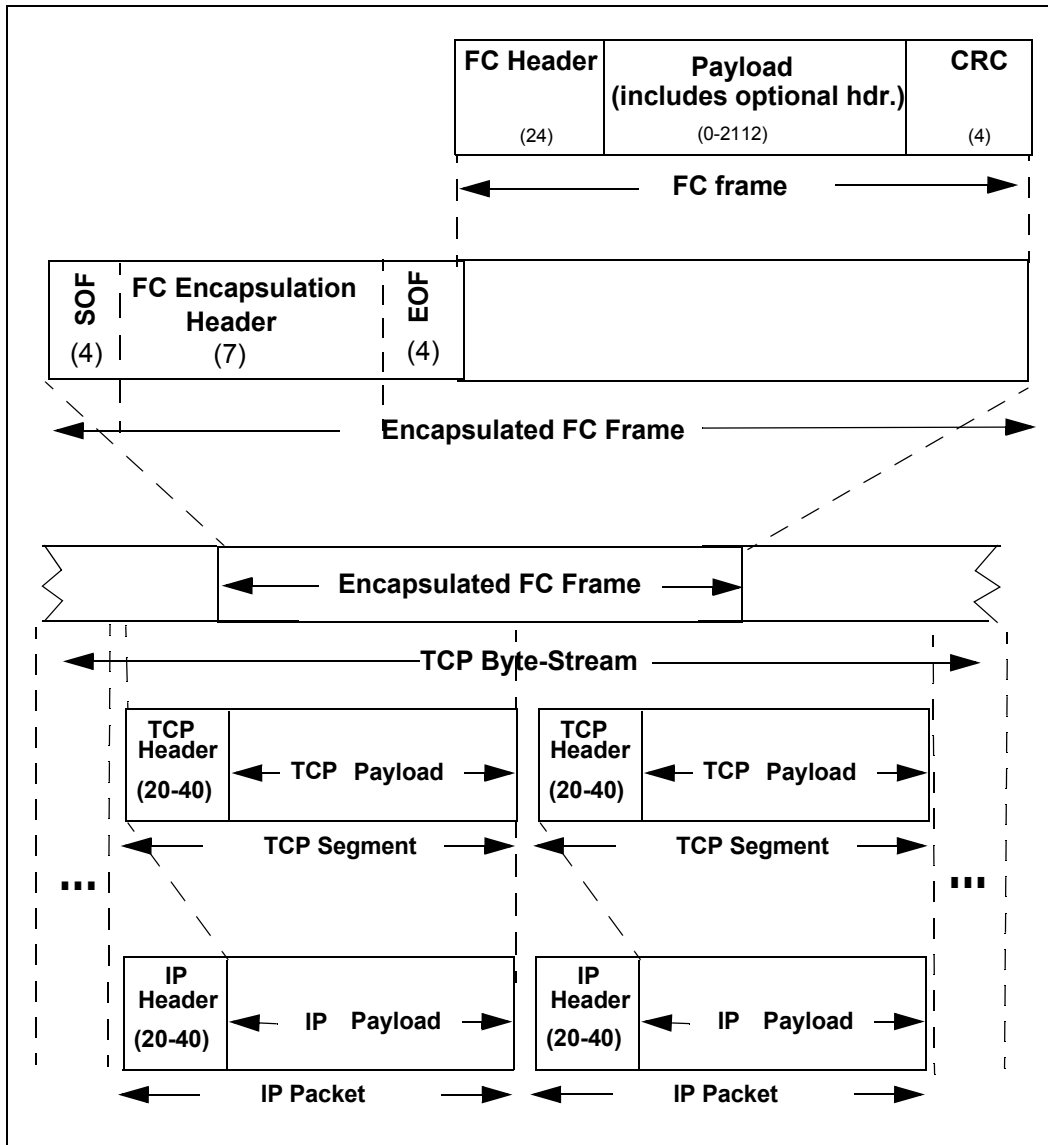


Figure 26 – TCP/IP Encapsulation of an Encapsulated FC Frame

15 The FC-BB-2_IP Protocol Procedures

15.1 Applicability

This Clause only applies to FC-BB-2_IP.

15.2 FC-BB-2_IP Protocol Procedures

This clause describes the FC-BB-2_IP protocol procedures for Platform Management (15.3), Connection Management (15.4), and Error Detection and Recovery (15.5). There are no specific procedures defined for housekeeping functions such as maintenance of error or event logs.

15.3 Procedures for Platform Management

15.3.1 Function

Platform Management has three main functions: Discovery, Security, and Time Synchronization.

15.3.2 Procedures for Discovery

Device Discovery is one of the functions of the Platform Management Module (PMM). Each FC-BB-2_IP device is statically or dynamically configured with a list of IP addresses and other identifiers (e.g., Port_Names) corresponding to participating FC/FCIP Entities. If dynamic discovery of participating FC-BB-2_IP devices is supported, the function is performed using the Service Location Protocol (SLPv2) [11].

FC/FCIP Entities themselves do not actively participate in the discovery of FC source and destination identifiers. Discovery of FC addresses (accessible via the FC/FCIP Entity) is provided by techniques and protocols within the FC architecture as described in FC-FS [8] and FC-SW-3 [6].

15.3.3 Procedures for Extending FC-SP Security

15.3.3.1 Authentication Mechanisms

The Platform Management Module (PMM) is responsible for extending security at the Fibre Channel Level.

Entity authentication occurs at the FCIP and Fibre Channel Levels as illustrated in Figure 21. Authentication mechanisms at the FCIP Level are defined in FCIP [10]. Authentication mechanisms at the Fibre Channel Level are defined in FC-SP [9].

During initialization of a Virtual ISL, each switch may authenticate the other switch with FC-SP authentication mechanisms. FC-BB-2_IP provides for extending the protection of FC-SP authentication to subsequently added TCP connections via either the ASF SW_ILS described in Clause 15.3.3.2 or vendor specific configuration information.

When a FCIP Entity receives a TCP Connect request for an additional TCP connection to an existing FCIP Link to which FC-SP authentication has been applied, the FCIP Entity generates a request to the FC Entity to authenticate the additional TCP connection including at least the following information:

NOTE 32– The unqualified use of the term Virtual ISL refers to both VE_Port Virtual ISL and B_Access Virtual ISL.

- a) Connection Nonce,
- b) Destination FC Fabric_Name,
- c) Connection Usage Flags, and
- d) Connection Usage Code.

If FC-SP authentication procedures are not being applied to the Virtual ISL, the FC Entity shall respond to the FCIP Entity indicating that the new TCP is authentic.

NOTE 33– If the first TCP connection in a Virtual ISL is not authenticated using the applicable FC-SP procedures, no security is gained by authenticating other TCP connections

NOTE 34– The preferred security mechanism for the Public Internet IP network is the success or failure of an ASF SW_ILS.

15.3.3.2 Authenticate Special Frame (ASF)

The Authenticate Special Frame (ASF) Switch Fabric Internal Link Service (SW_ILS) is used by a FC Entity to authenticate additional TCP connections on existing FCIP links. To authenticate a new TCP connection using the ASF SW_ILS, the FC Entity shall use the information provided by the FCIP Entity to transmit an ASF request on the Virtual ISL to which the new TCP connection is being added using a TCP connection in the Virtual ISL that has already been authenticated.

The FC Entity shall use the information from the (new) FSF request to populate the fields in the ASF request. The fields are the same as defined for FSF (see 14.2.2). The format of the ASF Request payload is shown in Table 34.

The FC Entity shall transmit the ASF over the previously authenticated TCP connection. This “piggy-backing” technique authenticates additional TCP connections by riding on the back of previously authenticated TCP connections.

A FC Entity that receives an ASF SW_ILS shall verify that the information in the request payload identifies a TCP connection initiated by that FC/FCIP Entity pair. If it verifies that this information is right then the FC Entity shall respond with a SW_ACC (see Table 35), otherwise it shall respond with a SW_RJT with a Reason Code of Unable To Perform Command Request and a Reason Code Explanation of Class F Service Parameter Error.

Protocol:

Authenticate Special Frame (ASF) Request Sequence

Reply Switch Fabric Internal Link Service Sequence

Format: FT_1

Addressing: The S_ID field shall be set to FFFFDh, indicating the Fabric Controller of the originating FC Entity. The D_ID field shall be set to FFFFDh, indicating the Fabric Controller of the receiving FC Entity.

Payload: The format of the ASF Request payload is shown in Table 34.

Table 34 – ASF Request Payload

Item	Size Bytes
28 03 00 00h	4
Destination FC Fabric_Name	8
Connection Nonce	8
Connection Usage Flags	1
Reserved	1
Connection Usage Code	2
Reserved	4

Destination FC Fabric_Name: This field is the Fabric_Name of the destination switch and is the Source FC Fabric_Name from the FSF frame.

Connection Nonce: This field is the Connection Nonce from the FSF request.

Connection Usage Flags: This field is the Connection Usage Flags from the FSF request and signifies the acceptance of these flags.

Connection Usage Code: This field is the Connection Usage Code from the FSF request and signifies the acceptance of these codes.

Reply Switch Fabric Internal Link Service Sequence:

Service Reject (SW_RJT)

Signifies the rejection of the ASF command

Accept (SW_ACC)

Signifies acceptance of the ASF Request.

-Accept payload

Payload: The format of the ASF Accept payload is shown in Table 35.

Table 35 – ASF Accept Response Payload

Item	Size Bytes
02 00 00 00h	4

15.4 Procedures for Connection Management

15.4.1 Function

The primary function of the Control and Services Module (CSM) is managing connections.

15.4.2 Procedures for Link Setup

In order to realize a Virtual ISL/FCIP Link between two FC-BB-2_IP endpoints, a FC-BB-2_IP device establishes TCP connection(s) with its peer FC-BB-2_IP device.

NOTE 35– A Virtual ISL exists between two VE_Ports or two B_Access portals and a FCIP Link exists between two FCIP_LEPs. Conceptually, the procedures for establishing these two are identical.

It may also be useful to assign a pool of connections for transmission of high priority and control frames (e.g., Class F) on connections so they do not encounter head of line blocking behind Class 2, Class 3 or Class 4 traffic. The use of multiple connections and policies for distributing frames on these connections is described in 15.4.5.

A Virtual ISL/FCIP Link and the two FC-BB-2_IP device endpoints that are involved become operational only after the first TCP connection is established. The sequence of operations performed in order to establish a (Virtual ISL/FCIP Link) is as follows.

- 1) The FC-BB-2_IP device initializes its local resources to enable it to listen to TCP connection requests.
- 2) The FC-BB-2_IP device discovers the FC-BB-2_IP device endpoints to which it is able to establish a Virtual ISL/FCIP Link. The result of the discovery shall be, at the minimum, the IP address and the TCP port of the peer endpoint. The discovery process may rely on administrative configuration or on services such as SLP as described in 15.3.2.

- 3) The processes defined by FCIP are used to establish TCP connections. Fibre Channel Level authentication of the first TCP connection is accomplished using the mechanisms and management controls described in FC-SP. To extend FC-SP authentication to additional TCP connections the mechanisms described in 15.3.3 shall be followed.
- 4) At this point, both endpoints have their respective VE_Port/FCIP_LEP pairs or B_Access/FCIP_LEP pairs established.
- 5) After connection establishment, the FC-BB-2_IP device constructs the Encapsulated FC Frames according to the methods described in 13.3.3.6.2.2 of this document.
- 6) At this point the Virtual ISL endpoints shall exchange Fibre Channel Virtual port initialization frames to enable and identify port operation. E_Port Port Mode Initialization State Machine is described in FC-SW-3 [6] and B_Access Portal Initialization State Machine is described in 13.4.3.3.2.1. Switch-to-Switch authentication shall use FC-SP authentication mechanisms.
- 7) A FC-BB-2_IP device operates in E_Port or B_Port mode. When operating in E_Port mode, normal FSPF messages are exchanged and the switch port becomes operational. When operating in B_Port mode, it is expected that the external E_Ports may exchange FSPF messages over the Virtual ISL which result in the link becoming operational.
- 8) Link costs are implementation defined.
- 9) In certain deployments, a single FC-BB-2_IP device may establish Virtual ISLs/FCIP Links with multiple FC-BB-2_IP device endpoints. In this situation, the FC-BB-2_IP device shall manage TCP operational parameters independently for each Virtual ISL or FCIP Link. Also, the FC-BB-2_IP device VE_Port may perform the E_Port initialization independently, for each Virtual ISL/FCIP Link. The B_Access also may perform initialization independently, for each Virtual ISL/FCIP Link.
- 10) The FC Entity may participate in determining allowed TCP Connections, TCP connection parameters, quality of service usage, and security usage by modifying interactions with the FCIP Entity that are modeled as a shared database. See FCIP [10].

15.4.3 Procedures for Data Transfer

The following procedures need to be followed for data transfer:

- 1) The sending FC Entity shall deliver FC frames to the correct FCIP_LEP/FCIP_DE in the correct FCIP Entity.
- 2) Each FC frame delivered to the FCIP_DE shall be accompanied by a time value synchronized with the clock maintained by the FC Entity at the other end of the FCIP Link (see 13.3.3.6.2.2).
- 3) When FC frames exit FCIP_DE(s) via the FC Transmitter Portal(s), the FC Entity should forward them to the FC Fabric. However, before forwarding the FC frame the FC Entity shall verify the end-to-end transit time as described in 13.3.3.6.2.3.

15.4.4 Procedures for FCIP Link Disconnection

The FC Entity may require the FCIP Entity to perform TCP close requests (e.g., to perform a controlled shutdown of a link or to respond to high link error rates). If the FC Entity requests the closure of all TCP connections in a FCIP Link, the FCIP Link is disconnected.

When the FCIP link is disconnected, notification of the disconnection shall be accomplished according to the procedures in 16.4.5.

15.4.5 Procedures for Multiple Connection Management

A pair of FC-BB-2_IP device endpoints may establish a number of TCP connections between them. Since a Virtual ISL potentially maps a fairly large number of FC flows (where a flow is defined as a pair of Fibre Channel S_ID, D_ID addresses), it may not be practical to establish a separate TCP connection for each Fibre Channel flow. However, once a FC flow is assigned to a FCIP_DE within the Virtual ISL, all FC frames of that flow shall be sent on that same FCIP_DE. This rule is in place to honor any in-order delivery guarantees that may have been made between the two end points of the Fibre Channel flow.

When a TCP connect request is received and that request would add a new TCP Connection to an existing FCIP_LEP, the procedures described in 15.3.3.1 shall be followed.

15.5 Procedures for Error Detection Recovery

15.5.1 Procedures for Handling Invalid FC frames

Data corruption is detected at two different levels: TCP checksum and FC frame encapsulation errors. Data corruption detected at the TCP level shall be recovered via TCP data recovery mechanisms. The recovery for FC frame errors is described below. The TCP and FC frame recovery operations are performed independently.

Fibre Channel frame errors and the expected resolution of those errors are described in FCIP [10] and summarized below:

NOTE 36– The behavior given below is that of the FCIP Entity.

- 1) All incoming frames on the FC Receiver Port are verified for correct header, proper format, valid length and valid CRC. A frame having an incorrect header or CRC shall be discarded or processed in accordance with the rules for the particular type of FC Port.
- 2) All frames transmitted by the Encapsulated Frame Transmitter are valid FC Encapsulations of valid FC frames with correct TCP check sums on the correct TCP/IP connection.
- 3) The FC frames contained in incoming encapsulated frames on the Encapsulated Frame Receiver Port are verified for a valid header, proper content, proper SOF and EOF values, and valid length. FC frames that are not valid according to those checks are managed according to the following rules.
 - a) The frame may be discarded.
 - b) The frame may be transmitted in whole or in part by the FC Transmitter Port and ended with an EOF indicating that the content of the frame is invalid.
- 4) If there is any discrepancy between statements in this sub clause and FCIP [10] then FCIP prevails.

15.5.2 Procedures for Error Recovery

The FC Entity shall recover from events that the FCIP Entity cannot handle, such as:

- a) loss of synchronization with FCIP frame headers from the Encapsulated Frame Receiver Portal requiring resetting the TCP Connection;
- b) recovering from FCIP frames that are discarded as a result of synchronization problems

The FC Entity may recover from connection failures.

Since FC Primitive Signals and Primitive Sequences are not exchanged between FCIP devices, there may be times when a FC frame is lost within the IP network. When this event occurs it is the responsibility of the communicating FC devices to detect and correct the errors based on the features defined in FC-FS [8].

In order to facilitate faster detection of loss of link connectivity, FC Entities shall make use of the Link-Keep-Alive (LKA) protocol. The LKA ELS is exchanged across the virtual ISL as shown in Figure 20 (E_Port implementation) or Figure 23 (B_Port implementation). The exact number of lost LKA heartbeats that will force the FC Entity to mark the link down is a configurable parameter with a default value of 2. Once the link has been marked down, the FC Entity shall attempt to re-establish the link via the FCIP Entity.

15.6 List of FC-BB-2_IP System Parameters

15.6.1 FC Timers

Fibre Channel has two important timeouts: E_D_TOV, and R_A_TOV.

E_D_TOV determines the life of an individual Fibre Channel frame in any particular fabric element. The effects of E_D_TOV on the fabric as a whole are typically cumulative since each fabric element contains its own E_D_TOV timers for any frame received.

R_A_TOV determines the life of an individual Fibre Channel frame in the fabric as a whole. For a fabric, R_A_TOV implies that no particular frame shall remain in (and thus be emitted from) the fabric after the timer expires.

K_A_TOV is a timer defined in this document that is used by the Link Keep Alive (LKA) ELS (see 13.3.3.4.2) as a trigger for issuing LKA. The LKA should be sent at least every K_A_TOV if no traffic has been sent and/or received on the connection. The default value for K_A_TOV is 1/2 E_D_TOV.

15.6.2 TCP Timers

Given the multitude of current and probable TCP implementations, IETF Requests For Comments related to TCP, applications network requirements, etc., it is impossible to provide even rudimentary guidance in suggesting values for the tunable parameters associated with TCP.

15.6.3 Maximum number of attempts to complete an Encapsulated FC Frame transmission

This is an unspecified parameter and is implementation specific.

15.6.4 Maximum number of outstanding Encapsulated FC Frames

This is an unspecified parameter and is implementation specific.

16 FC-BB-2_IP Service Considerations

16.1 Applicability

This clause only applies to FC-BB-2_IP.

16.2 Latency Delay

The time required for a frame to pass from one FC-BB-2_IP device to another across the IP network is variable and beyond the direct control of the FC/FCIP Entity pair. However, the IP network transit time affects the FC Entity's ability to meet Fibre Channel timeout requirements (e.g., the R_A_TOV requirements of the Fabric). Therefore, the FC Entity is required to use facilities provided by the FCIP Entity to compute the IP network transit time for frames. See 13.3.3.6.2.

Class F frames may be excepted from IP network transit time checking; however all other classes of frames shall have their IP network transit time computed and checked. If a frame is found to have an IP network transit time that would cause the frame's lifetime in the Fabric to exceed Fibre Channel requirements, the FC Entity shall discard the frame.

16.3 Throughput

16.3.1 How Timeouts affect Throughput

Both FC and TCP timeouts affect throughput:

- a) Small R_A_TOV values may cause Encapsulated FC Frames to be discarded frequently in the FCIP_DE necessitating FC end node retransmissions.
- b) Large TCP timeouts may result in Encapsulated FC Frames becoming stale in the IP network, leading the FCIP_DE to discard them again necessitating FC end node retransmissions.
- c) Discarding Encapsulated FC Frames due to improper settings of timeout values and errors in the IP network lowers the effective throughput

The FC/FCIP Entities have little or no control over TCP timeouts. The FC/FCIP Entities never initiate retransmissions, that is done either by TCP or by the FC end nodes.

16.3.2 How loss affects Throughput

TCP retransmissions occur due to loss (or corruption) of TCP Segments. If TCP retransmission result in causing the allowed transit time to exceed a threshold, then Encapsulated FC Frames shall be discarded. Either case is likely to cause the effective throughput to be reduced.

16.3.3 Other factors that affect Throughput

Throughput may be affected by a mismatch in the effective rates of data transfer across the FC and the IP network interfaces. This mismatch may occur due to differences in the physical line speeds at the FC network and the IP network interfaces or due to the fundamental difference in the two flow control mechanisms.

FC uses the Buffer-to-Buffer credit flow control and TCP uses the sliding window based flow control. FC-BB-2_IP does not specify the mechanism that aligns the two flow control schemes, although thought performance may be affected if this aspect is not considered. The FC-BB-2_IP device needs to ensure that the TCP connections are able to handle the frame arrival rate from the FC Fabric. The FC Entity shall work cooperatively with the FCIP Entity to manage flow control problems in either the IP network or FC Fabric.

In order to achieve better TCP aggregate throughput properties in the face of packet losses, a pair of

peer FC-BB-2_IP devices may use multiple DEs between them, and use appropriate policies for mapping FC frames to these connections.

16.4 Reliability

16.4.1 Loss of Connectivity

The FC-BB-2_IP device has the capability of detecting loss of connectivity with its remote partner. (See 15.5.2). Upon detecting a loss of connectivity, a FC-BB-2_IP device establishes a new connection, or uses an existing TCP connection to the same FC-BB-2_IP device endpoint. A FC-BB-2_IP device shall not retransmit an Encapsulated FC Frame on the new connection. This is to ensure exactly-once delivery semantics to the Fibre Channel endpoint.

The FC Entity may test for failed TCP Connections. Should such a test detect a failed TCP connection, the FC entity shall disconnect that connection following the procedures in 15.4.4.

16.4.2 Loss of Synchronization

The FC Entity shall recover from events that the FCIP Entity cannot handle, such as:

- a) loss of synchronization with FC-BB-2_IP Encapsulated FC Frame headers from the Encapsulated Frame Receiver Portal requiring resetting the TCP Connection;
- b) recovering from FC-BB-2_IP Encapsulated FC Frames that are discarded as a result of synchronization problems FCIP [10].

16.4.3 Loss or Corruption of TCP Segments

TCP Flow control and error control has mechanisms to detect lost or corrupted TCP segments. TCP retransmits the TCP segments that were lost or corrupted.

TCP Flow Control provides the ability to regulate the flow of data on the IP network interface based on the perceived IP network congestion conditions potentially avoiding large losses of data.

16.4.4 Loss or Corruption of FC frames

The FC interface of the FC-BB-2_IP device has no mechanisms to detect lost data but only corrupted frames. Corrupted frames detected prior to transmission into the IP network, are discarded and not sent over the IP network.

FC Buffer-to-Buffer Credit Control provides the ability to regulate the flow of data on the FC network interface with no loss.

16.4.5 FCIP Error Reporting

The FC Entity receives notifications from the FCIP Entity due to a number of errors detected by the FCIP Entity. As a result, the E_Port implementation of the FC Entity shall report those errors to the local FC switch element via the local VE_Port (see Figure 18). Similarly the B_Port implementation shall report the error to the local B_Access (see Figure 22). In addition, the FC Entity may pass these error reports to the local PMM for inclusion in a local event log.

The FC Entity shall convert the error message received from the FCIP Entity into a Registered Link Incident Report (FC-FS RLIR). It is the RLIR that is forwarded from the FC Entity to either the VE_Port (Figure 18) or B_Access (Figure 22).

On receipt of the message from the FC Entity, VE_Port or B_Access shall immediately forward the RLIR to the Domain Controller of the Switch.

As a minimum the FC Entity shall accept the following information from the FCIP Entity:

- a) Loss of FC frame synchronization (See Section 6.6.2.3 in FCIP [10])

- b) Failure to setup TCP connection (See Section 9.1.2.3 in FCIP [10])
- c) Duplicate connect request (See Section 9.1.3 in FCIP [10])
- d) TCP connect request timeout (See Section 9.1.3 in FCIP [10])
- e) Successful completion of FC Entity request to close TCP connection (See Section 9.2 in FCIP [10])
- f) Loss of TCP connectivity (See Section 9.4 in FCIP [10])
- g) Excessive number of dropped datagrams (See Section 10.4.3 in FCIP [10])
- h) Any confidentiality violations (See Section 10.4.3 in FCIP [10])
- i) SA parameter mis-match (See Section 10.4.4 in FCIP [10])

The FC Entity shall generate and forward an RLIR to the Management Server for the following:

- a) Loss of FC frame synchronization (See Section 6.6.2.3 in FCIP [10])
- b) Failure to setup additional TCP connection (See Section 9.1.2.3 in FCIP [10])
- c) Additional duplicate TCP connect request (See Section 9.1.3 in FCIP [10])

16.5 Quality of Service (QoS)

The FC-BB-2_IP protocol may use TCP/IP QoS features to support Fibre Channel capabilities.

16.6 Delivery Order

Each VE_Port/FCIP_LEP pair defines a separate FCIP Link. FCIP_DEs within a FCIP_LEP share the FCIP Link. Multiple FCIP_DEs between FCIP_LEPs introduce multiple traffic paths (e.g., Class F, Class 2/3/4). The order in which the FCIP_DEs are serviced on the FCIP Link is not specified. One possibility is providing different priority levels to each traffic path changing the overall delivery order.

The only delivery order guarantee provided by TCP is correctly ordered delivery of FC-BB-2_IP Encapsulated FC Frames between a pair of FCIP_DEs. The FC Entity is expected to specify and handle all other FC frame delivery ordering requirements.

NOTE 37– The order of the FC frames sent by the Encapsulated Frame Transmitter may not be the same as the order sent by the source FC End Node. This is due to the fact that some types of FC login allow FC frames to be re-ordered in the FC Fabric before reaching the FC Receiver Port.

16.7 IP Multicast and Broadcast

This standard does not make use of IP Multicast and Broadcast.

16.8 Security and Authentication

The IETF security standards referenced by FCIP [10] provide numerous mechanisms for securing TCP connections between FC/FCIP Entity pairs (e.g., IPsec packet authentication and confidentiality). It is important to note that the Public Internet IP network is subject to a large variety of security attacks, meaning that serious consideration should be given to enabling the full suite of security features described in FCIP whenever the Public Internet IP network is to be used to transit FCIP frames.

The TCP connection authentication mechanism described in 15.3.3.1 provides FC-BB-2 specific authentication for the second, third, etc. TCP connections in a FCIP Link and its associated Virtual ISL as long as the first TCP connection is authenticated using the mechanisms described in 15.3.3 and FC-SP.

Table A.2 – OS-Code Definition

OS-Code (hex)	Delimiter Name
42	EOFt
46	EOFdt
44	EOFrt
49	EOFni
4E	EOFdti
4F	EOFrti
50	EOFa

A.3 FC-BB-2 SOF OS codes

FC-BB-2 makes use of only a subset of SOF OS-Codes defined in Table A.2. The valid OS-codes for FC SOF delimiters are listed in Table A.3.

Table A.3 – FC-BB-2 SOF Codes

OS-Code (hex)	Delimiter Name	Class
28	SOFf	F
2D	SOFi2	2
35	SOFn2	2
2E	SOFi3	3
36	SOFn3	3
39	SOFc4	4
29	SOFi4	4
31	SOFn4	4

A.4 FC-BB-2 EOF OS codes

FC-BB-2 makes use of only a subset of EOF OS-Codes defined in Table A.4. The valid OS-codes for FC EOF delimiters are listed in Table A.4.

Table A.4 – FC-BB-2 EOF Codes

OS-Code (hex)	Delimiter Name	Class
41	EOFn	2, 3, 4, F
42	EOFt	2, 3, 4, F
49	EOFni	2, 3, 4, F
50	EOFa	2, 3, 4, F
46	EOFdt	4
44	EOFrt	4
4E	EOFdti	4
4F	EOFrti	4

Annex B (Informative)

ATM Traffic Management and Signaling

B.1 ATM Traffic Management

The first step in establishing an ATM connection is a signaled connection request setup message from an end user to the ATM network across the User Network Interface (UNI). The endpoint specifies its QoS parameters and the traffic parameters in this connection request setup message.

During this connection request, the ATM network determines if it has sufficient resources to support the QoS and traffic parameters specified by the user. This determination mechanism is called Connection Admission Control (CAC) and is an action taken by all the ATM switches in the connection's path from the source to the destination UNI. Only if CAC determines such resources are available does it admit the connection; otherwise, it rejects it.

Each ATM switch along the path of an ATM connection needs to determine if it has two essential resources: bandwidth and buffers. A Network Resource Management (NRM) function in each of these switches allocates the required bandwidth and buffers during the connection setup phase. CAC depends on NRM to report the availability of the resources. NRM manages the available bandwidth and buffer at each switch along the path and ensures the QoS and traffic requirements are met.

If the user connection request is accepted, a Traffic Contract is established between the user and the network. The traffic contract is a fundamental mechanism between the user and network where the user specifies the traffic characteristics of the source and the QoS requirements of the application at the ATM layer. The traffic contract of a connection is specified by the connection traffic descriptor and the set of QoS parameters for each direction. Once, a connection is established the endpoint is expected to conform to this contract and the network is required to guarantee the QoS specified in the contract.

End users may unintentionally or maliciously exceed their traffic specifications and this may adversely affect the QoS provided by the ATM network to other applications. The network therefore polices the traffic generated by each connection, to ensure that the traffic contract is observed by the end users. Usage Parameter Control (UPC) is the set of policing mechanisms implemented by the network at the UNI (network side) to monitor and control the traffic submitted by each end user. UPC implicitly requires the endpoint to provide a Traffic Shaping mechanism to conform to these parameters. UPC may drop or tag a cell that violates the contract. UPC is essentially an open loop control mechanism to avoid congested network conditions. A closed loop control mechanism, using a feedback to the end users is also implemented for only certain types of traffic, that use the Available Bit Rate Service (ABR).

B.2 ATM QoS Parameters

The ATM layer Quality of Service (QoS) is measured by a set of parameters characterizing the performance of an ATM layer connection. These QoS parameters quantify the end-to-end network performance at the ATM layer. The ATM Forum so far has defined six QoS parameters. Out of these, three may be negotiated between the end-systems and the network. The following parameters are negotiated:

- a) Peak-to-peak Cell Delay Variation (Peak-to-peak CDV)
- b) Maximum Cell Transfer Delay (maxCTD)
- c) Cell Loss Ratio (CLR)

The following parameters are not negotiated and are also called as Dependability Parameters:

- d) Cell Error Ratio (CER)

- e) Severely Errored Cell Block Ratio (SECBR)
- f) Cell Misinsertion Rate (CMR)

NOTE 38– QoS commitments are probabilistic in nature and are intended to be only a first order approximation of the performance that the network expects to offer over the duration of the connection. The actual QoS may vary during the duration of the connection. QoS commitments are a network objective and are not absolute guarantees.

The above QoS definitions are new to the networking industry and care needs to be taken in using them. It is important that the scope of these parameters is correctly understood before specifying it in this standard. Following is a brief description of these parameters. Details may be found in ([1] in B.8).

B.2.1 Maximum Cell Transfer Delay (maxCTD) - Negotiated Parameter

Maximum Cell Transfer Delay (maxCTD) specifies the time elapsed between the event when a cell exits a source UNI and the event when the same cell exits the destination UNI. This time delay consists of a fixed and a variable delay. Fixed delays are attributed to the propagation delay through the physical media and fixed minimum switching latency. Variable delays are due to other network conditions. Thus, the minimum CTD is the fixed delay.

The maxCTD is defined as the sum of the fixed delay and the Peak-to-peak CDV (defined below). A graph is often used to make these definitions clearer and is shown in Figure C.1. The maxCTD specified for a connection is defined as a statistical guarantee using a probability parameter α (where α is a very small number). The network has to guarantee the maximum CTD to be met with a probability of $(1-\alpha)$.

For real-time streaming traffic there is a maximum end-to-end delay requirement beyond which those late cells are treated as lost cells. In this case, α may be treated as the probability of cell loss that may be tolerated by the application due to late arrivals.

B.2.2 Peak-to-peak Cell Delay Variation (CDV)- Negotiated Parameter

Cell Delay Variation (CDV) specifies the variation in delay incurred in the ATM network and is induced by buffering and cell scheduling. Formally, this is defined as the Peak-to-peak CDV, that is the maximum difference in CTD among all the cells of the connection. In other words, Peak-to-peak CDV is simply the difference in value between the best (minimum CTD) and worst-case CTD (maxCTD) as shown graphically shown in Figure B.1.

Since the reference points of measurements are the source and destination UNIs for the CTD, this method is also called the 2-point CDV.

NOTE 39– CDV is also defined using only one reference point called a 1-point CDV. Using only one point of reference, a negative CDV indicates a late arrival and a positive CDV indicates an early arrival. The reference arrival time eliminates the effect of gaps and provides a measurement of cell dumping.

CDV is also defined as a statistical parameter that is satisfied by $(1-\beta)$ of all the cells transferred in the connection. A real-time streaming traffic application has a delay variation requirement beyond which the receiving buffer will either underflow or overflow. Hence, β is the cell loss probability due to receiving buffer overflow or underflow.

ATM networks have a finite ability to control Peak-to-peak CDV. Therefore, end-systems cannot expect to negotiate arbitrarily small values of Peak-to-peak CDV as their sole means of meeting jitter and wander tolerances. Path calculations are made during signaling to accumulate delay parameters as a call progresses. The accumulation algorithms inside ATM switches provide estimates of the end-to-end values of these parameters along a path. A simple accumulation based on worst-case assumptions is supported in ([1] in B.8).

An ATM switch receives the accumulated Peak-to-peak CDV and adds its contribution of the Peak-to-peak CDV to the accumulated Peak-to-peak CDV. Similarly, an ATM switch receives the accumulated maxCTD and adds its contribution of the maxCTD to the accumulated maxCTD.

NOTE 40– It is important to distinguish CDV from CDV Tolerance (CDVT), that is a traffic parameter that may be measured at the UNI to determine conformance as part of Usage Parameter Control (UPC) (defined later).

B.2.3 Cell Loss Ratio (CLR) - Negotiated Parameter

This is simply the ratio of lost cells to the total transmitted cells. Lost cells do not include Severely Errored Cell Blocks from the cell population. CLR is the value that the network agrees to offer as an objective over the lifetime of the connection. CLR objective applies to Cell Loss Priority (CLP) = 0 cell flow or the aggregate of CLP = 0 and CLP = 1 (denoted by CLP = 0+1).

B.2.4 Cell Error Ratio (CER) - Non-negotiated parameter

CER is defined as the ratio of Errored Cells to the total sum of Successfully Transferred Cells and Errored Cells. Successfully Transferred Cells and Errored Cells counted in Severely Errored Cells are excluded from the cell population.

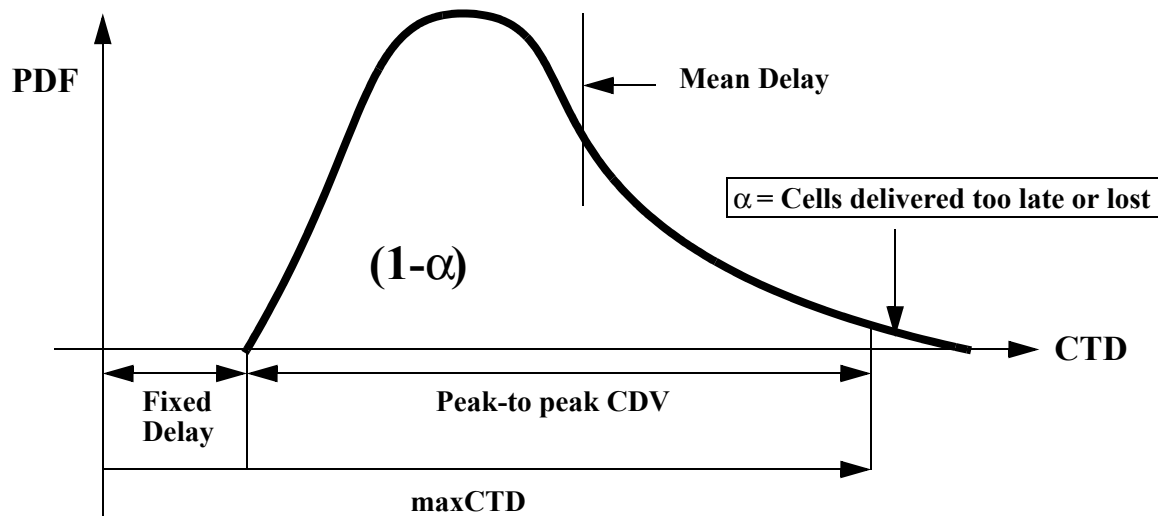


Figure B.1 – Cell Transfer Delay Distribution

B.2.5 Severely Errored Cell Block Rate (SECBR) - Non-negotiated parameter

SECBR is defined as the ratio of Severely Errored Cell Blocks (SECB) to the Total Transmitted Cell Blocks. A Cell Block is a sequence of N cells transmitted consecutively on a given connection. SECB outcome occurs when more than M errored cells, lost cells, misinserted cell outcomes are observed in a received cell block. Practically, this is measured between successive Operations Administration, and Maintenance (OAM) cells.

B.2.6 Cell Misinsertion Ratio (CMR) - Non-negotiated parameter

CMR is defined as the ratio of Misinserted cells over a given time interval. SECB is excluded from this cell population. CMR is caused by an undetected error in the header of a cell being transmitted on a different connection. CMR is independent of the number of cells transmitted.

The ITU-T ([2] in B.8) has defined some typical values for the QoS parameters for ITU defined QoS Classes. This is summarized Table B.1.

Table B.1 – I.356 defined QoS Parameters for different Traffic Classes

Traffic Requirement	QoS Parameter					
	CTD	CDV	CLR(0+1)	CLR(0)	CER	CMR
Constant Bit Rate	400 ms	3 ms	3×10^{-7}	None	None	None
RT-VBR	Unspecified	Unspecified	1×10^{-5}	None	4×10^{-6}	1 per day
NRT-VBR	Unspecified	Unspecified	Unspecified	1×10^{-5}	4×10^{-6}	1 per day

B.3 ATM Traffic Parameters

Traffic parameters describe traffic characteristics of a source at the ATM layer and are specified in the connection setup message. Once the connection is established, both the network and the user are bound by the traffic contract based on these traffic parameters. The endpoint is responsible for not exceeding these traffic parameters. The conformance to these traffic parameters is policed at the UNI by the UPC of the network. This implicitly requires the endpoint to provide a traffic shaping mechanism to conform to these parameters.

A Connection Traffic Descriptor specifies the traffic characteristics of the ATM connection as a whole; it includes the following:

- a) CDV Tolerance (CDVT)
- b) Conformance Definition
- c) A Source Traffic Descriptor

B.3.1 CDVT

The CDVT specifies the tolerance of an ATM network for incoming traffic that exceeds the PCR value. Cell clumping at an intermediate point in the ATM network may be caused beyond the control of the source leading to a nonzero CDV.

B.3.2 Conformance Definition

The Connection Traffic Descriptor contains sufficient information for Conformance testing the cells of the connection at the UNI on the network side. A leaky bucket algorithm, also known as the Generic Cell Rate Algorithm (GCRA), has been defined to measure the PCR performance ([1] in B.8). The leaky bucket has been modeled after an imaginary bucket of a fixed size K that leaks at the rate R. The traffic generated by a source is used to fill the leaky bucket draining at a rate R starting with an empty bucket. An arriving cell from the source is conforming to a PCR =R if it does not cause an overflow. K is specified as a tolerance to amount of cell clumping allowed for the cell stream, as described in the CDVT.

The conformance definition based on the GCRA is used to unambiguously specify the conforming cells of connection at the UNI for all services except ABR. ABR uses a different definition for conformance.

Even under ideal conditions, some cells may be non-conforming and therefore connection compliance does not imply that all cells associated with the connection are conforming

NOTE 41– Traffic contract parameters may be specified either explicitly or implicitly. A parameter value is explicitly specified when its value is assigned by the end-system using signaling for SVCs or Network Management Systems for PVCs. A parameter value specified at subscription (e.g., Service Provider) is also

considered explicit. A network may assign a default value when it is not explicitly specified. This is then called implicitly specified. In case there are no default rules, they are network specific. CAC and UPC procedures are network specific and take into account the knowledge of the specified traffic contract to operate efficiently.

B.3.3 Source Traffic Descriptor

A Source Traffic Descriptor includes the following parameters:

- a) Peak Cell Rate
- b) Sustainable Cell Rate
- c) Maximum Burst Size
- d) Minimum Cell Rate

Peak Cell Rate (PCR) specifies the peak bandwidth that may be sent by the source measured in cells per second.

Sustainable Cell Rate (SCR) specifies the average data rate in cells per second that may be sent by the source. SCR is also specified by the leaky bucket algorithm using the SCR parameter instead of PCR.

Maximum Burst Size (MBS) approximately specifies the maximum number of cells sent at the PCR rate and is proportional to the leaky bucket size associated with the SCR definition.

Minimum Cell Rate (MCR) is the minimum bandwidth guaranteed for an ABR service connection. It may be set to zero.

B.4 ATM Service Categories

The ATM Forum has defined 5 service categories with distinct attributes in traffic and QoS parameters: CBR, VBR-RT, VBR-NRT, UBR, and ABR ([1] in B.8).

NOTE 42– The service models defined by the ITU-T differ; this document shall follow the ATM Forum model.

ATM service categories are classified according to three main attributes: source traffic pattern (constant versus variable), end-to-end timing relationship requirement (real-time-based versus non-real-time-based), and connection mode (connection-oriented versus connectionless).

NOTE 43– Non-real-time-based may be a non-timed-based real-time application such as real-time stock, but without the need of strict CDV and delay requirements; the use of the term, *real-time*, is sometimes loose.

Requirements of real-time-based and non-real-time-based service requirements depend on whether the QoS delay objectives (CTD and CDV) are required to be satisfied. Constant Bit Rate (CBR) and Real-Time Variable Bit Rate (RT-VBR) are real-time services, while Non-Real-Time Variable Bit Rate (NRT-VBR), Unspecified Bit Rate (UBR), and Available Bit Rate (ABR) are non-real-time services.

B.4.1 Real-time Services

Real-time services provide real-time delivery of time-based information such as video, audio, voice, or animation and are primarily meant for support of streaming applications that are constant (using CBR) or variable (using RT-VBR).

Users of CBR service at the ATM layer are expected to send CBR traffic with a specified PCR (SCR equals PCR in this case) and CDVT. The QoS requirements for CBR service include CDV, maxCTD, and CLR. This service is most applicable to circuit emulation and real-time CBR video delivery.

The users of RT-VBR are expected to send VBR traffic with specified peak (PCR) and average (SCR) bandwidth, and associated traffic parameters (CDVT and MBS). The specified QoS objectives are peak-to-peak CDV, maximum CTD and CLR. This service is most applicable to real-time streaming applications that send at a variable bit rate.

B.4.2 Non-real-time Services

A non-real-time VBR service specifies the same parameters as the RT-VBR service: PCR, SCR, CDVT, and MBS. However, neither maxCTD nor the CDV is important and is not specified. NRT-VBR provides bandwidth guarantees at a PCR but no guarantees on delay bounds. Non-time-based Real-time block-transfer applications may be supported by this service because the guaranteed PCR may be used to satisfy the block transfer delay. CLR is specified because the data integrity is important here.

UBR service is used to support best-effort traffic. Hence, there are no QoS requirements. Traffic parameters PCR and CDVT may be specified but are not required. It is up to the network to decide whether UBR is subject to UPC. There are no CLR guarantees and makes it unsuitable for many non-real-time applications, that are sensitive to data loss.

ABR is an improvement over UBR to reduce CLR to an acceptable level and to provide more efficient use of the available network resources. ABR service attempts to share the bandwidth among all other ABR users and attempts to adapt to the changing network conditions due to congestion. The ABR user is expected to adapt the transmission rate according to this feedback and consequently should experience a low CLR and a fair share of the bandwidth. The CLR guarantees provided by the ABR service is based on this procedural commitment, instead of a quantitative one. If the ABR user does not behave according to the feedback mechanism and adjust for a lower rate during overload, the user may incur a significant cell loss.

The traffic parameters specified by ABR are PCR, CDVT, and MCR. MCR is the minimum cell rate that the user may send at all times.

ATM layer service categories and its Traffic and Qos attributes are summarized in Table B.2.

Table B.2 – Service Categories and its Traffic and QoS Attributes

ATM Layer Service Categories					
Attributes	CBR	RT-VBR	NRT-VBR	UBR	ABR
Traffic Parameters					
PCR CDVT	Specified			May not be subject to CAC & UPC	Actual rate subject to Flow Control
SCR MBS CDVT	N/A	Specified	N/A		
MCR	N/A				Specified
QoS Parameters					
Peak-to-peak CDV	Specified		Unspecified		
maxCTD	Specified		Unspecified		
CLR	Specified			Unspecified	Not signaled
Other					
Feedback Control	Unspecified				Specified

B.5 ATM Adaptation Layer (AAL) Types

The ITU-T ([3] in B.8) has defined 4 AAL types: AAL1, AAL2, AAL3/4, and AAL5. AAL2 is still being defined. The choice of an AAL type depends on the application requirements on error resiliency, packet efficiency, processing efficiency, and timing synchronization.

AAL1 and AAL5 are both capable of supporting CBR where timing is required.

AAL2, AAL3/4, and AAL5 are all capable of supporting VBR. In addition to VBR, when timing is required AAL2 is capable of supporting this type of application (e.g., compressed voice).

AAL5 has become the most popular AAL type because of its use in the ATM signaling protocol. This implies that all end users and switches (that use SVC) are required to support AAL5. AAL5 has become the most widely deployed AAL type to date. This is the main reason that applications such as MPEG video use this AAL type, although it is not best type for this application.

Although recently there has been industry wide consensus that AAL Types should be decoupled from the Service Class categories, the AAL types offer different data structures that make them more suitable for certain kinds of applications. If the data structure offered by a particular AAL type is satisfactory for the application at hand, then it may be used without penalty. This shall also be a guiding principle for using AAL5 as the AAL type to support the different classes of traffic.

B.6 ATM Multiplexing

Multiplexing of user traffic in ATM may occur in two ways: multiplexing within a single VC and multiplexing using different VCs.

Multiplexing within a single VC also referred to as VC Multiplexing is important when the available (active) number of VCs is a constraint. An identifier in the AAL-PDU distinguishes the different data streams multiplexed. Alternatively, this multiplexing may be provided outside the AAL at a higher layer protocol using an encapsulation scheme to self identify each PDU. In this case, the protocol multiplexing over the same VC would be transparent to the AAL.

VC Multiplexing also requires all multiplexed data streams to adhere to the same traffic contract. The biggest advantage of VC multiplexing could be the savings in signaling overhead that is incurred when a new VC connection is opened. If this is not a performance issue, then multiplexing using separate VCs may be preferable in cases where the different data streams have different QoS and traffic parameter specifications.

Connections using these two distinct service categories are forced to use different VCs. However, even while using the same service category it may be preferable to use separate VCs if the individual data streams have different QoS and traffic specifications. For example, the within the NRT-VBR different data streams may require different bandwidths and requests for such connections may arrive asynchronously.

NOTE 44– The present UNI signaling does not allow change of the traffic contract dynamically. Future, UNI versions may allow this capability at which time multiplexing within the same service category would be useful.

In addition, VC multiplexing within a VC may provide a very application-specific requirement to support a multipoint-to-point VC, using an appropriate AAL (AAL3/4) to differentiate between different senders.

What may be more appropriate is a single VPC (individual VCCs within this VPC) between a pair of locations. This shall force all VCCs to follow the same physical path and may provide a better synchronization between different VCC components by removing the differential delay between them. If, this level of synchronization is not required then in general completely different VCCs should be used.

B.7 Summary of ATM Services and Guarantees

Table B.3 summarizes the guarantees for delay, delay variance, bandwidth, and loss for all the ATM services classes. The table entries indicate whether a particular guarantee may be met by a service category. Items shown in parenthesis indicate the parameter mechanism used to achieve this guarantee. NRT-VBR service is best suited for non-time-based critical data, that require guarantees for loss and bandwidth. This service matches the requirements of FC-BB-2_ATM Control and FC Classes 2, 3, and F. CBR service is best suited for time-based applications that require guarantees for delay, delay variance, and bandwidth. This service most closely matches the requirements of FC Class 1. B.8 contains the UNI signaling procedures and connection negotiation of the QoS and traffic parameters.

Table B.3 – ATM Service Categories and Guarantees

Service Category	Guarantees				Other
	Loss	Delay Variance	Bandwidth	Delay	Flow Control
CBR	Yes (CLR)	Yes (pp-CDV)	Yes (PCR)	Yes (maxCTD)	No
RT-VBR	Yes (CLR)	Yes (pp-CDV)	Yes (SCR)	Yes (maxCTD)	No
NRT-VBR	Yes (CLR)	No	Yes (SCR)	No	No
ABR	Yes (qualitatively through Flow Control)	No	Yes (MCR)	No	Yes
UBR	No	No	No	No	No

B.8 ATM Signaling UNI Standard

Signaling is the mechanism used to request the ATM network to provide a Switched Virtual Connection (SVC) with the user specified desired parameters such as QoS and bandwidth. ATM signaling is used at the interface between an end user and the network and also within the ATM network. The end user is normally not concerned with the signaling that takes place inside the network.

NOTE 45– A different signaling protocol such as the Private Network Node Interface (P-NNI) operates between ATM switches inside the private ATM network cloud.

The ITU-T and the ATM Forum have both defined ATM signaling protocols used at the interface between the user and the immediate ATM device (e.g., switch). This interface is called the User Network Interface (UNI). The ITU-T calls this signaling protocol - Q.2931, while the ATM Forum has defined its own version of this protocol based on the Q.2931; several versions of this ATM Forum signaling currently exist (UNI 4.0, UNI 3.1), Encapsulated FC Frames each newer version adding functionality over the previous version. The BBW should at least support the ATM Forum's UNI 3.1.

NOTE 46– These are not always downward or upward compatible

There are actually 2 Types of UNI -Private and Public. The appropriate UNI is used depending on whether the user is interfacing with a Private or a Public ATM network. There are differences in the usage of the two types. The ATM addressing format depends on the type of ATM UNI used. Also, Address Registration is required at the Private UNI but is optional at the Public UNI.

NOTE 47– A Private UNI uses the private address format that is based on ISO NSAP format and has 3 forms: Data Country Code (DCC), International Code Designator (ICD), and E.164 Private Address. Part of this address is supplied by the network (suffix) and the other part by the End System. A Public UNI either uses the

E.164 or the Private Address format. Note that the private address format is optional when using Public UNI. There is a mechanism to tunnel the private address across a public network using a sub address field when the public UNI does not support the private address.

B.8.1 Switched Virtual Connections

BBWs may be connected by Permanent Virtual Connections (PVCs) or Switched Virtual Connections (SVCs). If PVCs are used then no ATM signaling is required and connections may be statically configured, manually or via Network Management. However, in general this arrangement is somewhat wasteful, inflexible, and does not scale well with changes and updates, as the number of BBWs grow. SVCs are dynamically provisioned via signaling and utilize resources more efficiently. SVCs are also administered automatically. This discussion of BBWs assumes Switched Virtual Connections (SVCs) as the primary mechanism that shall be used for data transport.

Using SVCs, two BBWs may be connected in a point-to-point topology.

NOTE 48– A number of BBWs may be connected via point-to-multipoint connection, with one BBW as the root and the other BBWs as leaves. Multicast connections may also be useful when exchanging BBW related control information such as routing table data and information that is shared amongst all BBWs.

A dedicated channel (Virtual Path Identifier (VPI) = 0 and Virtual Channel Identifier (VCI) =5) is reserved for signaling between the end user and the switch. When two BBWs are connected in a point-to-point connection, the Switch payload data may flow in one or both directions (unidirectional or I-directional) and the bandwidth may be the same or different in each direction.

B.8.2 Address resolution

The BBW's Call Handling logic shall initiate the ATM User Network Interface (UNI) Signaling after receiving a FC frame. Prior to this initiation, an Address Resolution function is required. This Address Resolution is between the 3-byte FC address (Domain-Area-Port ID) and the 20-byte ATM Address. The resolved address shall be referenced from a Routing Table prior to this UNI call request initiation. The exact mechanism(s) used in creating this table is outside the scope of this document.

B.8.3 SVC UNI messages

B.8.3.1 Message types

Figure B.2 pictorially shows the 3 normal phases associated with switched connections - Call Setup, Data Transfer, and Call Termination. The UNI Signaling uses distinct message types (SETUP, CONNECT, CALL PROCEEDING, RELEASE, etc.) that are meaningful during these phases. Multicast is supported with other messages such as ADD PARTY, DROP PARTY, etc.

B.8.3.2 Message information elements

Signaling messages carry additional information units called Information Elements (IE), with relevant information only to the particular message type, although more than one message type may share the same type of IE. Specialized IEs exist for many information categories. For example, IEs in the CONNECT request message carry the AAL IE among others. The information within each IE type varies. The AAL IE specifies the AAL Parameters such as the size of the Common Part Convergence Protocol (CPCS) - Service Data Unit (SDU). The Quality of Service is another IE that specifies the QoS Class. Another important IE is the Broadband Bearer Capability IE that specifies the end-to-end timing requirements, traffic type such as CBR or VBR, etc.

Not all information that is signaled is meant for the network. The Broadband Higher Layer and Lower Layer IE allows the end user to select useful information to be conveyed to the remote user transparent to the network. Such information, could for example, convey a vendor specific application identifier or the type of layer 2 or 3 protocol parameters (SNAP PID, OUI).

If the conditions (signaled parameters) of the ATM call request are acceptable to the ATM network (end-to-end), then a call shall be established. The network shall also signal back to the user its response to the requested call with information that it thinks it may meet. The initiator may then decide

to proceed or abort the request. All related FC-BB-2_ATM payload messages may now ride on this established call. The call shall last until it is explicitly closed. This scheme shall establish a switched FC connectivity on a call by call basis. Calls may be explicitly released by either user or the network.

NOTE 49– ATM call connection times could be in the order of several milliseconds and in large ATM networks this number could be in the 100s of milliseconds. Calls may therefore be kept open for longer periods of time when appropriate.

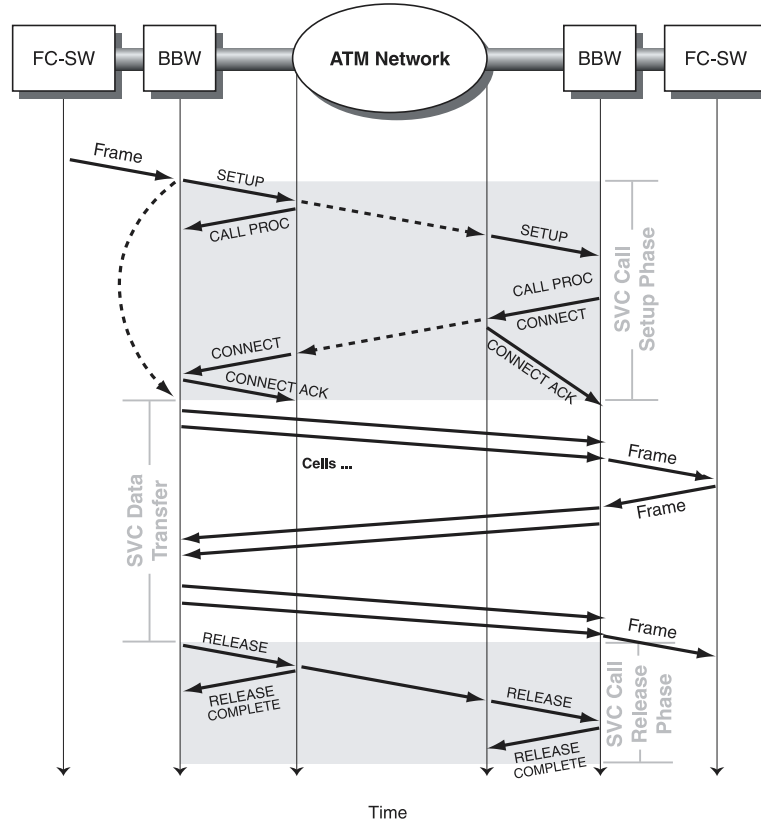


Figure B.2 – SVC Signaling at the UNI and Switched payload

B.9 References

- [1] ATM Forum (1996). *Traffic Management Specification 4.0*. AF-TM-0056.000
- [2] ITU-T1 (1995). Recommendation I.356, *B-ISDN ATM layer cell transfer performance*, Study Group 13, Geneva, 1995
- [3] ITU-T (1993). Recommendation I.363, *B-ISDN ATM Adaptation layer (AAL) Specification*, March, 1993

Annex C (Informative)

SR Protocol Parameter Guidelines and State Diagram

C.1 Assumptions

The following discussion provides guidelines for tuning the SR protocol parameters. Note that the order of discussion implies the order in which these values should be determined. The following suggests values that optimally deal with the case where traffic between BBW devices is unidirectional. In cases where there is sufficient (transmit frames at least T2 apart) bidirectional traffic the following calculated values become moot.

For a 1000 Kilometer (5 ms delay per 1000 kilometer), 155 Mbps link allowing 5 ms for packet processing and switch delay:

C.2 Calculated Ack Delay (T2) Timer variable

The timer T2 should be set to 1/4 the time it takes to transmit a Round Trip Delays worth of maximum length frames. Any less and the transmit window may close resulting in the least efficient use of the link. Any more and there shall be more ACKs returned than necessary:

$$T2 = ((2 * (((8 * 2136 \text{ bytes}) / 155520000) + 5 \text{ ms})) + 5 \text{ ms}) * .25 = 4 \text{ ms}$$

C.3 Calculated Window size

The Window Size W should be set to the time it takes to transmit a Round Trip Delays worth of maximum length frames plus the T2 timer value.

$$W = ((155520000 * (10\text{ms} + 4\text{ms})) / (8 * 2136 \text{ bytes})) = 127$$

C.4 Calculated Ack (T1) Timer variable

The timer T1 should be greater than the time it takes to send a windows worth of maximum length frames.

$$T1 = (((((127 * (8 * 2136 \text{ bytes})) / 155520000) + 5\text{ms}) + 5\text{ms})) = 24\text{ms} + 1 \text{ ms}$$

The window size is not negotiated dynamically between two SR end points.

C.5 SR Protocol State Diagram

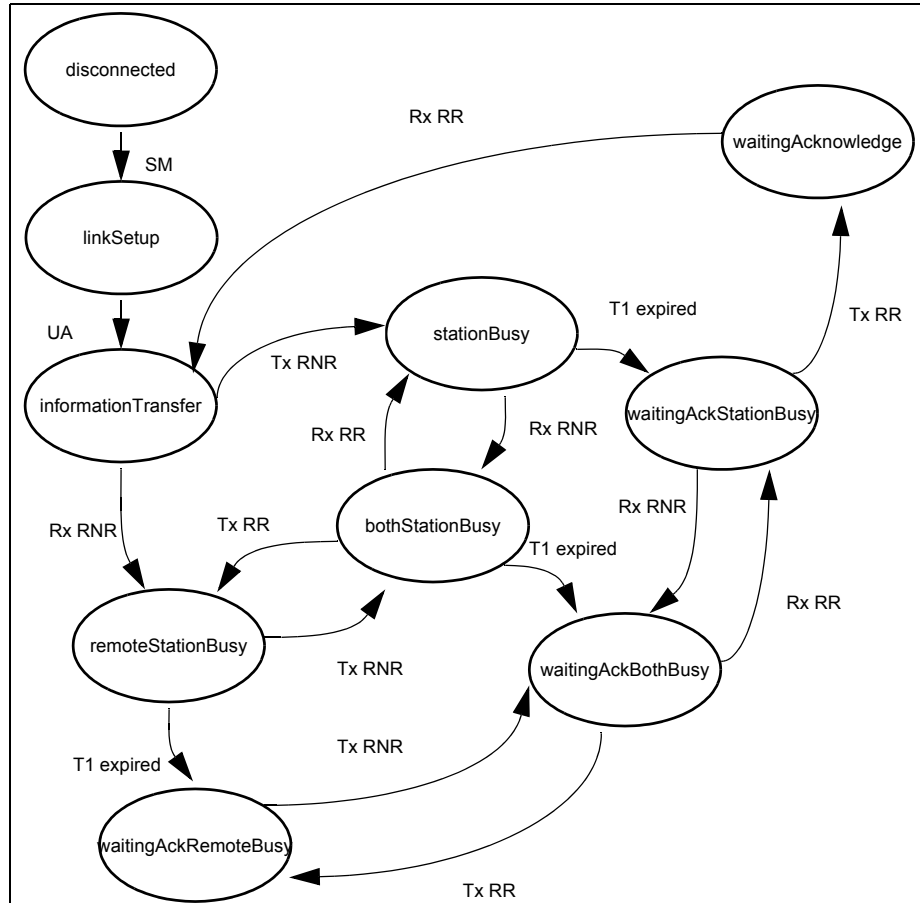


Figure C.1 – SR Protocol State Diagram

SM event indicates either the reception or transmission of a SM command frame (see Figure C.1). UA event indicates either the reception or transmission of an UA response frame. Note that if a SM is received the transition to informationTransfer requires that an UA is transmitted and if a SM is transmitted then an UA shall be received before a transition to informationTransfer.

disconnected: This is the initial state of the station; before, during and after initialization. The station shall return to this state whenever it receives a DISC.

linkSetup: The station shall transition to linkSetup when it receives a Set Mode (SM) frame.

informationTransfer: This is the normal interface state. The station reaches this state after either sending a Set Mode and receiving an Unnumbered Acknowledgement or receiving a Set Mode and transmitting an Unnumbered Acknowledgement.

stationBusy: When the station receives more frames than it is capable of processing, the station transitions to stationBusy state and transmits RNR.

waitingAckStationBusy: When the station is in stationBusy state and the T1 timer expires the station shall transition to waitingAckStationBusy state.

remoteStationBusy: When RNR is received from a remote station the station transitions to remote-StationBusy state.

waitingAckRemoteBusy: When the station is in remoteStationBusy state and the T1 timer expires the station shall transition to waitingAckRemoteBusy state.

bothStationsBusy: If RNR is received from the remote station and RNR has been transmitted to the remoteStation the station transitions to bothStationsBusy state.

The station returns to stationBusy state when a RR is received from the remote station.

The station returns to remoteStationBusy state when the station is ready to again receive frames from the remote station and transmits RR.

waitingAckBothBusy: When the station is in bothStationsBusy state and the T1 timer expires the station shall transition to waitingAckBothBusy state.

The station returns to waitingAckStationBusy state when a RR is received from the remote station.

The station returns to waitingAckRemoteBusy state when the station is ready to again receive frames from the remote station and transmits RR.

waitingAcknowledgement: When the station is in waitingAckStationBusy and is ready to again receive frames the station shall transition to waitingAcknowledgement and transmit RR.